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

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The development of autonomous agents increasingly relies on Multimodal Lan-
2 guage Models (MLMs) to perform tasks described in natural language with GUI
3 environments, such as websites, desktop computers, or mobile phones. Existing
4 benchmarks for MLM agents in interactive environments are limited by their focus
5 on a single environment, lack of detailed and generalized evaluation methods,
6 and the complexities of constructing tasks and evaluators. To overcome these
7 limitations, we introduce CRAB, the first agent benchmark framework designed to
8 support cross-environment tasks, incorporating a graph-based fine-grained evalua-
9 tion method and an efficient mechanism for task and evaluator construction. Our
10 framework supports multiple devices and can be easily extended to any environ-
11 ment with a Python interface. Leveraging CRAB, we developed a cross-platform
12 CRAB Benchmark-v0 comprising 100 tasks in computer desktop and mobile
13 phone environments. We evaluated four advanced MLMs using different single and
14 multi-agent system configurations on this benchmark. The experimental results
15 demonstrate that the single agent with GPT-4o achieves the best completion ratio
16 of 35.26%.

17 **1 Introduction**

18 The development of autonomous agents for human-centric interactive systems—such as desktop
19 OS [51], websites [56, 15], smartphones [52, 47], and games [38, 39]—has long been an impor-
20 tant goal of AI research, aiming to convert natural language instructions into concrete operations.
21 Traditionally, these challenges have been addressed using reinforcement learning [27]. Recently,
22 Large Language Models (LLMs) have demonstrated remarkable proficiency in natural language
23 understanding and commonsense reasoning, making them vital tools for developing autonomous
24 agents. This utility is further enhanced by Multimodal Language Models (MLMs), which improve
25 the ability to interpret visual information from GUIs [5].

26 To effectively develop MLM-based autonomous agents for real-world applications, it is essential to
27 create suitable benchmarks for standardized performance evaluation. However, existing benchmarks
28 still have limitations in terms of interaction methods, platform diversity, evaluation metrics, static
29 task dataset that prevent them from closely mirroring complex real-world applications. First, existing
30 benchmarks that interact with the environments through pre-collected observation data from system
31 environments [36, 26, 6] fail to capture the dynamic nature of real-world scenarios without interactive
32 exploration where data and conditions can change unpredictably. Second, existing benchmarks are
33 typically evaluated on a single platform, either Web, Android, or Desktop OS [34, 47, 46]. However,

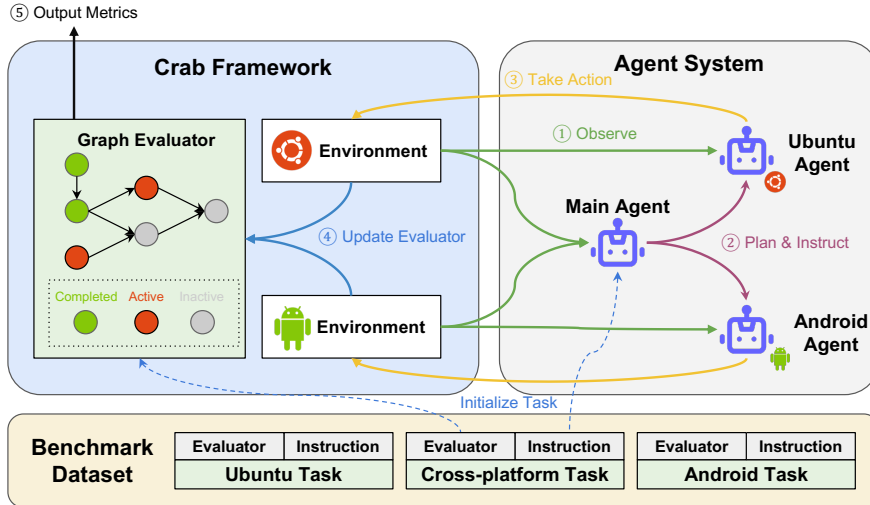


Figure 1: **Architecture of the Crab Framework demonstrating a benchmarking workflow for a multi-agent system.** A task is initialized by assigning instructions to the main agent and a graph evaluator inside the benchmark system. The workflow progresses through a cycle where the main agent observes, 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 throughout the workflow.

34 the practical applications usually involve tasks that span multiple platforms. For example, using a
 35 smartphone to take a photo and sending it to a desktop for editing with a graphics editor is a common
 36 real-world task across multiple platforms. Third, existing evaluation methods are generally either
 37 goal-based or trajectory-based [34, 47]. Goal-based methods typically employ coarse-grained binary
 38 metrics, solely evaluating whether the final system state aligns with the task’s goal. In contrast,
 39 trajectory-based methods can offer more fine-grained metrics by assessing the agent’s action trajectory
 40 against a ground truth trajectory yet ignore the possibility of multiple valid pathways to complete a
 41 task, making the evaluation results less fair. Lastly, task creation within these complex systems are
 42 not static and extensible with fixed templates [36, 46], which limits the diversity and scope of tasks.

43 We propose a benchmark that closely mirrors real-world situations and an evaluation method that more
 44 accurately reflects an agent’s performance on complex tasks. To this end, we introduce CRAB, a novel
 45 **C**Ross-environment **A**gent **B**enchmark framework. CRAB provides a comprehensive framework for
 46 evaluating cross-environment tasks in interactive environments, where the agent needs to operate
 47 simultaneously across various devices and platforms, adapting to varied system conditions to complete
 48 tasks efficiently. To the best of our knowledge, CRAB is the first autonomous agent benchmark
 49 framework that incorporates the **cross-environment tasks**. Moreover, we propose a novel evaluation
 50 method called **graph evaluator**. Unlike traditional goal-based and trajectory-based evaluation, our
 51 graph evaluator checks the intermediate procedures of completing a task by decomposing the task
 52 into multiple sub-goals. Each sub-goal is assigned a judge function to verify its completeness, and
 53 each is considered a node in the graph evaluator. The graph structure describes the sequential and
 54 parallel relationships between the sub-goals. Therefore, it offers fine-grained metrics similar to
 55 trajectory-based evaluations while accommodating multiple valid pathways to a solution, making it
 56 more suitable for evaluating tasks that involve various correct approaches. To solve the increasing
 57 complexity in cross-environment task construction. We also propose a highly extensible graph-based
 58 task construction method called **sub-task composition**. Combining multiple sub-tasks in a graph with
 59 task targets allows for efficient construction of various cross-environment tasks with corresponding
 60 graph evaluators. Table 1 compares CRAB with existing frameworks.

61 Based on CRAB framework, we develop a benchmark CRAB Benchmark-v0 with two collabora-
 62 ted environments that include an Android emulator and an Ubuntu desktop virtual machine. We

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; *Cross-platform* denotes support for multiple 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), 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 *Sub-task Composition* (composing multiple sub-tasks to construct tasks and evaluators).

	Interactive Environment	Multimodal Observation	Cross-platform	Evaluation	Task Construction
MINIWOB++ [34]	Web	✓	✗	Goal-based	Handmade
METAGUI [36]	✗	✗	✗	Trajectory-based	Handmade
GAIA [26]	✗	✗	✗	Goal-based	Handmade
MIND2WEB [6]	✗	✗	✗	Goal-based	LLM-inspired
AGENTBENCH [23]	Multi-isolated	✗	✗	Multiple	Handmade
INTERCODE [49]	Code	✗	✗	Goal-based	Handmade
WEBARENA [56]	Web	✓	✗	Goal-based	Template
WEBSHOP [50]	Web	✓	✗	Goal-based	Template
OMNIACT [12]	✗	✗	✗	Trajectory-based	Handmade
VWEBARENA [15]	Web	✓	✗	Goal-based	Template
ANDROIDARENA [47]	Android	✓	✗	Trajectory-based	LLM-inspired
OSWORLD [46]	Desktop OS	✓	✗	Goal-based	Template
CRAB	Desktop OS & Android	✓	✓	Graph-based	Sub-task Composition

63 have developed a total of 100 real-world tasks, encompassing both cross-environment and single-
64 environment tasks across multiple levels of difficulty. These tasks address a wide array of common
65 real-world applications and tools, including but not limited to calendars, email, maps, web browsers,
66 and terminals, and facilitate common collaboration between smartphones and desktops. Considerable
67 time has been invested in verifying the accuracy and comprehensiveness of the instructions for
68 sub-tasks, as well as the generalization and correctness of their evaluators. Most tasks are constructed
69 using a careful composition of sub-tasks, while some tasks are crafted manually to accommodate
70 specific multi-environment collaboration scenarios. We test 4 popular MLMs, including GPT-4 Turbo,
71 GPT-4o, Claude 3 Pro and Gemini 1.5 Pro, across different structures of single agent and multi-agent
72 systems, totaling 9 different agent settings in our benchmarks. The experimental results show that
73 the single agent with GPT-4o model achieves the best completion ratio of 35.26%, underscoring
74 the necessity for ongoing development of more effective autonomous agents. Our proposed metrics
75 successfully distinguish between different methods better than previous metrics. We further analyze
76 the different termination reasons that reflect the problems inherent in the function calling feature of
77 current models and communication within the multi-agent system.

78 2 Related Work

79 Leveraging LLMs as reasoning units has become an effective approach [42, 10, 45] for building
80 autonomous agents, including embodied agents [39, 35, 4], social simulations [30, 20], web naviga-
81 tion [24], game playing [16, 37], office assistants [18], and code generation [54], among others. With
82 common knowledge of Graphical User Interfaces (GUI) and operating systems, GUI agents [44, 41,
83 40, 55, 28] are becoming a productive research direction for developing autonomous agents capable
84 of operating systems with GUI interfaces to accomplish complex tasks. GUI agents can typically
85 operate multiple applications within a system, making them more versatile than the aforementioned
86 agents, which are often limited to a single application. Various benchmarks have been developed to
87 evaluate the performance of these GUI agents in interactive environments, which can generally be
88 categorized into three types: web, mobile phone, and desktop.

89 The web environment is one of the earliest environments used to benchmark agents due to its
90 simplicity, ease of reproduction, straightforward construction, and ease of parsing by agents. One of
91 the earliest examples is Miniwob++ [34], initially designed for evaluating reinforcement learning
92 agents. It quickly became a foundational benchmark for evaluating GUI agents. However, its web
93 page designs are overly simplistic and lack modern features, limiting its ability to assess agents’
94 performance on real-world websites. With the rise of LLMs as agent reasoning engines, more
95 complex web environments, such as WebShop [50] Mind2Web [6] and WebArena [56], have been
96 introduced for benchmarking language model agents, offering realistic and reproducible environments
97 and corresponding web-based tools to simulate sufficiently complex web tasks and cross-environment
98 interactions. Building on these works, Visual WebArena [15] focuses on evaluating multimodal
99 language model agents by incorporating tasks that require visual understanding. Although web
100 environments contain various real world scenarios, it is impossible to replace native applications for
101 complex tasks like multimedia editing, programming, etc.

102 Intelligent assistants have long been a commercial feature in mobile operating systems, making the
103 motivation to develop mobile agents clear. Additionally, mobile phone operations and observations
104 are generally simpler than those on personal computers, which has made mobile devices a popular
105 environment for benchmarking GUI agents. Several task datasets existed even before the rise of
106 GUI agents. MetaGUI [36] introduced a dataset that focuses on GUI-based task-oriented dialogue
107 systems (GUI-TOD), dividing mobile system control tasks into dialogues and GUI operation traces,
108 while AITW [32] builds the operation traces of challenging multi-step tasks on involving apps and
109 websites on mobile devices based on screenshots. Android Arena [47] underlines the collaboration
110 among android applications and expands simple android tasks into cross-App and constrained tasks,
111 which verifies the potential of LLM-based complicated android system control. AITZ [53] constructs
112 datasets with Chain-of-Thought (CoT) considerations, adding semantic annotations based on visual
113 models at each step and developing operational procedures for selected tasks. In addition, Mobile
114 Agent Bench [43] collects app event signals via Android accessibility services, builds a benchmark
115 with well-annotated operation trajectories, and organizes tasks into different levels of difficulty.

116 Desktop environments typically have a more complex action space, observation space, and operational
117 logic, making task creation and verification more difficult. Additionally, they are highly customizable
118 and lack generalized tools that can serve as a bridge for agents to interact with the system, which
119 complicates the creation of reproducible environments. OMNIACT [12] is a static benchmark that
120 captures data from multiple desktop operating systems, incorporating visual information from the OS
121 screen UI through segmentation and corresponding tagging. OSWorld [46] provides an interactive
122 and reproducible environment based on XML and screenshots with a standard format. However, both
123 of these works rely on the Python library PyAutoGUI¹ and code generation for operation, which
124 limits the generalizability.

125 While these benchmarks aim to evaluate an agent’s capacity across a wide range of applications,
126 they are built on human-annotated trajectories, which lack scalability. Most tasks are derived from
127 question-and-answer platforms like Stack Overflow or based on annotators’ daily usage. While these
128 resources are realistic, they may not effectively test the generalizability of the agent, as the texts
129 are highly likely to appear in the training data. Furthermore, the evaluation methods of previous
130 benchmarks often rely either on full task trajectories or only on the final goals, making it difficult to
131 capture the entire process or to account for partially completed tasks.

132 **3 Definitions**

133 **3.1 Problem Formulation**

134 Consider autonomous agents performing a task on a digital device (i.e. desktop computer). Such a
135 device typically has input devices (i.e. mouse and keyboard) for human interaction and output devices
136 (i.e. screen) to allow human observation of its state. In CRAB, we represent this type of device as an

¹<https://github.com/asweigart/pyautogui>

137 **environment.** Formally, this environment is defined as a reward-free Partially Observable Markov
 138 Decision Process (POMDP), denoted by the tuple $M := (S, \mathcal{A}, \mathcal{T}, \mathcal{O})$, where S represents the state
 139 space, \mathcal{A} the action space, $\mathcal{T} : S \times \mathcal{A} \rightarrow S$ the transition function, and \mathcal{O} the observation space.
 140 Considering the collaborative nature of multiple devices in real-world scenarios, we can combine
 141 multiple environments into a set $\mathbf{M} = M_1, M_2, \dots, M_n$, where n is the number of environments and
 142 each environment $M_j = (S_j, \mathcal{A}_j, \mathcal{T}_j, \mathcal{O}_j)$. We define a task that requires operations across multiple
 143 environments as a **cross-environment task**. This task is formalized as a tuple (\mathbf{M}, I, R) , in which
 144 \mathbf{M} is the environment set, I is the task objective in the form of natural language instructions, and
 145 R is the reward function of the task. An **agent system**, designed to complete a task represented
 146 by an instruction I , can be modeled as a policy $\pi((m, a) \mid (I, H, o_1, \dots, o_n))$, which defines
 147 the probability of taking action a in environment m when receiving observation (o_1, \dots, o_n) from
 148 environment (M_1, \dots, M_n) with a history action trajectory H . An **agent** within the agent system
 149 should have a fixed back-end MLM and system prompt, and retain its chat history. An agent system is
 150 composed of either a single agent responsible for planning, reasoning, and action-taking or multiple
 151 agents connected through a communication strategy to collaborate.

152 3.2 Graph of Task Decomposition

153 Decomposing a complex task into several simpler sub-tasks has been proved to be an effective
 154 prompting method for LLMs [13]. Some studies represent sub-tasks in a graph structure. For
 155 instance, PLaG [19] uses a graph-based structure to enhance plan reasoning within LLMs, while
 156 DyVal [57] employs directed acyclic graphs (DAGs) to facilitate dynamic evaluation of LLMs.
 157 By introducing this concept into cross-
 158 environment tasks, naturally, decompos-
 159 ing a cross-environment task into sub-tasks
 160 with in different environments that have
 161 both sequential and parallel connections
 162 forms a DAG. Therefore, we introduce
 163 the **Graph of Decomposed Tasks (GDT)**,
 164 where each node in the DAG is a sub-task,
 165 formalized as a tuple (m, i, r) , where m
 166 specifies the environment in which the sub-
 167 task is performed, i provides the subtask
 168 natural language instruction, and r repre-
 169 sents the reward function. The reward func-
 170 tion evaluates the state of m and returns a
 171 boolean value to determine if the sub-task
 172 is completed. The edges within GDT rep-
 173 resent the sequential relationship between
 174 sub-tasks. An example GDT is shown in Fig. 2.

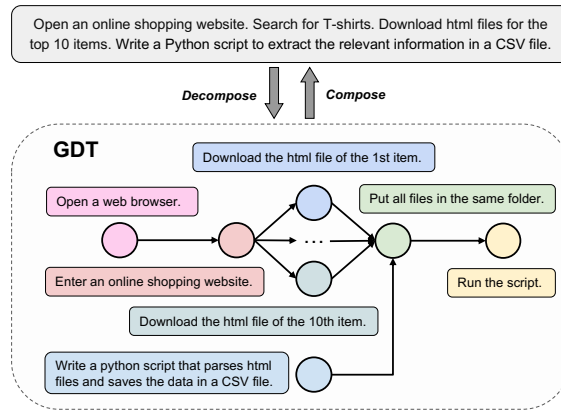


Figure 2: **Graph of Decomposed Tasks.**

175 4 The Crab Framework

176 4.1 Cross-environment Agent Interaction

177 Compared to single-environment tasks, cross-environment tasks offer three main advantages for
 178 benchmarking agents. First, cross-environment tasks reflect real-world scenarios where humans
 179 use multiple devices simultaneously to accomplish tasks. Second, these tasks require sophisticated
 180 message processing and information transfer between environments. Such tasks demand that the agent
 181 plan actions, construct outputs for each environment, and remember what needs to be transferred,
 182 showcasing a high-level understanding of environments and tasks. Lastly, role-playing multi-agent
 183 systems have proven to be effective in executing complex tasks [17, 9]. The underlying principle
 184 of their effectiveness is the division of responsibilities. Cross-environment tasks are suited to multi-
 185 agent, as they can be divided by distinct observation spaces, action spaces, and specialized knowledge

186 in each environment, as shown in Fig. 1. CRAB uses a unified interface for agents to operate in all
187 environments. Implementation details are in the Appendix A.2.

188 4.2 Graph Evaluator

189 Inspired by the "decomposing" idea from GDT (Sec. 3.2), we propose a novel integrated approach,
190 the *Graph Evaluator*, which provides fine-grained metrics and supports multiple valid paths. To build
191 a graph evaluator for a given task, we begin by decomposing the task into a GDT, where each sub-task
192 is associated with an intermediate environment state critical to completing the overall task. Nodes
193 in the graph evaluator activate when they either have no incoming edges or after all their preceding
194 tasks are completed, ensuring a sequential order of tasks. After an agent takes an action, the system
195 checks these active nodes to verify if the target state of each node is reached. A node completion
196 triggers successor nodes to activate and verify the state. This cycle repeats until no new nodes activate,
197 showing that the system's task sequence aligns with the current state of the environment. Unlike
198 trajectory-based methods, which compare sequences of agent actions, the Graph Evaluator does not
199 rely on the specific actions taken by the agent, allowing it the freedom to choose any path. Instead, it
200 concentrates on the key intermediate states of the environment necessary for reaching the final goal.

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

213 4.3 Task and Evaluator Construction

214 Despite the graph evaluator offers detailed evaluations, one challenge is the complexity in creating
215 each evaluator. Creating a graph evaluator requires: (1) adequately decomposing a task into multiple
216 sub-tasks, each with a well-defined graph structure; and (2) engaging an expert of the target platform
217 to carefully craft an evaluator for each sub-task. To efficiently create graph evaluators, we connect
218 sub-tasks as GDTs to formulate new tasks. There are two primary challenges in constructing GDT:
219 (1) Sub-tasks still require manual creation, necessitating a method to quickly generate them on a large
220 scale; (2) Properly modeling the sequential and parallel relationships between sub-tasks, ensuring
221 that the edges connecting sub-task nodes are semantically meaningful and systematically applicable.
222 A template-based approach is commonly used to address the first issue by generating a large number
223 of tasks efficiently. To tackle the second challenge, we employ the message transferring concept
224 (Sec. 4.1). Specifically, if a sub-task α produces an output message that serves as an input for another
225 sub-task β , then α can be considered a legitimate prerequisite of β , allowing us to connect α and β
226 with an directed edge in the GDT. To further refine our approach, we introduce a *sub-task template*
227 structure. Each sub-task is described using a natural language instruction template that includes
228 several replaceable input attributes and an optional output, where each input attributes and output
229 have a fixed type. To generate a GDT, input attributes can be filled with either a hand-crafted value
230 corresponding to their type or linked to a task with the same output type as the input type. From the
231 evaluator's perspective, each sub-task template is linked to an evaluator generator that uses the input
232 attribute value to generate evaluator subgraphs. Once a GDT is constructed, the graph evaluator is
233 created by interlinking each subgraph. The description for the composed task is initially generated by
234 GPT-4 using the sub-task descriptions as prompts and subsequently refined and polished by human
235 reviewers.

236 5 Experiments

237 5.1 Benchmark

238 We build an agent benchmark CRAB Benchmark-v0 featuring with cross-environment, graph evalua-
239 tor, and task generation through CRAB framework, including an Android smartphone emulator and a
240 Ubuntu Linux desktop virtual machine. Both environments are reproducible and standalone. Detailed
241 environment implementation, observation space and action space are provided in Appendix A.1. we
242 meticulously construct 16 sub-task templates for the Android environment and 19 sub-task templates
243 for the Ubuntu environment. The Ubuntu templates encompass a variety of tasks such as Command
244 Line Interface (CLI) operations, file system management, search engine usage, desktop configurations,
245 and map navigation. Conversely, the Android sub-task templates are primarily focused on the storage
246 and transmission of messages via various applications. Each sub-task template is linked to a graph
247 evaluator consisting of one to four nodes. Each sub-task is verified by at least two related field experts.
248 The dataset has 29 android tasks, 53 Ubuntu tasks and 18 cross-platform tasks. Besides, the sub-task
249 pool has 19 in Ubuntu and 17 in Android.

250 5.2 Baseline Agent System

251 At the core of MLM Agents are back-end Multimodal Language Models that provide natural language
252 and image understanding, basic device knowledge, task planning, and logical reasoning abilities. To
253 run in CRAB Benchmark-v0, the back-end model needs to support: (1) Accept multimodal mixed
254 input, as the system provides both screenshots and text instructions as prompts; (2) Handle multi-turn
255 conversations, as most tasks require the agent to take multiple actions, necessitating the storage of
256 history messages in its context; (3) Generate structured output through function calling, ensuring the
257 proper use of provided actions with type-correct parameters. We selected four MLMs that meet these
258 criteria for our experiments: GPT-4o (gpt-4o-2024-05-13) [29], GPT-4 Turbo (gpt-4-turbo-2024-04-
259 09) [1], Gemini 1.5 Pro (May 2024 version) [33], Claude 3 Opus (claude-3-opus-20240229) [2]. To
260 examine how different multi-agent structures impact performance, we design three agent system
261 structures. In the **single agent** structure, one agent manages all responsibilities, including observation
262 analysis, planning, reasoning, and format the output action. The **multi-agent by functionality**
263 structure splits tasks between a main agent, responsible for analysis and planning, and a tool agent
264 that translates instructions into actions without accessing environmental observations. This division
265 allows the main agent to concentrate on high-level tasks without managing functional call formats.
266 Meanwhile, in the **multi-agent by environment** setup, responsibilities are further distributed. A
267 main agent processes all environmental observations for high-level planning, while each environment-
268 specific sub-agent executes actions based on the main agent’s instructions, incorporating observations
269 from their respective environments.

270 For all agents, we utilized the default API parameters and retained two turns of historical messages.
271 The interaction turns are limited to 15 and the task will terminated because reaching max turns. The
272 agent can also terminate the task ahead if it thinks the task is completed. The screenshots do not
273 descale and passed through PNG format with the highest quality that the APIs provide. Detailed agent
274 and prompt designs are shown in Appendix B. In the experiment, we deployed four cloud machines
275 cloned from the same disk image to ensure a consistent environment for all agents. Running a single
276 agent setting in the benchmark requires at least 30 hours to complete on one machine. This duration
277 depends on the API call times and the necessity for manual resets in certain tasks.

278 5.3 Results

279 The primary outcomes are detailed in Table 2. The GPT-4o and GPT-4 Turbo models, developed by
280 OpenAI, achieve the highest average success rates and completion ratios among the tested models.
281 Specifically, GPT-4o slightly outperforms GPT-4 Turbo. This result suggests a tiny difference in their
282 underlying architectures or training data, but GPT-4o possibly be trained on more GUI data. Claude 3
283 outperforms Gemini 1.5 in all settings, according to CR. The multi-agent structures’ performances on

Table 2: **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 stops when it does not complete the task. *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	14.00	35.26	3.66	5.26×10^{-4}	7.00	59.00	20.00
GPT-4o	By Func	13.00	32.48	3.29	5.20×10^{-4}	12.00	54.00	21.00
GPT-4o	By Env	14.00	33.74	3.40	2.71×10^{-4}	8.00	49.00	29.00
GPT-4 TURBO	Single	11.00	31.52	3.60	6.45×10^{-4}	7.00	64.00	18.00
GPT-4 TURBO	By Func	13.00	29.99	3.53	4.79×10^{-4}	11.00	41.00	35.00
GEMINI 1.5 PRO	Single	6.00	17.19	1.69	\	3.00	55.00	36.00
GEMINI 1.5 PRO	By Func	6.00	14.53	1.50	\	10.00	33.00	51.00
CLAUDE 3 OPUS	Single	6.00	21.39	2.66	4.51×10^{-4}	7.00	53.00	34.00
CLAUDE 3 OPUS	By Func	5.00	18.79	1.90	3.31×10^{-4}	29.00	32.00	34.00

284 all back-end MLMs are slightly lower than the single agent, indicating that current autonomous agents
 285 mainly rely on back-end model performance. Regarding termination reason, multi-agent structures
 286 have higher possibility to take invalid action and incorrectly complete the task, this can caused by
 287 the hallucination when main agent generating the instruction messages or misunderstanding of the
 288 sub-agents when receiving these messages. We analyze the reasons for the poorer performance of
 289 multi-agent structures in Appendix C.2. In terms of execution efficiency, the GPT-4 series show
 290 strong performance. However, when evaluating cost efficiency, GPT-4 Turbo exhibited a lower CE
 291 value compared to GPT-4o, suggesting that GPT-4 Turbo is more cost-effective.

292 The completion ratio metric reveals a notable performance difference between models. For instance,
 293 even though Claude (single agent) and Gemini (multi-agent by functionality) have the same success
 294 rates, their completion ratios differ by up to 6.86%. This highlights the value of the completion
 295 ratio in assessing the effectiveness of different methods. We provide more detailed analyses and
 296 comparisons of agent configurations in Appendix C.

297 6 Conclusion

298 We propose the CRAB framework introducing cross-environment automatic task performing problem,
 299 featuring advanced graph-based task generation and evaluation methods, which reduce the manual
 300 effort in task step and provide a more dynamic and accurate agent assessments. Based on the
 301 framework, we propose CRAB Benchmark-v0, including a set of high quality cross-environment
 302 tasks for a smart phone and desktop, equipped with visual prompting strategy. We test various
 303 backend models and agent system structures on the dataset. The result reflects preference of different
 304 agent settings. Despite our work contributing to better cross-environment agent research, there are
 305 still some limitations. We build sub-tasks upon the original apps in the Ubuntu system and the
 306 Android system on Pixel, which cannot cover a wider range of applications. Moreover, the visual
 307 information is not used in the evaluation on the sub-tasks in Android System. Future works can focus
 308 on expanding datasets and environments and testing more models, prompts, structure of agents upon
 309 the benchmark.

310 Acknowledgement

311 We express our gratitude to Yuhui Wang for refining the expressions in our paper and providing
312 invaluable advice on writing. We would like to also thank Beichen Huang for the helpful discussions
313 on solving virtualization technology issues.

314 References

- 315 [1] Josh Achiam et al. *GPT-4 Technical Report*. Mar. 4, 2024. URL: [http://arxiv.org/abs/](http://arxiv.org/abs/2303.08774)
316 2303.08774. preprint.
- 317 [2] Anthropic. *The Claude 3 Model Family: Opus, Sonnet, Haiku*. [https://www-cdn.](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf)
318 [anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf)
319 [Claude_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf). Year.
- 320 [3] Fabrice Bellard. “QEMU, a fast and portable dynamic translator.” In: *USENIX annual technical*
321 *conference, FREENIX Track*. Vol. 41. 46. California, USA. 2005, pp. 10–5555.
- 322 [4] Yongchao Chen et al. “Autotamp: Autoregressive task and motion planning with llms as
323 translators and checkers”. In: *arXiv preprint arXiv:2306.06531* (2023).
- 324 [5] Kanzhi Cheng et al. *SeeClick: Harnessing GUI Grounding for Advanced Visual GUI Agents*.
325 Jan. 17, 2024. URL: <http://arxiv.org/abs/2401.10935>. preprint.
- 326 [6] Xiang Deng et al. *Mind2Web: Towards a Generalist Agent for the Web*. 2023. arXiv: 2306.
327 06070 [cs.CL].
- 328 [7] Johan Edstedt et al. “RoMa: Robust Dense Feature Matching”. In: 2024.
- 329 [8] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. “Exploring Network Structure, Dynam-
330 ics, and Function Using NetworkX”. In: *Proceedings of the 7th Python in Science Conference*.
331 Ed. by Gaël Varoquaux, Travis Vaught, and Jarrod Millman. Pasadena, CA USA, 2008, pp. 11–
332 15.
- 333 [9] Sirui Hong et al. “Metagpt: Meta programming for multi-agent collaborative framework”. In:
334 *arXiv preprint arXiv:2308.00352* (2023).
- 335 [10] Wenlong Huang et al. “Language models as zero-shot planners: Extracting actionable knowl-
336 edge for embodied agents”. In: *International Conference on Machine Learning*. PMLR. 2022,
337 pp. 9118–9147.
- 338 [11] Hanwen Jiang et al. “OmniGlue: Generalizable Feature Matching with Foundation Model
339 Guidance”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
340 *Recognition (CVPR)*. 2024.
- 341 [12] Raghav Kapoor et al. *OmniACT: A Dataset and Benchmark for Enabling Multimodal Generalist*
342 *Autonomous Agents for Desktop and Web*. Feb. 28, 2024. URL: [http://arxiv.org/abs/](http://arxiv.org/abs/2402.17553)
343 2402.17553. preprint.
- 344 [13] Tushar Khot et al. “Decomposed Prompting: A Modular Approach for Solving Complex
345 Tasks”. In: *The Eleventh International Conference on Learning Representations*. 2023. URL:
346 https://openreview.net/forum?id=_nGgzQjzaRy.
- 347 [14] Avi Kivity et al. “kvm: the Linux virtual machine monitor”. In: *Proceedings of the Linux*
348 *symposium*. Vol. 1. 8. Dttawa, Dntorio, Canada. 2007, pp. 225–230.
- 349 [15] Jing Yu Koh et al. *VisualWebArena: Evaluating Multimodal Agents on Realistic Visual Web*
350 *Tasks*. Jan. 24, 2024. URL: <http://arxiv.org/abs/2401.13649>. preprint.
- 351 [16] Yihuai Lan et al. “Llm-based agent society investigation: Collaboration and confrontation in
352 avalon gameplay”. In: *arXiv preprint arXiv:2310.14985* (2023).
- 353 [17] Guohao Li et al. “Camel: Communicative agents for "mind" exploration of large scale language
354 model society”. In: (2023).
- 355 [18] Hongxin Li et al. “SheetCopilot: Bringing Software Productivity to the Next Level through
356 Large Language Models”. In: *Advances in Neural Information Processing Systems 36* (2024).
- 357 [19] Fangru Lin et al. *Graph-Enhanced Large Language Models in Asynchronous Plan Reasoning*.
358 Feb. 5, 2024. URL: <http://arxiv.org/abs/2402.02805>. preprint.
- 359 [20] Jiaju Lin et al. “Agentsims: An open-source sandbox for large language model evaluation”. In:
360 *arXiv preprint arXiv:2308.04026* (2023).

- 361 [21] Nelson F. Liu et al. *Lost in the Middle: How Language Models Use Long Contexts*. Nov. 20,
362 2023. URL: <http://arxiv.org/abs/2307.03172>. Pre-published.
- 363 [22] Shilong Liu et al. *Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set*
364 *Object Detection*. arXiv.org. Mar. 9, 2023. URL: <https://arxiv.org/abs/2303.05499v4>.
- 365 [23] Xiao Liu et al. “AgentBench: Evaluating LLMs as Agents”. In: *The Twelfth International*
366 *Conference on Learning Representations*. 2024. URL: [https://openreview.net/forum?](https://openreview.net/forum?id=zAdUB0aCTQ)
367 [id=zAdUB0aCTQ](https://openreview.net/forum?id=zAdUB0aCTQ).
- 368 [24] Xing Han Lù, Zdeněk Kasner, and Siva Reddy. *WebLINX: Real-World Website Navigation*
369 *with Multi-Turn Dialogue*. arXiv.org. Feb. 8, 2024. URL: [https://arxiv.org/abs/2402.](https://arxiv.org/abs/2402.05930v1)
370 [05930v1](https://arxiv.org/abs/2402.05930v1).
- 371 [25] Michael M. McKerns et al. *Building a Framework for Predictive Science*. Feb. 6, 2012. URL:
372 <http://arxiv.org/abs/1202.1056>. preprint.
- 373 [26] Grégoire Mialon et al. *GAIA: A Benchmark for General AI Assistants*. Nov. 21, 2023. URL:
374 <http://arxiv.org/abs/2311.12983>. preprint.
- 375 [27] Volodymyr Mnih et al. “Human-Level Control through Deep Reinforcement Learning”. In:
376 *Nature* 518.7540 (Feb. 2015), pp. 529–533. ISSN: 1476-4687. URL: [https://www.nature.](https://www.nature.com/articles/nature14236)
377 [com/articles/nature14236](https://www.nature.com/articles/nature14236).
- 378 [28] Runliang Niu et al. *ScreenAgent: A Vision Language Model-Driven Computer Control Agent*.
379 Feb. 8, 2024. URL: <http://arxiv.org/abs/2402.07945>. preprint.
- 380 [29] OpenAI. *GPT-4 omni*. <https://openai.com/index/hello-gpt-4o/>. 2024.
- 381 [30] Joon Sung Park et al. “Generative agents: Interactive simulacra of human behavior”. In:
382 *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*.
383 2023, pp. 1–22.
- 384 [31] Guilherme Potje et al. “XFeat: Accelerated Features for Lightweight Image Matching”. In:
385 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.
- 386 [32] Christopher Rawles et al. *Android in the Wild: A Large-Scale Dataset for Android Device*
387 *Control*. 2023. arXiv: 2307.10088 [cs.LG].
- 388 [33] Machel Reid et al. *Gemini 1.5: Unlocking Multimodal Understanding across Millions of*
389 *Tokens of Context*. Apr. 25, 2024. URL: <http://arxiv.org/abs/2403.05530>. preprint.
- 390 [34] Tianlin Shi et al. “World of Bits: An Open-Domain Platform for Web-Based Agents”. In:
391 *Proceedings of the 34th International Conference on Machine Learning*. Ed. by Doina Precup
392 and Yee Whye Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, June 2017,
393 pp. 3135–3144. URL: <https://proceedings.mlr.press/v70/shi17a.html>.
- 394 [35] Chan Hee Song et al. “Llm-planner: Few-shot grounded planning for embodied agents with
395 large language models”. In: *Proceedings of the IEEE/CVF International Conference on Com-*
396 *puter Vision*. 2023, pp. 2998–3009.
- 397 [36] Liangtai Sun et al. *META-GUI: Towards Multi-Modal Conversational Agents on Mobile GUI*.
398 Nov. 24, 2022. URL: <http://arxiv.org/abs/2205.11029>. preprint.
- 399 [37] Weihao Tan et al. *Towards General Computer Control: A Multimodal Agent for Red Dead*
400 *Redemption II as a Case Study*. 2024. arXiv: 2403.03186 [cs.AI].
- 401 [38] Oriol Vinyals et al. “Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement
402 Learning”. In: *Nature* 575.7782 (Nov. 2019), pp. 350–354. ISSN: 1476-4687. URL: [https:](https://www.nature.com/articles/s41586-019-1724-z)
403 [//www.nature.com/articles/s41586-019-1724-z](https://www.nature.com/articles/s41586-019-1724-z).
- 404 [39] Guanzhi Wang et al. *Voyager: An Open-Ended Embodied Agent with Large Language Models*.
405 May 25, 2023. URL: <http://arxiv.org/abs/2305.16291>. preprint.
- 406 [40] Junyang Wang et al. *Mobile-Agent-v2: Mobile Device Operation Assistant with Effective*
407 *Navigation via Multi-Agent Collaboration*. 2024. arXiv: 2406.01014 [cs.CL]. URL: [https:](https://arxiv.org/abs/2406.01014)
408 [//arxiv.org/abs/2406.01014](https://arxiv.org/abs/2406.01014).
- 409 [41] Junyang Wang et al. *Mobile-Agent: Autonomous Multi-Modal Mobile Device Agent with Visual*
410 *Perception*. arXiv.org. Jan. 29, 2024. URL: <https://arxiv.org/abs/2401.16158v2>.
- 411 [42] Lei Wang et al. “A survey on large language model based autonomous agents”. In: *Frontiers*
412 *of Computer Science* 18.6 (2024), pp. 1–26.
- 413 [43] Luyuan Wang et al. *MobileAgentBench: An Efficient and User-Friendly Benchmark for Mobile*
414 *LLM Agents*. 2024. arXiv: 2406.08184 [cs.AI]. URL: [https://arxiv.org/abs/2406.](https://arxiv.org/abs/2406.08184)
415 [08184](https://arxiv.org/abs/2406.08184).

- 416 [44] Zhiyong Wu et al. *OS-Copilot: Towards Generalist Computer Agents with Self-Improvement*.
417 arXiv.org, Feb. 12, 2024. URL: <https://arxiv.org/abs/2402.07456v2>.
- 418 [45] Zhiheng Xi et al. *The Rise and Potential of Large Language Model Based Agents: A Survey*.
419 arXiv.org, Sept. 14, 2023. URL: <https://arxiv.org/abs/2309.07864v3>.
- 420 [46] Tianbao Xie et al. *OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in*
421 *Real Computer Environments*. Apr. 11, 2024. URL: <http://arxiv.org/abs/2404.07972>.
422 preprint.
- 423 [47] Mingzhe Xing et al. *Understanding the Weakness of Large Language Model Agents within a*
424 *Complex Android Environment*. Feb. 9, 2024. URL: <http://arxiv.org/abs/2402.06596>.
425 preprint.
- 426 [48] Jianwei Yang et al. *Set-of-Mark Prompting Unleashes Extraordinary Visual Grounding in*
427 *GPT-4V*. Nov. 6, 2023. URL: <http://arxiv.org/abs/2310.11441>. preprint.
- 428 [49] John Yang et al. *InterCode: Standardizing and Benchmarking Interactive Coding with Execu-*
429 *tion Feedback*. Oct. 30, 2023. URL: <http://arxiv.org/abs/2306.14898>. preprint.
- 430 [50] Shunyu Yao et al. “Webshop: Towards scalable real-world web interaction with grounded lan-
431 guage agents”. In: *Advances in Neural Information Processing Systems 35* (2022), pp. 20744–
432 20757.
- 433 [51] Chaoyun Zhang et al. *UFO: A UI-Focused Agent for Windows OS Interaction*. Mar. 1, 2024.
434 URL: <http://arxiv.org/abs/2402.07939>. preprint.
- 435 [52] Chi Zhang et al. *AppAgent: Multimodal Agents as Smartphone Users*. Dec. 21, 2023. URL:
436 <http://arxiv.org/abs/2312.13771>. preprint.
- 437 [53] Jiwen Zhang et al. *Android in the Zoo: Chain-of-Action-Thought for GUI Agents*. 2024. arXiv:
438 2403.02713 [cs.CL].
- 439 [54] Shun Zhang et al. “Planning with large language models for code generation”. In: *arXiv*
440 *preprint arXiv:2303.05510* (2023).
- 441 [55] Longtao Zheng et al. “Synapse: Trajectory-as-Exemplar Prompting with Memory for Computer
442 Control”. In: *The Twelfth International Conference on Learning Representations*. 2024. URL:
443 <https://openreview.net/forum?id=Pc8AU1aF5e>.
- 444 [56] Shuyan Zhou et al. *WebArena: A Realistic Web Environment for Building Autonomous Agents*.
445 Oct. 24, 2023. URL: <http://arxiv.org/abs/2307.13854>. preprint.
- 446 [57] Kaijie Zhu et al. “DyVal: Dynamic Evaluation of Large Language Models for Reasoning
447 Tasks”. In: *The Twelfth International Conference on Learning Representations*. 2024. URL:
448 <https://openreview.net/forum?id=gjfoL9z5Xr>.

449 **A Benchmark Detail**

450 Section A.1 shows the system design and implementation strategies of environments and evaluators.
 451 Section A.2 is the crab framework implementation details at code level. Section A.3 describes the our
 452 experiment settings in detail. Section A.4 describes the specific data format defined in our framework.
 453 Fig. 3 shows the structure of modules inside CRAB Benchmark-v0.

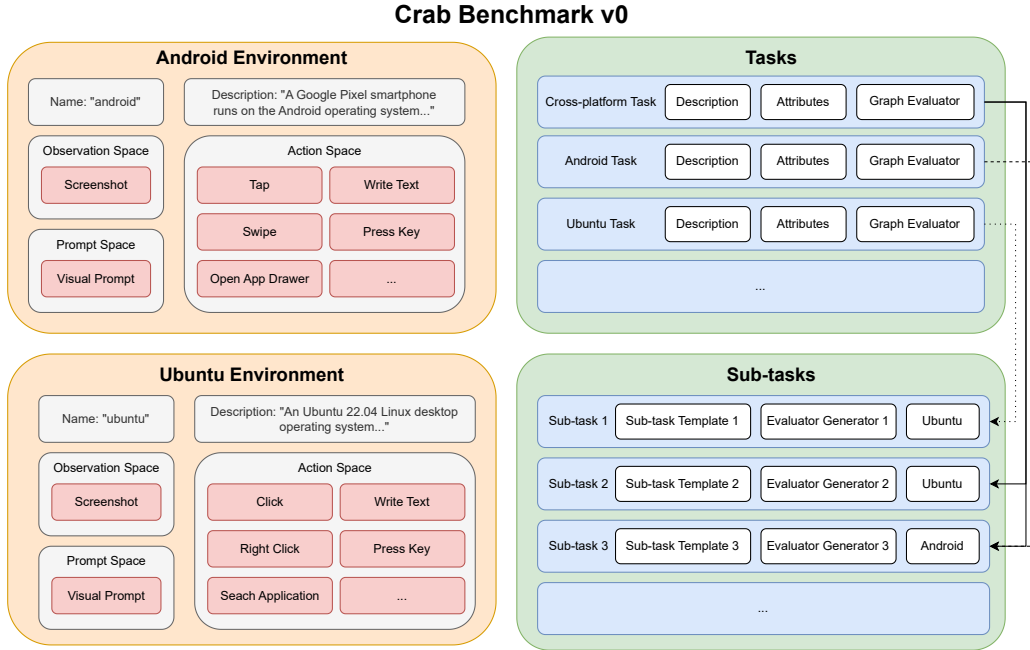


Figure 3: **Module Structure of CRAB Benchmark-v0.** The benchmark is divided into two primary sections: the left section, highlighted with warm hues, features two environments, while the right section, accentuated with cool hues, outlines various tasks. Each environment is defined by attributes including 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.

454 **A.1 Overview**

455 The Ubuntu environment is launched on a QEMU/KVM [3, 14] Virtual Machine, and the Android
 456 environment employs the Google Android Emulator². Both environments utilize snapshots to ensure
 457 a consistent state across all sessions. This allows each experiment to start from an identical state,
 458 providing a controlled setup for all test agents. Interaction with the Ubuntu environment is facilitated
 459 using PyAutoGUI³ and MSS⁴, which provide high-level commands for mouse and keyboard control
 460 and screen capture, respectively. For the Android environment, we use the Android Debug Bridge
 461 (ADB)⁵.

462 **Observation Space** The observation space consists solely of the current system screen for both
 463 environments, captured in image format at each step of the agent’s interaction. We employ the

²<https://developer.android.com/studio/run/emulator>

³<https://github.com/asweigart/pyautogui>

⁴<https://github.com/BoBoTiG/python-mss>

⁵<https://developer.android.com/tools/adb>

464 Set-of-Marks visual prompt method [48] to label each interactive element on the screen. Interactive
 465 elements are identified using the GroundingDINO [22] with `icon.logo. text` prompt to locate all
 466 interactive icons. Additionally, Optical Character Recognition (OCR) is utilized through EasyOCR⁶
 467 to detect and label interactive text elements. Each detected item is assigned a unique integer ID,
 468 facilitating reference within the action space.

469 **Action Space** The action spaces for Ubuntu and Android are distinct and designed to be close to the
 470 common interactions in the real devices. For Ubuntu, we define the following actions: mouse-based
 471 actions, keyboard-based actions and a shortcut action to search for applications. For Android, the
 472 action set includes tapping actions, a text action, a physical button action, and an action to open the
 473 app drawer. Additionally, we introduce two environment-irrelevant actions: completing the task and
 474 submitting an answer. Detailed descriptions for all actions are shown in Table 3.

Table 3: **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 <code>elem</code> .
<code>right_click(elem)</code>	Right-click on <code>elem</code> .
<code>write_text(text)</code>	Typing the specified <code>text</code> .
<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 <code>name</code> in the system.
<code>tap(elem)</code>	Tap on <code>elem</code> .
<code>long_tap(elem)</code>	Press and hold <code>elem</code> .
<code>swipe(elem,dire,dist)</code>	Swipe from <code>elem</code> in a specified <code>direction</code> and <code>distance</code> .
<code>write_text(text)</code>	Typing the specified <code>text</code> .
<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 <code>answer</code> if needed.
<code>complete()</code>	Tell system a task is completed.

475 **Evaluator Design** To assess the intermediate states of sub-tasks as described in Sec. 4.2, we have
 476 implemented a comprehensive suite of execution-based reward functions (evaluators) [46]. These
 477 evaluators retrieve and assess specific current states, such as the edited content of a file or a modified
 478 setting, thereby determining the successful completion of a sub-task. For each evaluator, input
 479 attributes are carefully selected to interpret software information or system settings relevant to the
 480 scenario defined for the sub-task. For instance, evaluators use file paths before and after edits as
 481 input parameters to verify the completion of file editing sub-tasks. Specifically, for sub-tasks on the
 482 Android platform, we incorporate XML-based evaluators [47]. We dump UI layout as XML path
 483 and verify whether the UI content matches the expected state. For the Ubuntu platform, we employ
 484 image matching techniques [31, 11, 7] and OCR to handle scenarios where acquiring necessary
 485 state information through conventional APIs is challenging. Image matching offers fine-grained
 486 visual correspondences by comparing keypoint features between images, allowing us to assess spatial
 487 relationships among visual elements. Using OCR and image matching, we can accurately evaluate
 488 tasks such as verifying whether an agent has successfully created a slide with specified images, text
 489 content, and layouts—tasks for which trivial evaluation methods are lacking. We utilize EasyOCR⁶
 490 and XFeat⁷ as our primary tools for OCR and image matching. For tasks with real-time characteristics
 491 that may change over time, we implement crawler scripts to capture dynamic values at the moment

⁶<https://github.com/JaidedAI/EasyOCR>

⁷https://github.com/verlab/accelerated_features

492 of evaluation. These values are then compared with the results achieved by the agent upon task
493 completion. We have a total of 59 evaluator functions.

494 A.2 Framework Design

495 CRAB offers a modular and extensible framework for evaluating agent performance in diverse tasks.
496 At the heart of the framework lies the *action*, a unit operation representing the fundamental operation
497 within the benchmark. The *action* is essentially an executable Python function that can be defined
498 with explicit typed parameters and a clear description. *actions* serve not only as building blocks but
499 also as interfaces through which agents interact with the environment. The *evaluator* is a specialized
500 *action* restricted to returning boolean values, signifying the success or failure of an agent’s task. It
501 enhances the *actions* by analyzing the state of the environment and the sequence of *actions* executed
502 by the agent, providing a decisive metric of task accomplishment. Additionally, multiple *evaluators*
503 can be interconnected to form a graph evaluator for complex tasks (Sec. 4.2).

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

510 CRAB utilizes Python functions to define all actions and evaluators, embodying a "code as configura-
511 tion" philosophy. Each function’s docstring outlines its description and parameter definitions, which
512 are then presented to the agent as structured prompts. Compared to traditional methods using data
513 interchange formats like JSON or YAML, Python code configurations provide a more structured
514 approach and fits in modern IDE.

515 By decoupling actions, environments, tasks, and evaluations, CRAB facilitates a plug-and-play archi-
516 tecture that can adapt to various scenarios. Such a system is scalable, maintainable and expandable,
517 allowing researchers and developers to introduce new tasks and environments without restructuring
518 the entire framework. Our implementation uses *networkx* [8] for building graph and *dill* [25] for
519 function serialization in our implementation.

520 A.3 Configuration Format by Modules

521 Building on the declarative and modular design of our framework, this section explains the configura-
522 tion and potential extensibility of each module.

523 **Environment** The environments in CRAB are a combination of multiple different uses of ac-
524 tions with some environment metadata, such as name and natural language description. In CRAB
525 Benchmark-v0, we use a computer desktop environment and a smartphone environment both based
526 on virtual machine technology. The computer desktop environment, named *Ubuntu*, is installed
527 from an ISO image of Ubuntu 22.04.4 LTS (Jammy Jellyfish) downloaded from the Ubuntu Official
528 website⁸. Necessary applications such as the LibreOffice suite (Writer, Calc, and Impress) and Slack
529 are installed later via snap and apt, according to the task dataset requirements. The smartphone
530 environment, named *Android*, is installed using pre-defined devices (Google Pixel 8 Pro with release
531 name *R*) provided in Google Android Studio⁹. We install additional required applications such as
532 *Keep Notes*, *Tasks*, and *Docs* from Google Play. The descriptions of the two environments in CRAB
533 Benchmark-v0, which are inserted in the agent prompts, are as follows:

- 534 • **Ubuntu:** An Ubuntu 22.04 Linux desktop operating system. The interface displays a current
535 screenshot at each step and primarily supports interaction via mouse and keyboard. You

⁸<https://releases.ubuntu.com/jammy/ubuntu-22.04.4-desktop-amd64.iso>

⁹<https://developer.android.com/studio>

536 must use searching functionality to open any application in the system. This device includes
537 system-related applications including Terminal, Files, Text Editor, Vim, and Settings. It also
538 features Firefox as the web browser, and the LibreOffice suite—Writer, Calc, and Impress.
539 For communication, Slack is available. The Google account is pre-logged in on Firefox,
540 synchronized with the same account used in the Android environment.

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

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

```
550 @action
551 def search_application(name: str) -> None:
552     """Search an application name.
553
554     For exmaple, if you want to open an application named "slack",
555     you can call search_application(name="slack"). You MUST use this
556     action to search for applications.
557
558     Args:
559         name: the application name.
560     """
561     pyautogui.hotkey("win", "a")
562     time.sleep(0.5)
563     pyautogui.write(name)
564     time.sleep(0.5)
```

Listing 1: Define "search_application" action.

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

566 • **Name:** The action name serves as the identifier for backend models. It should semantically
567 match the action’s behavior to improve the accuracy of the agent in executing the action.
568 The function name is extracted as the action name. In this example, `search_application`
569 is the assigned name.

570 • **Description:** The description provides a natural language explanation of the action to assist
571 the agent in understanding how to use it. The main body of the function’s docstring is used
572 as the description. For example, in this instance, the description outlines the basic usage of
573 the action: *Search an application name*, along with an example of its usage.

574 • **Parameters:** The parameters are the arguments that the functions accept, offering flexibility
575 for the agent to control the environment. Typically, a set of parameters is defined, each
576 consisting of a name, type, and a natural language description. Parameters are extracted
577 from the function’s parameters along with their type annotations. Additionally, parameter de-
578 scriptions are extracted from the `Args` section in the docstring. In this example, there is only
579 one parameter named `name`, with a type of `str`, and its description is the `application`
580 `name`.

581 • **Entry:** The entry represents the implementation of the function, defined within the function
582 body to specify how the action is executed. When the agent invokes the function, the entry
583 is executed with the provided parameters. In this example, we utilize the `pyautogui` package
584 for keyboard control. Initially, it presses a hotkey to enter the application search panel in
585 Ubuntu, then proceeds to type the application name provided by the parameters, finally
586 displaying the search results.

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

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

Listing 2: Define the "screenshot" observation action.

601 This action captures the screen's current view and encodes it in Base64 format. Additionally, visual
602 prompts are also defined by actions that utilize the output from an observation action as their input,
603 further processing it to generate a visual prompt for the agent.

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

```
608 @evaluator(env_name="ubuntu")
609 def check_text_in_current_window_name(text: str) -> bool:
610     try:
611         out = subprocess.check_output(
612             ["xdotool", "getwindowfocus", "getwindowname"], text=True
613         ).strip()
614     except subprocess.CalledProcessError:
615         return False
616     return text in out
```

Listing 3: Define "check_text_in_current_window_name" evaluator.

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

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

```
625 Task(
626     id="a3476778-e512-40ca-b1c0-d7aab0c7f18b",
627     description="Open \"Tasks\" app on Android, check the...",
628     evaluator=path_graph(
629         check_current_package_name("com.google.android.apps.tasks"),
630         check_current_window_process("gnome-control-center"),
631         check_color_scheme("prefer-dark"),
632     ),
633 )
```

Listing 4: Define a task.

634 In this model, each task is represented as an instance of the Task class, which is a subclass of
635 BaseModel in *Pydantic*¹⁰ package. Each task is uniquely identified by an ID and described by a
636 detailed description. The evaluator component is structured as a graph evaluator, which integrates
637 multiple evaluative functions into a directed graph using the *networkx*¹¹ package. Each evaluator
638 within this graph must be appropriately parameterized to assess specific conditions relevant to the
639 task. For example, the task demonstrated aims to open the "Tasks" app on Android and perform
640 a series of verifications: it checks whether the correct Android app is opened, whether the current
641 focused window's process name is `gnome-control-center`, and whether the color scheme is set to
642 dark.

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

```
645 SubTask(  
646     id="0f589bf9-9b26-4581-8b78-2961b115ab49",  
647     description="Open \{file_path}\" using vim in a terminal, write  
648     \{content}\", then save and exit vim.",  
649     attribute_dict={"file_path": "file_path", "content": "message"},  
650     output_type="file_path",  
651     evaluator_generator=lambda file_path, content: path_graph(  
652         check_current_window_process("gnome-terminal-server"),  
653         is_process_open("vim"),  
654         is_process_close("vim"),  
655         check_file_content(file_path, content),  
656     ),  
657 ),
```

Listing 5: Define a task.

658 In this sub-task model, each sub-task is defined using a similar approach to the main task. The
659 attributes of the sub-task are outlined in an `attribute_dict`, which details the types and roles of
660 each attribute used in the sub-task's operations. The `output_type` field specifies the expected type
661 of output from the sub-task. The types reflected in `attribute_dict` and `output_type`, play a
662 critical role in determining the compatibility and sequential logic of compose multiple sub-tasks.
663 The evaluator for the sub-task is dynamically generated using a lambda function, which crafts an
664 evaluator sub-graph based on the sub-task's attributes.

665 A.4 Task Dataset

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

670 B Agent system

671 B.1 Agent Implementation

672 In this section, we outline the implementation of the agents used in our experiments, which leverage
673 advanced multimodal language models from OpenAI, Anthropic, and Google. Each agent is designed
674 to function in multi-environment setups, interacting with various action spaces defined by different
675 environments.

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

¹⁰<https://pydantic.dev/>

¹¹<https://networkx.org/>

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

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

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

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

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

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

707 **B.2 Inter-agent Communication Strategies**

708 In this section we introduce the details of two multi-agent communications methods, which are
709 introduced in 5.2.

710 **Multi-agent Communication by Functionality** This setting involves two agents: a main agent
711 prompted with the task description and a tool agent with the entire action space. The main agent
712 generates the instruction for the next step and sends it to the tool agent. The tool agent chooses the
713 proper action with parameters and a target environment, then feeds it back to the system.

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

721 **B.3 Agent Prompt**

722 **B.3.1 Single 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. Your answer must be a least one function call. please do not output any other information. You must make sure all function calls get their required parameters.

723

724 **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.

725

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.

726

727 **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.

728

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.

729

Table 4: Evaluation results on Ubuntu tasks.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4O	Single	10.34	26.64	2.72	4.68×10^{-4}	5.17	60.34	24.14
GPT-4O	By Func	6.90	21.90	2.07	3.86×10^{-4}	6.90	60.34	25.86
GPT-4O	By Env	8.62	20.60	2.06	2.01×10^{-4}	3.45	48.28	39.66
GPT-4 TURBO	Single	12.07	28.36	3.82	8.79×10^{-4}	1.72	63.79	22.41
GPT-4 TURBO	By Func	10.34	24.45	3.10	4.74×10^{-4}	8.62	34.48	46.55
GEMINI 1.5 PRO	Single	1.72	7.61	0.54	\	0.00	46.55	51.72
GEMINI 1.5 PRO	By Func	1.72	3.30	0.30	\	0.00	20.69	77.59
CLAUDE 3 OPUS	Single	1.72	9.54	1.41	3.42×10^{-4}	5.17	56.90	36.21
CLAUDE 3 OPUS	By Func	1.72	6.75	0.65	2.81×10^{-4}	27.59	31.03	39.66

Table 5: 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	5.12	7.17×10^{-4}	13.79	58.62	3.45
GPT-4O	By Func	24.14	48.74	5.77	9.19×10^{-4}	24.14	37.93	13.79
GPT-4O	By Env	27.59	53.34	5.93	4.58×10^{-4}	13.79	44.83	13.79
GPT-4 TURBO	Single	10.34	30.53	2.84	3.36×10^{-4}	20.69	62.07	6.90
GPT-4 TURBO	By Func	20.69	37.01	4.32	5.92×10^{-4}	13.79	51.72	13.79
GEMINI 1.5 PRO	Single	17.24	34.52	4.09	\	10.34	65.52	6.90
GEMINI 1.5 PRO	By Func	17.24	35.99	3.88	\	31.03	41.38	10.34
CLAUDE 3 OPUS	Single	17.24	43.62	5.30	7.78×10^{-4}	13.79	51.72	17.24
CLAUDE 3 OPUS	By Func	13.79	42.30	4.20	5.07×10^{-4}	44.83	31.03	10.34

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.

730

731 C Further Result Analysis

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

735 **C.1 Result by Platforms**

736 Table 4, 5 and 6 show the experiment results on Ubuntu Tasks, Android Tasks, and cross-platform
737 Tasks, respectively.

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

744 In multi-agent system organized by environment, consistently yields better results in both Android
745 and cross-platform tasks. This configuration appears to enhance the agents’ ability to manage and
746 adapt to diverse tasks more effectively, leveraging environmental specifics to optimize performance.
747 This suggests that employing multiple agents that are either specialized or specifically configured to
748 operate within the same environment can significantly improve task handling and overall adaptability.

749 Cross-platform tasks present a greater challenge for all models, as evidenced by lower Success
750 Rates and Completion Ratios. These tasks, which necessitate functionality across different operating
751 systems or platforms, demand a broader capability range and more sophisticated agent coordination.
752 The importance of CR is especially critical in such environments, where it serves as a more reliable
753 metric for distinguishing between agent models than SR. Given the presence of all Gemini and
754 Claude agents’ SR is 0.0, indicating that Completion Ratio more effectively captures an agent model’s
755 capability, thereby better reflecting its robustness and adaptability to complex requirements.

756 Furthermore, analyzing the reasons for task termination offers additional insights into the operational
757 challenges these models encounter. False Completion is notably prevalent in Android tasks. Reach
758 Step Limit remains the most frequent cause of termination, particularly in cross-platform tasks. The
759 Claude model exhibits a significantly high Invalid Action ratio in cross-platform tasks, indicating its
760 difficulties in managing multi-environment scenarios effectively.

761 Overall, these findings underscore the necessity of selecting the appropriate agent model and con-
762 figuration based on specific platform and task needs. The variability in model performance across
763 different setups also highlights the ongoing need for development and refinement of multi-agent
764 systems to enhance their versatility and efficacy in increasingly diverse and complex operational
765 environments.

766 **C.2 Comparison between Single Agent and Multi-agent**

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

769 First, comparing in False Completion Rate, we attribute the lower Success Rate (SR) of Multi-agent
770 to a high False Completion Rate—where the agent incorrectly assumes that the task is complete. As
771 observed in failure cases (e.g., the Cross-platform Task case study in Appendix C.3), Sub-agents
772 often misinterpret the Main agent’s instructions. Despite being required to perform a final action,
773 the instructions lead Sub-agents to prematurely conclude that the task is complete, resulting in
774 incorrect “complete” actions. While this issue also occurs in Multi-Env, it happens less frequently.
775 We believe this is due to information loss during inter-agent communication. Natural language, while
776 effective for aligning with human understanding in LLM communication, is less suited for inter-agent
777 communication, leading to information loss during compression and interpretation, which weakens
778 the performance of multi-agent structures.

779 Next, comparing in Invalid Action Rate, we observe that in single-platform tasks, both Multi-Env and
780 Multi-Func suffer from similar inter-agent communication issues, as indicated by their high False
781 Completion and Invalid Action rates (Table 4 and 5). However, in cross-platform tasks (Table 6), the
782 Single agent’s Invalid Action rate is significantly higher than that of the Multi-agent structures. Cross-

Table 6: Evaluation results on cross-platform tasks.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4O	Single	7.69	45.53	4.54	3.57×10^{-4}	0.00	53.85	38.46
GPT-4O	By Func	15.38	43.41	3.19	2.25×10^{-4}	7.69	61.54	15.38
GPT-4O	By Env	7.69	48.61	3.69	1.68×10^{-4}	15.38	61.54	15.38
GPT-4 TURBO	Single	7.69	47.84	4.31	2.89×10^{-4}	0.00	69.23	23.08
GPT-4 TURBO	By Func	7.69	39.05	3.73	2.51×10^{-4}	15.38	46.15	30.77
GEMINI 1.5 PRO	Single	0.00	21.25	1.49	\	0.00	69.23	30.77
GEMINI 1.5 PRO	By Func	0.00	16.70	1.52	\	7.69	69.23	23.08
CLAUDE 3 OPUS	Single	0.00	24.69	2.33	1.62×10^{-4}	0.00	38.46	61.54
CLAUDE 3 OPUS	By Func	0.00	20.07	2.32	1.49×10^{-4}	0.00	38.46	61.54

783 platform tasks require frequent environment changes with varying action spaces, and if the model’s
784 performance output is inadequate, it often generates correct actions in the wrong environment, invalid
785 actions in the correct environment, or correct actions in correct environment but in the wrong format.
786 This phenomenon highlights the limitations of current general-purpose LLMs, where multi-agent
787 structures can be advantageous. By assigning each agent a specific responsibility and a limited action
788 space, multi-agent structures can mitigate these issues.

789 Last, when comparing different types of tasks, we observe that the Multi-Env structure significantly
790 outperforms the Single Agent in Android and cross-platform tasks but underperforms in Ubuntu tasks.
791 The key difference between the Single Agent and Multi-Env lies in the average context length each
792 agent processes. As demonstrated by Liu et al. [21], more context does not always lead to better
793 performance. The Single Agent is burdened with extensive knowledge across different fields, making
794 it challenging for the model to switch between multiple environments, particularly when managing
795 long history chat messages. In contrast, the sub-agents in the Multi-Env structure handle part of the
796 total prompt, which enhances their performance in more complex tasks. While different backend
797 models show varying performance across environments, resulting in some instability, the general
798 trend is that the more complex the task, the more advantageous the multi-agent approach becomes.

799 In summary, the performance difference between multi-agent and single-agent structures largely
800 depends on the task complexity. For tasks that are too complex for a single general-purpose agent, a
801 multi-agent structure may perform better. Conversely, for simpler tasks, multi-agent structures tend
802 to cause information loss during inter-agent communication, leading to misunderstandings among
803 downstream agents.

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

810 C.3 Case Study

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

814 C.3.1 Cross-platform Task

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

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

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

835 C.3.2 Ubuntu Task

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

838 This task can be approached through multiple methods. An agent may opt for a straightforward
839 strategy first using the `search_application` command to find the Terminal, then using Linux
840 commands to create the directory and copy the necessary files. Alternatively, the agent could employ
841 a GUI-based approach, manually creating the folder and selecting files through actions like `click`
842 and `right_click`. We evaluate various agent systems in a single-agent setting for this task. As
843 illustrated in Table 7–10, both GPT-4o and GPT-4 Turbo from OpenAI successfully interpret the task
844 instructions and employ a simpler solution using Terminal commands. These agents also demonstrate
845 superior capability in understanding the UI, selecting the correct commands, and accurately using the
846 Terminal application to fulfill the task requirements.

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

858 C.3.3 Android Task

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

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

867 **GPT-4o multi-agent by functionality** In steps 1-11, the agent tries to open the Contacts app but
868 mistakenly opens Google Assistant multiple times. In steps 12-14, the agent successfully enters the
869 Contacts app and finds the contact information. The agent then returns to the home page, and the
870 process is terminated due to the limitation of operation steps.

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

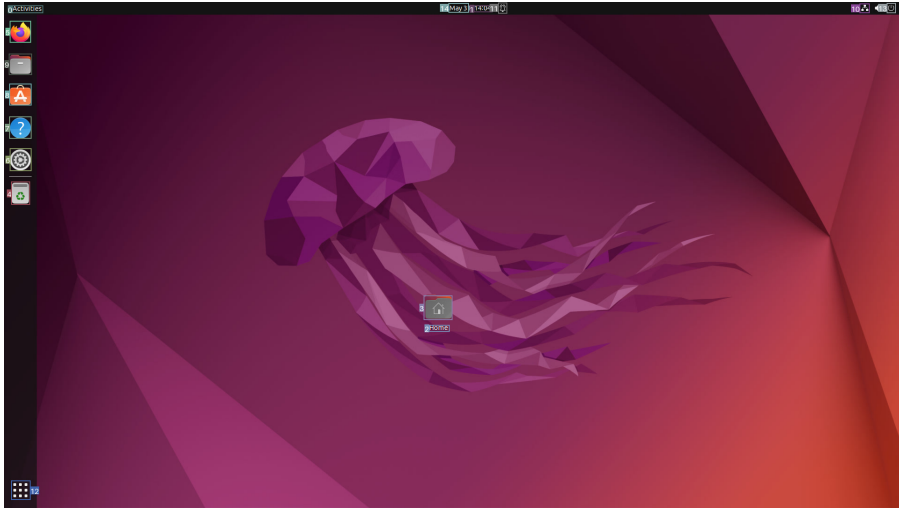
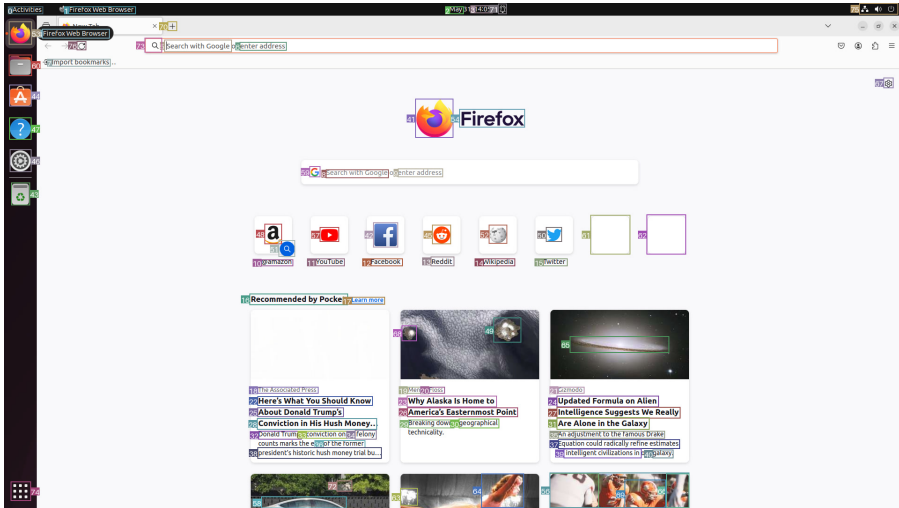
875 **Gemini 1.5 Pro multi-agent by functionality** In steps 1-2, the agent finds the Contacts app and
876 enters it. However, the agent misunderstands the instruction, gets lost in creating a new contact with
877 the given name, and cannot obtain the corresponding information.

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

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

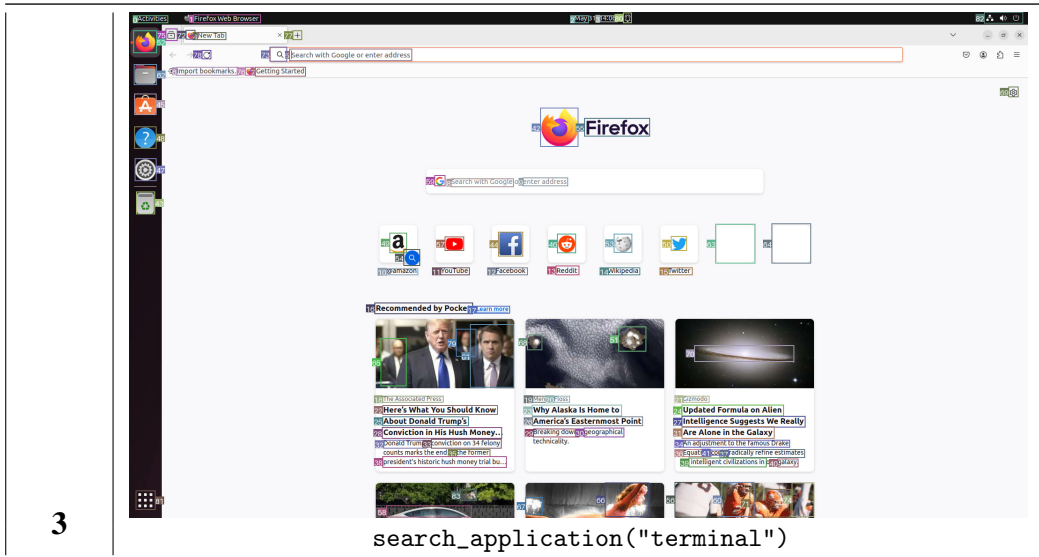
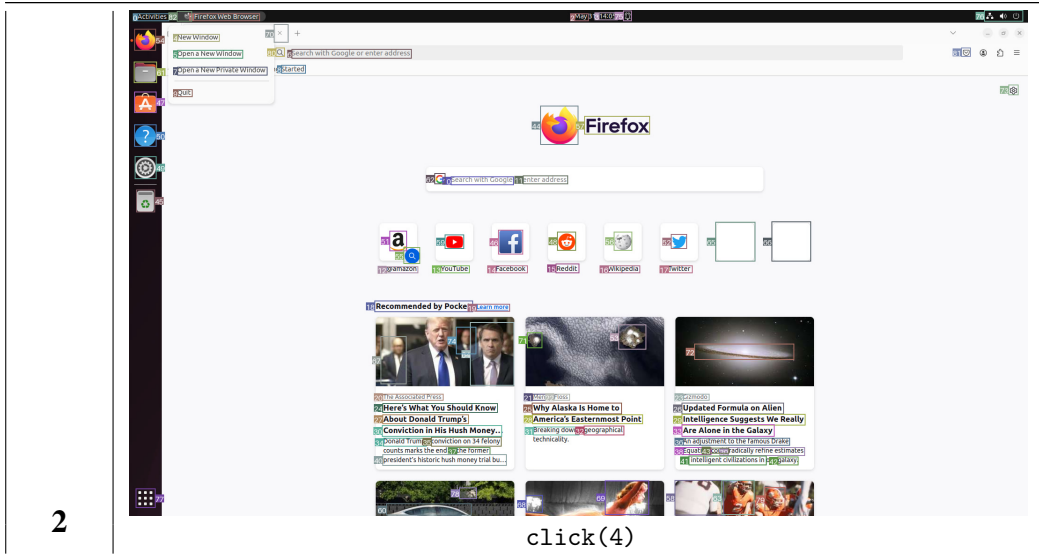
888 **Analysis** For the agents which are organized by functionality, Gemini 1.5 Pro struggles to complete
889 the first operation. Although it recognizes and opens the Contacts app as instructed, it fails to proceed
890 further. In contrast, Claude 3 Opus and GPT-4o successfully obtain the necessary information. In
891 the initial phase, the multiple agents agree that opening the Contacts app is the first step. However,
892 they often fail to find the correct position to tap, frequently opening incorrect apps such as Google
893 Assistant and Messages. Once the agents do open the correct app, they usually find the email address
894 of the contact quickly. Even when agents plan to go back home and open the Gmail app to send the
895 message, due to the limitation of operations, the system ended. As shown in steps 3-5, GPT-4 Turbo
896 quickly finishes the corresponding task after opening the correct app. However, similar to GPT-4o,
897 GPT4-Turbo agents get stuck as they can not open the correct apps in the following steps. Besides,
898 GPT-4o (multi-agent by environment) overcomes the issue encountered by GPT-4o (multi-agent by
899 functionality). Even affected by not being able to access the app drawer, the system could still find
900 and copy the corresponding information and change to the Gmail app for further operations.

Table 7: **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="808 867 915 898">click(5)</p>
1	 <p data-bbox="808 1419 915 1451">click(1)</p>

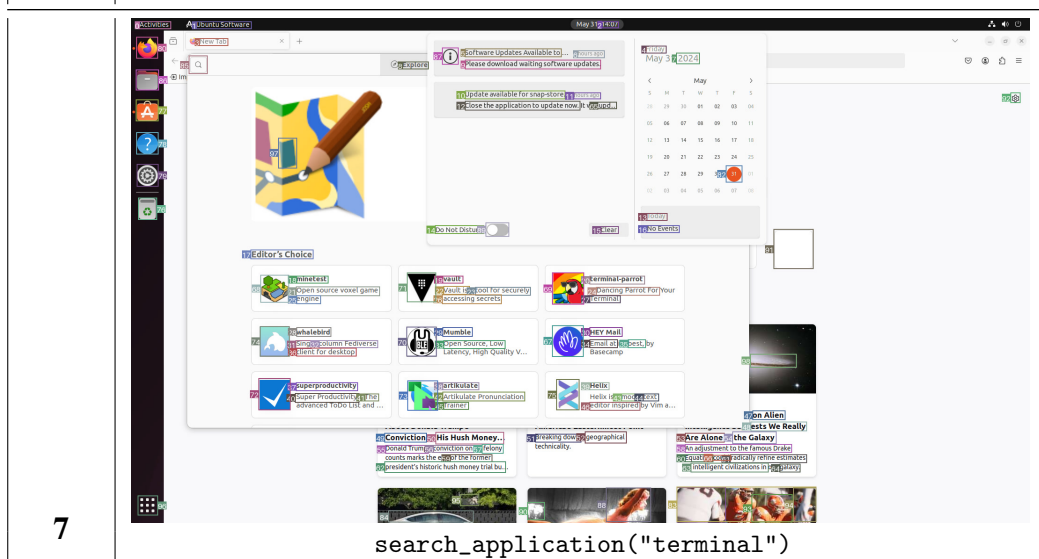
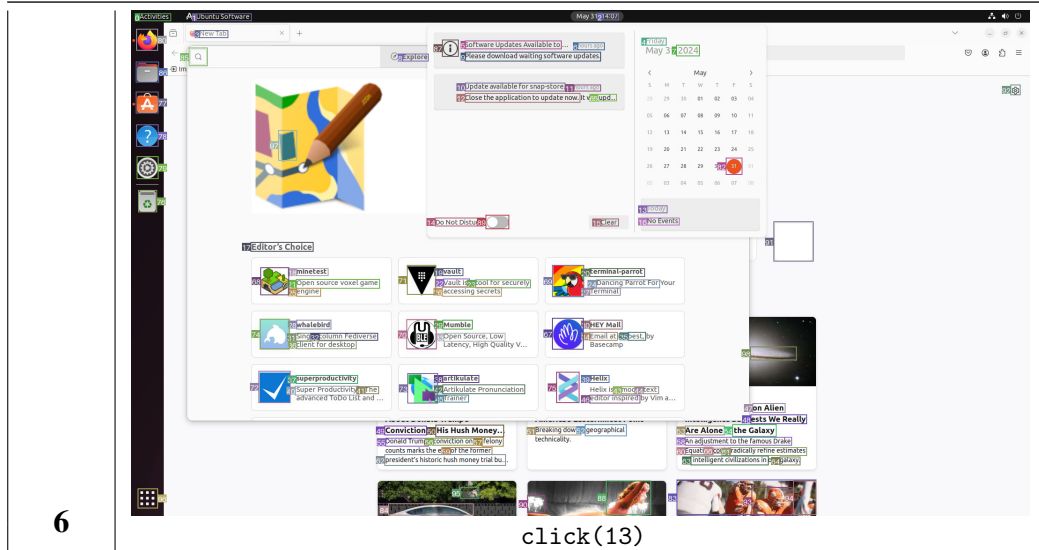
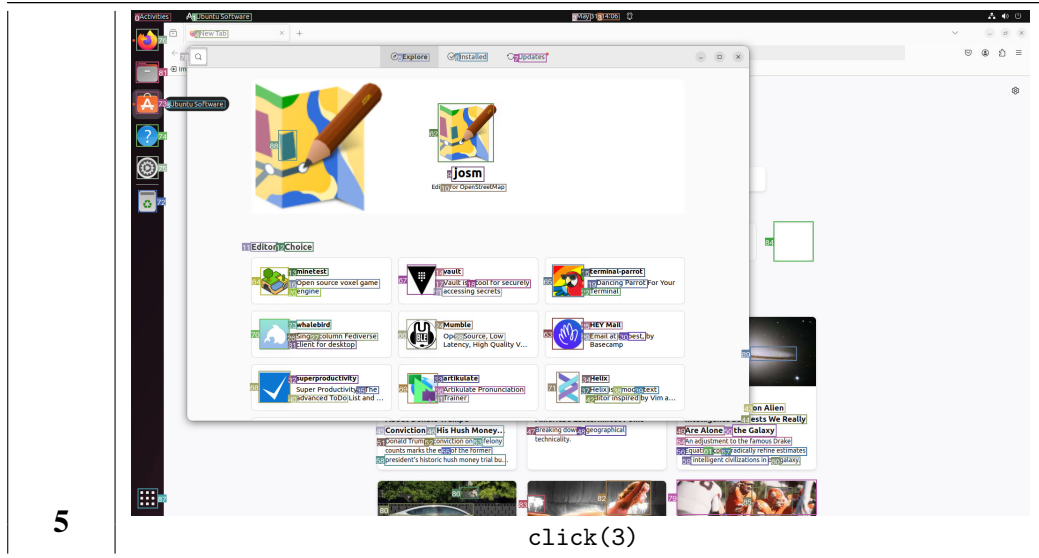
Continued on next page

Table 7 continued from previous page



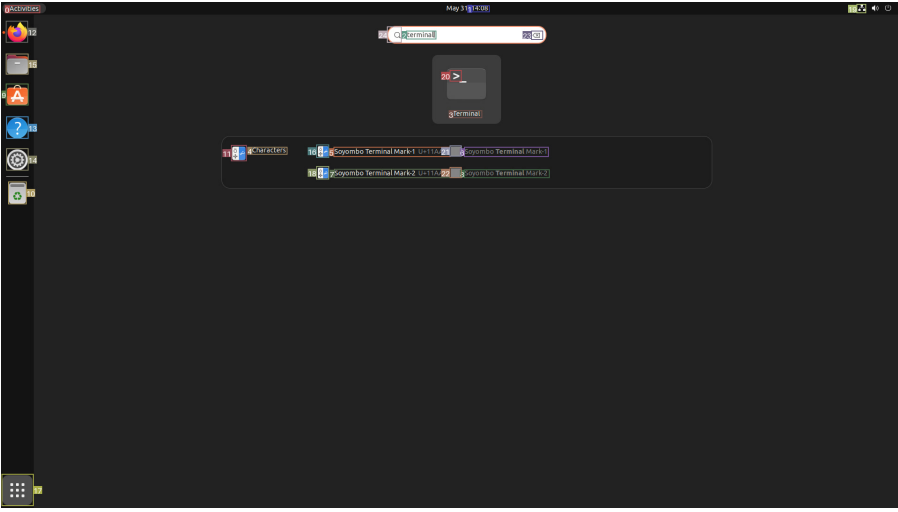
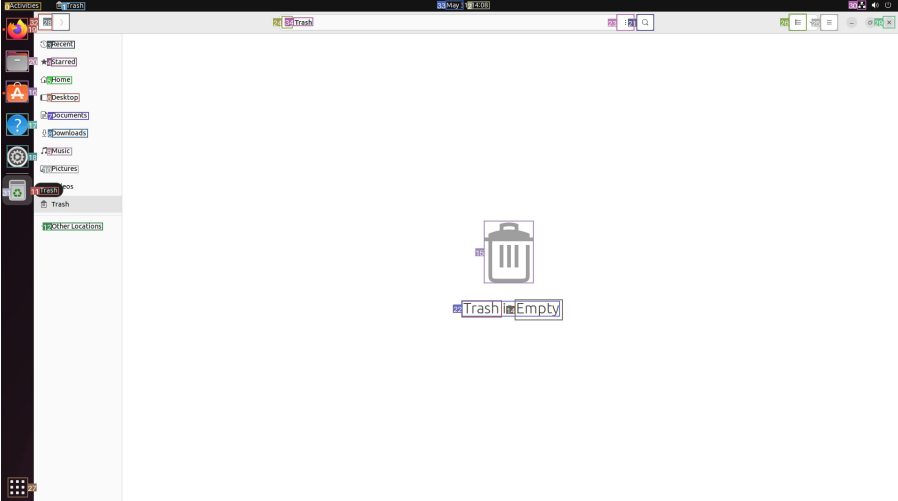
Continued on next page

Table 7 continued from previous page




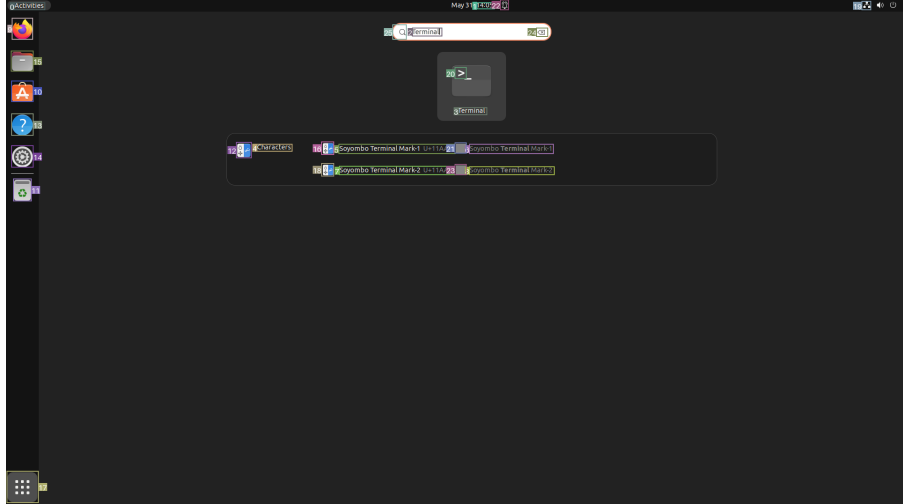
Continued on next page

Table 7 continued from previous page

8	 <p>click(10)</p>
9	 <p>click(6)</p>
10-14	The agent is stuck at this stage and keeps clicking useless elements.

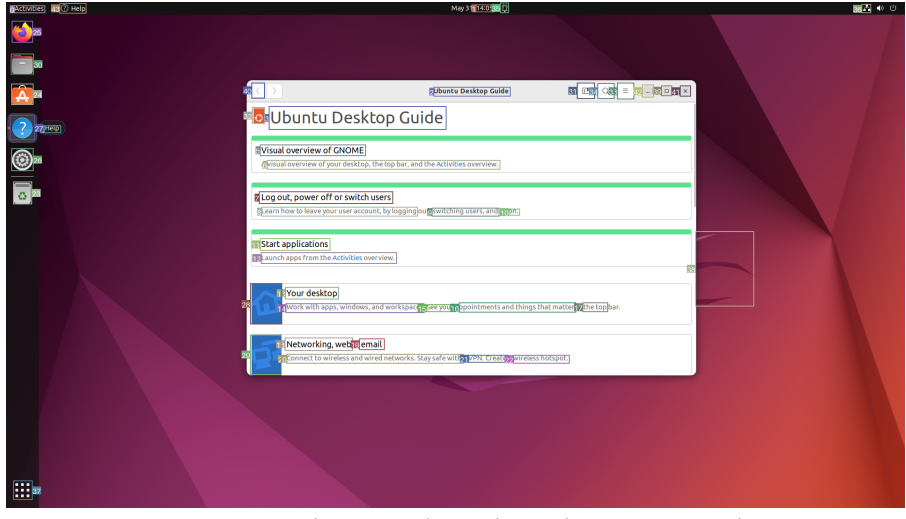
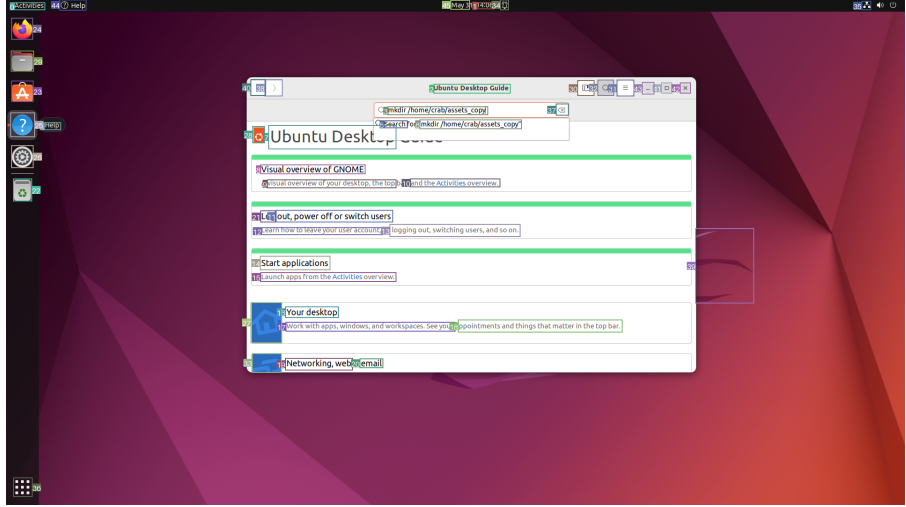
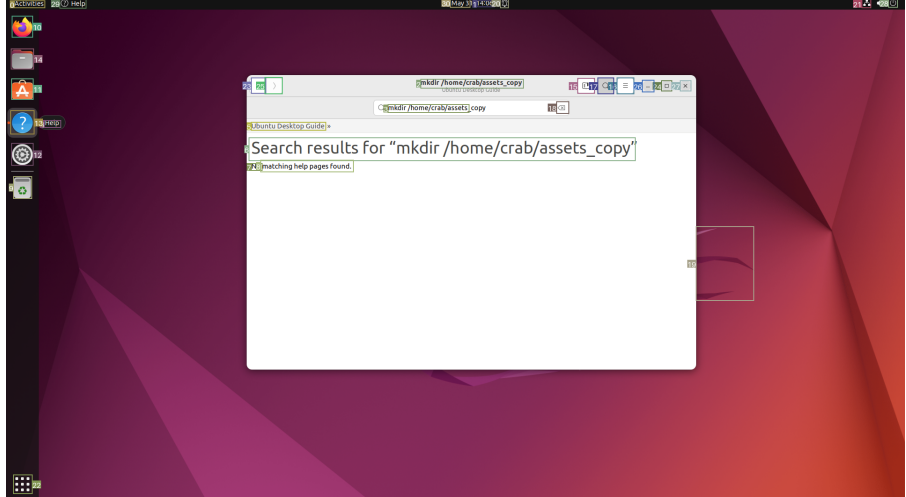
Completed Nodes / Total Nodes: 0 / 2 Termination Reason: Reach Step Limit

Table 8: **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>search_application("terminal")</p>
1	 <p>click(13)</p>

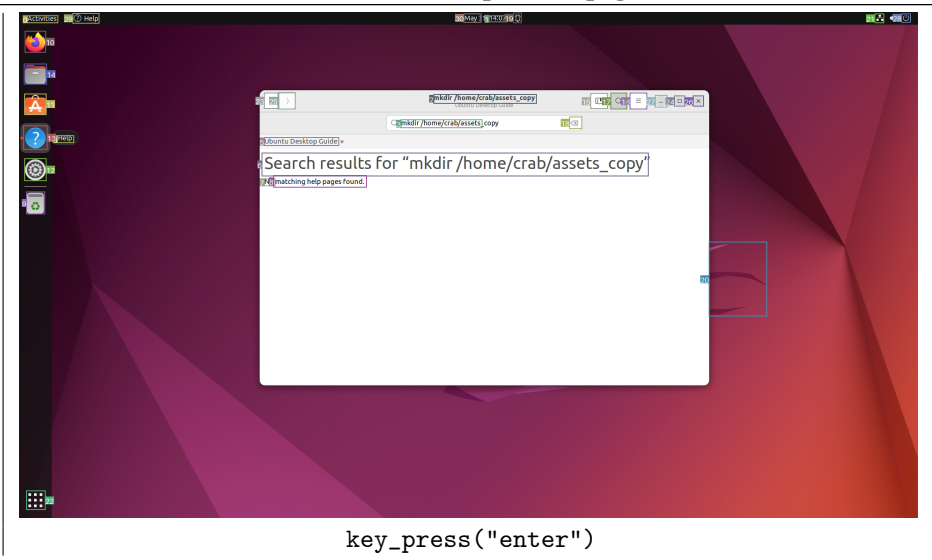
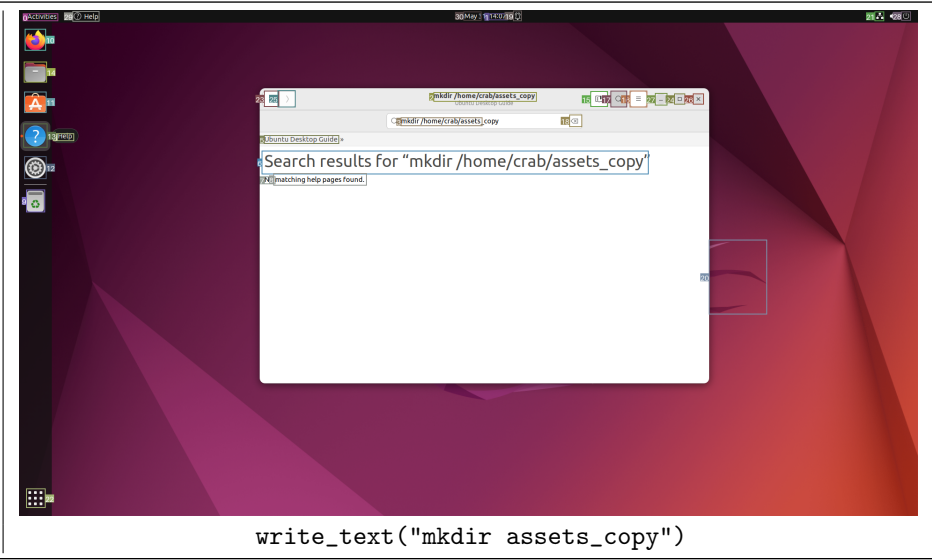
Continued on next page

Table 8 continued from previous page

2	 <p>The screenshot shows the Ubuntu Desktop Guide window open on a desktop. The window title is "Ubuntu Desktop Guide". The content includes sections for "Visual overview of GNOME", "Log out, power off or switch users", "Start applications", "Your desktop", and "Networking, web, email".</p>
3	 <p>The screenshot shows a terminal window with the command <code>mkdir /home/crab/assets_copy</code> entered. The terminal output shows the command being executed successfully.</p>
4	 <p>The screenshot shows a search results window titled "Search results for 'mkdir/home/crab/assets_copy'". The window displays the search results for the command entered in the terminal.</p>

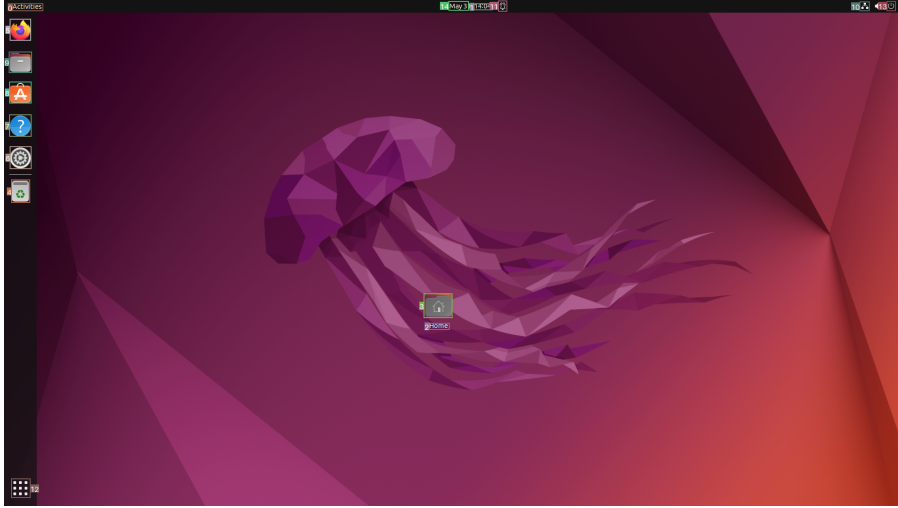

Continued on next page

Table 8 continued from previous page

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

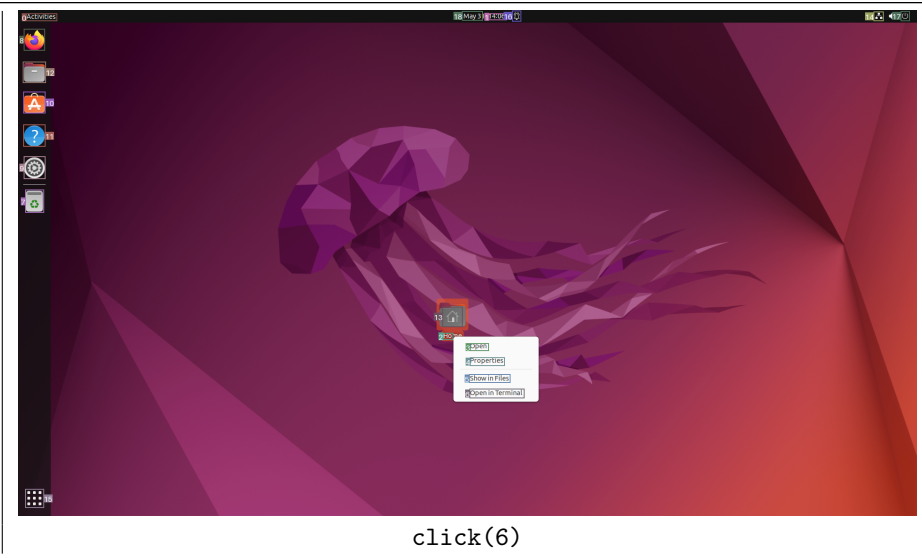
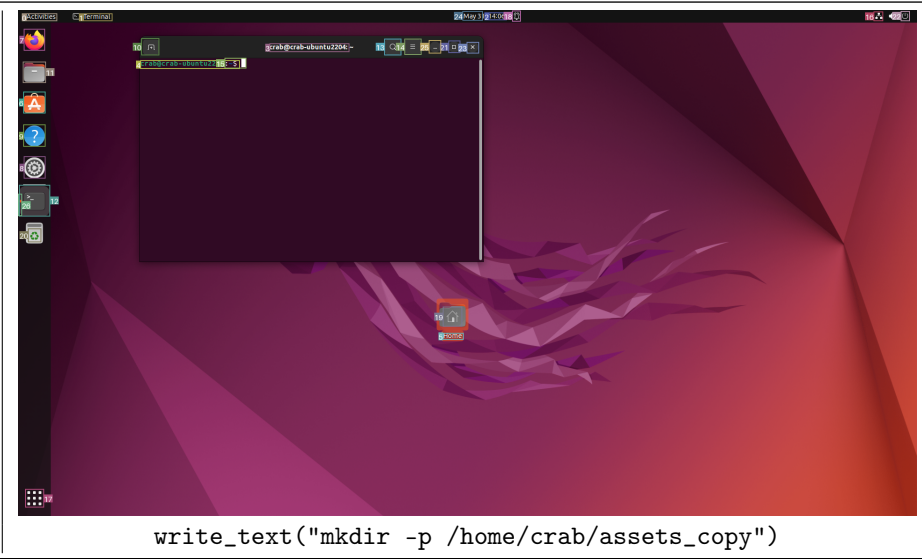
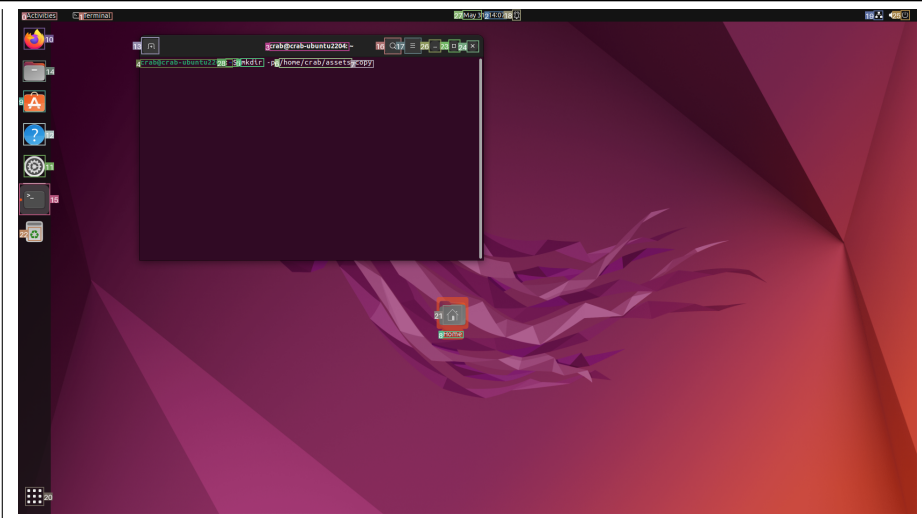
Completed Nodes / Total Nodes: 0 / 2 Termination Reason: Reach Step Limit

Table 9: **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>

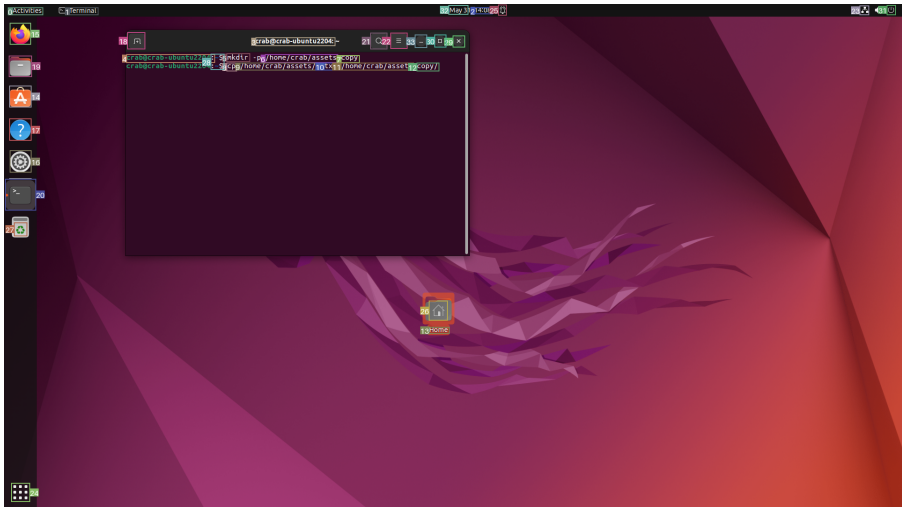
Continued on next page

Table 9 continued from previous page

3	 A screenshot of a Linux desktop environment with a purple and red geometric wallpaper. A context menu is open over a file icon on the desktop. The menu options are: Open, Properties, Show in Files, Show in Home, and Open in Terminal. The desktop has a dock on the left with various application icons.
4	 A screenshot of a terminal window on the same desktop. The terminal shows the command <code>mkdir -p /home/crab/assets_copy</code> being entered. The prompt is <code>crab@crab-ubuntu2204:~\$</code> . The terminal window is titled <code>crab@crab-ubuntu2204:~\$</code> .
5-6	 A screenshot of a terminal window on the same desktop. The terminal shows the directory path <code>/home/crab/assets_copy</code> being displayed. The prompt is <code>crab@crab-ubuntu2204:~\$</code> . The terminal window is titled <code>crab@crab-ubuntu2204:~\$</code> .

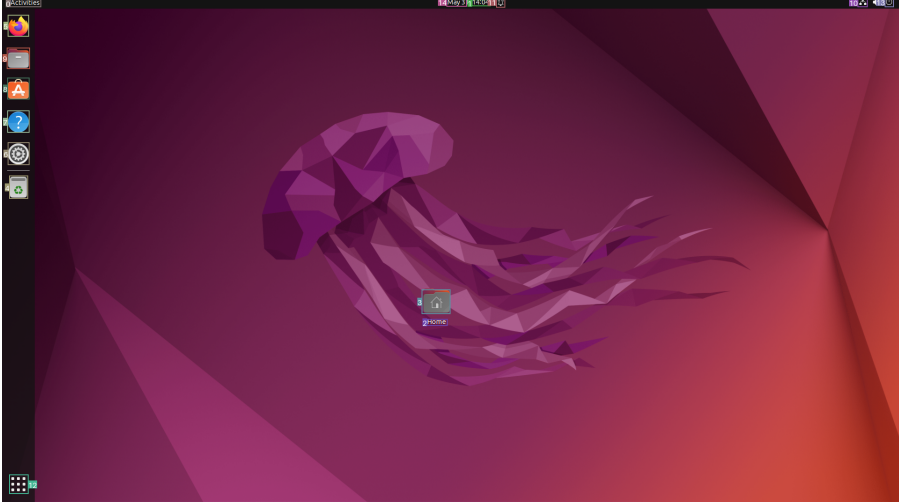
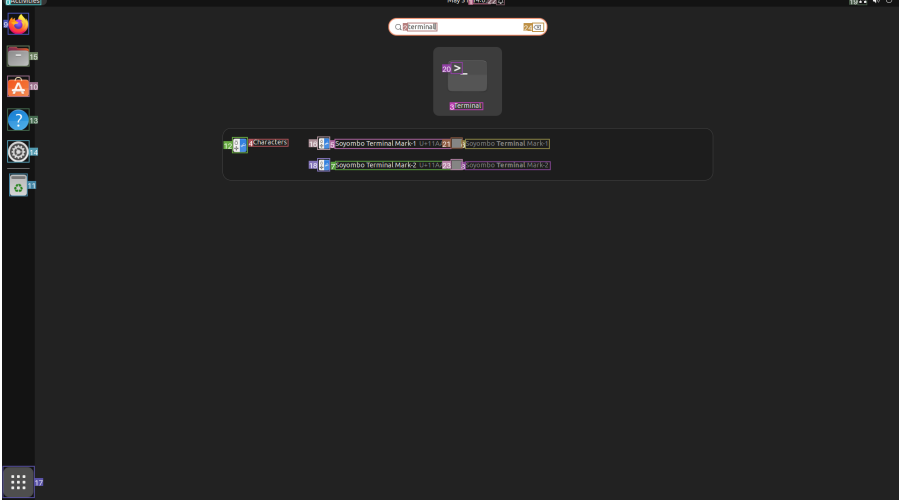
Continued on next page

Table 9 continued from previous page

5-6	<pre>key_press("enter") write_text("cp /home/crab/assets/*.txt /home/crab/assets_copy/")</pre>
7	 <p>The screenshot shows a Linux desktop environment with a dark purple and red geometric wallpaper. A terminal window is open in the center, displaying the command <code>cp /home/crab/assets/*.txt /home/crab/assets_copy/</code> and its execution. The desktop has a dock on the left with various application icons, and a system tray on the bottom right.</p>



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

Table 10: **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>The screenshot shows the Ubuntu desktop environment with a purple and red geometric background. A search application window is open, displaying a search bar and several application icons. The search bar contains the text "terminal".</p> <p><code>search_application("terminal")</code></p>
1	 <p>The screenshot shows the Ubuntu desktop environment with a dark background. A terminal application window is open, displaying a command prompt. The terminal window is titled "Terminal" and shows a cursor at the prompt.</p> <p><code>click(20)</code></p>

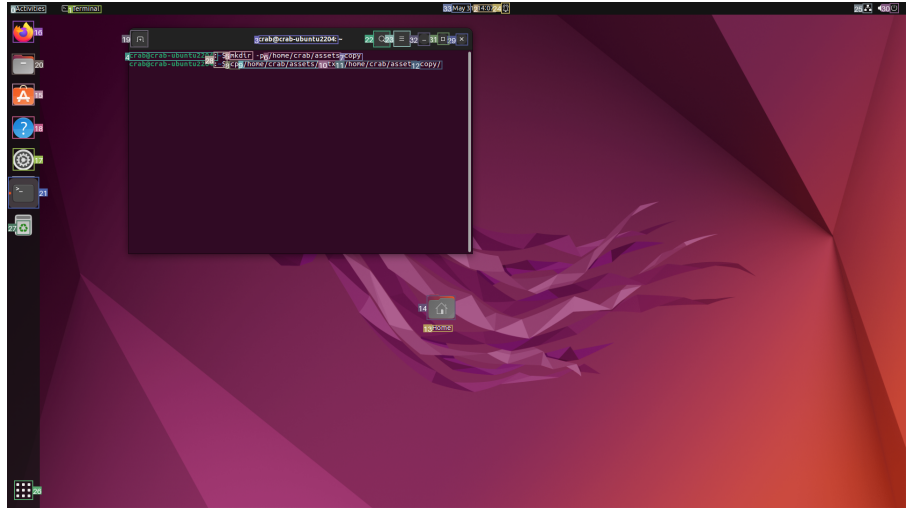
Continued on next page

Table 10 continued from previous page



Continued on next page

Table 10 continued from previous page

5	
<code>key_press("enter")</code>	

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