arXiv:2611.34232v1 [cs.CV] 25 Nov 2026

OS-Omni: Benchmarking Generalist Computer-Using Agents on Long-Horizon Workflows Across Platforms

OS-Omni Team

Website
os-copilot.github.io

Code
github.com/OS-Copilot/OS-Omni

Abstract

Computer-using agents are increasingly expected to complete realistic user goals rather than isolated GUI actions. However, existing benchmarks typically isolate web, mobile, or desktop environments, leaving it unclear whether a single generalist agent can execute long-horizon workflows that span multiple operating systems and persistent application states. We introduce OS-OMNI, a live cross-platform benchmark and execution environment for evaluating computer-using agents across Windows, Ubuntu, macOS, iOS, Android, and Web environments. OS-OMNI is released in versioned waves of 300 tasks each; this paper focuses on the first release, v1, which contains 48 tasks per platform and 12 cross-platform workflows that pair a desktop with an Android phone. OS-OMNI provides reproducible initialization, platform-specific adapters, trajectory logging, and executable final-state evaluators, and includes controlled service-backed applications for workflows involving account, cart, budget, calendar, message, and other persistent state. Human studies show that OS-OMNI tasks require 67.54 interaction steps, substantially exceeding OSWorld references of 15.0 steps. On v1, the strongest evaluated agent, Claude Opus 5.5, reaches 68.5% average success, trailing the 77.3% human reference by 8.8 points; the gap is uneven across platforms, with frontier models exceeding humans on Linux but trailing them by more than 25 points on iOS and macOS. These results indicate that current agents remain far from platform-robust, long-horizon computer use, especially when success requires sustained state tracking and durable final-state changes.

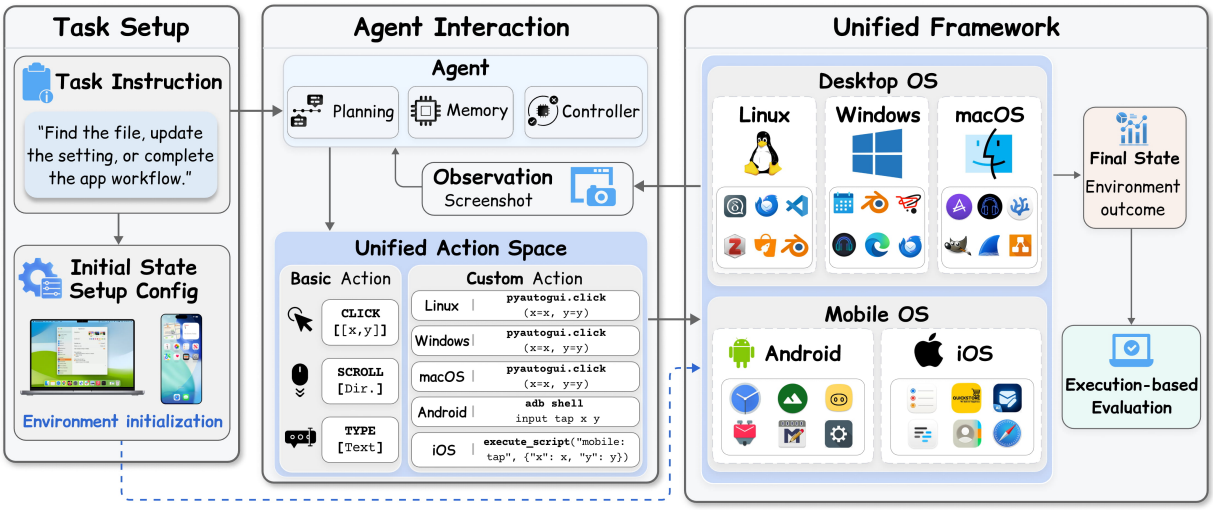

Figure 1: Overview of OS-Omni.

Author list

Hui Shen[*]
UMich, Touchstone AI

Yunta Hsieh[*]
UPenn, Touchstone AI

Jianing Ma[*]
Touchstone AI

Zhuoran Yin[*]
OpenAI

Zhuhan Bao[*]
Touchstone AI

Ziyuan Liu[*]
Touchstone AI

Qi Han[*]
Touchstone AI

Xiuqi Xu[*]
Touchstone AI

Xianbang Wang
MIT

Wenhao Chai
Princeton University

Sihan Xu
University of Michigan

Chenyang Zhao
SGLang Omni

Yanheng Shang
Touchstone AI

Dongfu Jiang
University of Waterloo

Rui Sun
MIT

Xingzhi Li
Touchstone AI

Chenqi Yin
Boston University

Zesen Zhao
University of Michigan

Boyuan Zheng
University of Michigan

Jisen Li
Together AI

Xiaoxia Wu
Together AI

Ruoyan Zhang
Touchstone AI

Ruiyao Liu
University of Pennsylvania

Jing Xiong
University of Hong Kong

He Xiao
University of Hong Kong

Nan Zheng
Touchstone AI

Hangxiang Chen
Touchstone AI

Shengyang Tao
New York University

Deyang Luorong
UW–Madison

Yifeng Liu
MIT

Jingxuan Zhang
The Ohio State University

Chaofan Tao
University of Hong Kong

Yuren Wang
Touchstone AI

Jiaqi Mo
UW–Madison

Rui Hu
Touchstone AI

Xing Wei
Touchstone AI

Xiaofeng Liu
Touchstone AI

Xin Xu
Touchstone AI

Han Wu
Touchstone AI

Mengyao Wen
Touchstone AI

Guobin Zhang
Touchstone AI

Shatong Zhu
Stanford University

Hei Ting (Una) Chan
University of Michigan

Sicheng Chen
The Ohio State University

Kai Zhang
The Ohio State University

Mengkang Hu
University of Hong Kong

Zhongwei Wan
The Ohio State University

Xin Wang
The Ohio State University

Chuanyang Zheng
Morgan Stanley

Ziheng Zhang
UT Austin

Hengyuan Zhang
University of Hong Kong

Zunhai Su
University of Hong Kong

Kanzhi Cheng

Dehan Kong
Georgia Institute of Technology

Yian Wang
NUS

Hanrong Zhang
University of Illinois Chicago

Jiawei Lu
UW–Madison

Yifan Zhang
University of Michigan

Junnan Li
SGLang Omni

Changyi Yang
Carnegie Mellon University

Jintao Qu
SGLang Omni

Wenyao Gao
SGLang Omni

Jiaxuan Luo
SGLang Omni

Jiaxin Deng
SGLang Omni

Yuhao Chen
SGLang Omni

Xuesong Ye
SGLang Omni

Jingwen Gu
SGLang Omni

Yichi Zhang
SGLang Omni

Qun Lou
SGLang Omni

Shen Yan
Bytedance Seed

Ben Athiwaratkun
Together AI

Qiushi Sun✉
University of Hong Kong

Mi Zhang
The Ohio State University

Ping Luo
University of Hong Kong

Wenhu Chen
University of Waterloo

Ngai Wong
University of Hong Kong

[*] Equal contribution. ✉ Corresponding author.
Zhuoran Yin contributed to this work independently and in his personal capacity, not as part of his role at OpenAI
Contact: Hui Shen (huishen@umich.edu), Qiushi Sun (qiushisun@connect.hku.hk).

1 Introduction

A useful computer-using agent [1–4] should be able to carry an ordinary user goal through the software environments where that goal actually lives. Consider a user preparing a small event: the agent may need to compare supplies in a controlled shopping service, add selected items to a cart, update a budget spreadsheet on a desktop, create follow-up calendar entries on a phone, export a short document, and send a message with the right details attached. None of these steps is remarkable in isolation, yet the full request crosses platform conventions, spans many intermediate decisions, and can only be judged by checking the resulting cart, records, calendar entries, exported file, and message state. Completing such a request requires far more than isolated widget grounding: the agent must maintain cross-application context and leave every underlying software state consistent [5].

This work asks: *can generalist computer-using agents reliably complete realistic, long-horizon, stateful workflows across heterogeneous software platforms?* Beyond sheer horizon length, agents face compounded challenges: tracking partially updated records across window switches, adapting to OS-specific UI conventions (e.g., desktop dialogs vs. mobile picker wheels), and committing changes that survive strict post-execution inspection. A benchmark for this capability therefore needs to evaluate both breadth across platforms and depth within workflows.

Existing benchmarks have made substantial progress by studying important slices of this problem. Browser-focused benchmarks evaluate web navigation, information seeking, and visual web interaction through browser-centric action and state spaces [6–11]. Enterprise-web benchmarks emphasize structured work over service-backed interfaces [12], while mobile benchmarks study dynamic interaction with phone applications [13–16]. Desktop benchmarks move evaluation into executable operating-system environments, where agents must operate real applications rather than predict isolated actions [17, 18]. These efforts establish that execution with verification [19] matters, but they still expose agents to one dominant software surface at a time, often with different platform scopes, action interfaces, task styles, and scoring protocols. As a result, strong performance on one family of environments can be hard to interpret as evidence of general computer use, rather than adaptation to a particular platform.

OSWorld was a major step toward general computer-using evaluation because it established execution-based evaluation in real desktop environments [17]. However, desktop-only evaluation leaves open whether agents can handle realistic workflows that span heterogeneous software surfaces, longer task horizons, and persistent application state. OS-OMNI extends this paradigm along two axes: **cross-platform evaluation** and **long-horizon workflows**, coupled with an execution-based evaluation protocol based on **final-state correctness**. Together, these features shift the benchmark from testing whether an agent can operate a particular interface toward testing whether it can complete a user goal across platforms. Human users require an average of **67.54** interaction steps and **200.97** seconds to complete OS-OMNI tasks, compared with OSWorld reference tasks requiring 15.0 steps and 111.94 seconds, indicating that OS-OMNI evaluates substantially longer and more involved workflows rather than only broader platform coverage.

These axes are coupled because failures often emerge at their intersections. Changing the platform changes the interaction bottleneck: a calendar update that is straightforward in a desktop dialog

may require picker wheels, modal sheets, and soft-keyboard input on a phone. Extending the horizon changes the burden of memory and coordination: adding one contact is easier than reconciling a directory with stale duplicates, and editing one field is easier than creating, deleting, reordering, and verifying a target set of records. Final-state evaluation also changes what counts as success: a document may look visually complete while its exported file or internal formatting remains wrong, and a shopping workflow may appear complete while the cart or budget record is inconsistent. Consequently, OS-OMNI scores tasks strictly by inspecting the persisted system state rather than intermediate visual progress.

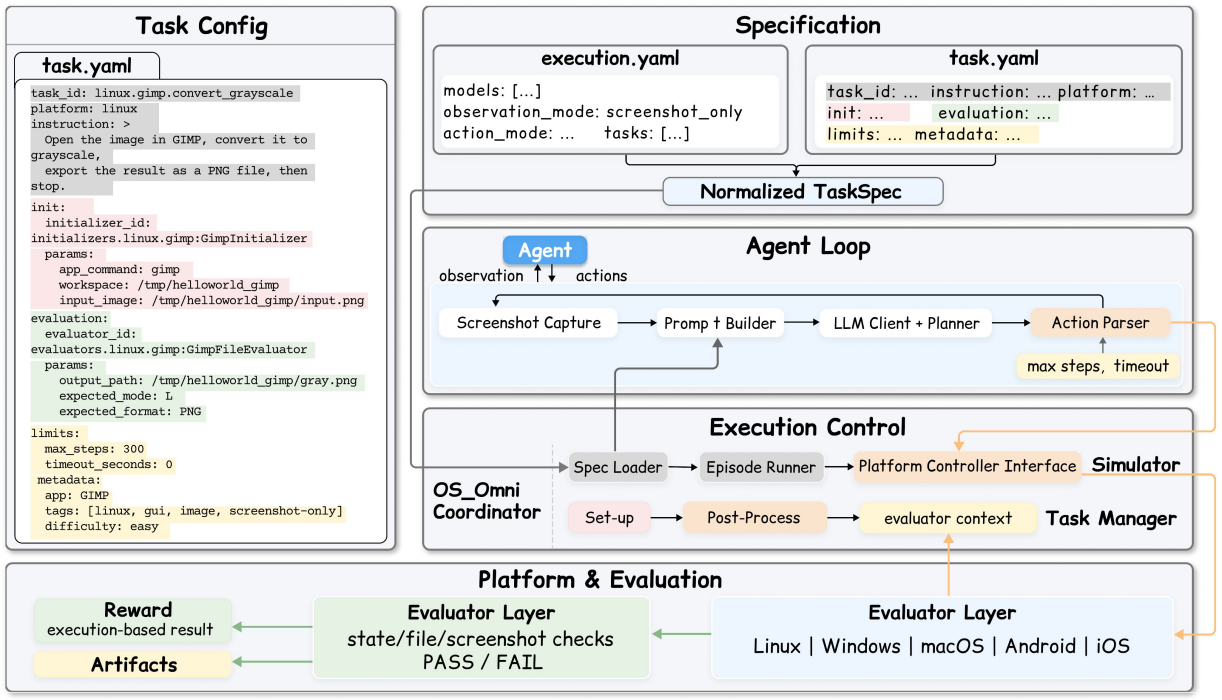

Figure 2: Overview of the OS-OMNI environment infrastructure. Execution specifications select tasks, models, observation modes, and action modes; platform runners instantiate the corresponding runtime, run the agent loop, and invoke task-specific evaluators.

To study this setting, we introduce OS-OMNI, a live cross-platform benchmark and execution environment for evaluating computer-using agents across Windows, Ubuntu, macOS, iOS, Android, and web environments. Rather than a fixed task pool, OS-OMNI is released in versioned waves of 300 tasks (v1, v2, ...), and each new version includes more mock applications and harder tasks (Section 3.3). This paper focuses on the first release, v1. Each task specifies a natural-language instruction, reproducible initialization, platform-specific adapters, trajectory logging, and an executable final-state evaluator. On v1, we evaluate nine models: GPT-6-Astra, Claude Opus 5.5, Gemini 4 Argon, GLM-5.3-753B, Kimi K3, Qwen 3.8 Max, Qwen 3.8 27B, Qwen 3.8 9B, and DeepSeek V4.1 Flash. The strongest model, Claude Opus 5.5, reaches 68.5% average success, 8.8 points below the human reference of 77.3%. Performance is highly uneven across platforms: all three frontier models (GPT-6-Astra, Claude Opus 5.5, and Gemini 4) exceed the human success rate on Linux, yet the best iOS and macOS results (41.7% and 54.2%) remain 33.5 and 25.6 points

below humans, and the smaller Qwen 3.8 27B and 9B models solve almost no tasks. These results suggest that current agents have not yet learned platform-robust, long-horizon computer use: they can solve many visually explicit tasks, but often fail when success requires sustained state tracking, platform-specific control, and durable final-state changes.

In summary, this paper makes the following contributions:

- We introduce OS-OMNI, a live benchmark for cross-platform computer-using agents released in versioned 300-task waves that progressively add mock applications and harder tasks; its first release, v1, covers six platforms and cross-platform workflows.
- We design long-horizon workflows centered on state tracking, cross-application coordination, and final-state correctness.
- We build a unified evaluation framework with reproducible initialization, native action backends, trajectory logging, and executable final-state evaluators.
- We evaluate nine models on v1, showing that even the best model remains 8.8 points behind humans on average, with gaps of more than 25 points on iOS and macOS.

Contents

1 Introduction **3**

2 OS-OMNI Environment **8**
2.1 Task Definition 8
2.2 Environment Infrastructure 8
2.3 Observation Space 9
2.4 Action Space 9

3 OS-OMNI Benchmark **10**
3.1 Operating System and Software Environments 10
3.2 Custom Software for Controlled Application Coverage 10
3.3 Live Benchmark and Versioned Releases 11
3.4 Tasks 12
3.5 Task Annotation Pipeline 12
3.6 Data Statistics 13
3.7 Comparison with existing benchmarks 14

4 Experiments **14**
4.1 VLM and Computer-Using Agent Baselines 14
4.2 Main Results 14
4.3 Analysis 15

5 Related Work **16**

6 Conclusion **17**

A Details of Human Experiments **21**

B Details of OS-OMNI Environment **21**
B.1 Environment Infrastructure 21
B.2 Observation Space 22
B.3 Action Space 23

C Details of the OS-OMNI Benchmark **25**
C.1 Platform Selection 25
C.2 Software Selection 26
C.3 The Open-Series Applications 27
C.4 Task Example Sources 28
C.5 Task Examples Collection 28
C.6 Initial State Setup Details 28
C.7 Evaluation Configuration Details 29
C.8 More Task Examples 35
C.9 Manual Verification for Ambiguous Cases 35

D Details of Agent and Evaluation Protocols **35**
D.1 Model and Execution Configuration 35
D.2 Prompt Details 35
D.3 Action Parsing and Episode Execution 46
D.4 Set-of-Mark Implementation Details 46
D.5 Result Reporting Protocol 46

E Qualitative Analysis **47**
E.1 Success and Failure Cases 47
E.2 Common Errors by Screenshot-only Agents 47
E.3 Human-Agent Difficulty Discrepancies 47
E.4 Comparative Analysis across Models 48
E.5 Analysis of SoM Effectiveness 48
E.6 Endpoint-Level Error Cases 48

2 OS-Omni Environment

This section specifies the execution environment behind OS-OMNI, including task specification, environment initialization, agent interaction, and final-state evaluation.

2.1 Task Definition

Each task in OS-OMNI specifies a stable identifier, a target platform, a curated instruction, optional initialization, evaluation logic, and descriptive attributes. We adopt a final-state evaluation protocol, where task success depends solely on the terminal environment state: $R(\tau) = g(s_T)$, where $\tau = (o_0, a_0, \ldots, o_T)$ denotes the interaction trajectory, s_T is the final environment state, and $g(\cdot)$ is the task-specific evaluator.

At each step, the agent receives the task instruction and the current observation o_t, then produces an action a_t, which is validated and executed by the platform controller. All cross-platform tasks follow a single-screen setup: an Ubuntu or Windows desktop displays a companion Android device through `scrcpy`. The agent interacts with both devices through the desktop interface using mouse and keyboard actions, including transferring artifacts between devices when required. This setup exposes a unified interaction surface without requiring separate platform endpoints or host-mediated transfers. An episode terminates when the agent emits `done` or `fail`, reaches the step or time limit, or otherwise stops. The final score is determined by the resulting environment state rather than similarity to a reference trajectory.

2.2 Environment Infrastructure

OS-OMNI is a scalable and executable environment for task design, evaluation, and interactive agent learning. It is organized as a shared orchestration layer plus platform-specific adapters covering most mainstream operating systems (*e.g.*, Windows, macOS, Ubuntu, Android, iOS).

2.2.1 Environment Initialization Real-world scenarios often begin with an open document, a seeded mailbox, or a transfer object that must move between machines. OS-OMNI aims to simulate these intermediate states to replicate real-world computer-using scenarios. Initialization procedures can be generic or application-specific. Generic desktop setup includes clearing stale processes, preparing assets, writing required data, launching applications, and arranging windows. Application-specific procedures seed app databases, configure local mail accounts, prepare mobile app storage, install or copy support assets, and connect benchmark apps to local services. The same task format can therefore describe both simple single-app starts and long multi-app workflow states.

State reset is implemented differently on each platform. Linux can revert a clean VM snapshot, Windows can restore golden VM storage, macOS can clone from a baseline Tart VM, iOS can clone a simulator, and Android can reset app and device state through ADB-backed preparation. These reset mechanisms are part of the environment contract because GUI agents routinely modify files, accounts, settings, and application caches. OS-OMNI first restores a platform baseline, then materializes task-specific data, and finally arranges the applications that the agent will observe.

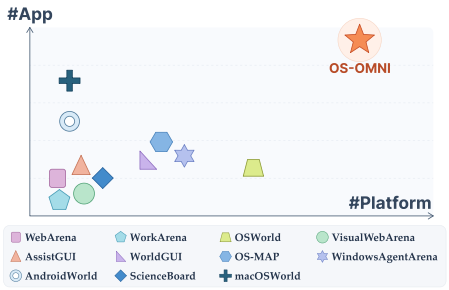

Benchmark	# Tasks	Platform						Cross-Platform	Cross-App	Curated Apps	# Apps / Sites
		Win.	Linux	macOS	iOS	Android	Web				
WorkArena	33	✗	✗	✗	✗	✗	✓	✗	✗	✗	1
WebArena	812	✗	✗	✗	✗	✗	✓	✗	✗	✗	6
VisualWebArena	910	✗	✗	✗	✗	✗	✓	✗	✗	✗	3
AndroidWorld	116	✗	✗	✗	✗	✓	✗	✗	✓	✗	20
WindowsAgentArena	154	✓	✗	✗	✗	✗	✓	✗	✓	✗	11
AssistGUI	100	✓	✗	✗	✗	✗	✗	✗	✗	✗	9
WorldGUI	611	✓	✗	✗	✗	✗	✓	✗	✗	✗	10
macOSWorld	202	✗	✗	✓	✗	✗	✗	✗	✓	✗	30
OSWorld	369	✓	✓	✓	✗	✗	✓	✗	✓	✗	9
ScienceBoard	169	✗	✓	✗	✗	✗	✗	✗	✓	✗	6
MMBench-GUI	967	✓	✓	✓	✗	✓	✓	✗	✓	✗	–
OS-MAP	416	✗	✓	✗	✗	✗	✓	✗	✓	✗	15
OS-OMNI	**300**	✓	✓	✓	✓	✓	✓	✓	✓	✓	**18**

Figure 3: #Platform indicates the number of supported platforms among Windows, Linux, macOS, iOS, Android, and Web, and #App indicates the number of covered applications or websites. Compared with existing benchmarks, OS-OMNI provides broader platform coverage, richer application diversity, cross-platform workflows, and curated applications.

2.2.2 Task Evaluation Every benchmark task declares an evaluator and its parameters. After the agent loop ends, the platform runner invokes this evaluator against the final environment state and records a structured result. This result includes success status and may include diagnostic information such as failure messages, extracted values, saved outputs, and step counts. As shown in Figure 2, OS-OMNI uses execution-based evaluation, which includes evaluators that parse office documents, inspect application databases, check external storage, and validate domain-specific application state. Explicit evaluators also make benchmark reliability auditable. OS-OMNI records trajectories and evaluation records, and also includes oracle variants of tasks for debugging.

2.3 Observation Space

OS-OMNI supports three observation modes through the execution specification: screenshot-only, accessibility tree-only, and combined observation. We adopt screenshot-only observation as the main experimental setting because screenshots are available across all supported platforms and provide the most uniform interface for generalist computer-using agents. This choice is also aligned with the recent trend of vision-centric CUA agents, where agents increasingly operate from screenshots to improve generality across heterogeneous applications and operating systems [20–23]. Accessibility trees, while useful when available, are platform- and application-dependent: their metadata can be incomplete or unavailable. Therefore, all experiments reported in this paper use screenshot-only observation unless otherwise stated. At the same time, OS-OMNI also supports accessibility tree-only and combined observation modes, enabling future studies to evaluate agents that rely on structured UI metadata or combine visual and accessibility-based information.

2.4 Action Space

In OS-OMNI, we introduce a two-level action space that combines a basic, platform-agnostic action interface with customized platform execution. The basic layer covers common human-computer operations, including pointer actions, keyboard shortcuts, text entry, scrolling, drag-

ging, touch-style gestures, waiting, and terminal decisions `DONE` and `FAIL`. This layer follows the widely used PyAutoGUI-style semantics for mouse and keyboard control, making the agent-facing syntax compact and familiar. The customized layer maps each high-level action intent to the most reliable native backend available on each endpoint. For example, a `click` action may be executed as a PyAutoGUI desktop event on Linux, a Windows or macOS runtime call, an ADB tap on Android, or an Appium/XCUITest action on iOS. This separation between action syntax and execution backend allows OS-OMNI to evaluate agents under a shared interface without forcing heterogeneous operating systems to expose the same automation API. As a result, the action space is broad enough to cover realistic GUI workflows across desktop and mobile platforms, while remaining structured enough for cross-platform comparison.

3 OS-Omni Benchmark

We introduce the OS-OMNI benchmark, a live benchmark of real-world computer use tasks executed across Linux, Windows, macOS, Android, iOS, Web, and cross-platform workflows (Section 3.3). * Each task includes a natural-language goal, a reproducible initialization procedure, and a final-state evaluator.

3.1 Operating System and Software Environments

OS-OMNI supports all mainstream operating systems, including ***desktop operating systems (Windows, macOS, Ubuntu)***, ***mobile operating systems (iOS, Android)***, **web (Chrome)**, and ***workflows that cross platforms*** (Figure 6). All platforms share the same specification format, while each platform retains the runtime mechanisms required by its own software stack. As shown in Figure 5, OS-OMNI covers the most common categories of applications used in everyday digital life.

3.2 Custom Software for Controlled Application Coverage

Real-world CUA workflows frequently involve commercial services or specialized applications that are difficult to reproduce in benchmark environments. We address these coverage gaps by developing a suite of 18 custom applications, collectively referred to as the `Open` series. Notably, these applications support **180 of the 300 tasks (60.0%)** in OS-OMNI v1, making controlled software development a central component of our benchmark design. We organize these applications into two categories; Appendix C.3 describes their design and deployment and lists the tasks that use each of them.

Controlled Cross-Platform Applications. The first category targets common consumer and content-oriented workflows whose real-world counterparts often depend on user accounts, payments, or unstable third-party services. We develop nine applications: `OpenNote` (social note sharing, including posting, liking, bookmarking, and commenting), `OpenShop` (e-commerce,

*Due to copyright issues, Windows tasks require activation by the user, while macOS and iOS tasks require Apple hardware.

including shopping carts, checkout, and reviews), `OpenTwitter` (social networking), `OpenBank` (online banking, transfers, and transaction disputes), `OpenDriver` (ride-hailing with passenger and driver interfaces), `OpenTrip` (flight and hotel booking), `OpenFood` (food delivery, coupons, and customer support), `OpenNavigation` (route planning, waypoints, and incident reporting), and `OpenPhotoEdit` (image editing and export). These applications provide realistic GUIs across multiple platforms while keeping task data and backend state under benchmark control, enabling reproducible evaluation without relying on live commercial services or real user accounts.

Web-Based Applications for Software Coverage Gaps. The second category targets specialized software domains that are unevenly supported across operating systems. We develop nine browser-based applications hosted in the Linux environment: `OpenDoc` (document editing), `OpenSheet` (spreadsheets), `OpenSlide` (presentations), `OpenMail` (email and contact management), `OpenFinance` (personal finance, budgeting, and portfolio management), `OpenCode` (code editing, debugging, and version control), `OpenBlender` (3D scene editing), `OpenWireshark` (network packet analysis), and `OpenZotero` (reference management). By reproducing these workflows in controlled browser environments, we extend task coverage beyond the software available in the underlying operating systems while maintaining a consistent evaluation interface.

Together, these applications enable OS-OMNI to evaluate a broader range of realistic CUA workflows under reproducible conditions, reducing the dependence of benchmark coverage on external services and platform-specific software availability.

3.3 Live Benchmark and Versioned Releases

OS-OMNI is a live benchmark rather than a fixed task pool. As agents improve, a static task set gradually saturates and its tasks may leak into training data, so a single release cannot measure frontier agents for long. OS-OMNI therefore releases tasks in versioned waves of 300 tasks each, starting with v1; later waves are published as v2, v3, and so on. Each release is a fixed task set, and results are reported together with the release version so that comparisons are always made on the same tasks.

Each new version is designed to be more demanding than the previous one in two ways. First, it includes more mock applications, i.e., the controlled applications described in Section 3.2, so that more service-backed workflows, such as banking, shopping, travel booking, and social media, can be evaluated without live commercial services or real user accounts. The first release, v1, already draws 180 of its 300 tasks from these Open-series applications. Second, each version makes the tasks harder, for example through longer horizons, more applications or devices per workflow, and more state that must be tracked and left correct at the end. Task specification, initialization, and final-state evaluation follow the same protocol across versions, so the benchmark can keep pace with agent progress without changing how tasks are run or scored. This paper introduces and evaluates the first release, v1.

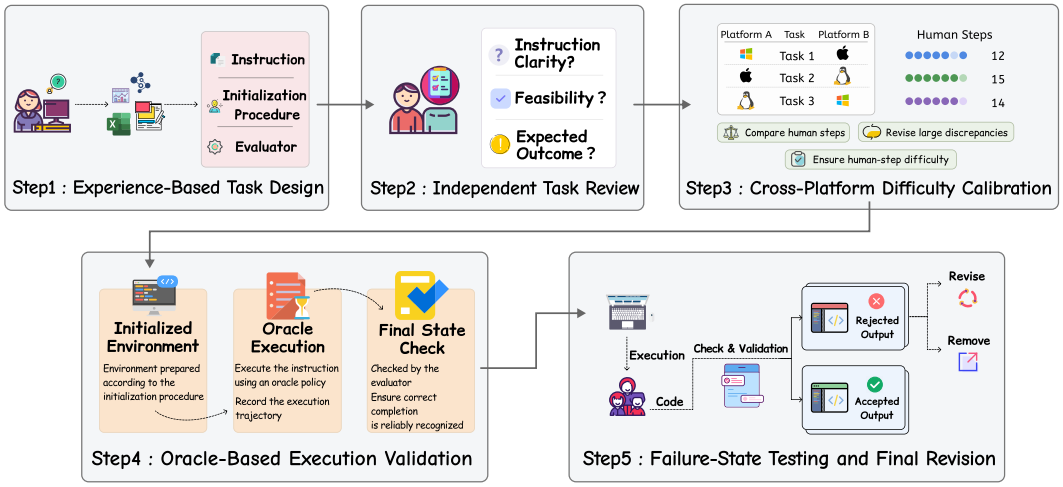

Figure 4: The five-stage task annotation pipeline for OS-OMNI.

3.4 Tasks

The first release, OS-OMNI v1, consists of 300 curated CUA tasks: 48 tasks on each of six platforms, ***macOS***, ***Windows***, ***Ubuntu***, ***iOS***, ***Android***, and ***Web***, plus 12 ***cross-platform*** workflows. Web tasks run in Chrome or Chromium on the Ubuntu VM. Each cross-platform workflow runs on a single screen: the Ubuntu or Windows desktop shows a companion Android phone in a scrcpy window, and the agent works on both devices, moves a file between them, and returns results to the phone (six Ubuntu + Android and six Windows + Android workflows). Each task is represented by a specification that names the platform, instruction, initialization, and evaluator. Overall, the dataset construction was completed by 30 computer science graduate students over two months, with approximately 2,400 person-hours devoted to annotation and an additional 400 person-hours devoted to double-checking.

3.5 Task Annotation Pipeline

As shown in Figure 4, we use a five-stage annotation pipeline:

(1) **Experience-Based Task Design.** Each task is constructed by an annotator with hands-on experience using the target software. The annotator designs a natural-language instruction together with the initialization procedure and evaluator. This stage ensures that the task is grounded in realistic software usage rather than being limited to basic operations. (2) **Independent Task Review.** A second annotator examines each task for feasibility and ambiguity. This review checks whether the instruction is understandable, whether the task can be completed under the initialized environment, and whether the expected outcome is sufficiently well-defined for reliable evaluation. (3) **Cross-Platform Difficulty Calibration.** Annotators then calibrate task difficulty across platforms by comparing the expected number of human operation steps required to complete each task. Tasks are revised when necessary to reduce large difficulty discrepancies, helping ensure that tasks from different platforms have comparable difficulty. (4) **Oracle-Based Execution Validation.** We use oracle trajectories to validate both the initializer and the evaluator. Specifically, the oracle execution is run from the initialized environment to verify that the task

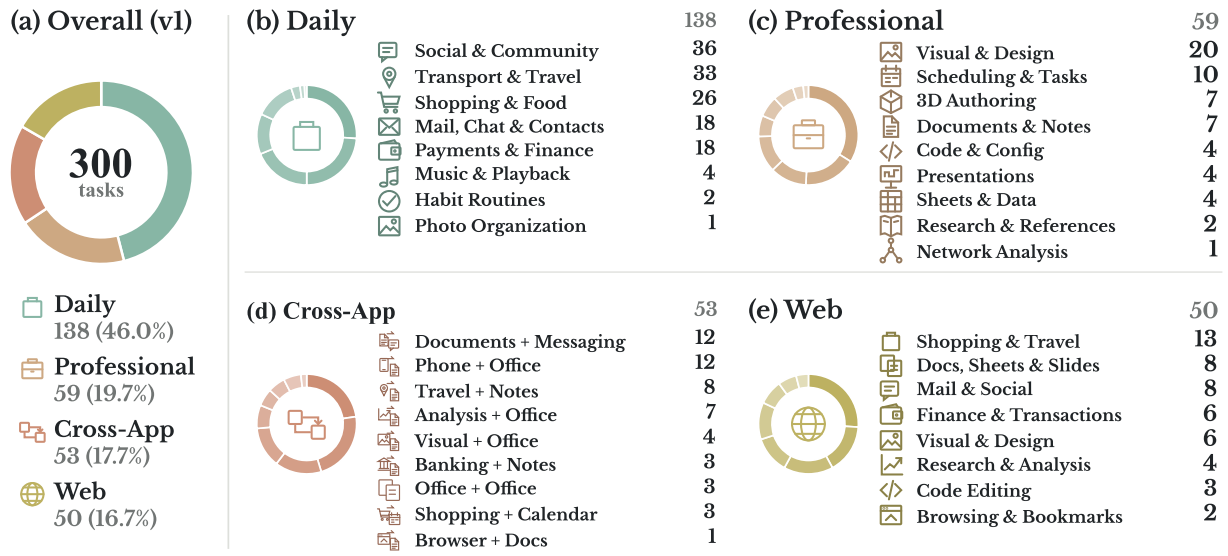

Figure 5: Task composition of OS-OMNI v1. (a) Overall distribution across task categories. (b–e) Secondary-category distributions within Daily, Professional, Cross-App, and Web tasks, respectively. Numbers denote task counts.

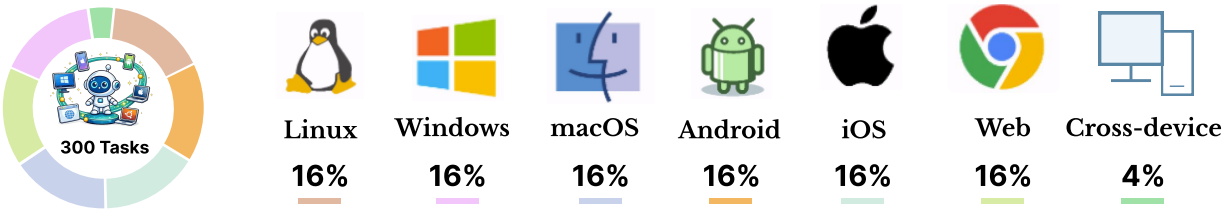

Figure 6: Distribution of the 300 tasks across evaluation tracks. Windows, macOS, Ubuntu, iOS, Android, and Web each contain 48 tasks (16%), while the cross-device track contains 12 workflows (4%).

is solvable under the provided state. The resulting final state is then checked by the evaluator to ensure that a correct completion is reliably recognized. (5) **Failure-State Testing and Final Revision.** To further test evaluator robustness, we examine representative failed cases and verify that the evaluator rejects incorrect solutions. Tasks whose initializer, instruction, difficulty level, or evaluator fails any stage of the checking process are revised through additional inspection or removed from the benchmark.

3.6 Data Statistics

The main statistics of OS-OMNI are summarized in Figures 5 and 6. As shown in Figure 6, OS-OMNI v1 contains 300 tasks: 48 tasks on each of the six platforms (16% each) and 12 cross-device workflows (4%). Figure 5 presents the task taxonomy, which comprises four primary categories: daily-life tasks (138; 46.0%), professional productivity tasks (59; 19.7%), cross-application workflows (53; 17.7%), and web tasks (50; 16.7%), further organized into 34 secondary categories.

Each task is assigned to exactly one primary category, with cross-application workflows taking precedence. The Web category includes the 48 Web-track tasks and two Safari browsing tasks from the iOS track.

3.7 Comparison with existing benchmarks

Figure 3 compares OS-OMNI with representative existing benchmarks in terms of platform coverage, task scale, application diversity, cross-platform workflows, cross-application workflows, and curated applications. WorkArena [12], WebArena [6], and VisualWebArena [7] evaluate agents in web environments. AndroidWorld [13] focuses on Android. OSWorld [17], WindowsAgentArena [18], AssistGUI [24], WorldGUI [25], macOSWorld [26], ScienceBoard [27], MMBench-GUI [28], and OS-MAP [29] extend evaluation to desktop or multi-platform GUI settings, but they do not provide the same level of coverage across desktop, mobile, web, and cross-platform workflows.

In contrast, OS-OMNI unifies Windows, Linux, macOS, iOS, Android, and Web environments under a single benchmark setting. It further introduces benchmark-controlled applications and services to support workflows involving persistent user state, backend records, shopping carts, orders, and payment-like options. Compared with benchmarks that rely on public websites, external accounts, or preloaded local profiles, this design enables reproducible initialization and reliable final-state evaluation without depending on live commercial platforms or real user accounts. As a result, OS-OMNI evaluates agents on realistic service-oriented workflows that span multiple applications and platforms, rather than isolated GUI operations alone.

4 Experiments

4.1 VLM and Computer-Using Agent Baselines

We conduct human evaluations with annotators who possess basic software usage skills but have not been exposed to the tasks before. For each human trajectory, we record the number of interaction steps and the wall-clock completion time, and aggregate the measurements by platform. Then, we evaluate nine models on OS-OMNI v1 (Section 3.3): GPT-6-Astra, Claude Opus 5.5, Gemini 4 Argon, GLM-5.3-753B, Kimi K3, Qwen 3.8 Max, Qwen 3.8 27B, Qwen 3.8 9B, and DeepSeek V4.1 Flash. By default, agents receive the current screenshot and the three most recent interaction steps and screenshots. Agents complete each task within the limits declared in its specification: a maximum of 300 steps and a per-step time limit of 180 seconds for single-platform tasks, and 300–650 steps with a per-step time limit of 120 seconds for cross-platform workflows. We present a summary of the results in Table 1.

4.2 Main Results

Even the strongest agents remain below humans, and most models are far behind. Claude Opus 5.5 obtains the best Macro-Avg. success rate in Table 1, 68.5%, 8.8 points below the 77.3% human average, followed by GPT-6-Astra (66.1%) and Gemini 4 Argon (58.6%). Under Micro-

Table 1: **Main results on OS-OMNI v1.** Each cell reports task success rate (%) on one platform or on the cross-platform workflows. **Macro-Avg.** is the mean over all seven columns (the six platforms and Cross-Platform). **Micro-Avg.** weights every task equally: it is the share of all 300 v1 tasks solved (288 single-platform tasks and 12 cross-platform workflows); for Human Performance it is the task-weighted mean of the column rates. Among models, the best result in each column is in **bold** and the second best is underlined.

Model / Agent	Success Rate (↑)								
	Windows	Linux	macOS	iOS	Android	Web	Cross-Platform	Macro-Avg.	Micro-Avg.
Human Performance	78.4	74.6	79.8	75.2	72.6	76.9	83.3	77.3	76.5
GPT-6-Astra	**66.7**	89.6	50.0	33.3	77.1	**87.5**	58.3	66.1	**67.0**
Claude Opus 5.5	62.5	85.4	**54.2**	**41.7**	**81.3**	70.8	**83.3**	**68.5**	66.7
Gemini 4 Argon	58.3	**95.8**	39.6	18.8	66.7	81.3	50.0	58.6	59.7
DeepSeek V4.1 Flash	22.9	62.5	16.7	12.5	41.7	22.9	33.3	30.4	30.0
GLM-5.3-753B	6.3	68.8	8.3	10.4	8.3	45.8	16.7	23.5	24.3
Kimi K3 2.8T	20.8	50.0	4.2	12.5	16.7	25.0	25.0	22.0	21.7
Qwen 3.8 Max	6.3	6.3	4.2	12.5	31.3	6.3	25.0	13.1	11.7
Qwen 3.8 27B	0.0	0.0	2.1	6.3	0.0	0.0	0.0	1.2	1.3
Qwen 3.8 9B	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Avg., which weights every task equally, GPT-6-Astra and Claude Opus 5.5 are nearly tied (67.0% and 66.7%, i.e., 201 and 200 solved tasks): the Macro-Avg. gives the 12 cross-platform workflows, where Claude Opus 5.5 solves three more tasks, the same weight as each 48-task platform. The remaining models fall far behind: DeepSeek V4.1 Flash reaches 30.4%, Qwen 3.8 Max 13.1%, and Qwen 3.8 27B and 9B solve almost no tasks (1.2% and 0.0%). These results indicate that current agents can solve a portion of visually explicit and short-horizon tasks, but still lack the robustness needed for long-horizon planning, precise GUI grounding, state tracking, and final-state persistence in real operating systems.

Humans are consistent across platforms, while agents are not. The human reference row is consistent across all platforms (72.6–83.3%), whereas agent performance varies sharply by operating system. Gemini 4 Argon reaches 95.8% on Linux but only 18.8% on iOS, and even Claude Opus 5.5 ranges from 85.4% on Linux to 41.7% on iOS. On Linux, all three frontier models exceed the human success rate of 74.6%, while on iOS and macOS the best model remains 33.5 and 25.6 points below humans. Similar drops appear for the other models, suggesting that the main challenge is not only performing well on a specific platform, but also succeeding consistently on every platform.

4.3 Analysis

Cross-application workflows are harder than single-application commands. Cross-application tasks add a second layer of difficulty: the agent must preserve information across window switches, identify which application owns each subgoal, and ensure that all destination states are updated before stopping. The failures are also qualitatively different: agents often complete a visually plausible intermediate action, but fail the final evaluator because a calendar event or follow-up message is missing or inconsistent. This makes cross-application evaluation closer to

real-world computer use, where user intent is rarely confined to a single GUI command.

Cross-platform workflows separate the strongest model from the rest. Cross-platform tasks require correct use of platform-specific interaction conventions and reliable transfer of artifacts between devices. In single-platform tasks, the agent can remain within one action backend and one set of visual conventions. The v1 cross-platform workflows instead require the agent to edit or build a file on the Android phone, move it to the desktop itself, build the deliverable in LibreOffice, and send results computed on the desktop back to the phone. On these workflows, Claude Opus 5.5 matches the human success rate of 83.3%, while GPT-6-Astra and Gemini 4 reach 58.3% and 50.0%, below their averages over the six single platforms, and Qwen 3.8 27B and 9B solve none. Rankings also shift: Qwen 3.8 Max and DeepSeek V4.1 Flash score higher on this split (25.0% and 33.3%) than on most single platforms. Because the split contains only 12 workflows, each task changes the success rate by 8.3 points, so these differences should be read with caution. Together with the platform sensitivity observed in Table **1**, this suggests that cross-platform CUA is not simply a larger task set, but a distinct generalization problem involving heterogeneous UI conventions and input primitives.

Error analysis. We further inspect trajectories to categorize common failure modes. The most frequent errors fall into four groups. First, *visual grounding errors* occur when the agent clicks a nearby control or misreads dense document and diagram interfaces. Second, *state-tracking errors* appear when the agent loses copied information across applications or forgets a required destination file. Third, *persistence errors* occur when the visible interface looks correct but the durable final state is not saved or committed in the format expected by the evaluator. Fourth, *platform and timing errors* arise from modal dialogs, delayed application loading or missing confirmation steps. These categories explain why screenshot-only agents can often produce plausible trajectories while still failing the task.

5 Related Work

Benchmarks for GUI Agents. GUI-agent benchmarks have progressed from static web navigation to executable environments. Mind2Web, WebArena, VisualWebArena, WebVoyager, and WorkArena evaluate browser or enterprise web workflows [6–9, 12]; AndroidWorld focuses on mobile-device interaction [13]; and OSWorld, WindowsAgentArena, WorldGUI, macOSWorld, OS-MAP, and MMBench-GUI study desktop GUI automation [17, 18, 25, 26, 28, 29]. These benchmarks establish the value of executable evaluation, but most still emphasize one platform family, one dominant interface assumption, or benchmark-specific runtime protocols.

Vision-language models for GUI Agents. Recent CUA agents are specialized for screen understanding and long-horizon interaction. AppAgent and MobileAgent study smartphone control, while SeeClick, ShowUI, and OmniParser improve visual grounding and screen parsing for GUI operation [20–23, 30]. Agent S explores open-ended desktop operation with explicit planning and recovery [31]. More recently, OpenCUA and ScaleCUA develop open computer-using agents and training data for cross-platform control [3, 4]. These systems show rapid progress, but their behavior remains sensitive to platform conventions and action interfaces, motivating unified

cross-platform evaluation in OS-OMNI.

6 Conclusion

We built OS-OMNI to study whether a single computer-using agent can reliably complete long-horizon workflows across desktop, mobile, and web environments. The first release contains 300 tasks, including 12 cross-device workflows, with success determined by the resulting environment state. Future releases will expand application coverage and task complexity in waves of 300 tasks.

Our results reveal two limitations of current computer-using agents. First, strong performance on one platform does not imply generality: all three frontier models outperform humans on Linux, yet the best results on iOS and macOS remain 33.5 and 25.6 percentage points below human performance, respectively (Section 4.2). Second, progress through an interface does not necessarily translate into task completion. In OpenBank, an agent submits the same transfer twice; in OpenTrip, an agent creates five confirmed bookings but duplicates one and omits another (Appendix E).

These findings suggest that generalist computer use requires more than proficiency with individual interfaces. Agents must maintain task state over extended interactions, adapt to different software environments, and verify the effects of their actions. We believe benchmarks for generalist computer use should measure both the breadth of environments an agent can operate in and the correctness of the state it leaves behind.

Acknowledgments

This paper is partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region (Project No. T41-517/25-N). We also acknowledge CompAI for their support of this project, conducted as part of a milestone.

References

[1] Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, and Yu Qiao. OS-ATLAS: Foundation action model for generalist GUI agents. In *The Thirteenth International Conference on Learning Representations*, 2025.

[2] Boyu Gou, Ruohan Wang, Boyuan Zheng, Yanan Xie, Cheng Chang, Yiheng Shu, Huan Sun, and Yu Su. Navigating the digital world as humans do: Universal visual grounding for GUI agents. In *The Thirteenth International Conference on Learning Representations*, 2025.

[3] Xinyuan Wang et al. OpenCUA: Open Foundations for Computer-Use Agents. arXiv preprint arXiv:2508.09123, 2025.

[4] Zhaoyang Liu et al. ScaleCUA: Scaling Open-Source Computer Use Agents with Cross-Platform Data. arXiv preprint arXiv:2509.15221, 2025.

[5] Qiushi Sun, Kanzhi Cheng, Yian Wang, Bowen Yang, Hang Yan, Liheng Chen, Fangzhi Xu, Zichen Ding, Nuo Chen, Jialin Cao, et al. Osreward: Instituting standardized evaluation for cross-platform computer-use reward models. *arXiv preprint arXiv:2607.28609*, 2026.

[6] Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*, 2023.

[7] Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov, and Daniel Fried. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 881–905, 2024.

[8] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114, 2023.

[9] Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. Webvoyager: Building an end-to-end web agent with large multimodal models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6864–6890, 2024.

[10] Qiushi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yian Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhoumianze Liu, et al. Os-genesis: Automating gui agent trajectory construction via reverse task synthesis. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5555–5579, 2025.

[11] Boyu Gou, Zanming Huang, Yuting Ning, Yu Gu, Michael Lin, Weijian Qi, Andrei Kopanev, Botao Yu, Bernal Jimenez Gutierrez, Yiheng Shu, Chan Hee Song, Jiaman Wu, Shijie Chen, Hanane Nour Moussa, TIANSHU ZHANG, Jian Xie, Yifei Li, Tianci Xue, Zeyi Liao, Kai

Zhang, Boyuan Zheng, Zhaowei Cai, Viktor Rozgic, Morteza Ziyadi, Huan Sun, and Yu Su. Mind2web 2: Evaluating agentic search with agent-as-a-judge. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2025.

[12] Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H Laradji, Manuel Del Verme, Tom Marty, Léo Boisvert, Megh Thakkar, Quentin Cappart, David Vazquez, et al. Workarena: How capable are web agents at solving common knowledge work tasks? *arXiv preprint arXiv:2403.07718*, 2024.

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

[14] Qiushi Sun, Mukai Li, Zhoumianze Liu, Zhihui Xie, Fangzhi Xu, Zhangyue Yin, Kanzhi Cheng, Zehao Li, Zichen Ding, Qi Liu, et al. Os-sentinel: Towards safety-enhanced mobile gui agents via hybrid validation in realistic workflows. *arXiv preprint arXiv:2510.24411*, 2025.

[15] Kanzhi Cheng, Zehao Li, Zheng Ma, Nuo Chen, Jialin Cao, Qiushi Sun, Zichen Ding, Fangzhi Xu, Hang Yan, Jiajun Chen, et al. Openmobile: Building open mobile agents with task and trajectory synthesis. *arXiv preprint arXiv:2604.15093*, 2026.

[16] Chenyang Yan, Yingying Zhang, Kanzhi Cheng, Qiushi Sun, Zhengyuan Pan, Zheng Ma, Hang Yan, Yian Wang, Nuo Chen, Jialin Cao, et al. Openmobile-2: Building versatile mobile agents with scalable environments and app-native tools. *arXiv preprint arXiv:2610.14157*, 2026.

[17] Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Advances in Neural Information Processing Systems*, 37:52040–52094, 2024.

[18] Rogerio Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Bucker, et al. Windows agent arena: Evaluating multi-modal os agents at scale. *arXiv preprint arXiv:2409.08264*, 2024.

[19] Qiushi Sun, Da Yu, Alexander Knop, Amer Sinha, Lingpeng Kong, Ben Kao, Ravi Kumar, and Pritish Kamath. Vericua: Verification as agentic exploration for computer-using agents. *arXiv preprint arXiv:2610.40578*, 2026.

[20] Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Li YanTao, Jianbing Zhang, and Zhiyong Wu. Seeclick: Harnessing gui grounding for advanced visual gui agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9313–9332, 2024.

[21] Junyang Wang, Haiyang Xu, Jiabo Ye, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. Mobile-agent: Autonomous multi-modal mobile device agent with visual perception. *arXiv preprint arXiv:2401.16158*, 2024.

[22] Yadong Lu, Jianwei Yang, Yelong Shen, and Ahmed Awadallah. Omniparser for pure vision based gui agent. *arXiv preprint arXiv:2408.00203*, 2024.

[23] Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Stan Weixian Lei, Lijuan Wang, and Mike Zheng Shou. Showui: One vision-language-action model for gui visual agent. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19498–19508, 2025.

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

[25] Henry Hengyuan Zhao, Kaiming Yang, Wendi Yu, Difei Gao, and Mike Zheng Shou. Worldgui: An interactive benchmark for desktop gui automation from any starting point. *arXiv preprint arXiv:2502.08047*, 2025.

[26] Pei Yang, Hai Ci, and Mike Zheng Shou. macosworld: A multilingual interactive benchmark for gui agents. *arXiv preprint arXiv:2506.04135*, 2025.

[27] Qiushi Sun, Zhoumianze Liu, Chang Ma, Zichen Ding, Fangzhi Xu, Zhangyue Yin, Haiteng Zhao, Zhenyu Wu, Kanzhi Cheng, Zhaoyang Liu, et al. Scienceboard: Evaluating multimodal autonomous agents in realistic scientific workflows. *arXiv preprint arXiv:2505.19897*, 2025.

[28] Xuehui Wang, Zhenyu Wu, JingJing Xie, Zichen Ding, Bowen Yang, Zehao Li, Zhaoyang Liu, Qingyun Li, Xuan Dong, Zhe Chen, et al. Mmbench-gui: Hierarchical multi-platform evaluation framework for gui agents. *arXiv preprint arXiv:2507.19478*, 2025.

[29] Xuetian Chen, Yinghao Chen, Xinfeng Yuan, Zhuo Peng, Lu Chen, Yuekeng Li, Zhoujia Zhang, Yingqian Huang, Leyan Huang, Jiaqing Liang, et al. Os-map: How far can computer-using agents go in breadth and depth? *arXiv preprint arXiv:2507.19132*, 2025.

[30] Chi Zhang, Zhao Yang, Jiaxuan Liu, Yanda Li, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. Appagent: Multimodal agents as smartphone users. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2025.

[31] Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, and Xin Eric Wang. Agent s: An open agentic framework that uses computers like a human. *arXiv preprint arXiv:2410.08164*, 2024.

A Details of Human Experiments

As shown in Figures 7a and 7b, human users require substantial effort across all OS-OMNI platforms. Averaged over Windows, Linux, macOS, iOS, and Android, a task takes 67.54 interaction steps and 200.97 seconds to complete. This is 1.80× the OSWorld reported median completion time of 111.94 seconds [17]. Even the fastest OS-OMNI platform, macOS, takes 165.41 seconds on average, or 1.48× the OSWorld reference; iOS, Linux, Windows, and Android cluster between 207.00 and 212.79 seconds, corresponding to 1.85–1.90× the OSWorld time.

The per-platform measurements also show that task burden is not determined by action count alone. Android has the fewest human steps (45.63) but still requires 207.00 seconds on average, while macOS has the most steps (89.43) but the shortest average time among OS-OMNI platforms. This gap indicates that human effort depends on platform-specific feedback latency, application startup and navigation cost, file or state inspection, and the need to verify persistent artifacts, rather than simply the number of GUI actions. OSWorld further reports 72.36% human accuracy, compared with 88% on WebArena [6], showing that executable desktop tasks are already harder than pure web tasks. OS-OMNI extends this difficulty to a broader cross-platform setting, where users must handle heterogeneous operating systems, application ecosystems, and state representations.

B Details of OS-Omni Environment

B.1 Environment Infrastructure

OS-OMNI supports both containerized and non-containerized deployment. The containerized path packages the host-side runner, dependencies, local services, and orchestration logic into a reproducible software environment, which is convenient for headless servers and shared evaluation machines. The non-containerized path runs the same components directly on the host system, which is useful for local development, debugging, and platform-specific runtime control. In both cases, deployment is separated from the environment seen by the agent.

OS-OMNI provides runtime support for desktop, mobile, and web endpoints through controlled

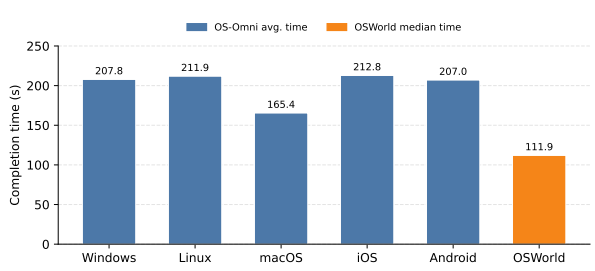

(a) Human completion time on OS-OMNI by platform. Blue bars report the average completion time for each OS-OMNI platform, and the orange bar shows the OSWorld reported median completion time.

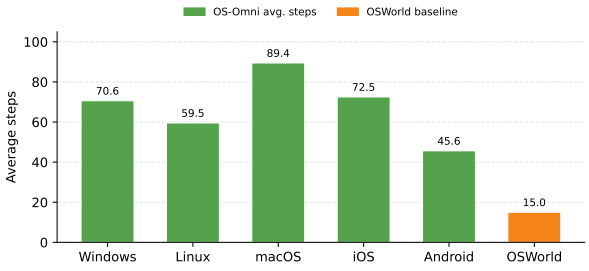

(b) Human interaction steps on OS-OMNI by platform. Bars report the average number of interaction steps needed by human annotators on each platform.

execution environments. The agent-facing environments are instantiated through platform-specific runtimes, with the host-side components deployed either inside or outside containers. Desktop endpoints are instantiated through managed virtual machines, mobile endpoints are instantiated through emulators or simulators, and web tasks are instantiated through controlled browser environments on the supported operating systems. Ubuntu and Windows are controlled through the desktop runtime interface, macOS through the Apple-platform runtime interface, Android through ADB, iOS through an XCUITest-based automation backend, and web tasks through browser-specific controllers and local benchmark services. This design preserves each platform's native UI stack, application state, and input behavior while keeping the host-side execution stack reproducible.

Across these backends, the runtime layer exposes the same episode contract. It restores or prepares the guest, applies the task initializer, captures observations, executes actions, retrieves artifacts, and invokes the final-state evaluator. In v1, cross-platform tasks run on a single screen: the Ubuntu or Windows guest shows a companion Android emulator in a scrcpy window, and the evaluator inspects both devices. The earlier cross-platform workflows, which ran two platform episodes in sequence and transferred the required artifact through the host, have been retired from the task set. Every guest is restored or reset from a clean baseline before each task so that initial state is reproducible, and multiple guests can run in parallel on a single host.

B.2 Observation Space

OS-OMNI supports three observation modes through the execution specification: screenshot-only, accessibility-tree-only, and combined screenshot-plus-tree observation. The primary protocol used in the main results is screenshot-only, because screenshots are naturally available across desktop, mobile, web, and cross-platform workflows without assuming that every application exposes clean structured UI metadata. Accessibility-tree-only and combined-observation modes are supported for alternative protocols, debugging, and ablation studies. Some runtime backends can also expose terminal output for command-line workflows, but terminal output is an auxiliary channel rather than a main benchmark observation mode. Trajectory artifacts (per-step screenshots, model responses, controller status, and the final evaluator output) are recorded for post-hoc inspection but are not surfaced to the agent at decision time. OS-OMNI supports observation refactoring and extension if needed, such as restricting the observation to a focused application or adding a scoped DOM extract.

B.2.1 Screenshot To align with the perception of a human user, we capture a screenshot of the entire device screen at every step. Including the mouse cursor also proves helpful in certain cases where mouse information is crucial. For desktop resolution, we default to 1920×1080, which offers a 16:9 aspect ratio and matches the most commonly used desktop resolution in worldwide statistics. On mobile, the screenshot is captured at the device's native resolution. OS-OMNI also supports modifying the resolution of virtual machines to avoid potential memorization of absolute pixel values and to assist studies on topics like generalization across different resolutions. Screenshots are persisted under per-episode run directories so that a trajectory can be replayed visually and the final state can be audited independently of the evaluator's structured checks.

B.2.2 Accessibility Tree An accessibility tree (or a11y tree) refers to an intricate structure generated by the OS or browser accessibility APIs that renders a representative model of the on-screen content, providing a means of interaction for assistive technologies. Each node within the accessibility tree hosts important information about a UI element. This could range from the nature of the object (a button, checkbox, list item, or paragraph of text), its current state (checked or unchecked, focused or hidden), and even its spatial orientation on the screen.

Different operating systems employ varied accessibility APIs to construct and manipulate the accessibility tree. These include UI Automation (UIA) and Microsoft Active Accessibility (MSAA) for Windows, the NSAccessibility / AX protocol for macOS, the Assistive Technology Service Provider Interface (ATSPI) for the GNOME desktop used on Ubuntu, the device view-hierarchy dump for Android, and XCUITest element trees for iOS. OS-OMNI normalizes these per-platform sources into a shared per-node schema that carries an identifier, a type, visible text, a selector, the enabled / focused / active flags, and the on-screen bounding rectangle, so that the same agent interface can address an element by either node identifier or pixel coordinate. Because the information returned by these accessibility APIs differs in coverage and noise level across platforms, we treat structured metadata as a diagnostic and ablation channel rather than as a privileged oracle.

B.2.3 Terminal Output OS-OMNI also supports terminal output as an optional observation channel for command-line workflows. When enabled, the runtime returns the text shown in the active terminal session or the output exposed by the platform command interface, depending on the backend. This channel is useful for tasks involving shell commands, script execution, logs, and text-based tools, where the relevant state may be textual and not fully recoverable from screenshots alone.

Terminal output is not part of the main screenshot-only protocol. We reserve it for compatible environments, debugging, and ablation studies so that the primary reported setting remains aligned with the visual information available across desktop, mobile, web, and cross-platform workflows.

B.3 Action Space

We implement two kinds of action space: a free-form PyAutoGUI action space, used as the main interface on desktop platforms, and a structured *computer-style* action schema that is exposed by every platform runner and shared across desktop, mobile, and web tasks. The PyAutoGUI action space saves tokens for describing the action space in prompting and matches the OSWorld baselines; the structured schema is finite and parameterized, which makes it more suitable for tool-calling APIs, constrained decoding, and reinforcement-learning-style training. Both interfaces share the same episode lifecycle and the same final-state evaluator, so they are experimental choices rather than separate benchmarks. `DONE` and `FAIL` are terminal actions shared by both interfaces. `WAIT` allows the interface to settle before the next observation without terminating the episode and counts toward the step budget. Episodes may also terminate upon budget exhaustion or an unrecoverable runtime condition.

B.3.1 PyAutoGUI `pyautogui` is an open-source, cross-platform Python module utilized for programmatically controlling the mouse and keyboard. It controls cursor movement, single and multi-button clicks, drags, scroll wheel events, and keyboard input. We adopt PyAutoGUI as our core desktop control component, and we use it as the main agent-facing action space on desktop platforms because it tokenizes compactly, generalizes across Ubuntu/Linux, Windows, and macOS, and matches the style of code used in prior baselines.

At each step, the agent is asked to first give a short reflection on the current screenshot and previous actions, and then to return one fenced Python block that contains the next action(s). Specifically, it may also return one of three special blocks: `WAIT` when it needs the application to settle, `FAIL` when it determines the task cannot be performed, or `DONE` when it determines that the task is complete. The on-guest sandbox only allows `import pyautogui`, `import time`, direct `pyautogui` calls, simple loops, and `print`; modules such as `os`, `subprocess`, or `tkinter` are not importable, and `pyautogui.screenshot()` and `pyautogui.locateCenterOnScreen` are explicitly disallowed. This isolates "the agent moves the mouse and keyboard" from "the agent edits files or runs shell commands," so that screenshot-only success cannot be accidentally laundered through an unrelated automation API.

A representative episode therefore alternates between short natural-language reflections and minimal PyAutoGUI snippets:

````
# step k reflection
The Calc workbook is open on Sheet1 and the cursor is in A1.
I need to enter the SUM formula, so I will click B10 first.

```python
import pyautogui, time
pyautogui.click(x=412, y=638)
time.sleep(0.3)
pyautogui.typewrite("=SUM(B2:B9)", interval=0.02)
pyautogui.press("enter")
```
````

B.3.2 Structured Action Schema To facilitate tool-calling APIs and potential reinforcement-learning applications, OS-OMNI also exposes a structured variant in which actions are wrapped into a finite class with parameterized enumeration. This structured approach supports constrained decoding by providing a distinct set of actions to be learned and optimized. The schema is shared across desktop, mobile, and web lanes so that the same agent interface covers heterogeneous endpoint types while preserving backend-specific execution. In every case, the `target` field of a pointing action accepts either an absolute pixel coordinate `[x, y]`, a normalized coordinate `{"x": 0.5, "y": 0.5}`, or, when the accessibility tree is exposed to the agent, a node identifier from that tree. Tables 2 and 3 summarize the main desktop and mobile action schemas.

The desktop schema retains the full mouse, keyboard, and shortcut surface that PyAutoGUI exposes, but presents it as a discrete, parameterized action space with explicit terminal symbols, making the desktop lane directly comparable to the mobile lanes during benchmark reporting.

Table 2: Action types and parameters defined in the desktop action schema (Ubuntu/Linux, Windows, and macOS). Each desktop guest exposes the same schema to the agent.

Action Type	Parameters	Note
`TAP`	`target`	Click at the specified element or position.
`DOUBLE_TAP`	`target`	Double-click at the specified element or position.
`RIGHT_CLICK`	`target`	Right-click at the specified element or position.
`DRAG`	`from_target, to_target, duration_ms`	Drag from one position to another with the left button held.
`SCROLL`	`direction, distance, target`	Scroll up, down, left, or right; optional anchor target.
`TYPE_TEXT`	`text, target, clear_first`	Type text, optionally after focusing the target field.
`KEYBOARD`	`action, count`	Press a single key (`enter`, `escape`, `backspace`, ...) one or more times.
`HOTKEY`	`keys`	Press a keyboard shortcut, e.g. `[ctrl, c]`.
`WAIT`	`ms`	Wait for the UI to settle.
`DONE`	`message`	Decide the task is done.
`FAIL`	`message`	Decide the task cannot be performed.

The mobile schemas extend the same conceptual interface with the additional intents that touch UIs require, including `LONG_PRESS`, `SWIPE`, `ALERT`, `SYSTEM`, and `LAUNCH_APP`. `DONE` and `FAIL` terminate the episode, whereas `WAIT` is non-terminal. Episodes may also terminate upon budget exhaustion or an unrecoverable runtime condition. After termination, the evaluator checks final-state correctness and any applicable execution constraints, rather than similarity to a reference trajectory.

C Details of the OS-Omni Benchmark

C.1 Platform Selection

OS-OMNI evaluates agents across Windows, Ubuntu, macOS, iOS, Android, and Web environments. The selection is intended to cover desktop, mobile, and browser-based computer use rather than optimize for market share alone. In StatCounter's March 2026 worldwide statistics, Windows accounts for 60.8% of desktop traffic, OS X/macOS for 14.76%, Linux for 3.16%, and ChromeOS for 1.62%.[†] We include Windows and macOS to cover major proprietary desktop ecosystems, and Ubuntu/Linux to provide an open, reproducible desktop stack with broad office, communication, graphics, networking, and developer workflows. The Web group is included as a first-class benchmark environment because many everyday computer tasks are performed through browsers and service-backed interfaces. The current release therefore centers

[†] https://gs.statcounter.com/os-market-share/desktop/worldwide, March 2026 data.

on controlled desktop, mobile, and web runtimes with comparable final-state evaluation.

For mobile environments, Android and iOS account for 67.46% and 32.27% of worldwide mobile operating-system traffic, respectively.[‡] They introduce interaction modes absent from desktop-only benchmarks, including touch gestures, software keyboards, permission dialogs, app switching, and device-specific state. Web tasks are instantiated through widely used browsers on supported operating systems, including Chrome, Edge, Firefox, and Safari; in v1, Web tasks run in Chrome or Chromium on the Ubuntu VM. Together, these environment groups require agents to operate under a unified screenshot-only protocol across pointer-and-keyboard desktops, touch-first mobile devices, and browser-based workflows. We additionally include cross-platform workflows to evaluate cases in which artifacts or information must be prepared in one environment and consumed in another.

C.2 Software Selection

We select applications to span common end-user workflows while preserving deterministic initialization and final-state evaluation. The suite is intentionally centered on everyday applications and task patterns, so that the benchmark covers the kinds of software needs that arise in ordinary desktop and mobile use. It emphasizes software with persistent, user-visible state, including documents, spreadsheets, presentations, messages, calendars, notes, budgets, shopping carts, media libraries, references, diagrams, images, 3D scenes, packet captures, files, and editor workspaces. Applications whose success depends primarily on remote availability or hidden server-side state are used only when their state can be seeded and evaluated through controlled artifacts.

On desktop platforms, OS-OMNI uses overlapping application families rather than an identical software stack. This supports cross-platform comparison while respecting native availability and user conventions. Representative families include office suites, communication tools, finance software, graphics and 3D editors, diagramming tools, packet-analysis tools, reference managers, media applications, and text or code editors. For shopping workflows, we use OpenShop, a controlled benchmark application, instead of real e-commerce services. This preserves the core interaction structure of shopping tasks, including search, cart editing, option selection, and checkout confirmation. At the same time, it avoids personal data collection, payment handling, real purchases, account security risks, and benchmark drift from external commercial systems.

On mobile platforms, the selected applications exercise touch-first interaction and app-level persistence through finance, notes, mail, calendar, reminders, contacts, browsing, file management, and system-utility tasks. Web tasks use browser-native interaction with benchmark-controlled services when external state would otherwise make evaluation unstable. Across all platforms, inclusion requires three properties: the environment can initialize meaningful state, the evaluator can inspect the final state, and the interface supports realistic visually grounded interaction, state tracking, and multi-step planning. Service-backed and platform-limited workflows are covered by the `Open`-series applications described next.

[‡] https://gs.statcounter.com/os-market-share/mobile/worldwide, March 2026 data.

C.3 The Open-Series Applications

Section 3.2 introduces the 18 controlled applications that we develop for OS-OMNI, collectively called the `Open` series. They are used by 180 of the 300 tasks in v1 (60.0%). Their real-world counterparts either depend on user accounts, payments, and live third-party data, or are not available on every operating system. The `Open` applications reproduce the interaction structure of these services while keeping every piece of task-relevant state under benchmark control. Table 4 lists the applications and the number of v1 tasks that use each of them on each platform.

Design principles. All `Open` applications follow the same principles, so that tasks built on them can be initialized and evaluated deterministically.

- **Controlled data.** Accounts, products, flights and hotels, posts, notes, transactions, and documents come from a fixed seed shipped with the application. Task logic needs no external service, and record identifiers are stable across runs.
- **Deterministic time.** Applications whose behavior depends on dates run on a fixed business date (for example, `OpenBank` treats 2026-06-01 as the current day), so relative dates in instructions and date-dependent states, such as whether a transfer is completed or scheduled, are reproducible.
- **Persistent, inspectable state.** Each application writes its business state to a local store: browser storage for the web builds, an embedded database or JSON store for the mobile builds, and a bundled local service for the desktop builds. Evaluators read the final records directly instead of inferring them from the final screen.
- **Business event records.** Applications also append durable business events, such as a submitted checkout, a posted review, or a created transfer. Evaluators use these records to attribute actions to the current episode and to detect duplicated actions. An order between actions is required only when the task instruction states it; browsing actions such as searching, opening a page, or switching tabs are never required or counted unless the instruction asks for them.
- **One data model across builds.** The nine cross-platform applications share their data model and seed across their builds, so the same kind of workflow can be posed on different endpoints and compared.

Deployment and initialization. On Windows and macOS, the applications are installed in the virtual-machine images as desktop builds. On Android and iOS, they are installed on the emulator or simulator as mobile builds, and their stores are read from the device. The nine browser-based applications, together with web builds of several cross-platform applications, are served inside the Ubuntu VM and opened in Chrome; their state lives in Chrome's local storage, which the harness flushes to disk before evaluation and the evaluator decodes with a LevelDB reader. For every task, the initializer resets the application store to the seed, adds the task-specific records (for example, previously placed orders, existing reviews, saved notes, or seeded messages), launches the application, and verifies the starting state before the agent receives control.

C.4 Task Example Sources

We construct task examples from sources that reflect ordinary computer use while remaining reproducible under execution-based evaluation. Public documentation, application tutorials, community Q&A, how-to resources, common productivity workflows, controlled benchmark applications, and cross-application scenarios are used to identify realistic user intents, application states, and evaluable final conditions. These sources seed tasks in office editing, email, chat and messaging, calendar management, finance, file handling, media, diagrams, packet analysis, reference management, and developer-style editing.

Not all realistic workflows can be used directly. Tasks are included only when the benchmark can initialize a meaningful state and evaluate the final result without relying on a hidden reference trajectory. For workflows involving external services, privacy-sensitive data, or unstable online state, we either use controlled local artifacts or a benchmark application with the same interaction structure. OpenShop is one such example: it lets us evaluate search, cart editing, option selection, and checkout confirmation without involving real purchases, payment information, or user accounts.

We also include tasks that cover common interaction conditions in real applications, such as dense visual layouts, small controls, modal dialogs, delayed saves, software keyboards, app switching, long forms, and information transfer across applications or platforms. The goal is not to imitate a single demonstration, but to define a user intent, an initial state, and a final-state check. Each task is therefore represented as a YAML specification containing the instruction, initializer, limits, metadata, and evaluator.

Task authoring follows the same validation discipline as the main benchmark pipeline. Annotators first design tasks from hands-on use cases, calibrate the difficulty so that completion requires meaningful planning rather than a single obvious click, and then execute the oracle version to confirm that the initializer, runtime, and evaluator support the intended goal. Failure-state checks are used to catch underspecified instructions, accidental shortcuts, and evaluators that accept incomplete results.

C.5 Task Examples Collection

OS-OMNI v1 contains 300 task specifications: 48 tasks for each of Windows, Ubuntu, macOS, iOS, Android, and Web, and 12 cross-device workflows. Each task is paired with an oracle specification used for validation, debugging, and sanity checks. The oracle versions are not treated as separate benchmark tasks; instead, they provide a controlled reference for checking task feasibility, initializer correctness, runtime stability, and evaluator behavior. Table 5 summarizes the task distribution and oracle coverage.

C.6 Initial State Setup Details

OS-OMNI tasks begin from realistic intermediate states rather than blank applications. Each initializer prepares the files, databases, app profiles, open windows or app views, local account fixtures, local services, and distractor artifacts needed by the task. This design mirrors ordinary

computer use, where a user typically resumes work from a partially configured environment, an existing inbox, an open document, a populated calendar, or a mobile app with saved state.

Initial state setup follows the same high-level contract across platforms. First, the runtime is restored or prepared from a clean baseline. Linux and Windows use VM restoration, macOS uses Tart clones, Android uses emulator reset with ADB-controlled seeding, and iOS uses prepared Simulator state. Second, task-specific artifacts are placed into the environment, such as documents, spreadsheets, images, notes, packet captures, contact entries, or shopping records. Third, optional setup actions open the target application or navigate to a stable starting view. These setup operations are part of the benchmark specification and are completed before the first agent observation, regardless of the observation mode used in the run.

C.7 Evaluation Configuration Details

OS-OMNI evaluates task completion from the final application state rather than from the action trajectory. After the agent stops or reaches the step limit, the runner dynamically loads the evaluator named in the task file, builds a platform-specific evaluation context, and records the returned `PASS`, `FAIL`, `UNDETERMINED`, or `ERROR` result together with diagnostic metadata. This design is important for open-ended GUI tasks: there are often many valid ways to complete the same instruction, but the completed environment should contain the requested user-visible result.

Evaluator interface. Each evaluator is a Python class loaded from an import path such as `evaluators.linux.OpenShop:OpenShopApiEvaluator`. All evaluators implement the shared `BaseEvaluator` interface and return an `EvalResult` containing a status, an optional message, and an open metadata dictionary. The context always includes the task parameters, the runtime controller, and the final observation. Desktop runners additionally pass the run directory and trajectory path when available, so evaluators can save retrieved artifacts or inspect episode logs. Windows also supports supplemental file and PowerShell checks that run after the primary evaluator passes.

Execution limits. The v1 task specifications use a long-horizon interaction budget: each single-platform task sets `limits.max_steps: 300` and `limits.step_timeout_seconds: 180`, and each cross-device workflow sets 300–650 steps with a 120-second step timeout. The step timeout is a per-step execution budget rather than a whole-episode wall-clock limit; it bounds the model/action turn and prevents a hung action or tool call from blocking the run indefinitely. The evaluator is invoked after the agent emits `DONE` or `FAIL`, after the episode reaches its step budget, or after another terminal runtime condition. If `step_timeout_seconds` is omitted, the shared task loader falls back to 120 seconds.

Evaluation pipeline. Most evaluators follow a three-stage pipeline. First, a getter retrieves task-relevant evidence from the final environment: a saved file, application database, profile folder, local service response, UI tree, simulator container, or device storage path. Second, the evaluator

normalizes this evidence into structured records, such as workbook cells, document paragraphs, diagram labels, email messages, calendar events, shopping cart items, budget transactions, or browser bookmarks. Third, task-specific predicates compare the structured record against the specification in the YAML file. Failures therefore produce concrete mismatch messages and metadata, such as missing rows, unexpected files, unmatched events, observed cart contents, copied artifacts, command output, or parsed database rows.

C.7.1 Ubuntu/Linux

LibreOffice. Calc, Writer, and Impress tasks are evaluated from the saved workbook, document, or presentation rather than from the clicks used to edit it. The evaluator first resolves the result file: if the path is already available on the host it is read directly, otherwise it is fetched from the guest with `controller.get_file`. Spreadsheet checks use workbook-level parsers to compare cells, formulas, sheets, sheet order, and rule-based properties such as filtered ranges or derived tables. Writer and Impress checks parse document and presentation artifacts to verify text, tables, page breaks, slide contents, and transitions. The metadata records the resolved actual and expected paths and the specific mismatch when comparison fails.

Thunderbird and Delta Chat. Thunderbird evaluation uses both profile-level evidence and final UI evidence. Some tasks read files from the Thunderbird profile, such as preference files, filter definitions, or SQLite-backed state; others run post-processing commands, cache command output, or inspect the final accessibility tree when a compose window should remain open rather than send a message. The checker supports list, CSV, JSON, SQLite, preference, filter, and accessibility-tree rules. Delta Chat tasks inspect persisted conversation state through local database queries and task-specific message predicates, including sent messages, replies, and group or coordination state.

Actual Budget and OpenShop. Actual Budget tasks are evaluated by reading the application's local budget store and SQLite database. The checker searches for accounts, starting-balance transactions, categories, budget amounts, ordinary transactions, and transfer pairs while allowing the internal money representation used by the app. OpenShop tasks query the controlled local shopping app through its local API. The evaluator checks session state, cart item counts, product identifiers, quantities, selected specifications, order totals, address labels, payment options, and order notes. Because OpenShop is benchmark-controlled, these checks avoid real purchases, payment data, and external service drift.

Blender, GIMP, draw.io, Marble, Wireshark, Zotero, Notepadqq, VS Code, and Spotube. Creative, analytical, reference, development, and media tasks use application-specific evidence. Blender evaluators run guest-side scripts to extract scene properties and compare object placement, render settings, and camera or lighting requirements. GIMP evaluators retrieve exported

images and inspect dimensions, mode, and pixel-level differences with image-processing libraries. draw.io evaluators decode the saved diagram XML, including compressed diagram payloads, then compare pages and cell labels. Marble tasks parse bookmark or configuration state. Wireshark tasks verify exported captures or protocol reports. Zotero checks query the local library database for collections, items, tags, creators, and notes. Notepadqq and VS Code tasks compare edited workspace files or settings JSON. Spotube tasks inspect persistent profile, market, queue, and playback-related state.

C.7.2 Windows

LibreOffice. Calc, Writer, and Impress tasks are evaluated from saved files and structured document properties. Windows tasks often use checked-in PowerShell scripts that force-save when necessary, retrieve the target file, and run Python checks with libraries such as `openpyxl`, `python-docx`, `python-pptx`, XML parsing, or ZIP inspection. These scripts verify spreadsheet formulas and charts, document formatting and structure, page breaks, images, table-of-contents state, slide shapes, notes, animations, and transitions.

Thunderbird, Calendar, and Delta Chat. Thunderbird and Calendar tasks check saved messages, folders, filters, contacts, and events by combining profile inspection, SQLite queries, and guest-side scripts. Delta Chat tasks verify sent messages, group creation, and shared synchronization artifacts through local database or runtime checks. The Windows runner can also attach supplemental file checks or inline PowerShell checks to the primary evaluator, which is useful for tasks whose success depends on both an app state and an auxiliary artifact.

Actual Budget and OpenShop. Actual Budget tasks use PowerShell-launched Python checks over the application's SQLite state to validate accounts, categories, transactions, and setup details. OpenShop uses the same controlled benchmark application as the Linux lane but queries it from the Windows guest with Python. Its evaluator polls the local API until the state stabilizes, then checks the session, cart, order, and optional UI evidence such as process/window visibility.

Blender, GIMP, draw.io, Wireshark, Zotero, VS Code, Notepad, Spotube, and PowerShell checks. These tasks are evaluated through saved project files, exported artifacts, packet captures, library metadata, source files, plain-text files, persistent application settings, or command-level evidence. Several Windows evaluators are thin wrappers over named PowerShell scripts in `evaluators/windows/checks`. The script output and exit code become structured metadata, and selected guest artifacts can be copied into `evaluation_artifacts` under the run directory for later inspection.

C.7.3 macOS

LibreOffice. Calc, Writer, and Impress tasks check saved spreadsheet, document, and presentation artifacts retrieved from the Tart guest. The evaluators parse the ZIP/XML structure of LibreOffice-compatible documents to verify cell contents, document paragraphs, title formatting, tables, footer or heading requirements, and slide contents. This avoids depending on the visible final screen, which may not show all required document state.

Thunderbird Email, Thunderbird Calendar, and Delta Chat. Email and calendar tasks inspect saved messages, folders, filters, contacts, and events using guest-side Python scripts and local profile or SQLite state. Delta Chat tasks run guest-side database queries and return JSON summaries of conversations, recipients, and messages. The evaluator then checks exact recipients, subjects, message bodies, and required reply or coordination state.

Actual Budget, OpenShop, and Spotube. Actual Budget tasks scan `global-store.json`, locate the active `db.sqlite`, copy sidecar WAL/SHM files when present, and query accounts, categories, budgets, transactions, and transfers. OpenShop tasks check the controlled shopping application's local API or persistent state. Spotube tasks inspect the local SQLite and settings state used by the application to validate profile, market, plugin, queue, and playback-related requirements.

Blender, GIMP, draw.io, Wireshark, Zotero, VSCodium, and TextEdit. These tasks are evaluated from saved scene files, images, diagrams, packet evidence, reference library entries, workspace files, and edited text files. Blender and Wireshark evaluators run guest-side scripts that return JSON payloads. GIMP checks compare exported images with expected image properties or fixtures. draw.io decodes diagram XML and compares page/cell semantics. Zotero copies and queries `zotero.sqlite` and sidecar files to check collections, item titles, creators, tags, and notes. VSCodium and TextEdit evaluators inspect workspace files, settings, and edited text content.

C.7.4 Android

K-9 Mail and Delta Chat. Mail and chat tasks inspect sent messages, forwarded content, replies, archived state, contacts, and status changes. Evaluators use ADB-backed access to app state, local mail service evidence, or task-specific API records where available. Delta Chat evaluators validate exact outbound messages, forwarded content, search-and-reply behavior, archive state, and created contact state. Cross-application mail tasks additionally verify that information extracted from one app is correctly written or sent in another.

Fossify Calendar, Contacts, and Tasks.org. Scheduling and personal-information tasks check persisted event, contact, reminder, due-date, note, and subtask records. Contacts are validated through Android content-provider or database-style ADB queries. Tasks.org evaluators inspect persisted task fields, due dates, notes, parent tasks, and subtasks while accounting for the device timezone.

Markor, LibreOffice, Material Files, and Fossify Gallery. Note, document, file, and gallery tasks are evaluated through external storage and document artifacts. Generic filesystem evaluators check expected paths, forbidden paths, file/directory types, and exact directory entries on shared storage. Markor and text-file tasks read final note contents. LibreOffice Viewer tasks inspect modified office documents by pulling files from device storage and parsing the document package. Gallery tasks check renamed, moved, or deleted photo artifacts while preserving distractor files.

Oinkoin and OpenShop. Oinkoin tasks pull the app's `movements.db` through ADB root access, copy it to a temporary host location, and query records for matching title, note, category, sign, amount, and category type. OpenShop Android tasks read the WebView localStorage-backed state through the benchmark store helper. They poll until the state stabilizes and then check login status, cart size, required or forbidden cart items, latest order contents, address/payment choices, and order notes.

Spotube, Calculator, Clock, and Settings. Media and utility tasks check persistent app or device state. Spotube evaluators inspect XML/shared-preference or database state for profile, queue, playback, and search outcomes. Calculator and Clock tasks use UI-tree, screenshot-difference, activity, or device-state evidence when the final state is primarily visible rather than stored in a file. Settings tasks use ADB-readable device state, such as Bluetooth status.

Cross application Android tasks. Android also includes workflows combining Fossify Calendar with OpenShop, K-9 Mail with Markor or LibreOffice, and Fossify Calendar or Organic Maps with Delta Chat. These evaluators compose the relevant child checks and report child metadata separately, so a failure can be attributed to the source app, target app, or transfer step.

C.7.5 iOS

Calendar and Reminders. Calendar tasks are backed by the simulator's Calendar database. Evaluators locate the simulator by UDID when a direct database path is not provided, query `Calendar.sqlitedb`, normalize times with the configured timezone, and check required or forbidden event summaries, start/end times, locations, and notes. Reminders tasks combine visible UI evidence with app state: they check the foreground bundle, the selected list view, expected titles, forbidden titles, and exact-list constraints when the task requires a particular visible reminder list.

Contacts, Delta Chat, and Maily. Contacts tasks inspect `AddressBook.sqlitedb` from the simulator container and validate saved names, phone numbers, and related fields. Delta Chat and Maily tasks evaluate sent messages, forwarded content, deleted or retained messages, recipients,

subjects, and message bodies, using local app state and controlled mail/message service evidence rather than the final screenshot alone.

FSNotes, Flow, and Habit. FSNotes tasks inspect the note root inside the simulator data container to verify created, moved, renamed, archived, or edited Markdown notes. Flow tasks inspect the app support directory and ObjectBox-backed state for expenses, transfers, income records, and budget updates. Habit tasks inspect the local habit database for created, removed, or modified routines. These evaluators poll briefly and include observed summaries in metadata to reduce false negatives from asynchronous app persistence.

OpenShop and Safari. OpenShop iOS tasks use the same controlled shopping-state abstraction as Android, but resolve the local store from the simulator container or explicit database path. They check login state, cart contents, forbidden items, latest order state, and checkout details. Safari tasks query simulator history, bookmarks, folders, tabs, active-tab state, and optional foreground-bundle evidence, so they can distinguish opening the target page from merely typing a URL into a visible field.

C.7.6 Web

Browser and controlled-service tasks. Web tasks are instantiated through Chrome, Edge, Firefox, and Safari on supported operating systems. Evaluators inspect the browser-visible final state and, when applicable, the state of benchmark-controlled services. The checked evidence includes sessions, form submissions, shopping carts, orders, bookmarks, tabs, history entries, downloads, and local storage or database records. Service-backed workflows keep state under benchmark control so that final-state checks do not depend on live commercial services, personal accounts, or unstable remote data.

C.7.7 Cross-Application and Cross-Platform Workflows Cross-application tasks combine the app-level checks above when a user goal spans multiple applications on the same endpoint, such as Writer with Thunderbird, Thunderbird Calendar with OpenShop, K-9 Mail with Markor or LibreOffice, Calendar with OpenShop, or Calendar, Maps, and Delta Chat. The cross-app evaluator does not reward an intermediate handoff by itself; it checks the final state of every required destination artifact or application record. In v1, cross-device tasks run on one screen: the Ubuntu or Windows desktop shows a companion Android phone in a scrcpy window. The agent edits or builds a file in Markor on the phone, moves it to the desktop itself (for example, through the synced clipboard), builds the deliverable in LibreOffice, and sends results computed on the desktop back to the phone as a Markor note, Tasks items, Calendar events, or the exported file. The evaluator checks the phone's source file and returned records (exactly the expected ones and no others), checks that the copies on the phone match the desktop files, and checks the desktop document. Any agent action that invokes or types `adb` or `scrcpy` fails the task.

These workflows replace the earlier cross-VM workflows (Android-to-Windows and Windows-to-Linux), which ran source and target stages on different endpoints. When a deterministic checker cannot represent all acceptable outcomes, the case is routed to the manual verification protocol described in Section C.9.

C.8 More Task Examples

Additional task examples from the screenshot-only setting are provided at the end of the appendix. They include Ubuntu/Linux examples in Figures 13–18, Windows examples in Figures 19–25, macOS examples in Figures 26–31, Android examples in Figures 32–33, and iOS examples in Figures 34–35. Each example summarizes the related applications, task instruction, screenshots, and abilities needed to complete the task.

C.9 Manual Verification for Ambiguous Cases

Manual verification is reserved for cases where a deterministic evaluator cannot fully capture acceptable user intent. It is not used to repair ordinary model failures. Reviewers inspect the saved instruction, screenshots, action trace, final artifacts, and evaluator diagnostics. Each manual override records the accepted final state and the reason the automatic check was insufficient.

D Details of Agent and Evaluation Protocols

D.1 Model and Execution Configuration

OS-OMNI runs agents through execution specifications that bind together a task list, model identifier, observation mode, action mode, and history window. Model endpoints, provider-specific identifiers, authentication variables, and generation parameters are specified through a shared model registry, implemented in the release as `configs/models.yaml`. This separation keeps model selection independent from task initialization and evaluation, so the same task specifications can be reused across providers without changing the benchmark contract.

Unless otherwise stated, the main protocol uses screenshot-only observations and the platform action backend selected by the execution specification. The default history window contains the three most recent interaction steps. Tasks use the long-horizon limits declared in the task specification: in v1, `limits.max_steps:300` and `limits.step_timeout_seconds:180` for single-platform tasks, and 300–650 steps with a 120-second step timeout for cross-device workflows. The step timeout bounds one model and action turn rather than the whole episode.

D.2 Prompt Details

Prompts are constructed from the task instruction, the platform-specific control interface, the current observation, and a bounded history of recent steps. In the main screenshot-only protocol, the agent is instructed to ground actions in the visible screen, use the coordinate frame of the current screenshot, avoid hidden application state or evaluator internals, and return one executable action at a time. Before the action, the agent provides a short reflection on the current

screen and recent progress.

Across platforms, the agent follows the same high-level contract: it first gives a brief reflection, then returns an executable action block or one of the episode-control blocks `WAIT`, `DONE`, and `FAIL`. The platform runner validates the action block and dispatches it through the corresponding desktop, mobile, or browser controller. This design matches the shared action-space abstraction in the main paper while preserving the native execution channel used by each runtime.

Representative screenshot-only system prompts for the OS-specific runners are shown in Figures 8–12. Web tasks reuse the same benchmark-facing contract with browser-specific runtime instructions.

Table 3: Action types and parameters defined in the mobile action schema. A and I indicate availability on Android and iOS, respectively.

Action Type	Platform	Parameters	Note
`TAP`	A, I	`target`	Tap at the specified element or position.
`DOUBLE_TAP`	I	`target`	Double tap at the specified element or position.
`LONG_PRESS`	A, I	`target`, `duration_ms`	Press and hold the target.
`SWIPE`	A	`from_target`, `to_target`, `duration_ms`	Swipe from one position to another.
`DRAG`	I	`from`, `to`, `duration_ms`	Drag from one element/position to another, e.g. to reorder.
`SCROLL`	A, I	`direction`, `distance`, `target`, `duration_ms`	Scroll the screen or specified container.
`TYPE_TEXT`	A, I	`text`, `target`, `clear_first`	Type text, optionally after tapping a target field.
`KEYBOARD`	A, I	`action`, `count`	Send a keyboard or system key (`enter`, `back`, `home`, `escape`, ...).
`CLIPBOARD`	I	`action`, `text`	Set, get, or paste clipboard content.
`ALERT`	I	`action`	Accept or dismiss a system alert / popup.
`SYSTEM`	I	`action`	System action (`home`, `hideKeyboard`).
`LAUNCH_APP`	A, I	`name`	Launch an app by package or bundle identifier.
`WAIT`	A, I	`ms`	Wait for the UI to settle.
`DONE`	A, I	`message`	Decide the task is done.
`FAIL`	A, I	`message`	Decide the task cannot be performed.

Table 4: `Open`**-series applications in OS-Omni v1.** Counts are the number of v1 tasks that use each application on each platform. Cross-application tasks are counted under their platform; three of them combine `OpenShop` with Thunderbird Calendar (one on Windows, two on Linux). In total, 180 of the 300 tasks use at least one `Open` application.

Application	Domain and main operations	Win	Linux	macOS	iOS	Android	Web	Total
Controlled cross-platform applications								
`OpenShop`	E-commerce: search, product options, cart, checkout, order notes, reviews	5	2	4	4	6	2	23
`OpenNote`	Social note sharing: posting, liking, bookmarking, commenting, following, profile	4	–	4	4	6	3	21
`OpenTwitter`	Social networking: posts, replies, reposts, likes, saves, direct messages, profile	4	–	4	4	6	2	20
`OpenBank`	Online banking: recipients, transfers, bill payments, disputes, alerts, statements	4	–	3	3	6	3	19
`OpenDriver`	Ride-hailing with rider and driver interfaces: rides, documents, support tickets, payouts	4	–	4	4	4	3	19
`OpenTrip`	Flight and hotel booking: search, checkout, vouchers, reviews, rescheduling	4	–	4	4	4	2	18
`OpenPhotoEdit`	Image editing: adjustments, filters, calibrated exports	–	–	4	4	4	3	15
`OpenFood`	Food delivery: ordering, coupons, merchant and courier roles, customer support	4	–	–	–	4	3	11
`OpenNavigation`	Route planning: destinations, waypoints, travel modes, incident reports	4	–	–	–	3	3	10
Web-based applications for software coverage gaps (served in the Ubuntu VM)								
`OpenDoc`	Document editing, comments, and export	–	–	–	–	–	3	3
`OpenSheet`	Spreadsheets: formulas, filters, and export	–	–	–	–	–	3	3
`OpenSlide`	Presentations: layouts, shapes, notes, comments, slide show	–	–	–	–	–	2	2
`OpenMail`	Email, attachments, and contact management	–	–	–	–	–	3	3
`OpenFinance`	Personal finance, budgeting, and portfolio management	–	–	–	–	–	3	3
`OpenCode`	Code editing, debugging, and version control	–	–	–	–	–	3	3
`OpenBlender`	3D scene editing and animation	–	–	–	–	–	3	3
`OpenWireshark`	Network packet analysis	–	–	–	–	–	2	2
`OpenZotero`	Reference management	–	–	–	–	–	2	2
Total		33	2	27	27	43	48	**180**

Table 5: Task distribution and oracle coverage in OS-OMNI v1.

Platform group	Tasks	Oracle coverage	Representative task families
Windows	48	48/48	Office editing, communication, shopping, diagrams, packet analysis, reference management, and development workflows
Ubuntu	48	48/48	Office editing, email and chat, finance, graphics, 3D editing, packet analysis, references, media, and file workflows
macOS	48	48/48	Text and code editing, office workflows, email and calendar tasks, references, media, graphics, and productivity applications
iOS	48	48/48	Notes, browsing, mail, calendar, reminders, contacts, finance, shopping, and mobile productivity workflows
Android	48	48/48	Mail, notes, chat, calendar, finance, shopping, file management, media management, contacts, and system utilities
Web	48	48/48	Browser-based workflows in Chrome or Chromium on the Ubuntu VM
Cross-device	12	12/12	Phone-to-desktop round trips: Markor on Android, LibreOffice on the Ubuntu or Windows desktop, results returned to the phone
Total	300	300/300	v1 task collection

Table 6: Representative final-state signals used by OS-OMNI evaluators.

Signal	Access path	Examples
Office/text artifacts	Guest file retrieval; host-side parsers	XLSX, DOCX, PPTX, plain text
App stores and profiles	SQLite, JSON, XML, plist, profile files	mail, calendar, finance, Zotero
Controlled services	Local API or localStorage	OpenShop session, cart, order
Creative/analysis files	Project parsing and app scripts	GIMP, Blender, draw.io, Wireshark
Mobile state	ADB, device storage, simulator containers, UI state	notes, contacts, Safari, reminders
Cross-application state	Child evaluators plus transfer artifacts	attachments and follow-ups

Table 7: Default execution parameters used in the OS-OMNI v1 experiments.

Parameter	Setting
Observation mode	`screenshot_only`
Action mode	Platform-specific controller selected by the execution specification
History window	3 recent steps
Maximum steps	300 per single-platform task; 300–650 per cross-device workflow
Step timeout	180 seconds per step (single-platform); 120 seconds (cross-device)

```
You are an agent that follows the instruction and performs desktop computer tasks
↳ by controlling the mouse and keyboard on a computer.
↳ Assume good computer knowledge and internet connection.

For each step, you receive an observation image:
↳ the screenshot of the current computer screen.
↳ Predict the next computer action based on the screenshot.

Use `pyautogui` to perform actions grounded in the observation.
↳ Do NOT use `pyautogui.locateCenterOnScreen`
  to locate UI elements.
↳ Do NOT use `pyautogui.screenshot()`
  to make screenshots.
↳ Return one line or multiple lines of Python code
  to perform the action efficiently.

The code sandbox only allows:
↳ `import pyautogui`
↳ `import time`
↳ direct pyautogui calls
↳ simple loops
↳ `print`

Do NOT import or use:
↳ tkinter, os, subprocess, pyperclip, bpy, bmesh,
  or any other module.
↳ clipboard-based automation.
↳ application-internal scripting unless reachable through
  normal keyboard and mouse actions.

Linux desktop restrictions:
↳ Do NOT press Alt+F4 or close windows unless the task
  explicitly asks to close that exact app window.
↳ Complete the requested work through the visible desktop UI
  using normal keyboard and mouse actions.

When predicting multiple lines of code:
↳ add small sleeps such as `time.sleep(0.5)`
  so that the machine can respond.
↳ each prediction must be a complete code snippet.
↳ no variables or functions can be shared from previous steps.
↳ specify coordinates by yourself based on the current screenshot.
↳ ensure the coordinates are correct.

Before returning the action:
↳ briefly reflect on the current screenshot and previous actions.
↳ the reflection should be short.

Return the executable response as Python code:
```python
# your code here
```

Figure 8: System prompt used for Ubuntu/Linux screenshot-only runs.
```

```
You are an agent that follows the instruction and performs desktop computer tasks
↳ on a Windows computer by controlling the mouse and keyboard.
↳ Assume good computer knowledge and internet connection.

For each step, you receive an observation image:
↳ the screenshot of the current computer screen.
↳ Predict the next computer action based on the screenshot.

Use `pyautogui` to perform actions grounded in the observation.
↳ Do NOT use `pyautogui.locateCenterOnScreen`
  to locate UI elements.
↳ Do NOT use `pyautogui.screenshot()`
  to make screenshots.
↳ Return one line or multiple lines of Python code
  to perform the action efficiently.

The code sandbox only allows:
↳ `import pyautogui`
↳ `import time`
↳ direct pyautogui calls
↳ simple loops
↳ `print`

Do NOT import or use:
↳ tkinter, os, subprocess, pyperclip, bpy, bmesh,
  or any other module.
↳ clipboard-based automation.
↳ application-internal scripting unless reachable through
  normal keyboard and mouse actions.

Windows-specific restrictions:
↳ Do NOT press Alt+F4 or close windows unless the task
  explicitly asks to close that exact app window.
↳ Never open or use cmd.exe, PowerShell, pwsh,
  Windows Terminal, or any terminal / shell window.
↳ Do NOT use Win+R to launch shell commands.
↳ Complete the requested work inside the target application UI.

When predicting multiple lines of code:
↳ add small sleeps such as `time.sleep(0.5)`
  so that the machine can respond.
↳ each prediction must be a complete code snippet.
↳ no variables or functions can be shared from previous steps.
↳ specify coordinates by yourself based on the current screenshot.
↳ ensure the coordinates are correct.

Before returning the action:
↳ briefly reflect on the current screenshot and previous actions.
↳ the reflection should be short.

Return the executable response as Python code:
```python
# your code here
```

Figure 9: System prompt used for Windows screenshot-only runs.
```

```
You are an agent that follows the instruction and performs desktop computer tasks
↳ by controlling the mouse and keyboard on a computer.
↳ Assume good computer knowledge and internet connection.

For each step, you receive an observation image:
↳ the screenshot of the current computer screen.
↳ Predict the next computer action based on the screenshot.

Use `pyautogui` to perform actions grounded in the observation.
↳ Do NOT use `pyautogui.locateCenterOnScreen`
  to locate UI elements.
↳ Do NOT use `pyautogui.screenshot()`
  to make screenshots.
↳ Return one line or multiple lines of Python code
  to perform the action efficiently.

The code sandbox only allows:
↳ `import pyautogui`
↳ `import time`
↳ direct pyautogui calls
↳ simple loops
↳ `print`

Do NOT import or use:
↳ tkinter, os, subprocess, pyperclip, bpy, bmesh,
  or any other module.
↳ clipboard-based automation.
↳ application-internal scripting unless reachable through
  normal keyboard and mouse actions.

macOS desktop restrictions:
↳ Do NOT press Command-Q or close windows unless the task
  explicitly asks to close that exact app window.
↳ Complete the requested work through the visible desktop UI
  using normal keyboard and mouse actions.

When predicting multiple lines of code:
↳ add small sleeps such as `time.sleep(0.5)`
  so that the machine can respond.
↳ each prediction must be a complete code snippet.
↳ no variables or functions can be shared from previous steps.
↳ specify coordinates by yourself based on the current screenshot.
↳ ensure the coordinates are correct.

Before returning the action:
↳ briefly reflect on the current screenshot and previous actions.
↳ the reflection should be short.

Return the executable response as Python code:
```python
# your code here
```

Figure 10: System prompt used for macOS screenshot-only runs.
```

````
You are an Android GUI agent.
↳ For each step, you receive an Android screenshot.
↳ Predict the next Android action grounded in the current observation.

Before returning the action:
↳ briefly reflect on the current screenshot and previous actions.
↳ the reflection should be short.

Return exactly one fenced adb command block
↳ or one fenced special code block.
↳ Do not include anything else.

Use adb commands to perform actions grounded in the observation.
↳ Return one line or multiple lines of adb commands each time.
↳ Be time efficient.

When predicting multiple lines of commands:
↳ add small sleeps such as `sleep 0.5`
so that the device can respond.
↳ each prediction must be a complete adb command block.
↳ no variables or state can be shared from previous steps.

...

Android interaction rules:
↳ use `input swipe` for drag or scroll.
↳ type into a field with tap, short sleep, then input text.
↳ press Enter / Back / Delete / Space with input keyevent.
↳ you do not need to tap individual soft-keyboard keys.

Targeting rule:
↳ when a visible button, tab, or floating action button is clearly shown,
aim for the geometric center of the whole control.
↳ do not click the edge or only the inner icon glyph.

Coordinate rule:
↳ coordinates are Android screenshot pixel coordinates.

Special outputs are also allowed:
↳ return ```WAIT``` when the interface needs more time.
↳ return ```DONE``` when the task is finished.
↳ return ```FAIL``` when the task cannot be completed.

Do NOT return:
↳ JSON.
↳ pyautogui code.
↳ explanatory bullets.
↳ commands outside the allowed list.

Task:
↳ {goal}

Final output rule:
↳ Return exactly two parts:
first a brief plain-text reflection,
then the fenced adb command block or fenced special code.
↳ Do not include anything else.
````

Figure 11: System prompt used for Android screenshot-only runs.

```
You are an iOS GUI agent.
↳ For each step, you receive an iOS screenshot.
↳ Predict the next iOS action grounded in the visible screenshot only.

Before returning the action:
↳ give a short reflection in plain text
  about the current screen and recent progress.
↳ the reflection should be brief.

Return exactly one fenced JSON action block
↳ or one fenced special code block.
↳ Do not include extra prose after the fenced block.

Use only the iOS action schema below.
↳ Do NOT invent new action types.
↳ Do NOT invent new parameters.

Pure screenshot mode:
↳ do not reference accessibility node ids.
↳ do not reference hidden tree metadata.
↳ choose targets from the visible screenshot only.

Coordinate rule:
↳ use normalized coordinates such as
  `{"x": 0.42, "y": 0.77}`.
↳ pixel coordinate arrays may also be used
  when choosing targets from the screenshot.

Action rules:
↳ `ACTION_RULES = {action_rules}`

Return the action as a single JSON object
↳ inside a fenced code block.

Example:
```json
{"type":"tap","target":{"x":0.42,"y":0.77}}
```
```

Figure 12: System prompt used for iOS screenshot-only runs.

D.3 Action Parsing and Episode Execution

At each step, the runner sends the task instruction, the current screenshot observation, and the bounded recent history to the model. The model response is parsed into two parts: a short reflection and an executable action block. The reflection is stored in the trajectory for later analysis but is not executed. The action block is passed to the platform runner, which validates the format and dispatches the corresponding mouse, keyboard, touch, or device-control operation.

For desktop platforms, the action block is executed under a restricted PyAutoGUI-style sandbox or the corresponding desktop runtime controller. The sandbox permits direct mouse and keyboard control, short waits, and simple control flow, while disallowing hidden automation channels such as shell commands, screenshot APIs, image-location APIs, clipboard-based shortcuts, or application-internal scripting. For mobile platforms, the runner maps the same action intent to ADB-backed Android control or Appium/XCUITest-backed iOS control. For web tasks, browser-facing actions are dispatched through the runtime backend associated with the selected browser and host platform. This keeps the model-facing interaction protocol unified while allowing each environment to use its native execution backend.

After each executed action, the environment waits for the interface to settle, captures the next screenshot, and appends the response, action, controller status, and observation artifact to the trajectory. An episode terminates when the agent emits `DONE` or `FAIL`, when the task reaches the maximum step budget, or when a runtime condition such as a step timeout prevents further execution. The evaluator is invoked only after termination and checks the final environment state rather than the intermediate action sequence.

D.4 Set-of-Mark Implementation Details

Set-of-Mark (SoM) overlays are treated as optional visual grounding aids, not as a default benchmark observation. A SoM run annotates candidate regions on the screenshot and exposes mark identifiers to the agent. This can help isolate whether a model fails because it cannot localize a target or because it cannot plan the workflow after localization. Since SoM changes the visual input, SoM results should be reported as an ablation rather than mixed with screenshot-only results.

D.5 Result Reporting Protocol

Cross-platform result tables are valid only when platform lanes are run under a declared task inventory, model registry, observation protocol, action mode, and failure-accounting policy. The report should distinguish task failure from invalid runs caused by runtime setup, model-service errors, or evaluator exceptions. This convention follows the execution-based protocol in the main paper: a model succeeds only when the final environment state satisfies the task evaluator, while infrastructure failures are accounted for separately.

Recommended reporting fields. For each model and platform lane, reports include the release version (for example, v1), task count, success rate, invalid-run rate, average steps for successful

tasks, timeout rate, observation mode, action mode, and the execution specification used for the run. Aggregate tables separate screenshot-only, accessibility-tree-only, combined-observation, and SoM settings. This prevents a visually grounded run from being compared against a metadata-assisted run as if the agent received the same information.

E Qualitative Analysis

Qualitative analysis in OS-OMNI is based on saved trajectories, runtime metadata, and final artifacts rather than isolated screenshots. Each inspected run records the platform lane, task identifier, model, observation mode, action mode, episode status, evaluator result, and the first step at which the trajectory diverges from a plausible solution. This makes qualitative inspection useful for separating model errors from benchmark-system failures such as initialization drift, runtime instability, or evaluator mismatch.

E.1 Success and Failure Cases

Successful episodes typically share three properties. The agent maintains the user goal across application transitions, grounds actions in the current observation, and leaves the target artifact or application record in a verifiably saved state. Failed episodes often diverge before the final evaluator is invoked: the agent clicks an adjacent control, overlooks a modal dialog, types into the wrong field, loses track of transferred information, or stops after reaching a visually plausible but incomplete state. These cases illustrate why OS-OMNI scores final environment state rather than trajectory similarity.

E.2 Common Errors by Screenshot-only Agents

See Table 8.

E.3 Human-Agent Difficulty Discrepancies

Human difficulty and agent difficulty can diverge sharply in GUI benchmarks. Humans are robust to small UI variations, hidden save states, and modal dialogs, but may be slower on repetitive data-entry tasks. Agents can be efficient when the interface is visually simple and the target state is explicit, but they remain brittle under dense layouts, long-horizon memory requirements, platform-specific shortcuts, and tasks where success depends on confirming a saved or exported artifact.

For this reason, human studies should use the same initializer and final-state evaluator as model runs, while recording completion time and ambiguity notes. Infrastructure failures such as VM restore errors, simulator crashes, app-service outages, or evaluator exceptions should be excluded from human-agent comparisons and reported as benchmark-system failures.

Table 8: Failure taxonomy for screenshot-only OS-OMNI runs.

Failure type	Typical symptom	Where it is especially visible
Visual grounding	Selects a neighboring button, wrong cell, hidden tab, or stale menu item	Dense office files, diagrams, packet views, mobile settings
Semantic interpretation	Misreads visible text, UI labels, document content, or task-relevant state	Forms, calendars, references, packet views, canvas-heavy apps
Long-horizon state tracking	Forgets copied values, loses schedule details, or repeats a completed subgoal	Multi-app desktop tasks; Android and iOS coordination tasks
Persistence	Performs the visible edit but fails to save, export, send, or confirm it	Office suites, image editors, mail clients, notes, shopping carts
Timing or modal handling	Acts before loading finishes, ignores a permission prompt, or types into a dialog	Mobile apps, browser tasks, app launches, restored desktop sessions
Input or tool syntax	Enters text in the wrong field, hits keyboard occlusion, or uses a platform-specific shortcut incorrectly	Android, iOS, macOS, and Windows text-entry tasks
Transfer	Produces an incomplete source artifact or consumes the wrong transferred object	Cross-platform workflows
Benchmark system	Initializer drift, runtime failure, or evaluator mismatch prevents a clean judgment	Reported separately from model failure

E.4 Comparative Analysis across Models

Model comparisons should be organized by failure mechanism rather than by anecdotal task outcome. For example, two VLMs may both fail a spreadsheet task, but one may misread the screen while the other may identify the correct cell and then fail to persist the edit. Similarly, a text-oriented model with accessibility metadata may perform well on labeled forms but fail on canvas-heavy image or diagram tasks. We therefore group inspected trajectories by observation mode, platform lane, task family, and first-failure cause.

E.5 Analysis of SoM Effectiveness

SoM can improve localization when a model can reason over discrete visual marks, but it can also introduce clutter or bias the agent toward marked regions that are not semantically relevant. In OS-OMNI, SoM tests a specific hypothesis about visual grounding. It is reported as an ablation rather than a replacement for screenshot-only evaluation, because the benchmark's core cross-platform question is whether agents can operate from the same visual signal across desktop, mobile, and web environments.

E.6 Endpoint-Level Error Cases

We further provide representative error cases from the Ubuntu/Linux, Windows, macOS, Android, and iOS endpoints. The cases are selected to illustrate distinct failure mechanisms rather than repeated instances of the same error type. The full set is shown in Figures 36–41.

Linux-Actual Budget

Task
task_id: linux.actual_budget.household_budget_setup_hard
Platform: Linux
App: Actual Budget

Task instruction
Actual Budget is already open. If no budget file is open yet, create a new local budget called `Bench Household Budget` in the default documents location. Then set up a household budget with:
⋯ - Twelve payment transactions across the created accounts ⋯
Finish only after all twelve transfers are recorded.

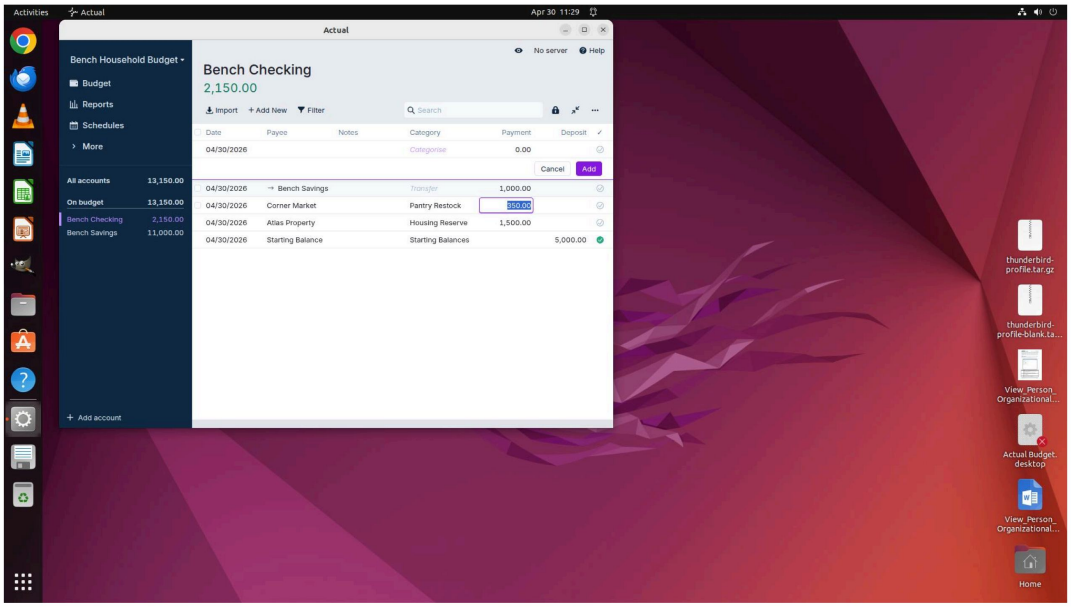

Linux-Blender

Task
task_id: linux.blender.prepare_social_render
Platform: Linux
App: Blender

Task instruction
In Blender, open the prepared product stage scene and update the render setup for a social preview. Set the render engine to EEVEE, set the output resolution to 1800 by 1000 at 100% scale, ...derivatives=social_review_closeup.blend|social_archive_wide.blend|social_story_crop.blend, shared_workspace=blender_social_render, handoff_status=ready.

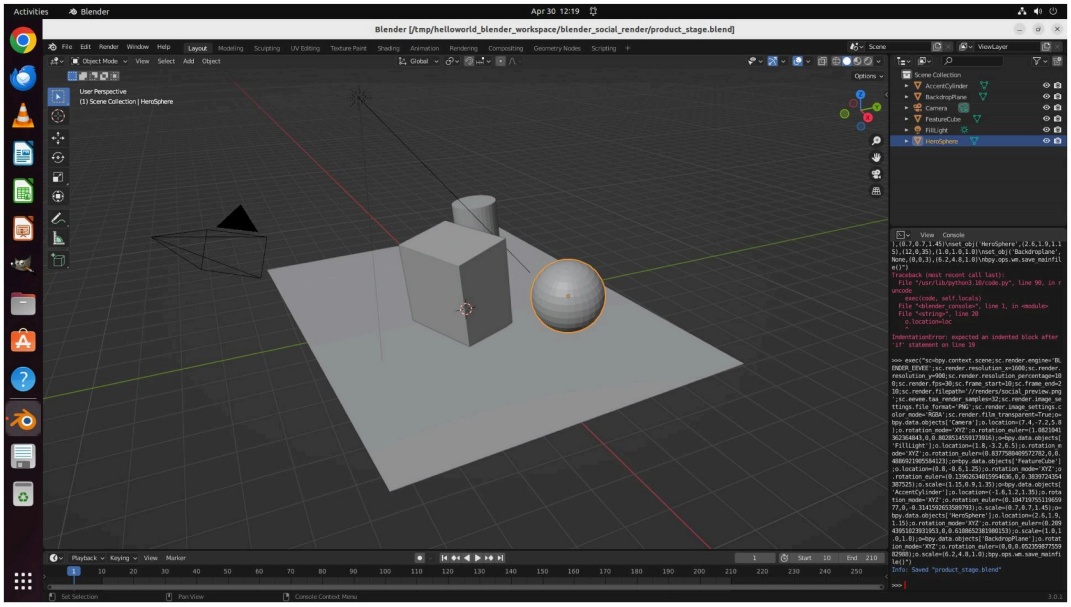

Figure 13: Representative Ubuntu/Linux task examples.

Back to list of figures

Linux-LibreOffice Calc

Task
task_id:
linux.calc.freeze_header_range
Platform: Linux
App: LibreOffice Calc

Task instruction
The spreadsheet is already open in LibreOffice Calc. Complete the workbook setup and audit delivery:
- Freeze the range `A1:B1` so the headers remain visible.
……
- Set exact row heights for rows `1:12`, including row 1 `24 pt`……
- Create the exact audit delivery files under `/tmp/helloworld_office/`: ……
- Save the workbook when finished.

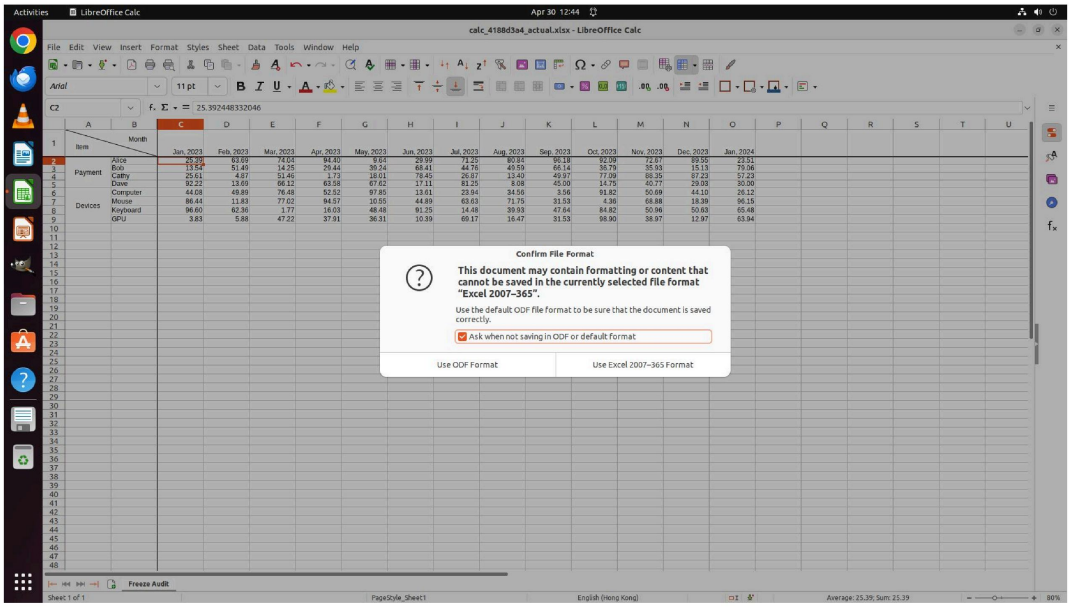

Linux-Thunderbird Calendar

Task
task_id:linux.cross_app.thunderbird_calendar_add_coffee_cart_event_easy
Platform: Linux
App: Thunderbird Calendar

Task instruction
`In Thunderbird Calendar, create a new event titled "Bench Coffee Cart Review" in the Home calendar on May 1, 2026, from 4:30 PM to 5:00 PM. Ensure ... and navigate to the corresponding week. Leave Thunderbird Calendar open with the newly created event visible.`

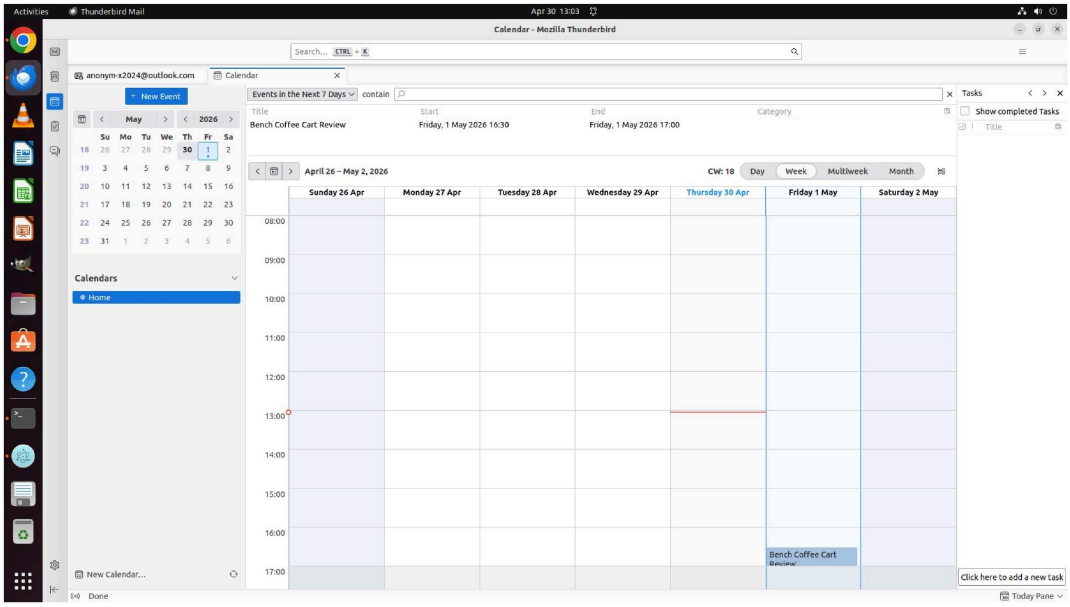

Figure 14: Representative Ubuntu/Linux task examples.

Back to list of figures

Linux-OpenShop

Task

task_id: linux.OpenShop.add_headphones_to_cart
Platform: Linux
App: OpenShop

Task instruction

In **OpenShop**, sign in and build a six-item cart through exact catalog operations:
.....
- Sign in with phone number `13800138000` and password `spring2026`.
.....
- Stop only when all six products are still listed together with quantity `1`.

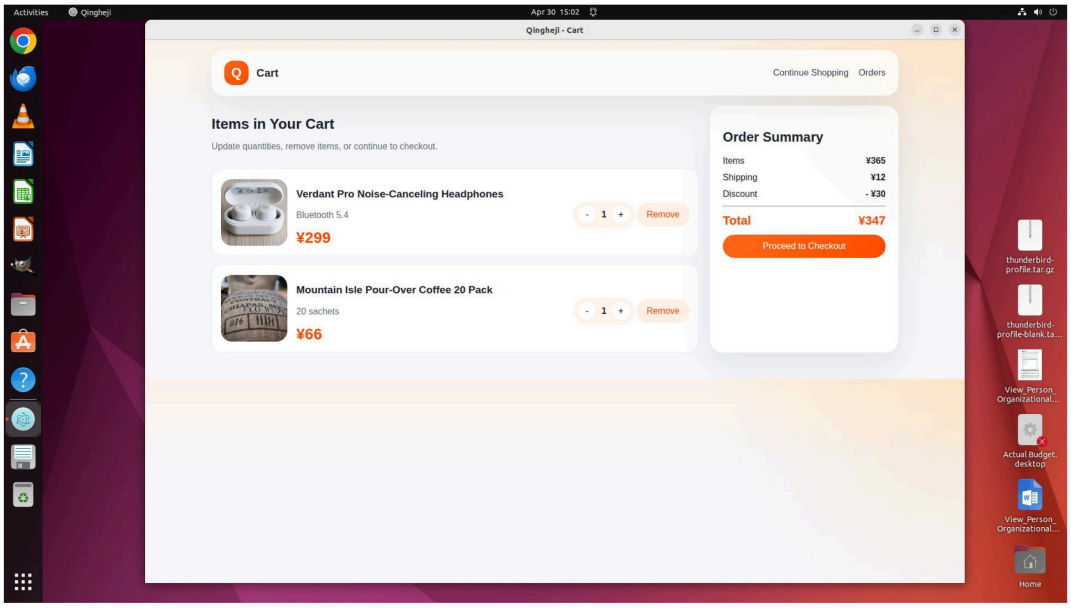

Linux-Rhythmbox

Task

task_id: linux.rhythmbox.import_and_play_track
Platform: Linux
App: Rhythmbox

Task instruction

In Rhythmbox, import and play the target local music track:
.....
- Import the folder `Music > BenchVault > North Africa Sampler > Sahara Set`.
- Use search to find `sahara_dawn`.
- Play `sahara_dawn`.

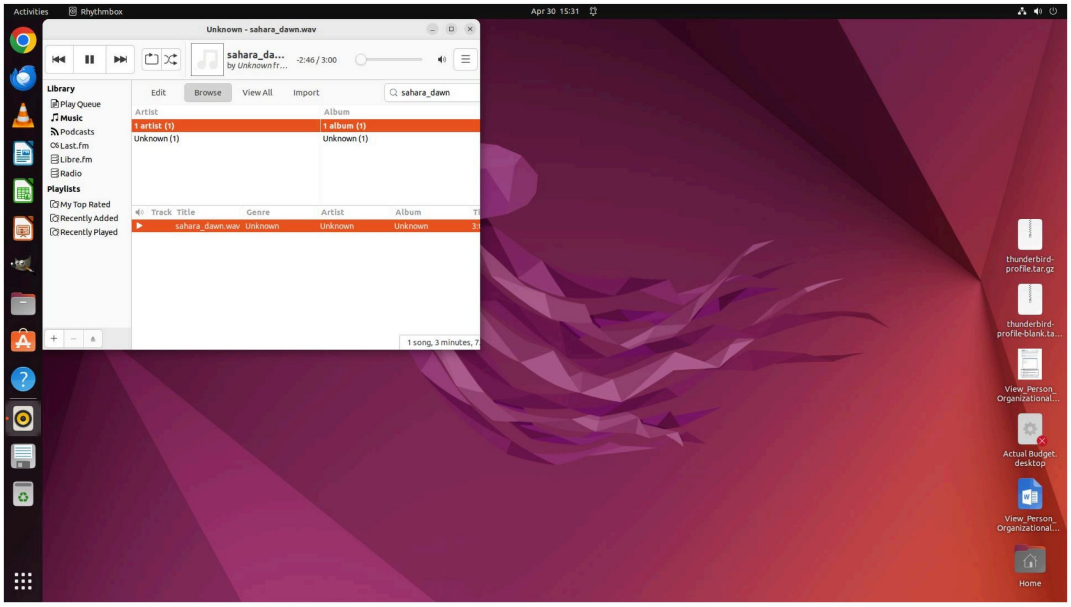

Figure 15: Representative Ubuntu/Linux task examples.

Back to list of figures

Linux-Thunderbird

Task
task_id:linux.thunderbird.thunderbird_enable_dark_mode
Platform: Linux
App: Thunderbird

Task instruction
In Thunderbird, enable full dark mode for late-night use:
- Open Thunderbird.
- Navigate to Thunderbird's theme/add-ons or appearance controls.
- Enable a complete dark mode theme.
- Ensure Thunderbird is using a dark theme persistently.

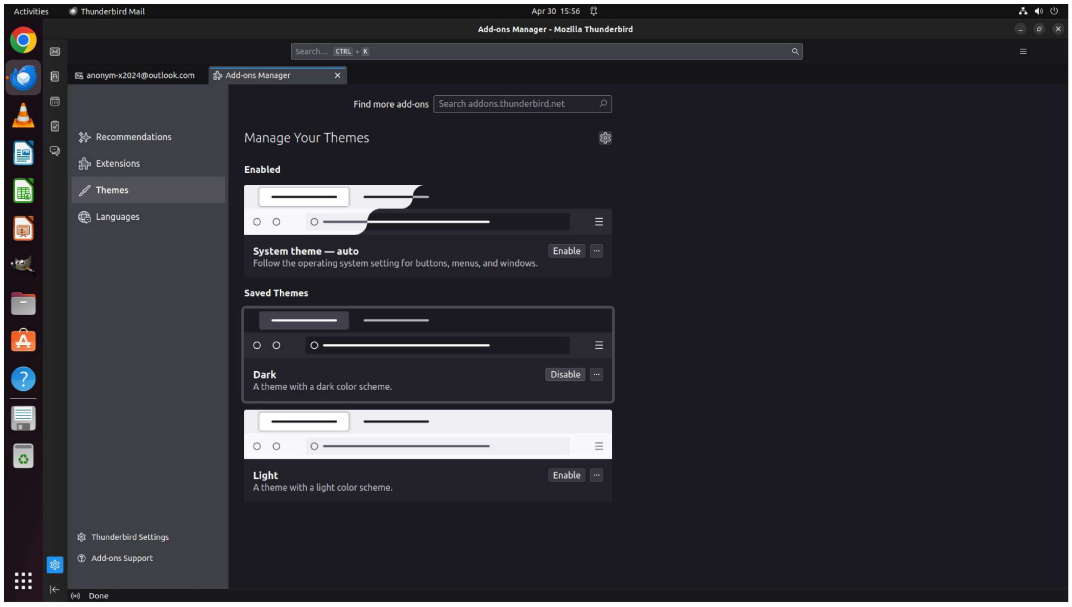

Linux-Zotero

Task
task_id:linux.zotero.research_collection_with_notes
Platform: Linux
App: Zotero

Task instruction
In Zotero, build a research library with collections, manual references, tags, notes, and handoff files:
- Create `22` collections under My Library, including `Research`, ·····
- In each collection, add the requested manual item, such as: ·····
Krizhevsky`
·····
- Finish after the required files exist.

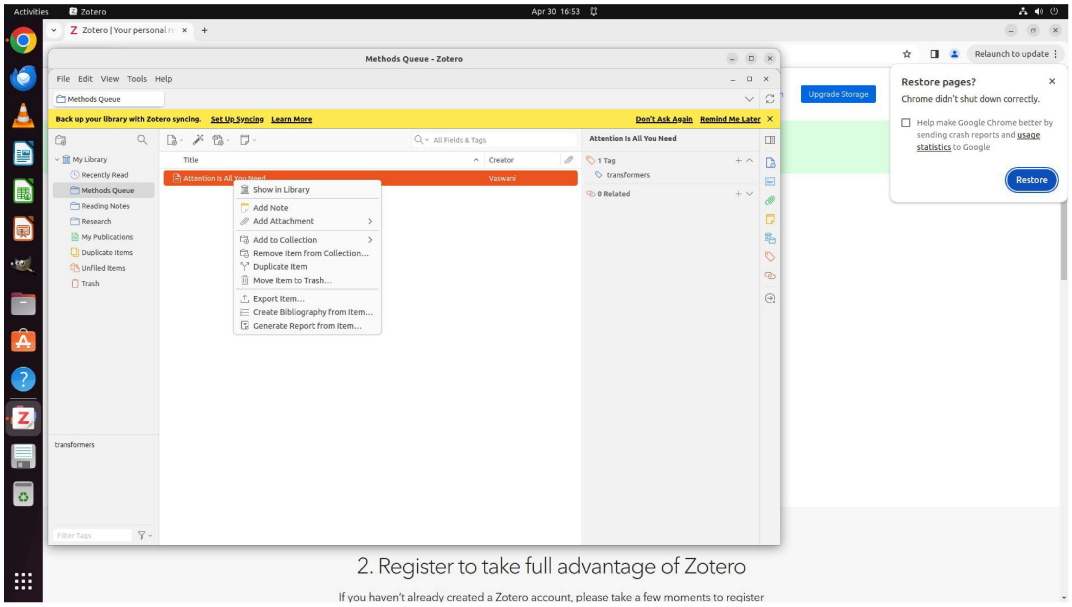

Figure 16: Representative Ubuntu/Linux task examples.

Back to list of figures

Linux-LibreOffice Writer

Task

task_id:linux.cross_app.writer_save_copy_export_attach_draft_hard
Platform: Linux
App: LibreOffice Writer, Thunderbird

Task instruction

In LibreOffice Writer, the document `Vendor Review Notes` is already open. Build the supplier summary attachment package and leave the Thunderbird draft unsent:
.....
- Save a Desktop copy named `Bench_Supplier_Summary.docx`.
.....
- Leave the Thunderbird draft open without sending it.

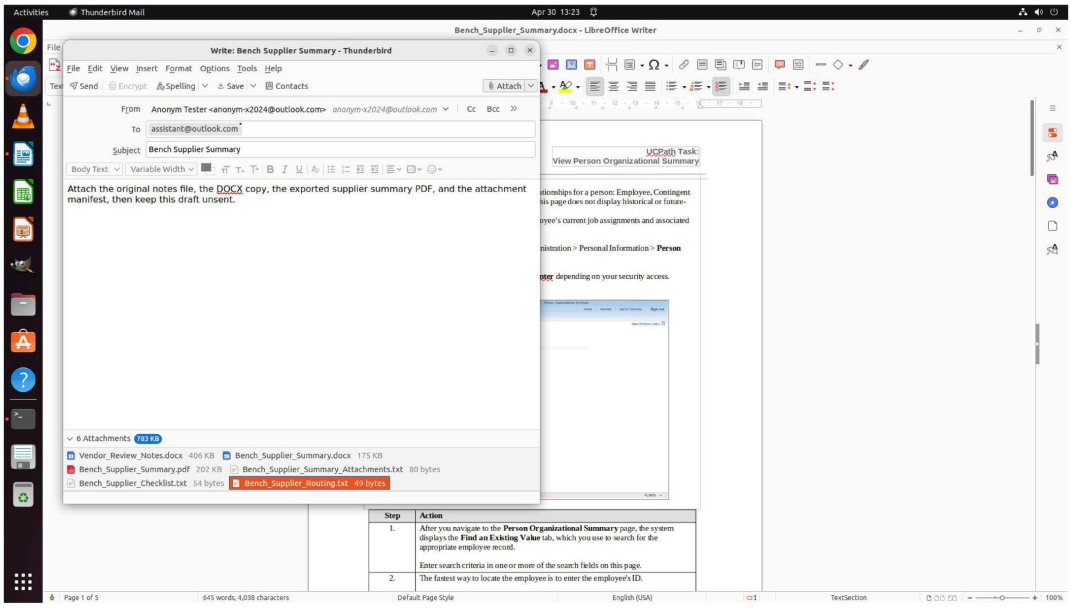

Linux-draw.io

Task

task_id:
linux.drawio.launch_badge
Platform: Linux
App: draw.io

Task instruction

In draw.io, the blue-green launch badge diagram is already open. Update the diagram and save it:
- Locate the large center label currently reading `Draft Label`.
.....
- Save the diagram after making the label change.

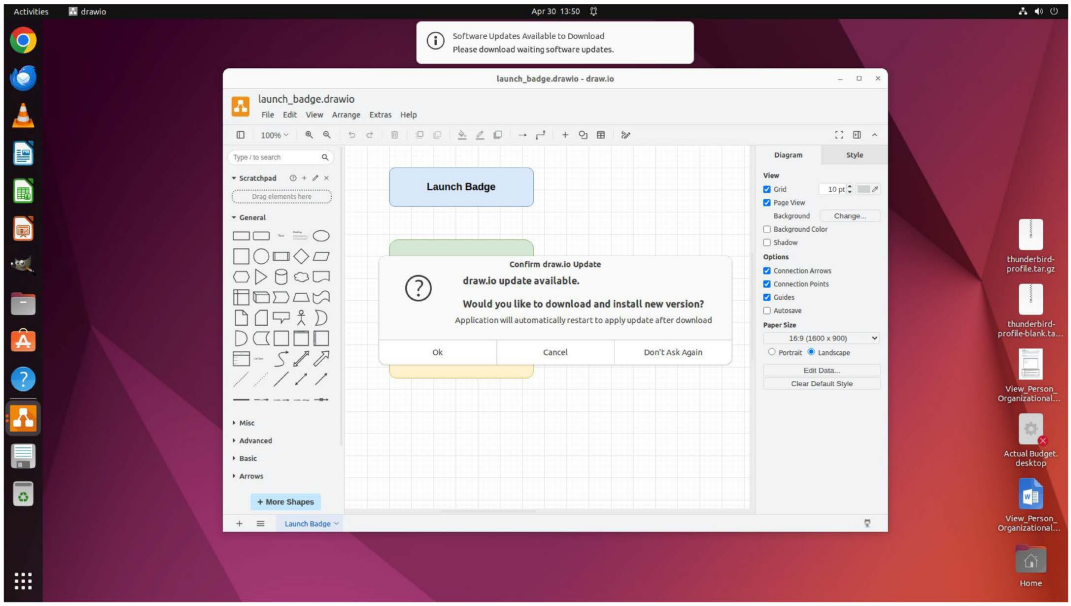

Figure 17: Representative Ubuntu/Linux task examples.

Back to list of figures

Linux-Wireshark

Task
task_id:linux.wireshark.export_protocol_bundle
Platform: Linux
App: Wireshark

Task instruction
In Wireshark, use the seeded retail capture to create a protocol export bundle:
- Export displayed DNS packets to `/tmp/helloworld_wireshark/dns_only.pcapng`.
……
- Save every export as `pcapng`.
……
- Preserve the `/checkout` request URI and the `200` response code in the combined HTTP export.

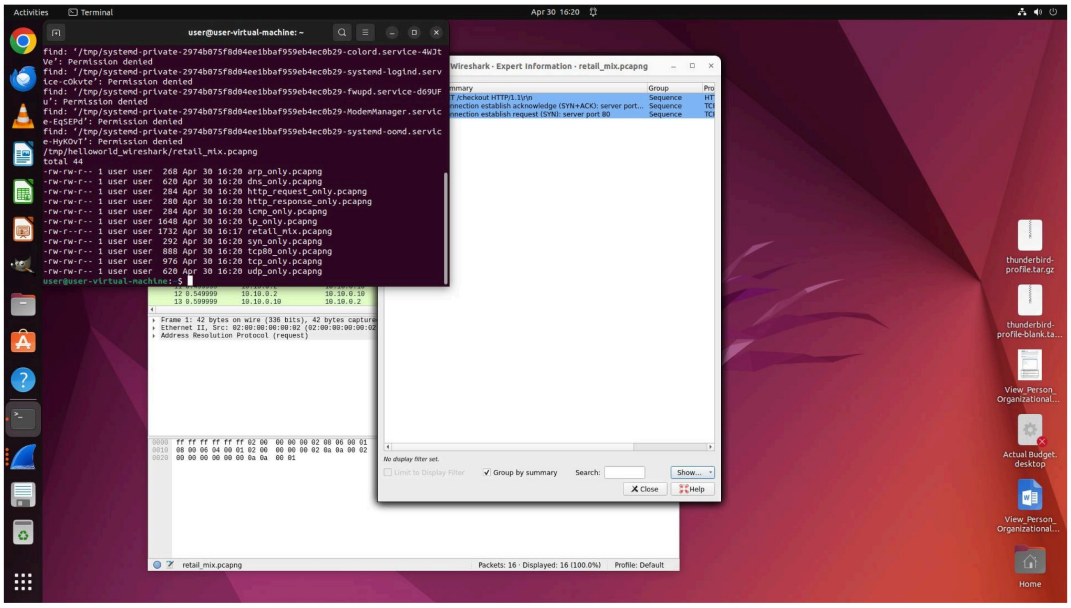

Linux-Gedit

Task
task_id:
linux.Gedit.launch_note
Platform: Linux
App: Gedit

Task instruction
In Gedit, edit the prepared launch note:
…
- Open `notes/launch-note.txt`.
- Change the first line from `Launch note draft` to `Launch note final`.
- Leave the remaining note content unchanged.
- Save the file.

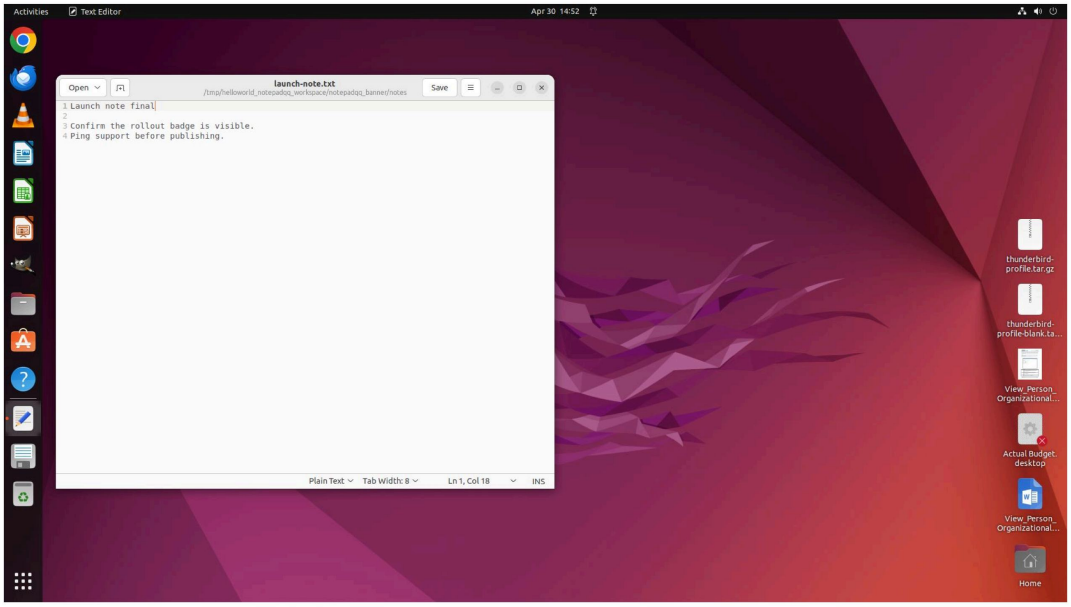

Figure 18: Representative Ubuntu/Linux task examples.

Back to list of figures

Windows-OpenShop

Task

task_id: OpenShop.login_20260503
Platform: Windows
App: OpenShop

Task instruction

Open **OpenShop** and sign in with phone number 13800138000 and password spring2026. After signing in, build a two-item starter cart. First return to the catalog, search "Mountain Isle Pour-Over Coffee", choose Grocery and Price Low to High, open Mounta ... 2, 1, 3, 1; phone stand 2, 1. Stop only when the signed-in cart contains exactly these two products, each at quantity 1.

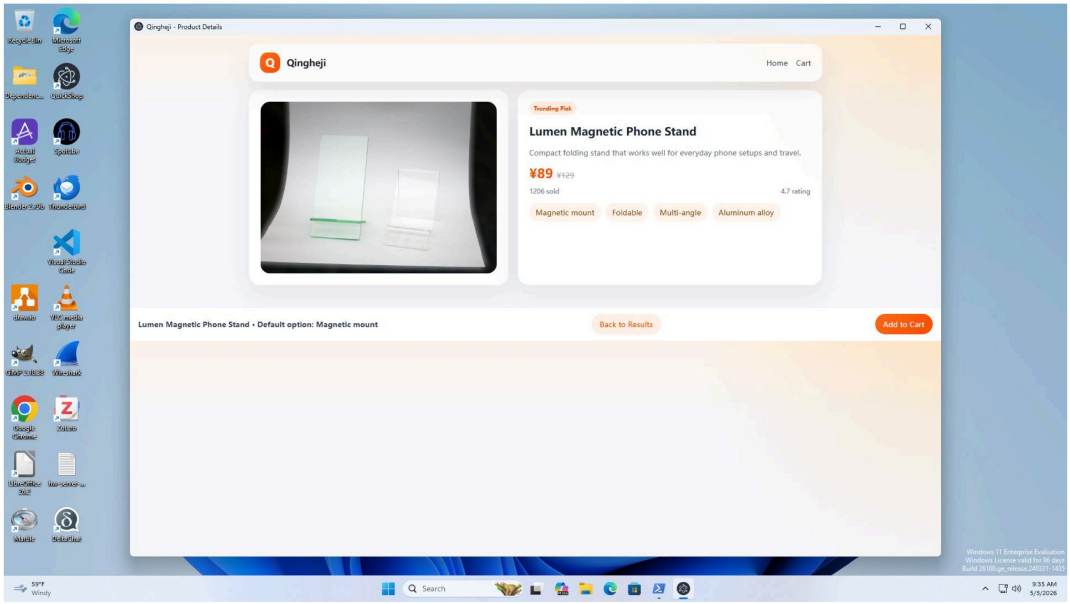

Windows-Spotube

Task

task_id: spotube.get_started_market_canada_disable_endless
Platform: Windows
App: Spotube

Task instruction

In Spotube, the getting-started flow is already open on the Language & Region step. Change Marketplace Region to Canada (CA). Then go to the next step, turn Endless Playback off, continue to the metadata provider/plugins page, and leave Spotube open there.

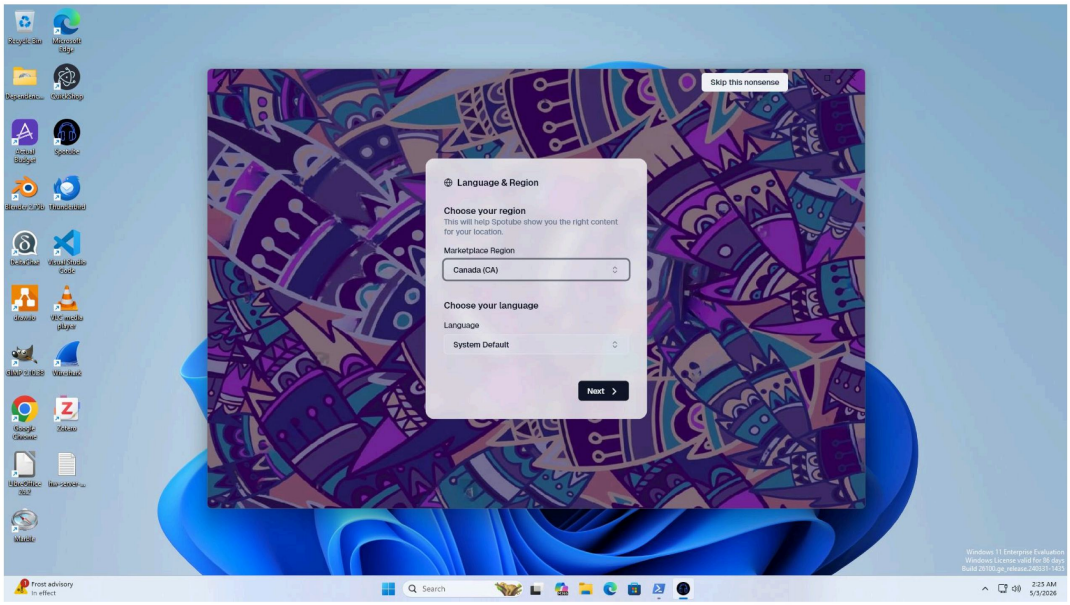

Figure 19: Representative Windows task examples.

Back to list of figures

Windows-Thunderbird

Task
task_id: thunderbird.create_contact_send_by_nickname_20260503
Platform: Windows
App: Thunderbird

Task instruction
In Thunderbird, create a new contact with the full name "Anonymous Helloworld", nickname "hello2026", email address "anonym-x2026@outlook.com", and address "70 Washington Square S, New York, NY 10012, USA". Then compose and send an email to hello2026. When addressing the email, use only the nickname "hello2026" and do not type the email address. Use the subject "Morning" and the exact body "Good Morning".

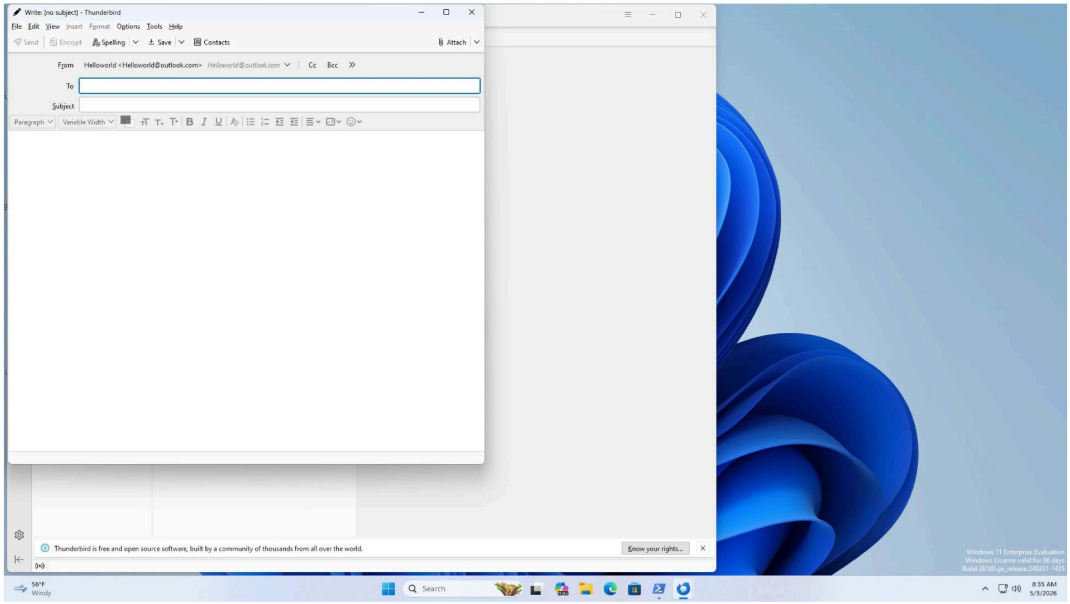

Windows-Zotero

Task
task_id: zotero.add_item_by_identifier_20260503
Platform: Windows
App: Zotero

Task instruction
In Zotero, use the "Add Item by Identifier" feature (the magic wand icon in the toolbar) to add a new item by entering the DOI "10.1038/nature14539". Confirm the lookup and verify the new entry appears in My Library.

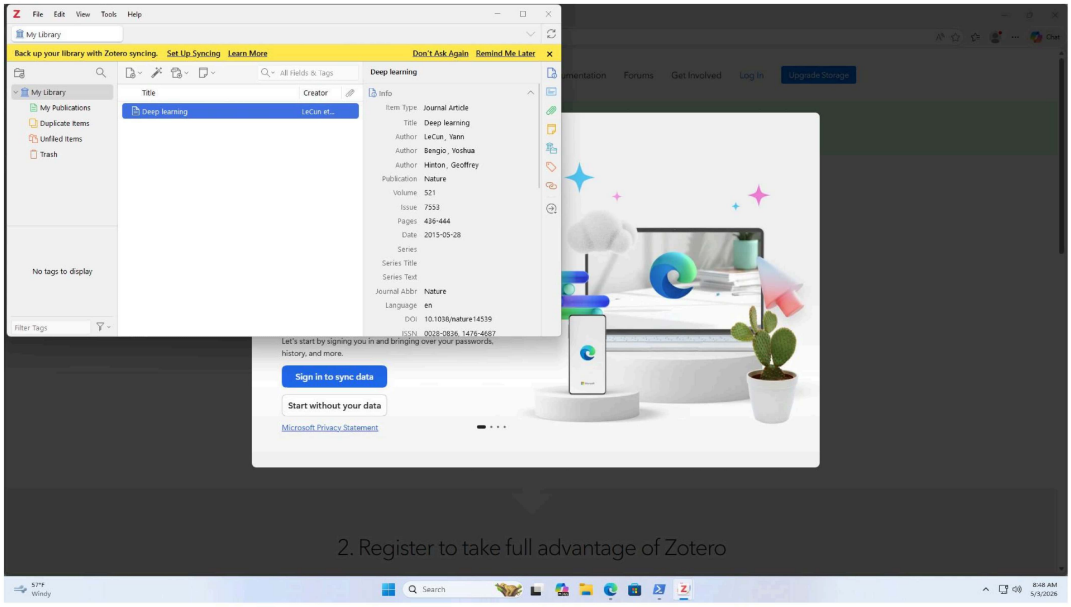

Figure 20: Representative Windows task examples.

Back to list of figures

Windows-Actual Budget

Task
task_id:
windows.actual_budget.create_account
Platform: Windows
App: Actual Budget

Task instruction
In Actual Budget, create a new bank account:
- Open Actual Budget.
- Create a new bank account called \"Checking Account\" with an initial balance of $5,000.

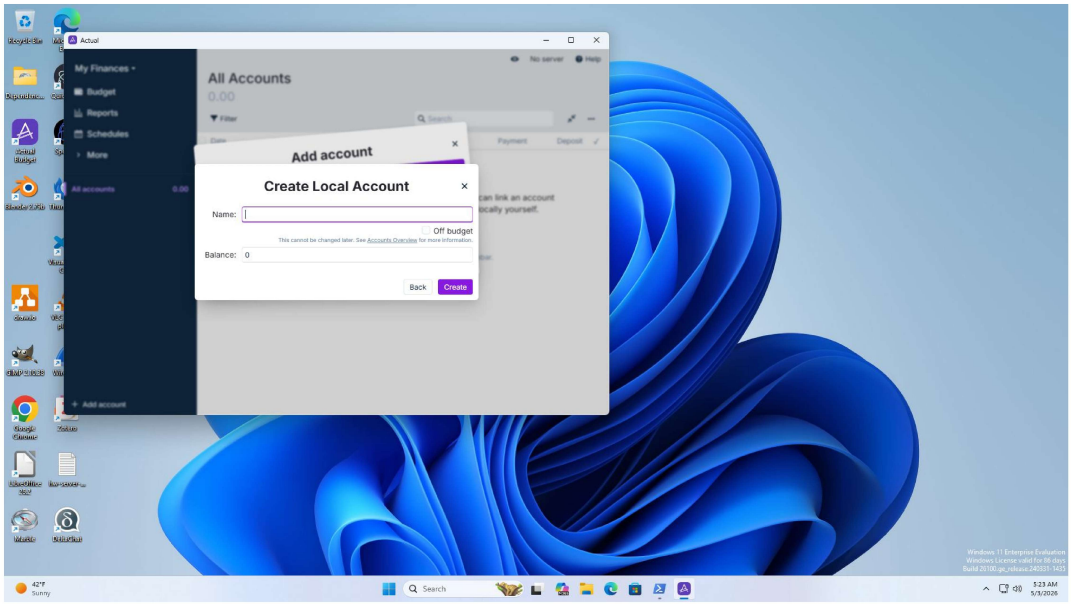

Windows-Blender

Task
task_id:
blender.marker_cube_scene
Platform: Windows
App: Blender

Task instruction
In Blender, starting from the default startup scene:
- Dismiss the splash screen.
- Rename the default cube to "MarkerCube".
- Set MarkerCube location to X = 2, Y = -1, Z = 1.
- Save the currently opened scene.

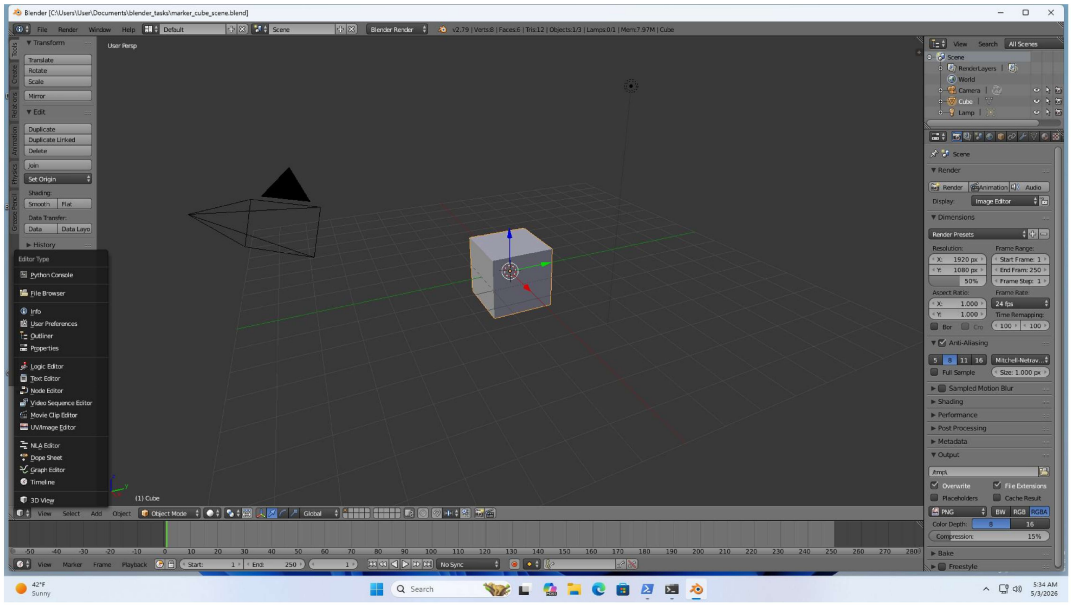

Figure 21: Representative Windows task examples.

Back to list of figures

Windows-Calc

Task
task_id:
windows.calc.freeze_filter
Platform: Windows
App: Calc

Task instruction
A sales data spreadsheet is open in LibreOffice Calc:
- Freeze the header row so it stays visible when scrolling, using A2 as the freeze point so row 1 remains frozen.
- Enable auto-filter on the data columns.
- Save when done.

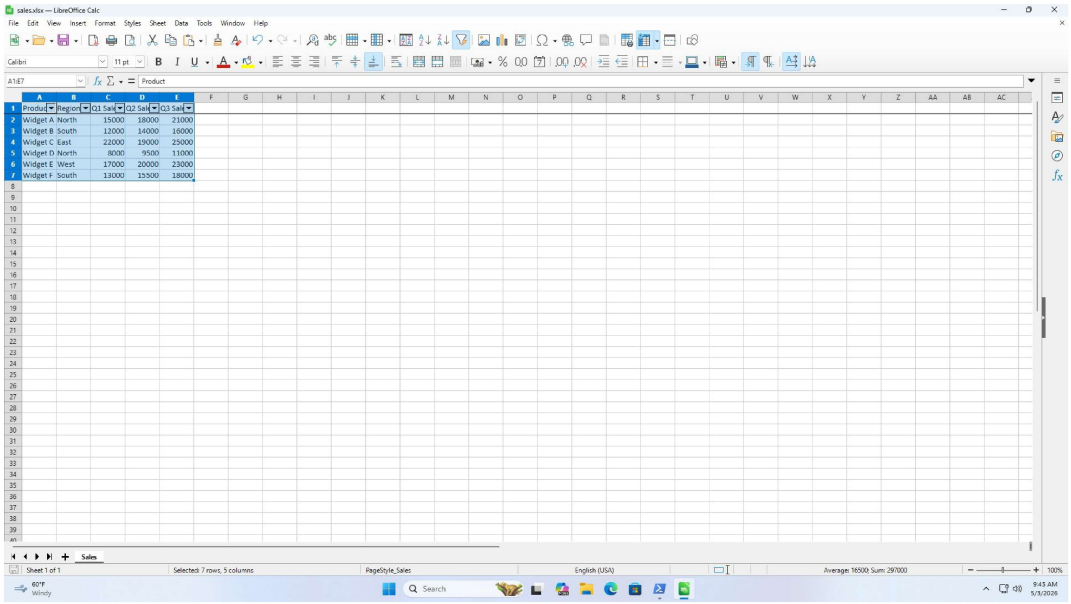

Windows-Calendar

Task
task_id:
windows.calendar.create_event
Platform: Windows
App: Calendar

Task instruction
Open Thunderbird and switch to the Calendar tab:
- Create a new event called "Team Meeting" scheduled for tomorrow at 2:00 PM, lasting 1 hour.
- Save the event.

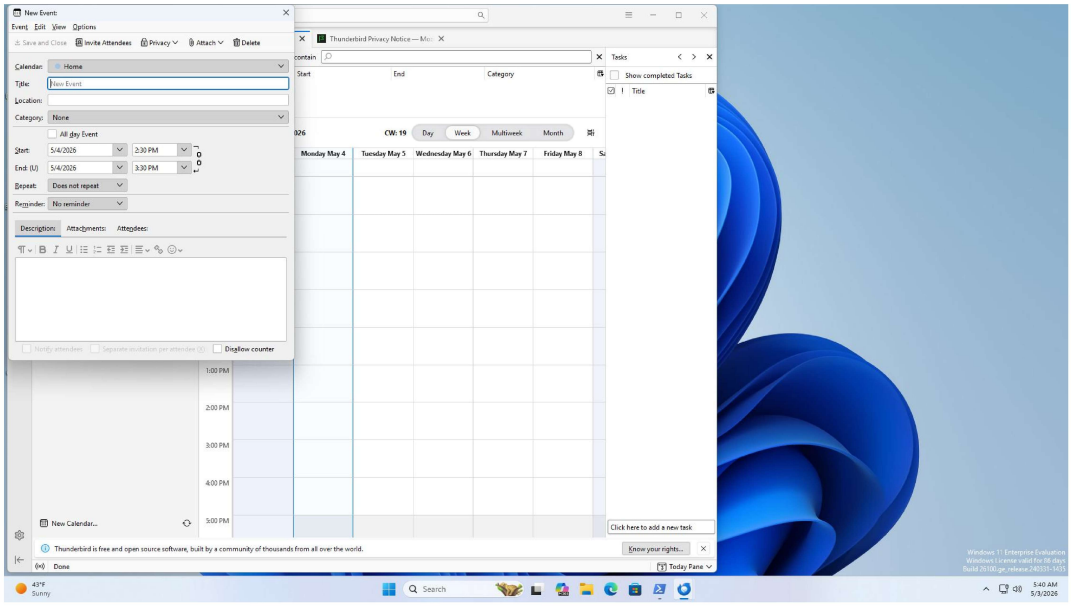

Figure 22: Representative Windows task examples.

Back to list of figures

Windows-Wireshark

Task

task_id:wireshark.open_and_filter_20260503
Platform: Windows
App: Wireshark

Task instruction

In Wireshark (already showing "C:\Users\User\Downloads\sample.pcapng"), apply the display filter "http" and save the displayed (filtered) packets to a new file called "http_only.pcapng" in the same Downloads folder.

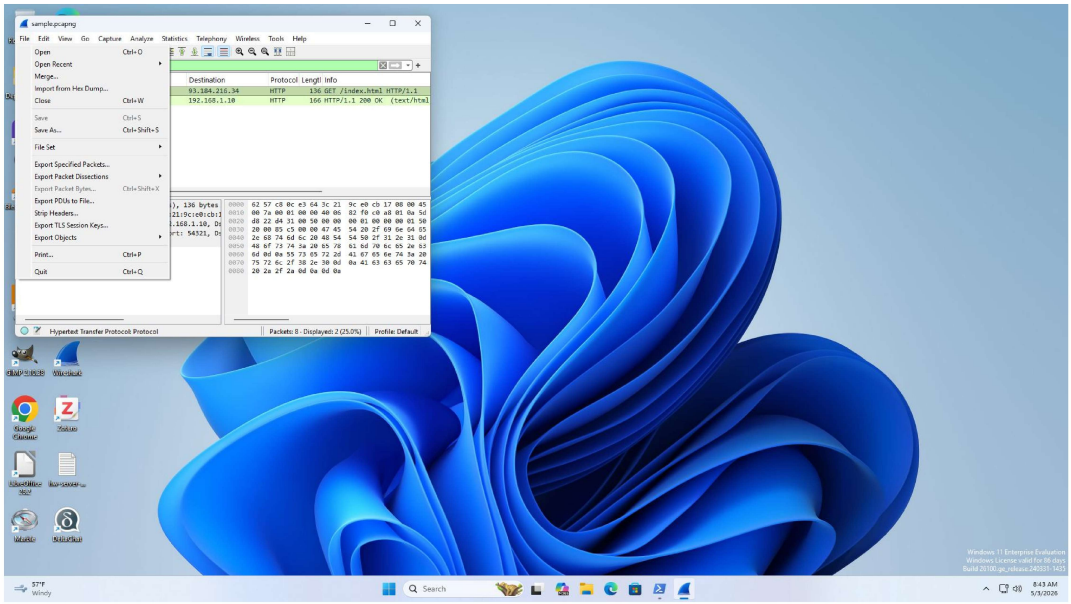

Windows-VS Code

Task

task_id:vscode.disable_minimap_enable_autosave_20260503
Platform: Windows
App: VS Code

Task instruction

In Visual Studio Code, open the user settings JSON and update it so that the editor minimap is disabled and Auto Save is set to afterDelay. Save the settings file and leave VS Code open. You must use Visual Studio Code only for all file creation and editing; do not use Notepad or any other editor.

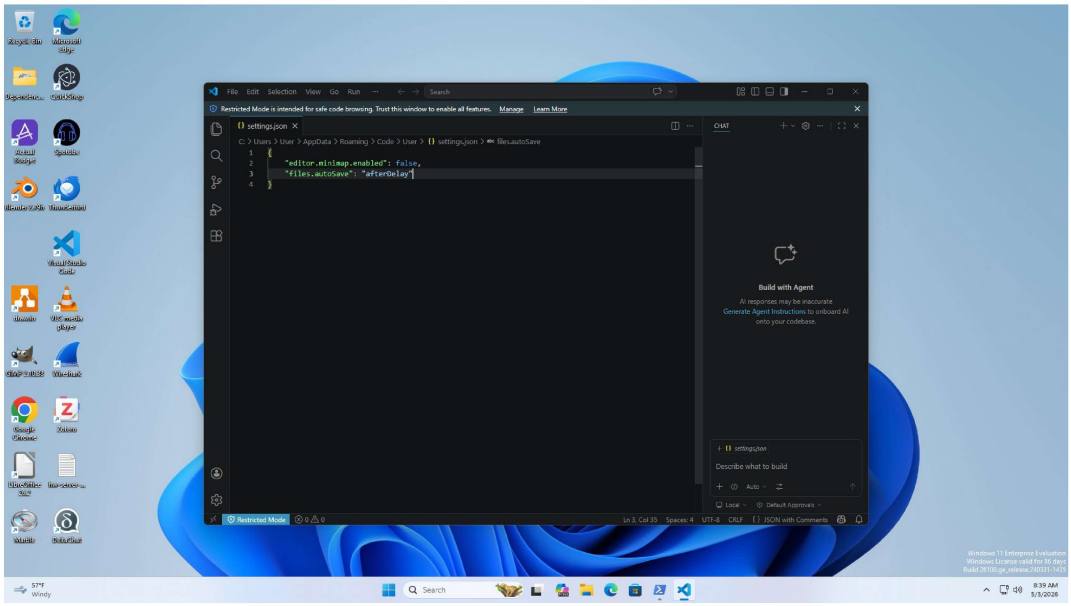

Figure 23: Representative Windows task examples.

Back to list of figures

Windows-Drawio

Task
task_id:windows.drawio.square_center_helloworld
Platform: Windows
App: Drawio

Task instruction
In the draw.io desktop app, use the blank diagram that is already opened from C:\\Users\\User\\Documents\\drawio_square_center_helloworld.drawio:
- Add exactly one square.
- Insert the text \"Helloworld\" and make sure the text is centered in the square.
- Save the currently opened diagram and leave the finished file open.

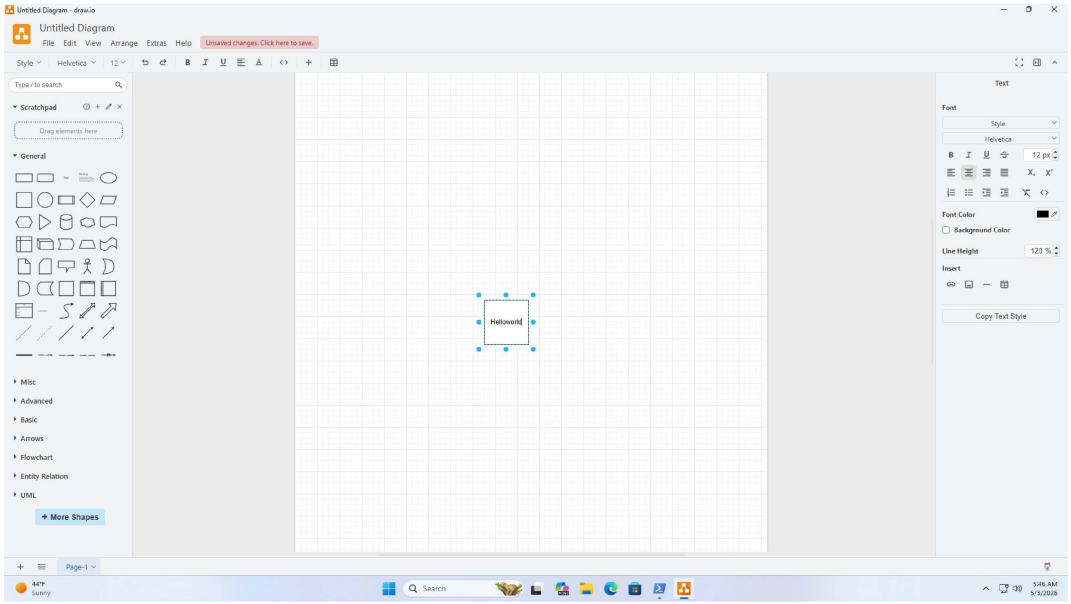

Windows-Gimp

Task
task_id:
windows.gimp.center_helloworld
Platform: Windows
App: Gimp

Task instruction
In GIMP, create a new default image.:
- Add the text "Helloworld" and center it on the canvas.
- Export the finished image to C:\Users\User\Downloads\gimp_center_helloworld.png, and leave the finished image open.

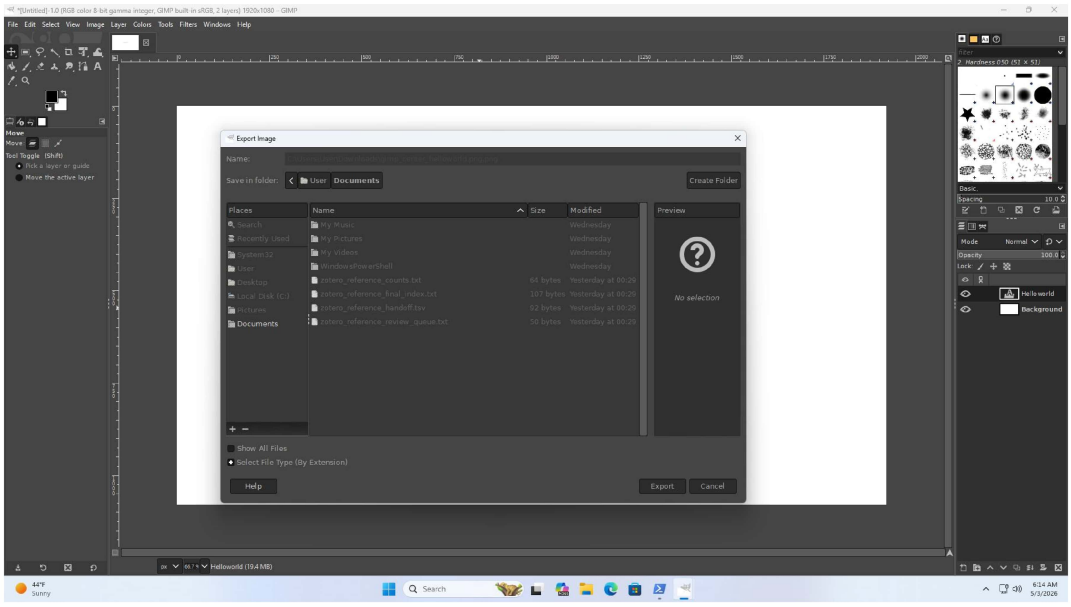

Figure 24: Representative Windows task examples.

Back to list of figures

Windows-Impress

Task
task_id:
windows.impress.title_notes
Platform: Windows
App: Impress

Task instruction
A presentation is open in LibreOffice Impress:
- Change the title on slide 1 from "Kickoff" to "Project Kickoff Meeting".
- Add a speaker note on slide 2 with the text "Remember to demo the live system" .
- Save when done.

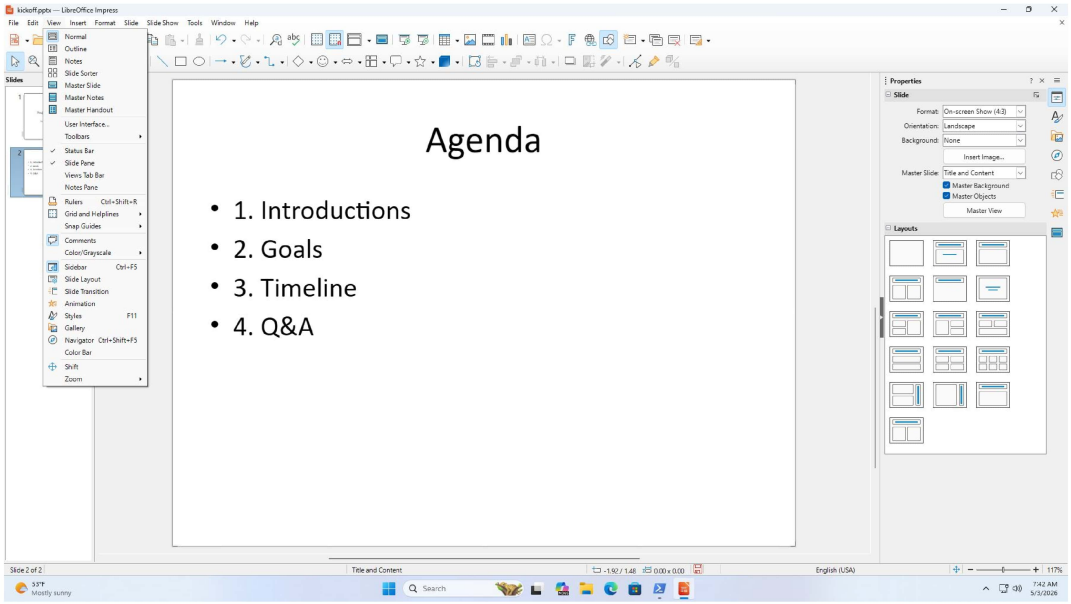

Windows-Notepad

Task
task_id:
windows.notepad.find_replace
Platform: Windows
App: Notepad

Task instruction
A server list named hosts.txt is open in Notepad:
- Use the Replace feature to change all occurrences of the word "production" to "staging".
- Navigate to the end of the file and manually append 8 new server entries on separate lines, from "10.0.0.21 node21" up to "10.0.0.28 node28".
- Add these 2 final footer lines in this exact order: "# migration complete" and "# verified-by: ops".
- Save when done.

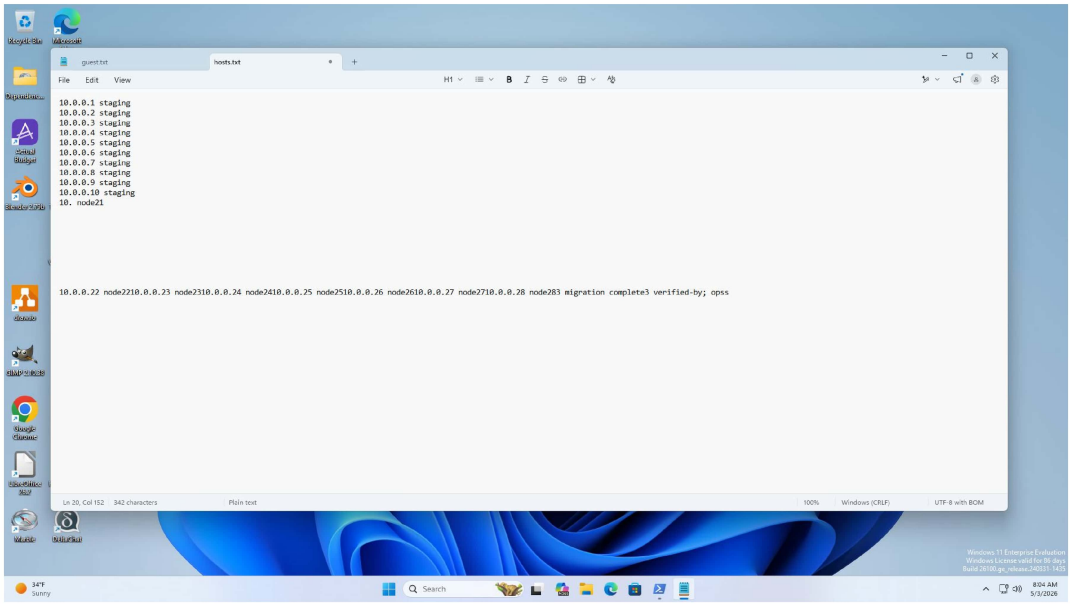

Figure 25: Representative Windows task examples.

Back to list of figures

Macos-Libreoffice.writer

Task
task_id: macos.libreoffice.writer.insert_table
Platform: Macos
App: Libreoffice.writer

Task instruction
I have a team roster document open in LibreOffice Writer.
Can you add a table at the end listing our team members?
...
The columns should be Name, Department, and Extension.
Please bold the header row. Save when done.

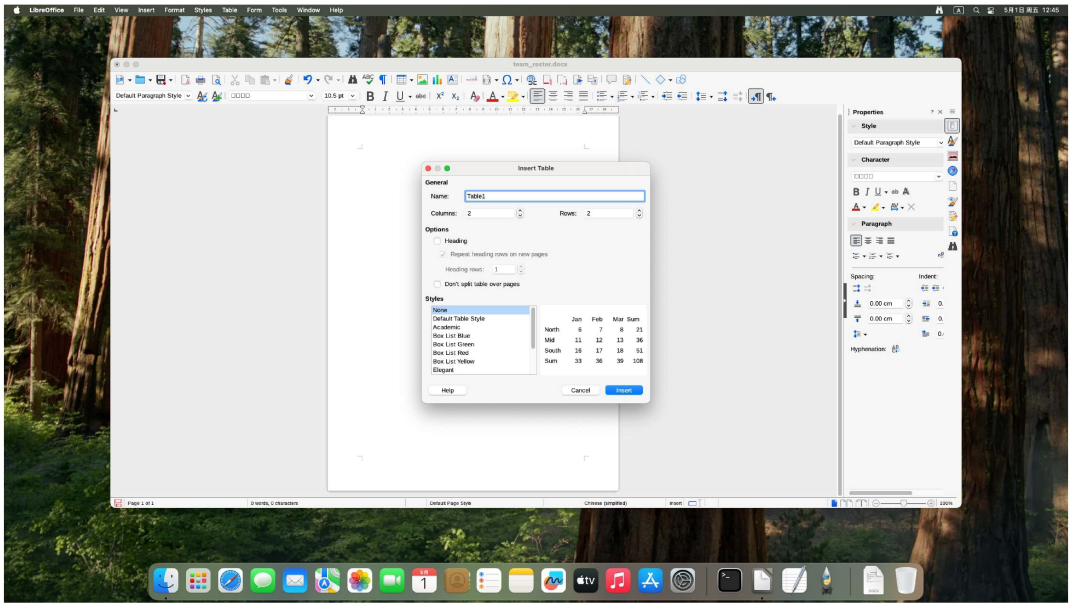

Macos-Spotube

Task
task_id: macos.spotube.light_connect_japan_profile
Platform: Macos
App: Spotube

Task instruction
Open Spotube and prepare an English light profile for Japan without using login or media search.
Complete these changes in order:
...
11. Leave Spotube open when you finish.

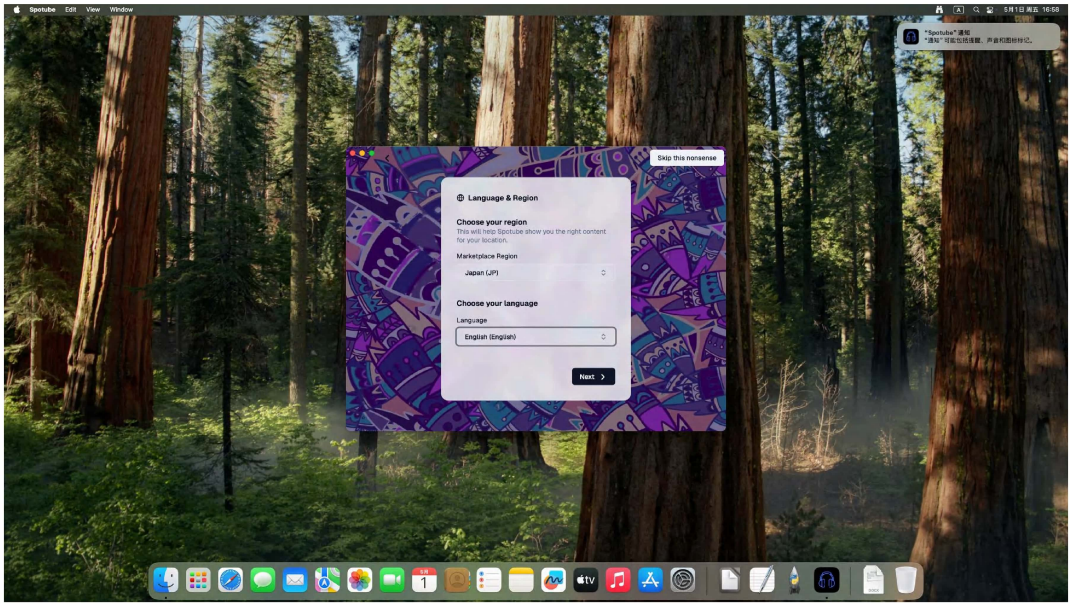

Figure 26: Representative macOS task examples.

Back to list of figures

Macos-textedit

Task
task_id:
macos.textedit.config_rewrite
Platform: Macos
App: textedit

Task instruction
TextEdit is already open with `/Users/bench/Downloads/app.ini`.
Edit the file with these exact changes:
- Delete the entire second and third lines.
...
Save when done.

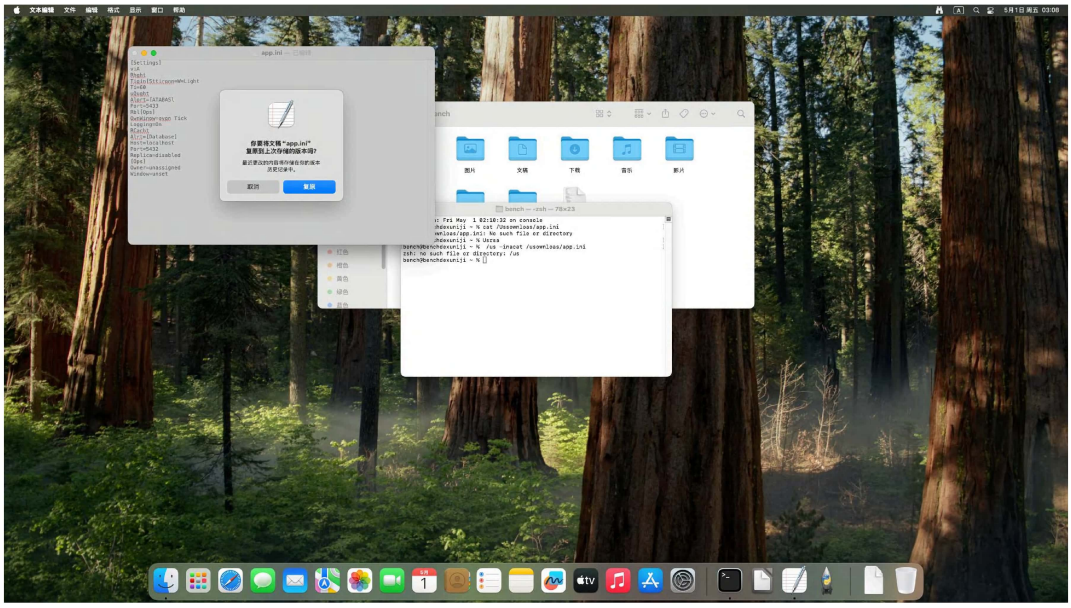

Macos-GIMP

Task
task_id:
macos.gimp.rotate_flip_export
Platform: Macos
App: GIMP

Task instruction
Apply the following edits in this exact order:
1. Rotate the image **90 degrees clockwise**.
...
8. Export the final result as a PNG file to:
`/Users/bench/Downloads/gimp_transform_input.png`

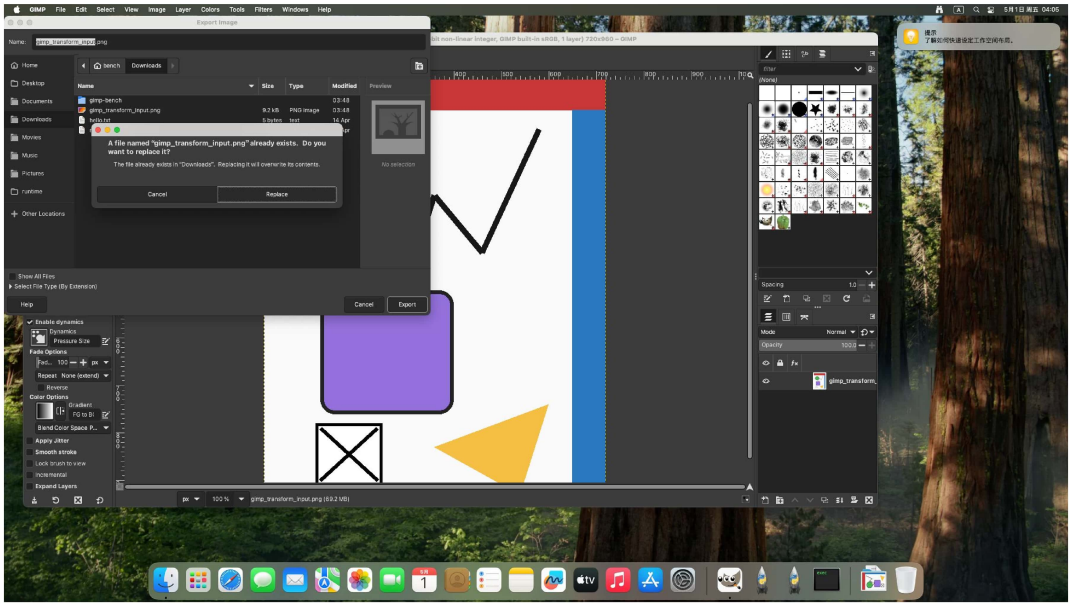

Figure 27: Representative macOS task examples.

Back to list of figures

Macos-Actual Budget

Task
task_id:macos.actual_budget.create_emergency_fund_account
Platform: Macos
App: Actual Budget

Task instruction
Open Actual Budget.
Create a local account named "Emergency Fund" with an opening balance of $8,000.
Leave the app showing the created account when you finish.

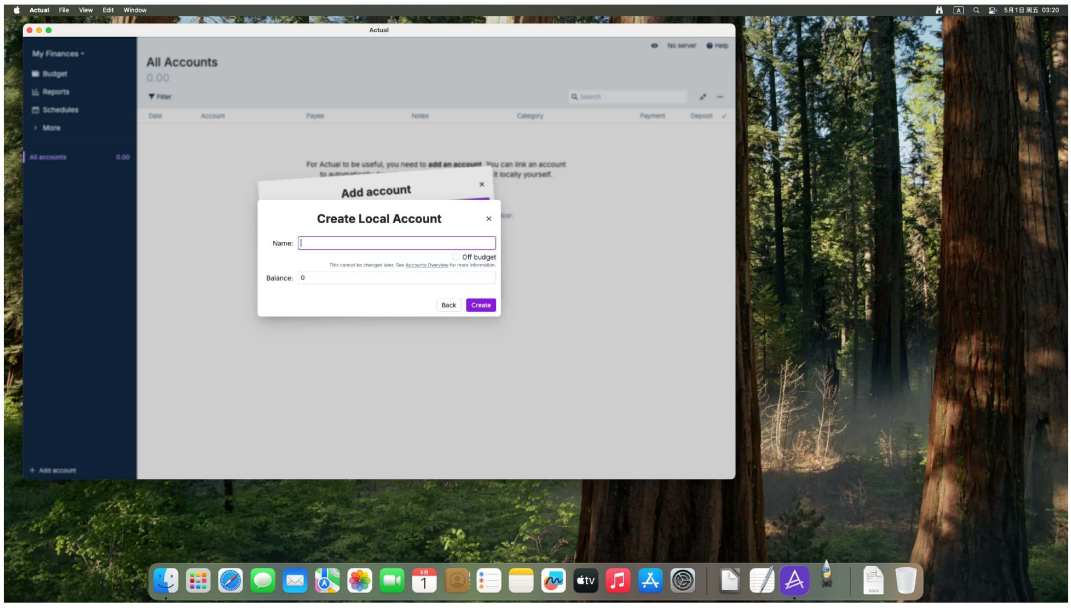

Macos-Delta Chat

Task
task_id:
macos.deltachat.support_escalation
Platform: Macos
App: Delta Chat

Task instruction
Replace each placeholder with the correct values from the conversations.
...
3. In `Ops Sprint`,send: `Escalation update: {ticket_id}, {error_code}, recovered after second launch, owner {owner}, window {deploy_window}.`

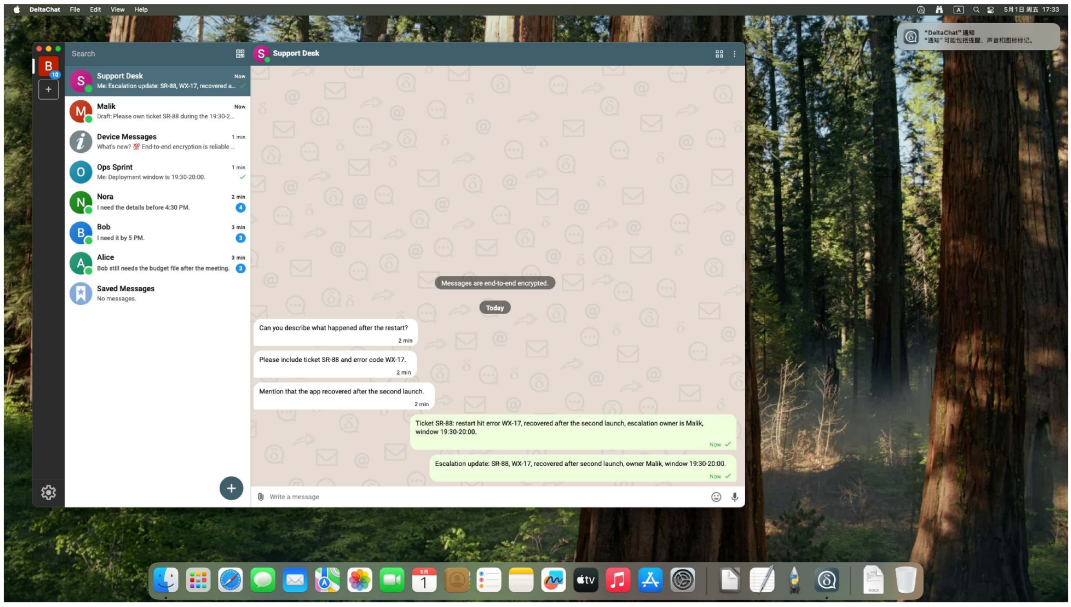

Figure 28: Representative macOS task examples.

Back to list of figures

Macos-0penShop

Task
task_id:
macos. OpenShop.checkout_two_items
Platform: Macos
App: OpenShop

Task instruction
In **OpenShop**, first sign in with phone number 13800138000 and password spring2026.
On the home page, set the category to Electronics and the sort to Price: High to Low.
...
... and place the order.

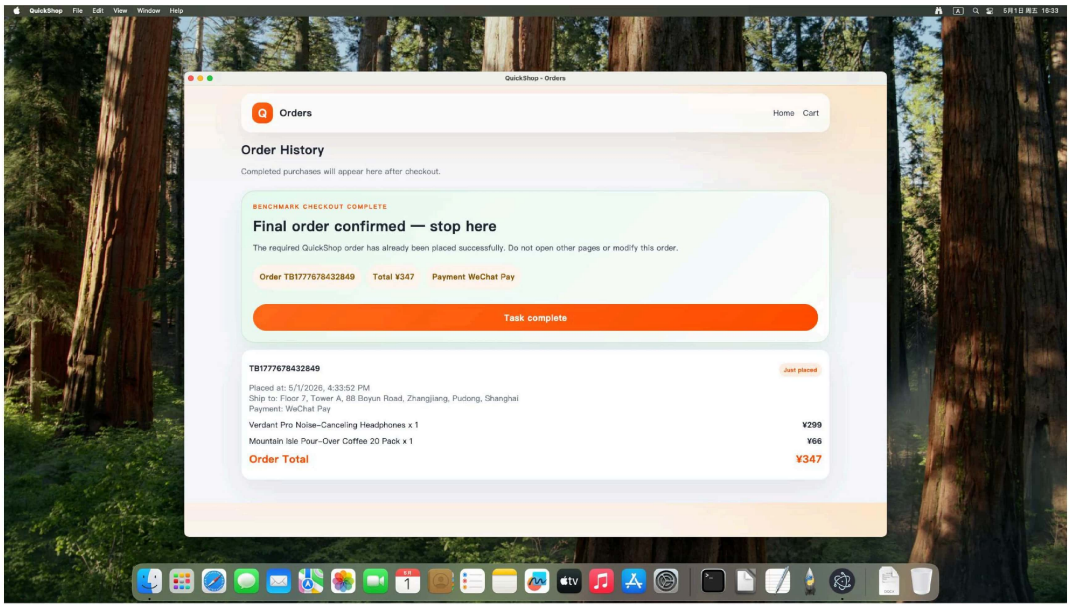

Macos-Libreoffice.calc

Task
task_id:
macos.libreoffice.calc.fill_totals
Platform: Macos
App: Libreoffice.calc

Task instruction
I opened an inventory spreadsheet in LibreOffice Calc.
...
Then complete the Grand Total row by putting the total quantity in cell B7 and the total value in cell D7.
Save the spreadsheet when finished.

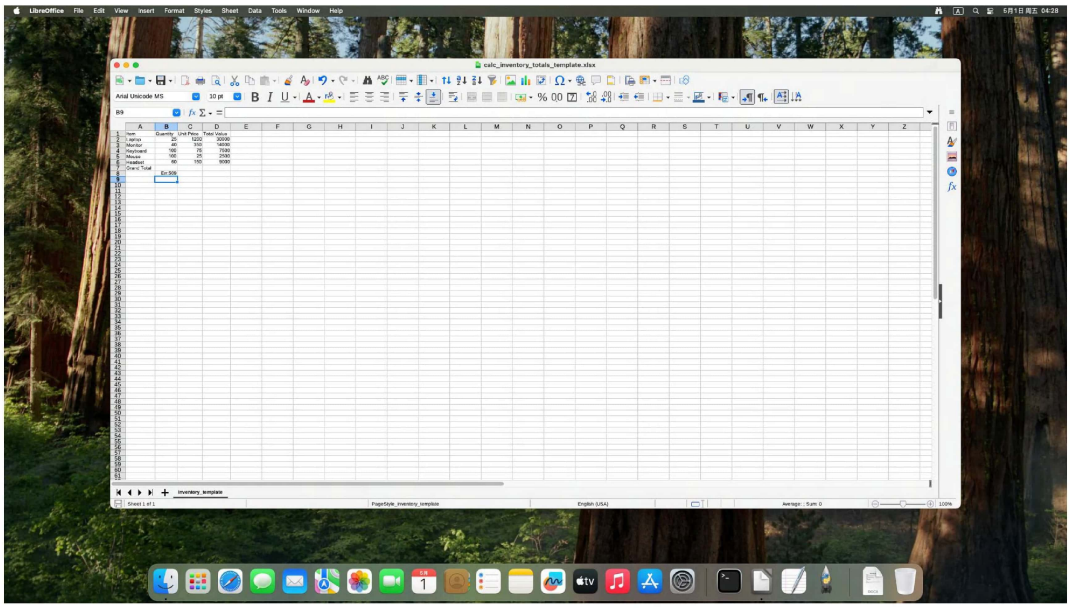

Figure 29: Representative macOS task examples.

Back to list of figures

Macos-draw.io

Task
task_id:
macos.drawio.release_board_refresh
Platform: Macos
App: draw.io

Task instruction
There are 6 rounded rectangles visible in a 2-row by 3-column grid.
Replace the text inside the rectangles exactly as follows:
- `Card 1` → `Inbox`
- `Card 4` → `QA Queue`
...then save the file.

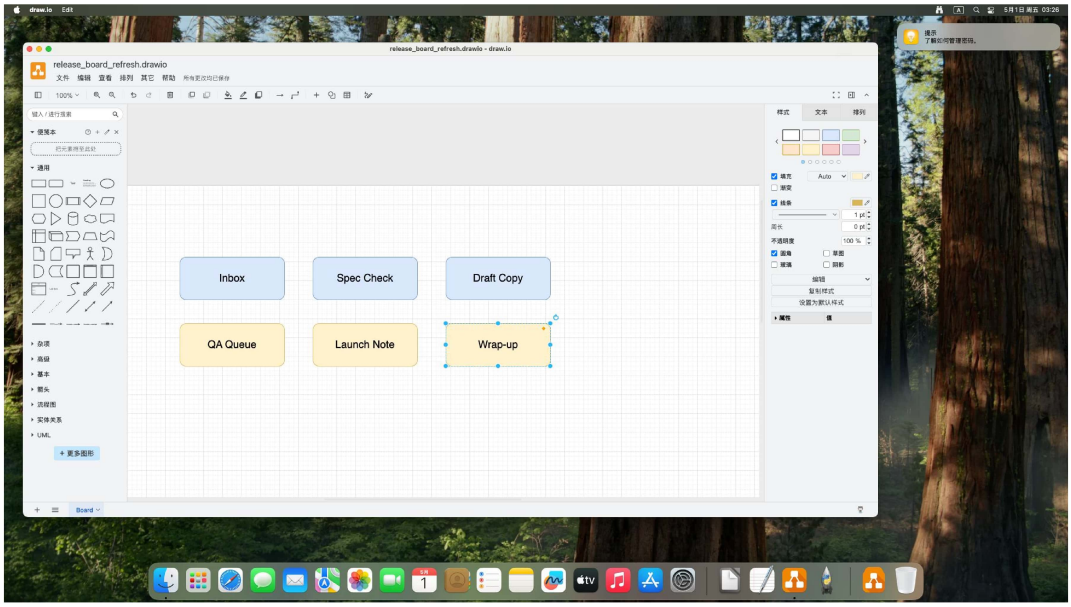

Macos-Thunderbird

Task
task_id:macos.thunderbird.calendar.
create_recurring_event
Platform: Macos
App: Thunderbird

Task instruction
Open Thunderbird and switch to the Calendar view.
Create a recurring event called "Weekly Standup" every Monday at 9:00 AM, lasting 1 hour.
...
... so the recurrence is anchored to Monday rather than the default day.

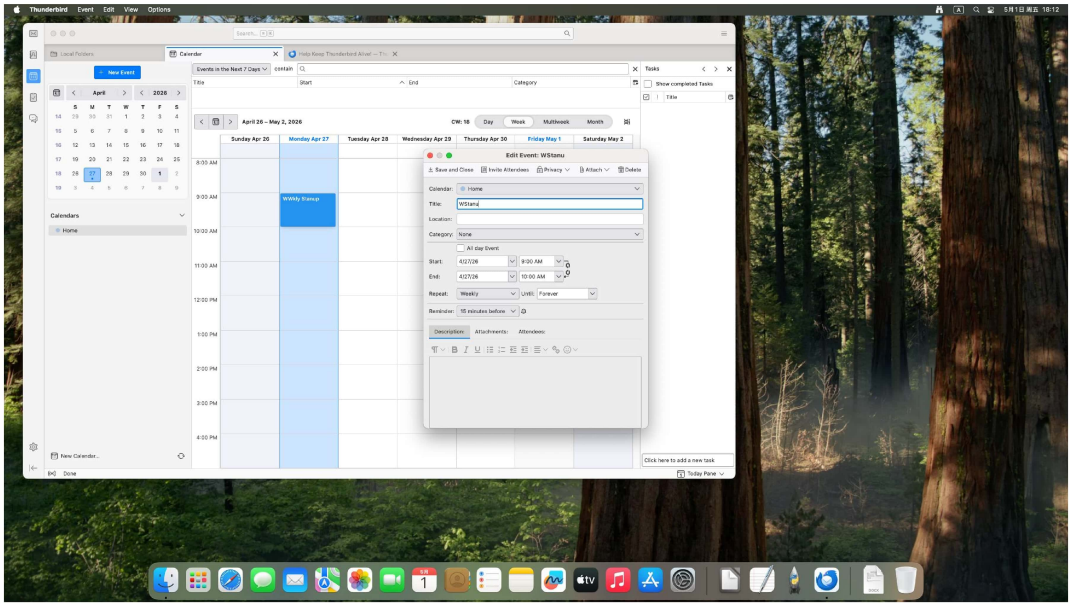

Figure 30: Representative macOS task examples.

Back to list of figures

Macos-Thunderbird.Mail

Task

task_id: macos.thunderbird.email.collect_st_club_progress_local

Platform: macOS

App: Thunderbird

Task instruction

Open Thunderbird in Mail view. Compose and send three emails to qa-local@example.com that summarize ST Club progress. The subjects and body openings are intentionally similar, so send them one by one in the exact order shown below. Send exactly three emails, do not resend any subject, and keep the blank lines exactly as written.

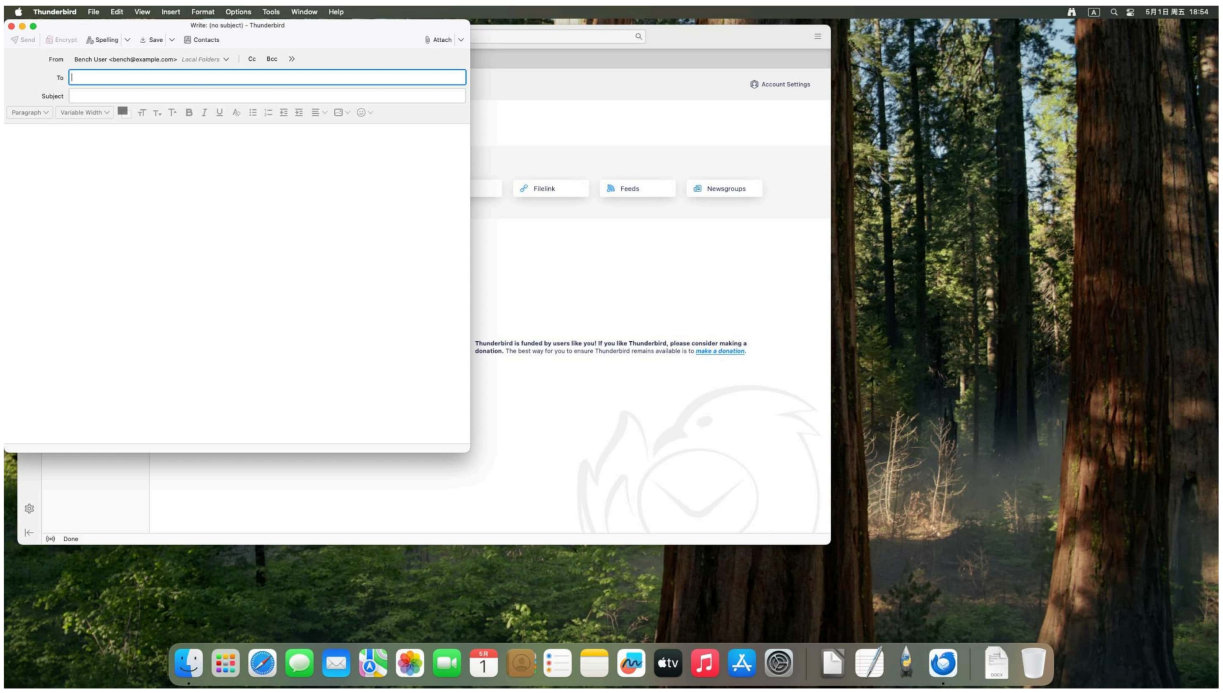

Figure 31: Representative macOS task examples.

Back to list of figures

Android OpenShop : Add Coffee to Cart

Task

task_id:
android.OpenShop.add_coffee_to_cart

Platform: Android

App: OpenShop

Task instruction

```
In  OpenShop , search for
“Mountain Isle Pour-Over
Coffee 20 Pack”.
Identify the exact target
product among nearby shopping
options,
add exactly one unit to the
cart, and ensure that no
unrelated
products are added before
stopping.
```

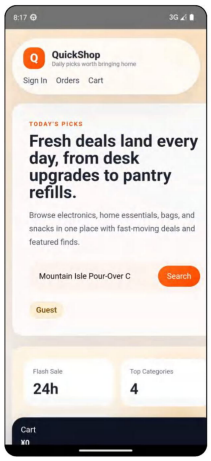

Android K9: forward meeting notes

Task

task_id:
android.k9.forward_meeting_notes

Platform: Android

App: K-9 Mail

Task instruction

```
```text
Use K-9 Mail to read the 10
inbox messages whose subjects
start with BMK Manual Mail
Packet. Each message contains
one line in the form "field:
value". Send one synthesized
email to ……test on the local
test mail server 10.0.2.2; ……
and outgoing mail is available
via SMTP on port 3025 with no
encryption.
```
```

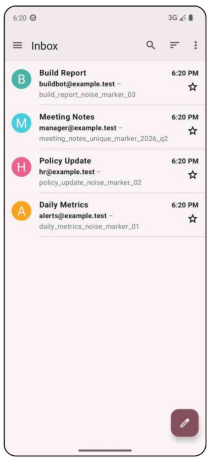

Android DeltaChat: create_contact_send_message

Task

task_id:
android.deltachat.create_contact_send_message

Platform: Android

App: DeltaChat

Task instruction

```
Open the existing Delta Chat
conversation with
benchrecv@example.test. The
incoming message is an
instruction card with DRAFT,
DECOY, and FINAL sections.
Reply only with the exact
three-line text after
FINAL_RESPONSE=, ……. Do not
start a new contact and do
not use an email app.
```

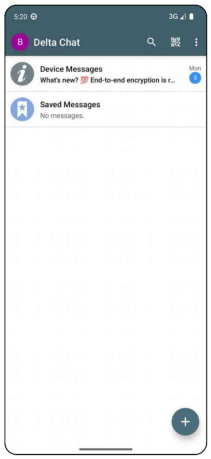

Android Fossify Calendar: Update and Follow-up

Task

task_id:
android.fossify_calendar.update_and_followup_hard

Platform: Android

App: Fossify Calendar

Task instruction

```
In Fossify Calendar, open BMK
Draft Planning and use the f-
-inal launch plan stored in
its description. Apply that
plan exactly, delete the obs-
-olete backup draft, ... lau-
-nch metrics. Add BMK Launch
Risk Sweep on 2026-05-01
16:05-16:35 at Risk Desk with
description Resolve launch
blockers. Leave BMK Reference
Hold unchanged.
```

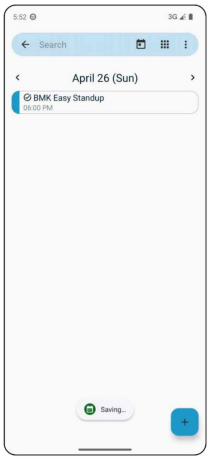

Android Fossify Gallery: Delete Photo

Task

task_id:
android.fossify_gallery.delete_photo

Platform: Android

App: Fossify Gallery

Task instruction

```
In Fossify Gallery, delete the
blue photo from the camera
album while keeping the red
camera photo and the green
receipts photo.
```

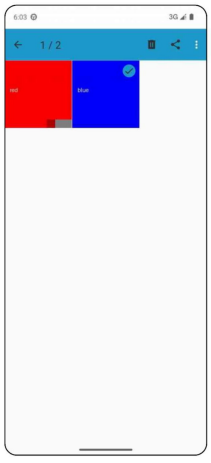

Android Libreoffice: notes append

Task

task_id:
android.libreoffice.notes_append_hard

Platform: Android

App: Libreoffice

Task instruction

```
Open notes.odt in LibreOffice,
go to the very end, and append
only the FINAL checklist, not
the archival sections. The
exact twelve new lines must be
Final 01 - Task A completed,
……
scheduled, Final 08 - Device
handoff confirmed, Final 09 -
Metrics snapshot attached,
Final 10 - Rollback owner Mei,
…… and make sure each final
line appears exactly once.
```



Figure 32: Representative Android task examples.

Back to list of figures

Android Spotube: Queue Naina Zara

Task

task_id:
android.spotube.queue_naina_zara

Platform: Android

App: Spotube

Task instruction

```
In Spotube, make Jiosaavn
Audio the selected audio
source, set Language to
English and Marketplace Region
to Austria, then build one
queue in this exact order as
closely as the app allows:
Naina, Zara Zara, ……. Do not
stop after only Naina and Zara
Zara, do not use a YouTube-
only source, and do not leave
the final screen on Search or
Settings.
```

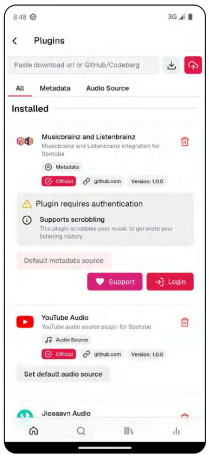

Android Tasks Org. : Create Task With Subtasks

Task

task_id:android.tasks_org.create_task_with_subtasks

Platform: Android

App: Tasks Org.

Task instruction

```
In Tasks.org, the list already
contains seeded weekend-errand
decoys. Delete exactly the
stale top-level task Weekend
errands - duplicate charger
draft.
……
bench_weekend_errands_note_v5
/ route=blue owner=family
window=Sat-09:30 / ……
Do not create duplicate
subtasks, reminders, completed
tasks, or any other top-level
task.
```

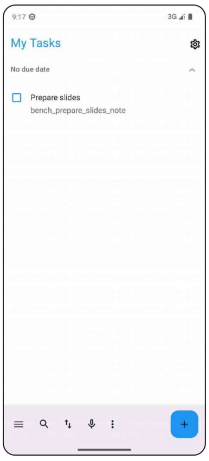

Android Material Files: Rename

Task

task_id:
android.material_files.rename_file

Platform: Android

App: Material Files

Task instruction

```
In Material Files, open
incoming/
rename_instructions.txt inside
the benchmark files workspace
and apply only the FINAL BLUE
batch. Ignore the DRAFT, HOLD,
BACKUP, ……
 keep the remaining final
files inside incoming, and
delete rename_instructions.txt
when finished. Do not delete,
rename, or move any hold/
backup/draft file.
```

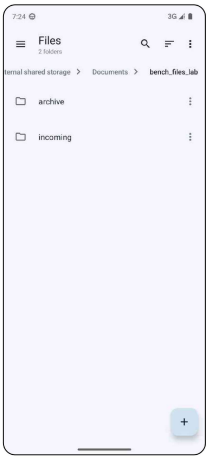

Android Fossify Calendar: Reorganize schedule

Task

task_id:android.fossify_calendar.reorganize_schedule_hard

Platform: Android

App: Fossify Calendar

Task instruction

```
Please clean up the BMK
schedule in Fossify Calendar.
The event BMK Candidate Alpha
contains the final routing
plan in its description. Apply
that plan exactly, delete the
candidate/draft BMK events
that the plan says are
obsolete, preserve the hold
event it says to preserve, and
create every final BMK event
named in the plan. Ignore any
conflicting candidate event
titles or old draft details.
```

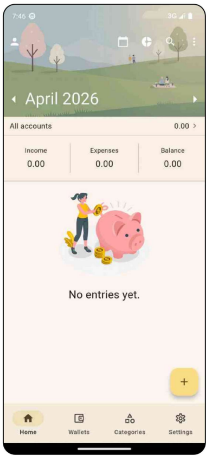

Figure 33: Representative Android task examples.

Back to list of figures

IOS Calendar:Create Today Event

Task

task_id:
ios.calendar.create_today_event

Platform: IOS

App: Calendar

Task instruction

In Calendar, update today's event Bench Volunteer Briefing Draft so the title becomes Bench Volunteer Briefing Final, the time becomes 18:15-19:00…… Create a new event for tomorrow titled Bench Volunteer Debrief from 09:30-10:00 at Cedar Room. Then switch to Reminders, …… Send Bench Volunteer Briefing Final agenda…… and Review Bench Volunteer Debrief notes.

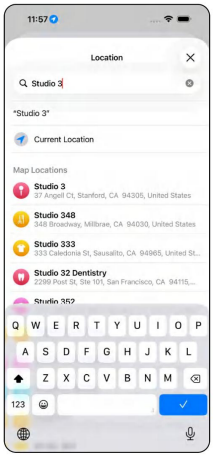

IOS Contacts:Reconcile Seeded Contacts

Task

task_id:
ios.contacts.reconcile_seeded_contacts

Platform: IOS

App: Contacts

Task instruction

Clean up the existing contacts list. Update Bench Catering so its company is Bench Catering Team, its phone number is 555-010-3888, ……. Delete the outdated contact Bench Shuttle. Then create a new contact named Bench Venue with company Bay Hall, phone number 555-010-3999, and email venue@bench-example.test. Leave Bench Office unchanged.

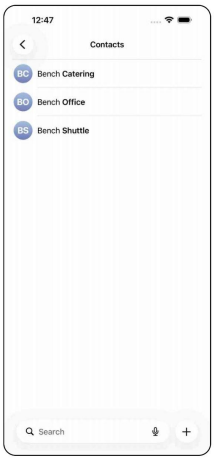

IOS Deltachat:Send Message

Task

task_id:
ios.deltachat.send_message

Platform: IOS

App: DeltaChat

Task instruction

In Delta Chat, create a new contact with email address benchrecv@example.test and name "Bench Courier". …….test and forward the message "courier_packet_ready_2047" to Bench Courier. After that, open the conversation with Bench Courier and send "Please confirm when you receive the forwarded packet.…….test and send "I forwarded the packet to Bench Courier and notified Bench Supervisor."

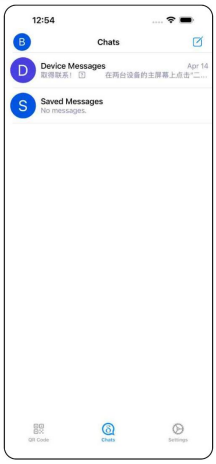

IOS OpenShop : Add Coffee to Cart

Task

task_id:
ios. OpenShop.add_coffee_to_cart

Platform: IOS

App: OpenShop

Task instruction

In OpenShop , you are already signed in. Search for coffee and add Mountain Isle Pour-Over Coffee 20 Pack to the cart. …… Open the cart, increase Polar Day Slim Business Backpack to quantity 2, ……
note OpenShop order: backpack, headphones, and coffee for the offsite setup. Leave the existing transaction Bench Router Cable unchanged.

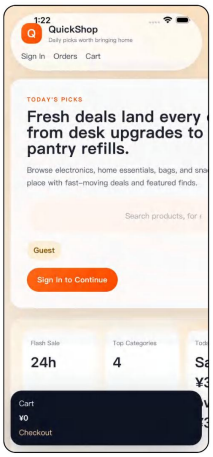

IOS flow:Rework Budget with Transfer and Income

Task

task_id:
ios.flow.rework_budget_with_transfer_and_income

Platform: IOS
App: flow

Task instruction

In Flow, find the transaction titled Bench Taxi Draft and edit it. Change the title to Bench Taxi Receipt, ……, and change the note to Late train replacement ride. Then create a transfer of 60.00 from Bench Checking to Bench Cash titled Bench Cash Refill with the note Refill field wallet. After that, create an income transaction on Bench Checking titled Bench Vendor Rebate...

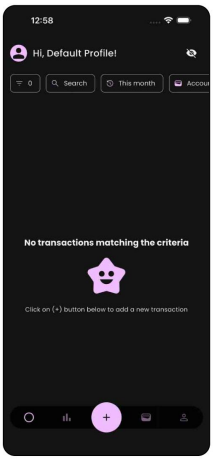

IOS Fsnotes:Archive Note and Create Followup

Task

task_id:
ios.fsnotes.archive_note_and_create_followup

Platform: IOS

App: Fsnotes

Task instruction

In FSNotes, create a folder named Bench Archive. Move the note Launch Retrospective Draft into that folder, ……and append this exact section at the end of the note: ## Decisions - Send recap by noon - Confirm postmortem room Then create a new note in the root named Bench Follow-up Log with these two lines exactly: Owner: Ops Status: Waiting on room reply Leave the other notes unchanged.

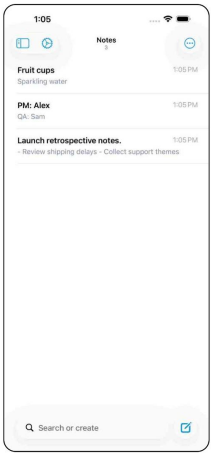

Figure 34: Representative iOS task examples.

Back to list of figures

IOS Habit:Rework Seeded Routine

Task

task_id:
ios.habit.rework_seeded_routine

Platform: IOS

App: Habit

Task instruction

```
In Habit, make these changes
to the seeded habits. Rename
Morning Tidy to Studio Reset,
change its motivation to Reset
the tables and wipe the
counters before lunch, ……
Rename Stretch Break to
Mobility Reset, change its
motivation to Two rounds of
shoulder and hip stretches,
and change its color to
purple. Delete Inbox Zero. Add
a new habit named Pack
Tomorrow Bag...
```

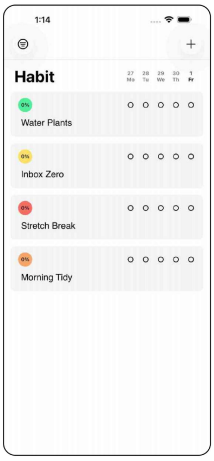

IOS Maily:Forward Delete and Send Launch Notice

Task

task_id:
ios.maily.forward_delete_and_send_launch_notice

Platform: IOS

App: Maily

Task instruction

```
In Maily, open the email with
subject \"Bench Launch
Packet\" ……  Then go back to
the inbox and delete the
emails with subjects \"Bench
Launch Packet Draft\" and
\"Bench Launch Packet Working
Notes\". After that, compose a
new email to
benchrecv@example.test with
subject \"Bench Launch Packet
Routed\" and body \"The final
launch packet was forwarded to
Bench Ops...
```

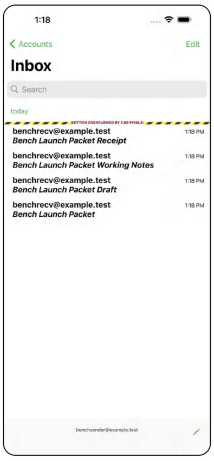

IOS Safari:Organize Tabs and Bookmarks

Task

task_id:
ios.safari.organize_tabs_and_bookmarks

Platform: IOS

App: Safari

Task instruction

```
In Safari, open each of these
pages in its own tab: https://
example.com/?bench=safari-
alpha, https://example.org/?
bench=safari-beta, ……
 Create a bookmark folder
named Bench Reading Queue.
Save the alpha page in that
folder with the title Bench
Alpha Note, save the beta page
in that folder with the title
Bench Beta Note, do not
bookmark the gamma page, close
the gamma tab...
```

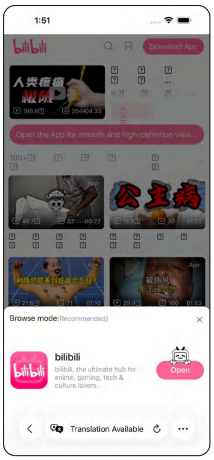

IOS Reminders:Manage Seeded List

Task

task_id:
ios.reminders.manage_seeded_list

Platform: IOS

App: Reminders

Task instruction

```
In the Reminders app, use the
already open list \"Benchmark
Tasks\" and update it to the
final state below. Rename
\"Plan team lunch\" to \"Plan
Friday team lunch\", rename
\"Send vendor ping\" to \"Send
final vendor ping\", delete
\"Old lunch draft\" and \"Book
extra chairs\", add four new
reminders titled \"Send venue
options\", \"Book projector\",
\"Confirm headcount\", and
\"Pack receipt envelope\"...
```

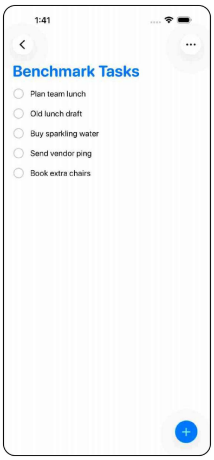

Figure 35: Representative iOS task examples.

Back to list of figures

Error case Linux Blender: Configure and Save a Product Gallery Scene

task_id	linux.blender.product_gallery_review
Platform	Linux
App	Blender
Human_Steps	75

Task instruction

In Blender, open `/tmp/helloworld_blender_workspace/blender_product_gallery/product_stage_seed.blend`, ... Make ReviewBlock the selected active object in the Outliner, then save to the exact path above so the saved file keeps ReviewBlock active.

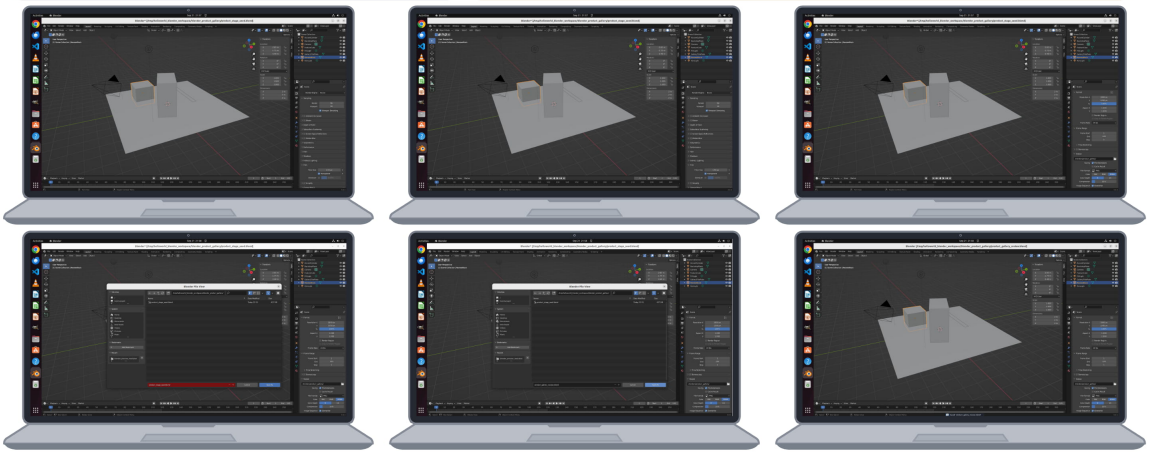

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	29

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Saved-artifact verification failure

The saved Blender file retained a 50 mm camera lens instead of the required 52 mm. The agent declared completion after saving, without detecting the mismatch. The evaluator reported this error; it does not establish that all other requirements passed.

Error case Linux LibreOffice Calc: Build Demographic Percentage Pivot

task_id	linux.calc.demographic_profile_pivots
Platform	Linux
App	LibreOffice Calc
Human_Steps	85

Task instruction

Please create a new sheet. Keep its sheet name as "Sheet2". Merge cells A1:C1 in the new sheet ... Attainment. They should be stacked one by one in Sheet2, each separated with a blank line. The spreadsheet should already be open in LibreOffice Calc. Save the workbook when finished.

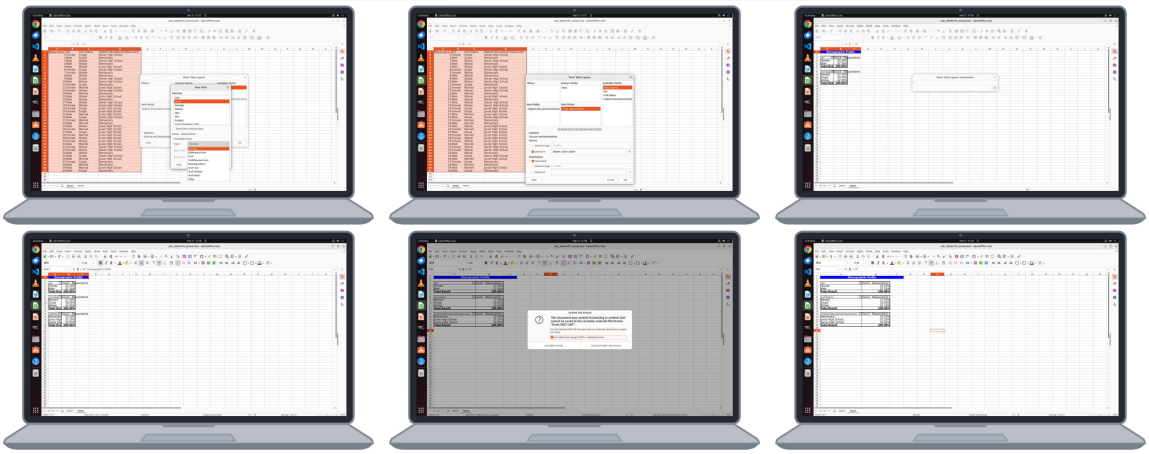

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	41

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Spreadsheet structure verification failure

The agent declared the three percentage pivot tables complete and saved. The workbook comparison failed because Sheet2!A2 was empty instead of the expected Sex label. This establishes a saved-workbook structure mismatch, not that all pivot tables were absent.

Figure 36: Representative Ubuntu/Linux error case.

Back to error-case list

Error case Windows OpenBank: Create Recipients, Transfers, Payments, and Disputes

task_id	windows.openbank.task01_551454
Platform	Windows
App	OpenBank
Human_Steps	78

Task instruction

In the OpenBank application, create three new transfer recipients, submit three new transfers, schedule two bill payments, and file two transaction disputes. The three transfer ... contacte--d.". OpenBank must remain the foreground application when all tasks are complete.

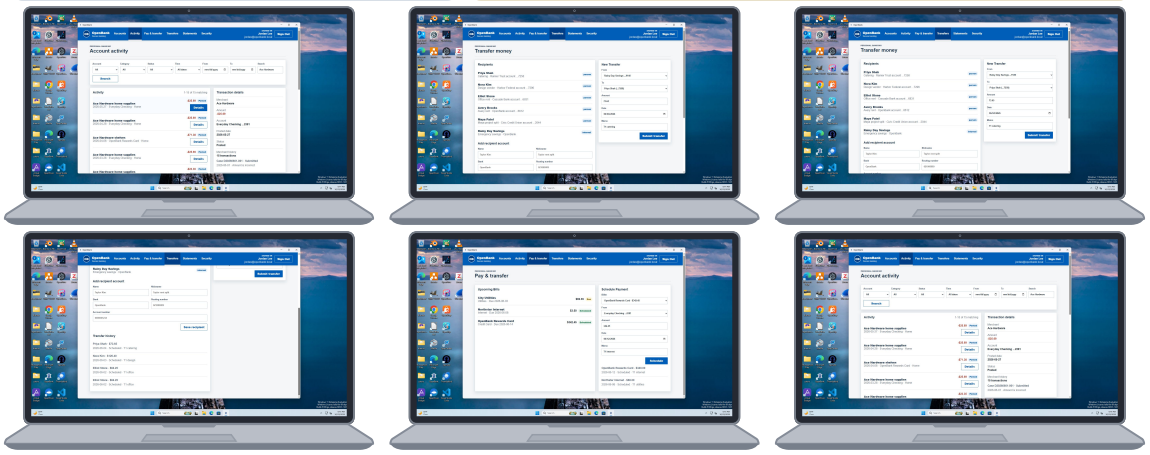

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	99

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Duplicate transaction submission

Two distinct transfers matched the requested Elliot Stone transfer: 84.25 from savings-9140 on 2026-06-02 with memo T1 office. The transfer had to exist exactly once. The uniqueness and business-order checkpoints failed although the agent declared completion.

Error case Windows OpenTrip: Create Five Distinct Confirmed Bookings

task_id	windows.opentrip.v2_multi_booking_906787
Platform	Windows
App	OpenTrip
Human_Steps	80

Task instruction

Create five confirmed bookings in OpenTrip under the logged-in account for Lucas Lin with phone number ... ES6821 from Shanghai (SHA) to Sanya (SYX) on 2026-06-12, with Maya Chen · Passport as the traveler, Business cabin, and the note V1SanyaBaseB.

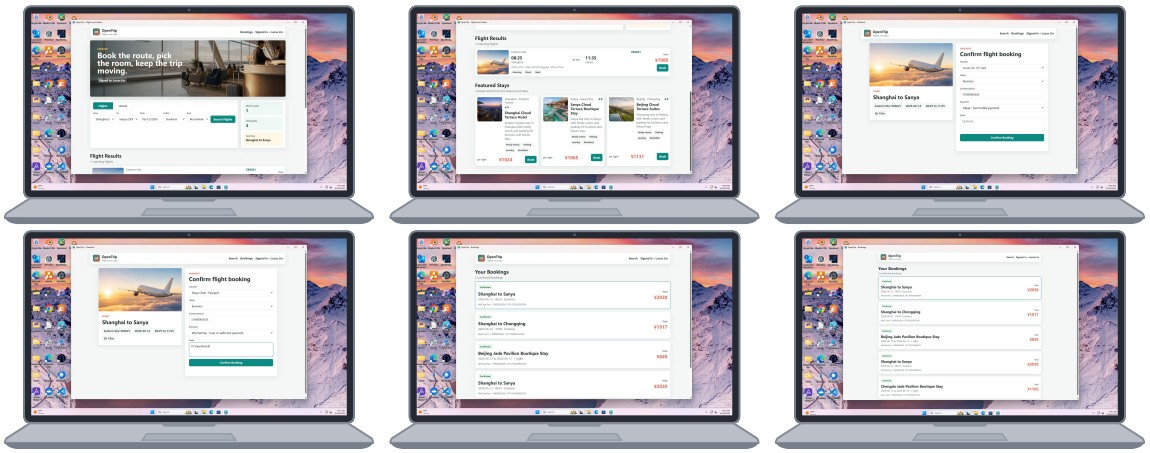

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	41

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Confusion between similar subtasks

The agent created two matching V1SanyaExecB bookings and no V1Sanya--BaseB booking. At trajectory step 38 it reused the executive booking note and WeChat Pay selection. It then declared completion based on five vis--ible Confirmed records without checking the distinct booking identities.

Figure 37: Representative Windows error case.

Back to error-case list

Error case **macOS DeltaChat: Compile the Nova Command Clearance Record**

task_id	macos.deltachat.nova_command_event_clearance
Platform	macOS
App	DeltaChat
Human_Steps	75

Task instruction

Compile the Nova Command event clearance record in DeltaChat by forwarding three specific source messages to ... chat ends with "Owner is Morgan" as the latest self-sent text, and the Nova Command chat ends with "Nova event cleared" as the latest message.

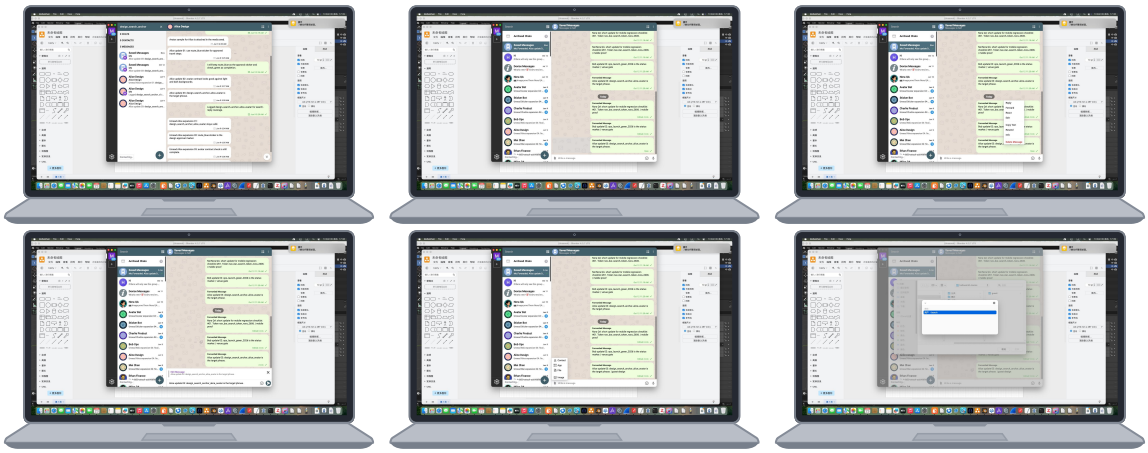

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	42

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Cross-message ordering failure

The evaluator reported one failed final-state check: the three source-derived messages in Saved Messages were out of instruction order. Their matched message IDs were 548, 546, and 552 for the required Nora-Bob-Alice sequence. The remaining final-state checks passed.

Error case **macOS Thunderbird: Create and Update a Budget Release Schedule**

task_id	macos.thunderbird.calendar.budget_release_schedule
Platform	macOS
App	Thunderbird
Human_Steps	75

Task instruction

Create five calendar entries in the Thunderbird Calendar, consisting of four single events and one... and a reminder set to `15 minutes before` for all occurrences. ... , with the `Release Sync` series fully updated, and the Thunderbird Calendar must be displayed in the Month view.

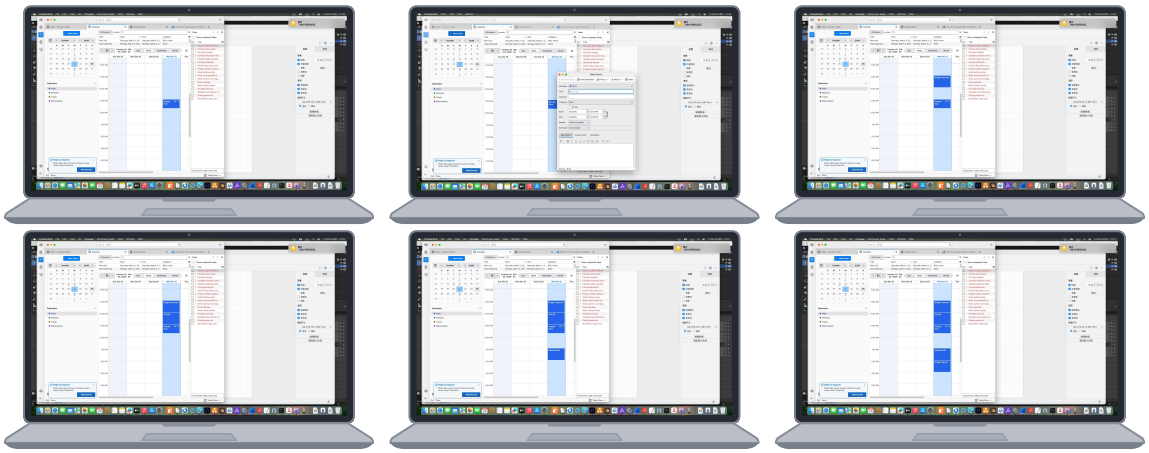

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	41

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Duplicate creation during recovery

The calendar contained 4 Budget Approval events, 3 Budget Prep events, and 2 Budget Signoff events instead of exactly one of each. Release Brief also had the wrong date, duration, and reminder. The agent declared completion despite these persisted errors.

Figure 38: Representative macOS error case.

Back to error-case list

Error case **Android DeltaChat: Compile and Share a Mobility Safety Review**

task_id	android.deltachat.mobility_safety_actual_input_1787395000000_739204
Platform	Android
App	Delta Chat
Human_Steps	74

Task instruction

Create a new Delta Chat group named `Mobility Safety Cell` with the description `Bike car ...` and **`Cafe: Safe fallback`**. Save the note and share the file with the `Mobility Safety Cell` group in Delta Chat. The final state must be the Delta Chat chat list.

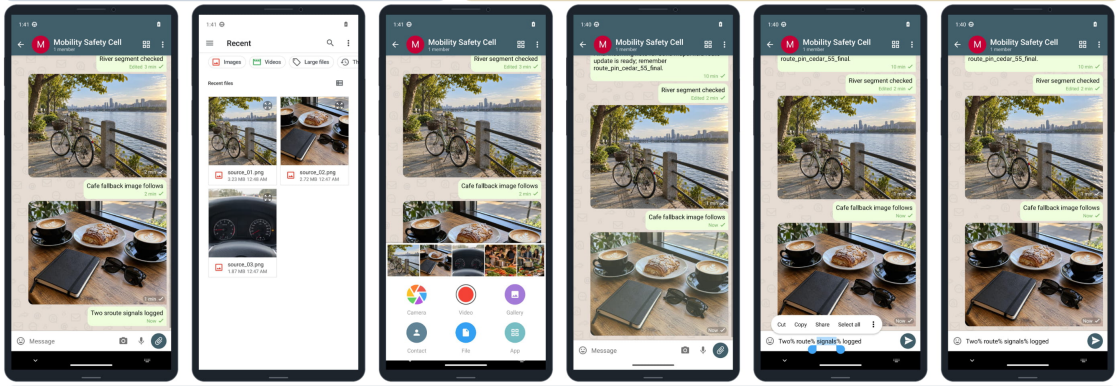

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	108

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Text-input corruption and incomplete correction

The required message Two route signals logged was sent as Two sroute si--gnals logged. The Markor evidence line also retained percent signs betw--een words. The agent shared the note and declared completion, but both exact-text checks failed.

Error case **Android OpenFood: Complete Reorders, Coupons, and Merchant Updates**

task_id	android.openfood.dumpling_biryani_coupon_merchant_clean_20260731
Platform	Android
App	OpenFood
Human_Steps	73

Task instruction

Produce a completed OpenFood session for Lucas Lin that includes three Dumpling Room No. 20 San Jose CA orders, ... and all specified orders, the support ticket, the claimed voucher, and the merchant availability changes in place.

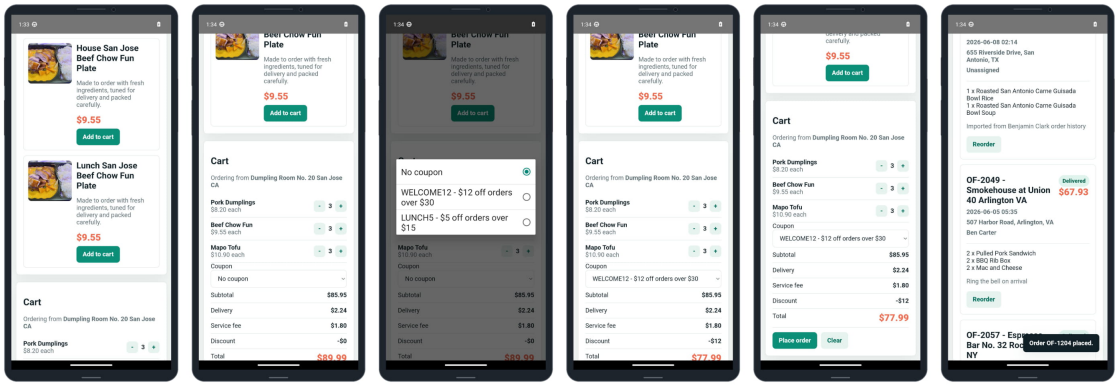

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	149

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Missed quantity update after reordering

The third Dumpling Room order kept Pork Dumplings, Beef Chow Fun, and Mapo Tofu at 3 each instead of 4. The agent applied WELCOME12 and submitted order OF-1204 without correcting the quantities. The other order, support, coupon, and merchant checks passed.

Figure 39: Representative Android error case.

Back to error-case list

Error case **iOS OpenShop: Place Two Orders, Then Submit Reviews**

task_id	ios.openshop.home_four_item_cart_checkout_368575
Platform	iOS
App	OpenShop
Human_Steps	77

Task instruction

Place two separate orders in the OpenShop app and then submit two product reviews, ensuring the final state reflects the exact items, quantities, ... in the first order must have a rating of 4 / 5 and the text "Stable patio set". No reviews may be submitted for any other products.

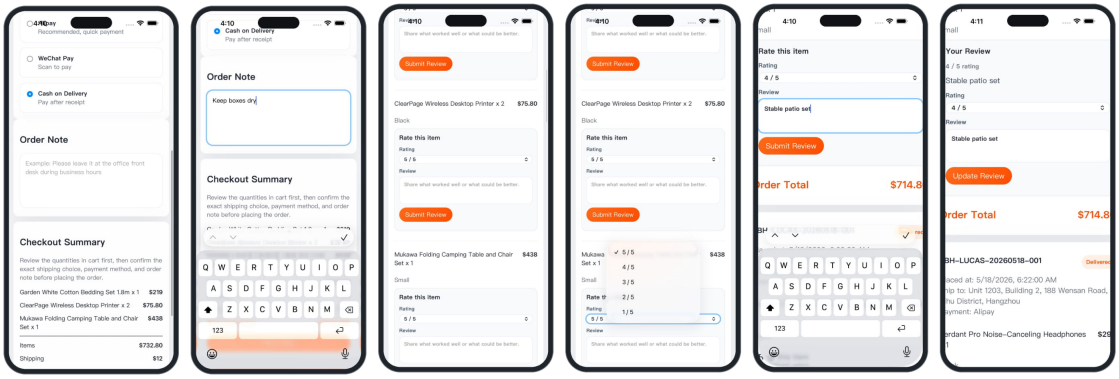

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	64

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Temporal dependency violation

The agent submitted the Mukawa review after placing the Office order but before placing the Home order. The final orders and reviews passed their object checks, but the required order-before-review dependency failed.

Error case **iOS FSNotes: Organize Notes and Create a Follow-up Checklist**

task_id	ios.fsnotes.move_rename_and_update_note
Platform	iOS
App	FSNotes
Human_Steps	75

Task instruction

In FSNotes, create folders named Release Archive, Meeting Prep, and Carrier Updates. Move the note Sprint Review Draft into Release ... lines exactly: Owner: Alex Room: Cedar Status: Pending confirmation Bring: Printed agenda Leave the other notes unchanged.

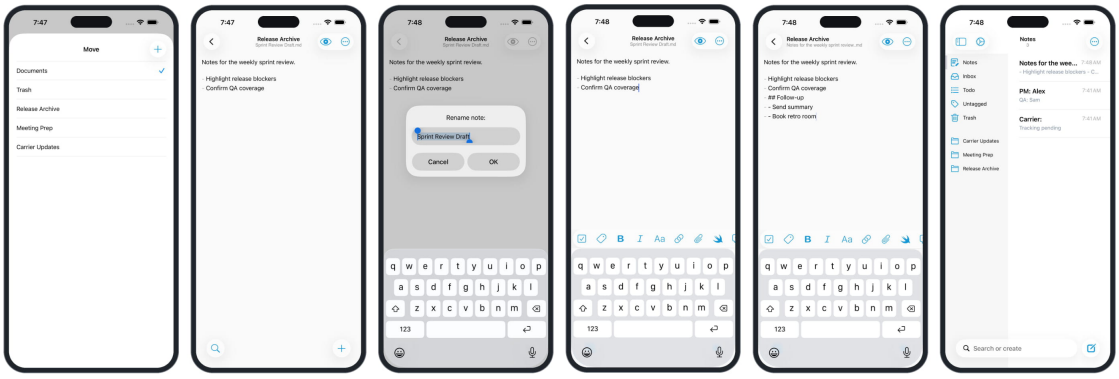

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	24

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Incorrect saved state and incomplete workflow

The final state contained 3 notes instead of 4. The release note was saved as Release Archive/Notes for the weekly sprint review..md instead of Re--lease Archive/Sprint Review Final.md. The new Retro Checklist was miss--ing, and the required updates to the moved notes did not pass verification.

Figure 40: Representative iOS error case.

Back to error-case list

Error case Web OpenSlide: Build a QA Review Deck

task_id	linux.openslide.qa70_review_199367
Platform	Web
App	OpenSlide
Human_Steps	72

Task instruction

Create a new presentation using the Business review template and set the presentation title to ... the slide's content area, along with a comment containing the text quarterly decision matrix owner review and the notes field populated with Decision owners locked for Friday review.

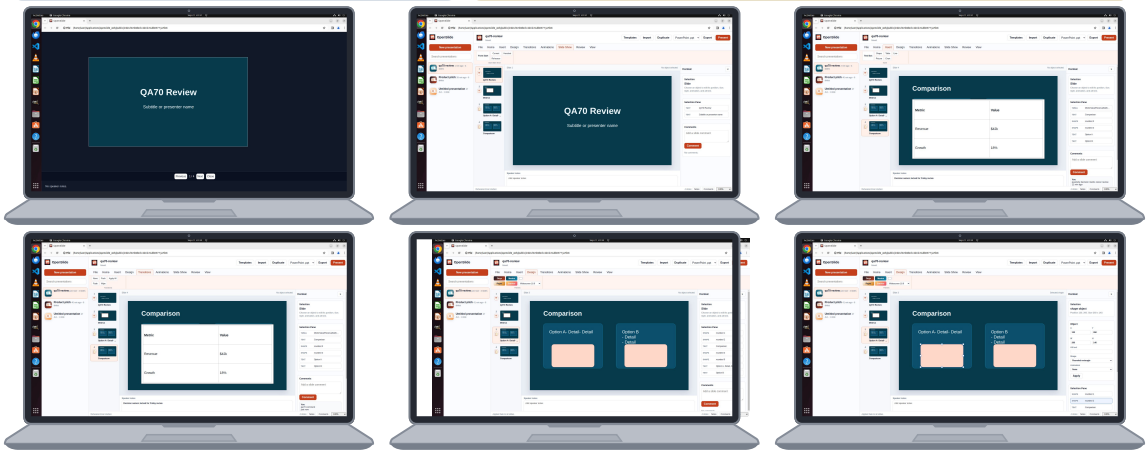

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	300

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Repeated non-idempotent actions

Slide 4 contained 7 comments instead of the required 2, with repeated texts. The run exhausted 300 steps. The main evaluator verdict named slide4.comments; nested rehearsal evidence included a failed attempt, not a separate top-level verdict.

Error case Web OpenTwitter: Complete a Social Control Trail

task_id	linux.opentwitter.control_trail
Platform	Web
App	OpenTwitter
Human_Steps	80

Task instruction

Sign in to OpenTwitter as Alex Chen with the username alex and password chirp2026 to establish ... control is present in the profile editor, it must be activated before the profile changes are saved. All notifications must be marked as read.

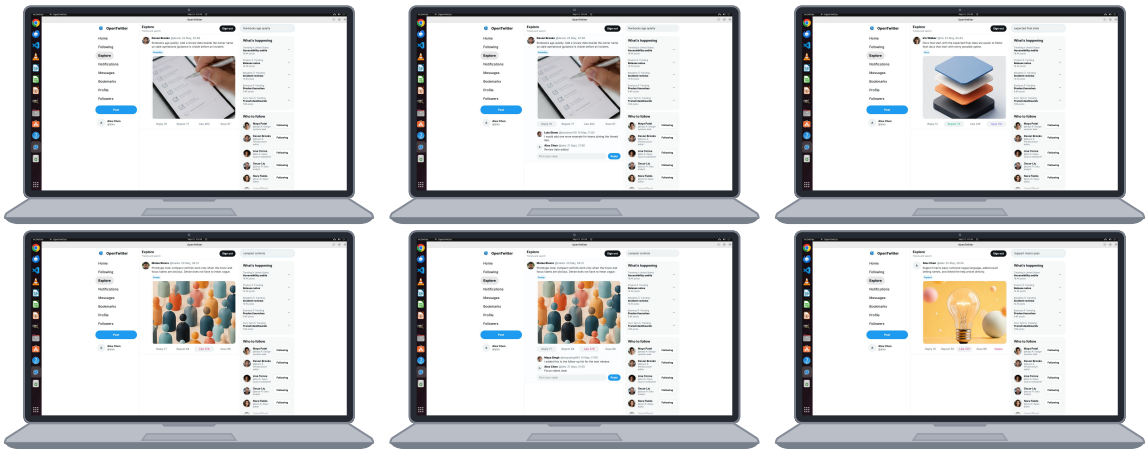

Run Result

Model	GPT-6 Astra
Action mode	code_adapter
Observation mode	screenshot_only
Steps	300

Initializer	PASS
Evaluation	FAIL
Overall status	FAIL

Failure reason

Progress-tracking failure and incomplete messaging

The required outbound message to Mateo Rivera was missing. Other posts, interactions, and profile changes had been completed, but late trajectory steps continued searching and checking posts until the 300-step budget was exhausted.

Figure 41: Representative Web error case.

Back to error-case list