



SWEeper-Bench: Can Agents Discover Bugs in Interactive Software?

Yang Yao^{1*} Haozhe Chen^{1*} Bingyi Kang Karthik R Narasimhan¹ Zhuang Liu¹

¹Princeton University

Website Code Data Agent Traces




Figure 1. From waiting for bug reports to sweeping for bugs. (a) Existing benchmarks start after a user reports a bug, which comes only after a slow, costly deployment. (b) In SWEeper-Bench, the agent must discover the problem itself by exploring the code and the running application, then diagnose and fix it before users encounter it.

Abstract

Coding agents increasingly interact with users by generating interactive software on the fly rather than replying in text, yet many bugs in such software reveal themselves only after multi-step interactions. Today, users must discover these bugs and wait for fixes, making the development loop slow and expensive. Existing benchmarks typically assume the bugs have already been identified, leaving agents’ ability to find interaction-dependent failures largely unevaluated. We introduce SWEeper-Bench, 200 tasks built from real bug fixes in 200 widely used open-source web applications. Each task gives the agent only the codebase and an open-ended request such as “Find and fix issues in Focalboard’s board sharing and membership.” Because these bugs are hard to detect from code alone, agents must design interactive tests to uncover them. An agentic verifier then judges each repair by what a user would observe in the browser. The best of 15 frontier agents passes only 59.0% of tasks. Through stage-by-stage analysis of agent trajectories, we show that the biggest bottleneck is identifying the buggy behavior: agents often reach the relevant feature but never trigger the failure or fail to recognize it. SWEeper-Bench measures progress toward agents that close this loop themselves, delivering software that works without waiting for user feedback.

* Equal Contribution




Figure 2. SWEeper-Bench is far from solved. Pass rate versus mean prediction cost per task for 15 agents. Each agent pairs a model with a harness. A task passes only if both its target and preservation behavior tests pass. Nine agents pass fewer than half of the tasks, and agents with similar pass rates differ in cost by more than 30 times. Table 2 gives the full results.

1 Introduction

Coding agents such as Codex (OpenAI, 2025), Claude Code (Anthropic, 2026a), and Cursor (Cursor, 2026) can now build interactive software from short requests, but such software is hard to debug. Its bugs often appear only after a specific sequence of user interactions: a board shared with a colleague never shows up in the colleague’s sidebar, or a coupon accepted in the cart is rejected at checkout. Because these failures depend on how people use the software, they are hard to spot from code alone and usually surface only when people use it in the wild. Repair therefore stays reactive: users run into problems and report them, and agents must wait for this feedback before they can improve the software (Figure 1a).

This loop matters more as coding agents increasingly generate interactive software on the fly instead of replying in text. Imagine asking an assistant to organize a team offsite. Instead of replying in text, it builds a small app with a sign-up form, a shared schedule, and a budget tracker, and shares it with twenty colleagues within minutes. No testing team ever sees this app, so if some sign-ups silently vanish, the colleagues discover the bug themselves and wait for a fix. Routing every bug through users is slow and expensive, and users become the bottleneck of the whole loop. Agents must instead close the loop on their own and sweep out as many bugs as possible before presenting their software. We therefore ask: *can agents discover and fix the problems users will encounter in interactive software on their own?*

Existing benchmarks typically assume the bug has already been identified, leaving agents’ ability to find interaction-dependent failures largely unevaluated. SWE-bench (Jimenez et al., 2024) and its successors (Yang et al., 2025; Rashid et al., 2025; Deng et al., 2025; Zan et al., 2025; Zhang et al., 2025) start from an issue that a user has already reported, so they measure diagnosis and repair but not discovery. They also grade repairs with hidden unit tests (Jimenez et al., 2024), which check internal state rather than what a user sees and can reject valid fixes that take a different approach.




Figure 3. Task execution and verification. Given only the buggy codebase and an open-ended prompt, the agent explores the code and the running application in a sandbox, discovers the bug, and produces a patch. An agentic verifier applies the patch to a fresh copy of the application and runs predefined behavior tests through the browser. It judges only what a user would observe.

We introduce SWEeper-Bench, 200 tasks that test whether agents can sweep the bugs out of real web applications (Figure 1b). Each task comes from a real bug fix in a different open-source application, and its bug appears only after a sequence of user interactions. The agent receives the codebase, the running application with a browser, and an open-ended request that names only a product area. To succeed, it must understand how the application is used, anticipate what could go wrong, and design interactive tests that expose the bug before fixing it. An agentic verifier then operates the repaired application in the browser and judges only what a user would observe, so any correct implementation passes.

Across 15 frontier agents, the best passes only 59.0% of tasks, and nine pass fewer than half (Figure 2). Agents can fix these bugs once someone finds them: when we replace the open-ended request with a detailed bug description, all five agents we test pass about 96% of tasks. Through stage-by-stage analysis of agent trajectories, we show that the biggest bottleneck is identifying the buggy behavior. Agents reach the feature containing the bug in 75.4% of trajectories but identify the buggy behavior in only 50.0%: they either never trigger the failure or do not recognize it. They also rarely test through the browser unless told to.

SWEeper-Bench moves the evaluation of coding agents from repairing bugs to discovering them, a step toward agents that close the loop of building, testing, and fixing software on their own.

2 Constructing SWEeper-Bench

Each SWEeper-Bench task pairs a buggy application and an open-ended prompt with two behavior tests that an agentic verifier runs in the browser (Figure 3). We describe how we curate tasks (Section 2.1), what agents receive (Section 2.2), and how we verify repairs (Section 2.3).

2.1 Curating Tasks

We build each task from a merged pull request that fixes a user-facing bug in an open-source web application: the parent commit is the buggy version, and the merged fix is the reference repair. Each application contributes one task. The 200 tasks thus span 200 applications in seven domains (Figure 4). The median repository has 4.6k GitHub stars.

Verifier	Buggy version		Reference fix		Avg. cost
	Target ↓	Preservation ↑	Target ↑	Preservation ↑	
🌀 GPT-5.6 Luna, Codex	0.2 ± 0.3	99.2 ± 0.8	99.8 ± 0.3	99.3 ± 0.3	\$0.160
🌀 GPT-5.6 Luna, Browser Use	0.8 ± 0.3	99.8 ± 0.3	99.8 ± 0.3	99.7 ± 0.3	\$0.635
🌟 Claude Sonnet 5, Browser Use	0.2 ± 0.3	99.8 ± 0.3	99.8 ± 0.3	100.0 ± 0.0	\$3.840

Table 1. The verifier is accurate and consistent. Pass rate (%) of each behavior test on the buggy version and the reference fix, mean $\pm$ standard deviation over three runs of 200 tasks. Arrows mark the ideal: 0% for the target test on the buggy version, 100% otherwise. Avg. cost is USD per task. Verifiers use GPT-5.6 Luna (OpenAI, 2026b) or Claude Sonnet 5 (Anthropic, 2026e) with Codex (OpenAI, 2025) or Browser Use (Müller and Žunič, 2024).

We select bugs that are hard to find but obvious once seen: they are reproducible through normal use, and any user would agree the behavior needs fixing. An agent may fix several issues in the named area, but we evaluate only the target bug. A complete sweep of the area should include it.

Figure 5 shows a typical task. Focalboard is a project board tool. When a user shares a board, it never appears in the recipient’s sidebar because adding a member saves the membership but never files the board under the member’s default sidebar category. Nothing looks wrong to the user who shares the board, and the recipient can still open it through a direct link. The failure appears only after nine steps across two accounts. The fix is a single function call.




Figure 4. 200 tasks from 200 real-world web applications. Each task comes from a different open-source application, grouped here into seven domains. Examples are the two most-starred repositories in each domain.

We reproduce every bug by hand in the browser, confirm that the reference fix resolves it, and walk through both behavior tests. The agent never sees the original issue, pull request, reference patch, repository history, or behavior tests, and outbound network access is blocked.

2.2 Open-Ended Prompt

Each task gives the agent a one-sentence request that names a product area, such as “Find and fix issues in Focalboard’s board sharing and membership.” It does not say what is broken, where the code lives, or how to reproduce the problem. This mirrors asking a colleague to sweep a feature before release. The agent works in an isolated sandbox with the buggy source code, a running instance that reflects its edits, a command interface for tests and builds, and a Chromium browser with Playwright. It decides for itself what to read, which scenarios to test, and when to stop. Our default prompt also asks the agent to test the application in the browser like a real user (Appendix A), and Section 4.1 shows why this matters. When the agent stops, we evaluate only its code changes.

2.3 Agentic Verifier

We verify repairs the way a user would notice them. We apply the agent’s code changes to a fresh copy of the buggy application, and a verifier agent then uses it through the browser without seeing the source code. Each task has two behavior tests, written as step-by-step user instructions with explicit pass and fail criteria (Appendix B). The target behavior test reproduces the interactions that expose the bug and checks that the intended behavior is restored; the preservation behavior test checks that nearby features still work. A task passes only if both tests pass.




Figure 5. A bug that takes nine steps and two accounts to see. In Focalboard, the owner creates and names a board, shares it with a second user, and logs out (1 to 7). The second user then logs in and refreshes (8). With the bug, the shared board is missing from the recipient’s sidebar (Fail); with the fix, it appears (Pass). The root cause is that adding a member saves the membership but never adds the board to the member’s default sidebar category.

To measure reliability, we run both tests on the buggy version and on the reference fix with three verifier configurations, three times each. An accurate verifier passes the target test on none of the buggy versions and all of the reference fixes, and passes the preservation test on both. All configurations come within one percentage point of these ideal rates (Table 1). Across runs, the standard deviation is at most 0.8 points. We use the cheapest configuration, GPT-5.6 Luna with Codex, as the default verifier.

3 How Do Agents Perform?

3.1 Benchmark Results

We evaluate 15 agents (Table 2). Each pairs a frontier model with its native harness if one exists and with Codex otherwise. We list the models (xAI, 2026; Qwen Team, 2026; Anthropic, 2026c,d; OpenAI, 2026c,b; DeepSeek, 2026c,b; Moonshot AI, 2026b; Meta, 2026b; Z.ai, 2026) and harnesses (OpenAI, 2025; Anthropic, 2026a; Cursor, 2026; Moonshot AI, 2026a; Meta, 2026a; DeepSeek, 2026a). All runs use the browser instruction and an 8-hour limit per task.

SWEeper-Bench is far from solved. Grok 4.6 with Cursor passes 59.0% of tasks, followed by Qwen 3.8 Flash at 57.5% and three agents at 56.5%; nine of fifteen pass fewer than half. Nearly all failures come from the target behavior test, while preservation pass rates stay between 97.5% and 100%. Agents rarely break working features but often miss the bug they were asked to find. Similar pass rates also hide very different time and cost: Grok 4.6 averages 12 minutes per task and Qwen 3.8 Flash 102 minutes, and Claude Opus 5 matches Claude Fable 5.1 at more than twice the cost.

Model	Harness	Reasoning	Pass $\uparrow$	Target $\uparrow$	Pres. $\uparrow$	Time (min)	Cost (\$)
🔗 Grok 4.6	Cursor	xhigh	59.0	59.0	100.0	12.3 $\pm$ 3.9	6.41
🔗 Qwen 3.8 Flash	Codex	xhigh	57.5	58.0	97.5	101.8 $\pm$ 46.6	0.76
🌟 Claude Fable 5.1	Claude Code	xhigh	56.5	56.5	100.0	28.2 $\pm$ 15.8	11.76
🌟 Claude Opus 5	Claude Code	xhigh	56.5	57.5	98.5	42.8 $\pm$ 23.2	25.01
🌀 GPT-6 Astra	Codex	xhigh	56.5	56.5	100.0	13.7 $\pm$ 3.4	5.73
🔗 DeepSeek V4.1 Flash	DeepSeek	xhigh	51.5	51.5	99.0	41.8 $\pm$ 20.0	0.35
🌀 Kimi K3	Kimi Code	high	48.5	49.0	99.5	90.1 $\pm$ 63.1	10.19
🔗 DeepSeek V4 Pro	DeepSeek	xhigh	47.0	47.0	99.5	58.2 $\pm$ 25.8	1.38
🌀 GPT-5.6 Sol	Codex	xhigh	47.0	47.0	98.0	16.8 $\pm$ 5.8	4.28
🌀 Muse Spark 1.3	Muse Code	xhigh	47.0	47.5	99.5	25.6 $\pm$ 20.3	4.75
🔗 GLM 5.3	Codex	max	43.5	43.5	99.0	20.3 $\pm$ 10.9	2.53
🌀 GPT-5.6 Luna	Codex	xhigh	42.0	43.5	98.5	16.5 $\pm$ 6.0	0.15
🔗 GLM 5.3 Flash	Codex	max	39.5	39.5	99.0	22.9 $\pm$ 11.7	0.26
🌀 GPT-5.6 Terra	Codex	xhigh	35.0	35.5	99.0	22.7 $\pm$ 7.0	1.64
🔗 Qwen 3.8 Max	Codex	none	28.0	28.0	99.0	24.7 $\pm$ 13.5	2.51

Table 2. Results on SWEeper-Bench. Pass rates are percentages of the 200 tasks. Pass requires both behavior tests to pass; Target and Pres. (preservation) report each test separately. Time is mean $\pm$ standard deviation of agent work time per task; cost is mean prediction cost per task.

3.2 Where Do Agents Stumble?

Failing agents often get close to the bug (Figure 6). To locate where they stop, we split each trajectory into five cumulative stages (Figure 7). GPT-5.6 Luna reads each trajectory and, using a task-specific description written from the reference fix, judges whether the agent identifies the target feature, the buggy behavior, and the root cause. The last two stages come from the behavior tests. Averaged over 14 agents, the share of trajectories falls from 75.4% at the first stage to 40.8% at the last.

Finding the target feature. Agents reach the feature that contains the bug in 75.4% of trajectories (63.5% to 91.4% across agents). In the AbanteCart task, Grok 4.6 fixes phone numbers, addresses, and notifications on the account profile page but never tries the ID photo upload where the bug lives.

Identifying the buggy behavior. This is the largest drop, to 50.0%. Agents often reach the feature but never run the interaction that exposes the bug: in the Answer task, Grok 4.6 accepts and reverts answers but never opens a question through its short-ID link. Even agents that trigger the failure may not recognize it (Figure 6). On the Focalboard task, all four agents we inspect reach the empty recipient sidebar. GPT-6 Astra moves on without comment, and Claude Opus 5 concludes the behavior is intended because the existing unit tests expect no sidebar update. Existing tests can thus reinforce the wrong expectation.


Diagnosing the root cause. This stage keeps 43.8% of trajectories. In the SearXNG task, GLM 5.3 Flash sees ampersands escaped in search suggestions but blames client-side query encoding instead of server-side HTML escaping. In the Audiobookshelf task, Qwen 3.8 Flash reproduces the failed author merge but attributes it to case-sensitive lookup rather than duplicate insertion.

Fixing and maintaining behavior. Once the cause is found, agents usually fix it: 41.0% of trajectories pass the target test, and 40.8% also keep nearby features working. The remaining failures are rare regressions and incomplete fixes. For example, GLM 5.3 Flash corrects the EverShop variant editor’s default stock but omits the API field that supplies the quantity.


3.3 Behavioral Tendencies

Longer work does not mean better results. Agents differ in work time by nearly an order of magnitude, and this time does not track pass rate (Figure 8). Grok 4.6 and GPT-6 Astra average about 13 minutes per task.


☀️ **Claude Fable 5.1** treated the empty sidebar as a bug and **fixed** it.




🔍 **GLM 5.3** flagged the empty sidebar as a bug right away and **fixed** it.



☀️ **Claude Opus 5** saw the empty sidebar but trusted the unit tests, so it left it **unfixed**.



🌀 **GPT-6 Astra** saw the empty sidebar but pursued other bugs, so it left it **unfixed**.



👉 Browser action (Playwright code) 🖨️ Browser output (page text in terminal) 💬 Agent message or reasoning ✍️ Code-edit tool call

Figure 6. Seeing a bug is not the same as recognizing it. Four agents on the Focalboard task all reach the empty recipient sidebar (shaded column) but respond differently. Claude Fable 5.1 and GLM 5.3 treat it as a bug and fix it; Claude Opus 5 trusts the existing unit tests and judges it intended; GPT-6 Astra pursues other bugs.

Qwen 3.8 Flash and Kimi K3 average 102 and 90 minutes for similar or lower pass rates. Long runs can fail to converge: on three tasks, Qwen 3.8 Flash works for nearly the full 8 hours without producing a patch. Short runs risk missing the bug: Grok 4.6 stops after 7.8 minutes on the AbanteCart task.

More browser use does not mean better results either. Among agents whose harness records tool calls, browser actions make up 17% to 62% of turns, or 371 to 1,082 actions per task (Figure 9). Neither tracks pass rate. Qwen 3.8 Max issues 517 browser actions per task and passes 28.0%. Claude Fable 5.1 uses the browser in only 25% of its turns and passes 56.5%. What matters is which scenarios an agent tests and whether it recognizes a failure when it appears (Section 3.2).




Figure 7. The largest drop is at identifying the buggy behavior. Share of trajectories reaching each stage. Each stage requires all previous ones.




Figure 8. Working longer does not lead to higher pass rates. Pass rate versus mean agent work time per task. Numbers are agent ranks from Figure 2 and Table 2.

4 What Affects Performance?

4.1 Does Browser Prompting Help?

Our default prompt asks agents to test in the browser like a real user (Appendix A gives the full template and the lines we remove). To see whether they would do so on their own, we remove this instruction for GPT-5.6 Luna and GLM 5.3 and keep everything else unchanged. The browser remains installed. The agents then almost never open the browser: GPT-5.6 Luna drives one on 2 of 200 tasks and GLM 5.3 on 6, compared with all 200 when instructed (Table 3). Instead, they read code, run shell commands, and send HTTP requests, and pass rates drop from 42.0% to 39.5% and from 43.5% to 31.5%. Interactive testing is thus a capability agents do not yet reach for on their own. We keep the instruction so the benchmark measures how well agents test rather than whether they think to.

Model	Instruction	Browser	Pass
GPT-5.6 Luna	No	2/200	39.5
	Yes	200/200	42.0
GLM 5.3	No	6/200	31.5
	Yes	200/200	43.5

Table 3. Agents rarely test in the browser unless asked. Browser: tasks where the agent drives a real browser.

4.2 Is Finding Harder Than Fixing?

To separate finding a bug from fixing it, we give five agents from the main evaluation a detailed prompt that describes the target bug and its expected behavior (Appendix A). This prompt plays the role of a bug report. All five then pass between 96.0% and 96.5% of tasks instead of 35.0% to 51.5%, and the gap between agents almost disappears (Table 4). Trajectories also narrow: for GPT-5.6 Luna, the median number of browser launches falls from 26 to 14, and the median number of changed lines from 49 to 13. The difficulty of SWEeper-Bench lies in producing the bug report, not in acting on it.

Model	Open-ended	Detailed
DeepSeek V4.1 Flash	51.5	96.5
GPT-5.6 Sol	47.0	96.5
GLM 5.3	43.5	96.0
GPT-5.6 Luna	42.0	96.0
GPT-5.6 Terra	35.0	96.0

Table 4. Agents fix bugs reliably once told what the bug is. Pass rate (%) by prompt. GPT-5.6 Sol has 199 detailed-prompt tasks.




Figure 9. How agents use the browser. (a) Percentage of agent turns that use the browser. (b) Mean browser actions per task, grouped by type.

4.3 Does the Harness Matter?

Each agent in Table 2 pairs a model with a harness. To isolate the effects of harnesses, we run GPT-5.6 Luna under five harnesses with all else held constant (Table 5). Pass rates range from 35.0% to 42.5%, and Luna does best in Codex, its native harness. A well-matched harness helps the model use its capabilities, but the model itself accounts for most of the variation across agents.

4.4 Does Working Longer Help?

Section 3.3 compares agents that work for different lengths of time. To test whether more time helps the same agent, we give GPT-5.6 Luna, GPT-5.6 Terra, and GLM 5.3 budgets of 20, 40, 60, and 80 minutes on the same 40 tasks. Each agent is told to keep working until its deadline. Pass rates rise from 20 to 40 minutes but not beyond (Figure 10). They peak at 50.0%, 52.5%, and 62.0%, respectively. With 40 tasks these differences are not statistically significant, but they give no sign that more time alone closes the gap.

5 Related Work

Software engineering agents and benchmarks. Repository-level agents improve through better interfaces, code localization, and structured repair workflows (Zhang et al., 2024; Yang et al., 2024; Xia et al., 2025; Wang et al., 2025; Pan et al., 2025; Applis et al., 2026; Lian et al., 2026; Tian et al., 2026; Ganhotra et al., 2026; He and Roy, 2026; Kuang et al., 2026). SWE-bench (Jimenez et al., 2024) evaluates them on real GitHub issues, and its successors raise difficulty (Deng et al., 2025; Huang et al., 2026; Ehrenberg et al., 2026; Zheng et al., 2026), add visual issues and more languages (Yang et al., 2025; Rashid et al., 2025; Zan et al., 2025; Shi et al., 2026), refresh tasks (Zhang et al., 2025), and cover freelance work, broader engineering tasks, user feedback, and rebuilding programs (Miserendino et al., 2025; Raghavendra et al., 2026a,b; Yang et al., 2026). Web, desktop, and command-line benchmarks likewise specify the goal (Zhou et al., 2024; Xie et al., 2024; Merrill et al., 2026). In all of these, the task states what to fix or achieve, so they evaluate the stages after a bug report. SWEeper-Bench names only a product area and evaluates the step that produces the report.

Testing and verification. SWT-Bench (Mündler et al., 2024) asks agents to write tests that reproduce a

Model	Harness	Pass
GPT-5.6 Luna	 Codex	42.5
	 Cursor	40.0
	 Kimi Code	37.5
	 Claude Code	35.0
	 Muse Code	35.0

Table 5. GPT-5.6 Luna does best in its native harness, Codex. Pass rate (%) by harness.




Figure 10. More time helps only up to a point. Pass rate (%) by time budget on 40 tasks.

reported issue, so the bug is known in advance. WebProber (Ye et al., 2025) lets an agent explore websites like a user and report usability issues, but does not evaluate repairs. In SWEeper-Bench, agents must design tests before any bug is known and then fix what they find. To grade repairs, most benchmarks run hidden unit tests written with the reference fix (Jimenez et al., 2024; Rashid et al., 2025; Zan et al., 2025), and these tests can reject valid alternative fixes. Others use engineer-written browser tests (Miserendino et al., 2025) or scripts in the task environment (Huang et al., 2026; Merrill et al., 2026). Browser agents (Müller and Žunič, 2024) make it possible to check behavior through the interface; our agentic verifier follows step-by-step behavior tests and accepts any fix that restores what a user would observe.

AI code review. Finding bugs before users do is increasingly delegated to code review agents. Early work learned to review code from data (Li et al., 2022; Frömmgen et al., 2024). Tools such as Cursor Bugbot (Cursor, 2025), GitHub Copilot code review (GitHub, 2025), CodeRabbit (CodeRabbit, 2026), Greptile (Greptile, 2026), Codex code review (OpenAI, 2026a), Claude Code Review (Anthropic, 2026b), and Gemini Code Assist (Google, 2026) now review pull requests for bugs as they are opened, and benchmarks score their comments on real pull requests (Zeng et al., 2026; Zhang et al., 2026). Like SWEeper-Bench, these tools look for problems no one has reported, but they mostly analyze the diff and surrounding code. SWEeper-Bench evaluates the harder case where bugs appear only when the running application goes through a sequence of user actions.

6 Conclusion and Discussion

Current agents fix bugs reliably once given a bug report, yet pass at most 59.0% of SWEeper-Bench tasks when they must find the bug themselves. The gap lies in testing, not repair. Agents must choose which scenarios to try, run them through the interface, and recognize a failure when they see one. Recognition is often where they fall short. On the Focalboard task, all four agents we inspect see the failure, but only two treat it as a bug.

This testing gap is likely to matter more as agents build software on the fly for individual users. Such software rarely passes through a testing team, so if the agent does not find a bug, the user will. Our results suggest that closing the gap requires agents that use the application the way a user would and judge its behavior by what it is meant to do, not only by what its tests expect. Verifiers that check user-observable behavior, like ours, could also reward this kind of testing during training. As agents take on more of building software, it is worth asking more of them than fixing the bugs we report: can they find those bugs before we do?

SWEeper-Bench has limitations. It covers only web applications and scores one known bug per task, so it measures recall rather than precision. Fixes to other real issues earn no credit, and unnecessary changes go unpenalized unless they break a preservation test. Its behavior tests check functional correctness only, not visual design, ease of use, or responsiveness. Extending the benchmark to other interactive software and to these aspects of user experience is a natural next step.

Ethics statement

SWEeper-Bench is constructed from publicly available open-source software projects in accordance with their licenses. All agent experiments run in isolated environments and do not affect the original repositories, deployed products, or real user data. The benchmark focuses on bugs triggered through normal user interactions rather than the discovery or exploitation of security vulnerabilities, and is not intended to facilitate offensive use. This work does not involve human subjects or private user data.

Acknowledgments

We thank Modal for providing the cloud sandboxes in which we ran all agent and verifier experiments.

References

- Anthropic. Claude Code, 2026a. URL <https://code.claude.com/docs/en/overview>. Software documentation; accessed September 26, 2026.
- Anthropic. Code review. <https://code.claude.com/docs/en/code-review>, 2026b. Accessed September 26, 2026.
- Anthropic. Claude Fable 5.1, 2026c. URL <https://www.anthropic.com/claude/fable>. Accessed September 26, 2026.
- Anthropic. Claude Opus 5, 2026d. URL <https://www.anthropic.com/news/claude-opus-5>. Accessed September 26, 2026.
- Anthropic. Claude Sonnet 5, 2026e. URL <https://www.anthropic.com/news/claude-sonnet-5>. Accessed September 26, 2026.
- Leonhard Applis, Yuntong Zhang, Shanchao Liang, Nan Jiang, Lin Tan, and Abhik Roychoudhury. Unified software engineering agent as AI software engineer. In *ICSE*, 2026. URL <https://www.cs.purdue.edu/homes/lintan/publications/USEagent-icse26.pdf>.
- CodeRabbit. CodeRabbit: AI code reviews. <https://www.coderabbit.ai>, 2026. Accessed September 26, 2026.
- Cursor. Bugbot. <https://cursor.com/docs/bugbot>, 2025. Accessed September 26, 2026.
- Cursor. Cursor, 2026. URL <https://cursor.com/docs/cli/overview>. Software documentation; accessed September 26, 2026.
- DeepSeek. DeepSeek coding agent, 2026a. URL <https://www.deepseek.com/harness/en/>. Accessed September 26, 2026.
- DeepSeek. DeepSeek-V4, 2026b. URL <https://deepseek.com/en/news/v4-preview/>. Accessed September 26, 2026.
- DeepSeek. DeepSeek-V4.1-Flash, 2026c. URL <https://www.deepseek.com/en/news/deepseek-v4-1-flash/>. Accessed September 26, 2026.
- Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, Chetan Rane, et al. SWE-bench pro: Can AI agents solve long-horizon software engineering tasks? *arXiv preprint arXiv:2509.16941*, 2025. URL <https://arxiv.org/abs/2509.16941>.
- Henry Kiss Ehrenberg, Vincent Sunn Chen, Austin W. Hanjie, Gabriel Orlanski, Frederic Sala, and Karthik R. Narasimhan. Senior SWE-bench: evaluating agents on their ability to act as senior engineers, 2026. URL <https://snorkel.ai/leaderboard/senior-swe-bench/>.

- Alexander Frömmgen, Jacob Austin, Peter Choy, Nimesh Ghelani, Lera Kharatyan, Gabriela Surita, Elena Khrapko, Pascal Lamblin, Pierre-Antoine Manzagol, Marcus Revaj, Maxim Tabachnyk, Daniel Tarlow, Kevin Villela, Daniel Zheng, Satish Chandra, and Petros Maniatis. Resolving code review comments with machine learning. In *ICSE-SEIP*, 2024. URL <https://doi.org/10.1145/3639477.3639746>.
- Jatin Ganhotra, Sami Serhan, Antonio Abu Nassar, Avraham Shinnar, Ziv Nevo, and Martin Hirzel. Resolving Java code repository issues with iSWE agent. *arXiv preprint arXiv:2603.11356*, 2026. URL <https://arxiv.org/abs/2603.11356>.
- GitHub. Copilot code review now generally available. <https://github.blog/changelog/2025-04-04-copilot-code-review-now-generally-available/>, 2025.
- Google. Review GitHub code using Gemini Code Assist. <https://developers.google.com/gemini-code-assist/docs/review-repo-code>, 2026. Accessed September 26, 2026.
- Greptile. Greptile: AI code review. <https://www.greptile.com>, 2026. Accessed September 26, 2026.
- Kang He and Kaushik Roy. SWE-Adept: An LLM-based agentic framework for deep codebase analysis and structured issue resolution. *arXiv preprint arXiv:2603.01327*, 2026. URL <https://arxiv.org/abs/2603.01327>. Accepted at AACL-IJCNLP 2026 (Main).
- Wenqi Huang, Charley Lee, Leonard Tng, and Serena Ge. DeepSWE: Measuring frontier coding agents on original, long-horizon engineering tasks. *arXiv preprint arXiv:2607.07946*, 2026. URL <https://arxiv.org/abs/2607.07946>.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *ICLR*, 2024. URL <https://openreview.net/forum?id=VTF8yNQM66>.
- Shiqi Kuang, Zhao Tian, Kaiwei Lin, Chaofan Tao, Shaowei Wang, Haoli Bai, Lifeng Shang, and Junjie Chen. REAgent: Requirement-driven LLM agents for software issue resolution. *arXiv preprint arXiv:2604.06861*, 2026. URL <https://arxiv.org/abs/2604.06861>.
- Zhiyu Li, Shuai Lu, Daya Guo, Nan Duan, Shailesh Jannu, Grant Jenks, Deep Majumder, Jared Green, Alexey Svyatkovskiy, Shengyu Fu, and Neel Sundaresan. Automating code review activities by large-scale pre-training. In *ESEC/FSE*, 2022. URL <https://doi.org/10.1145/3540250.3549081>.
- Shuquan Lian, Juncheng Liu, Yazhe Chen, Yuhong Chen, and Hui Li. SWE-AGILE: A software agent framework for efficiently managing dynamic reasoning context. In *Findings of ACL*, 2026. URL <https://aclanthology.org/2026.findings-acl.868/>.
- Mike Merrill, Alexander Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Shin, Thomas Walshe, E Kelly Buchanan, et al. Terminal-Bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *ICLR*, 2026. URL <https://openreview.net/pdf?id=a7Qa4CcHak>.
- Meta. Muse Code, 2026a. URL <https://dev.meta.ai/docs/muse-code>. Accessed September 26, 2026.
- Meta. Muse Spark 1.3, 2026b. URL <https://research.meta.ai/blog/introducing-muse-spark-1-3>. Accessed September 26, 2026.
- Samuel Miserendino, Michele Wang, Tejal Patwardhan, and Johannes Heidecke. SWE-Lancer: Can frontier LLMs earn \$1 million from real-world freelance software engineering? In *ICML*, 2025. URL <https://proceedings.mlr.press/v267/miserendino25a.html>.
- Moonshot AI. Kimi Code, 2026a. URL <https://www.kimi.com/code/en>. Accessed September 26, 2026.
- Moonshot AI. Kimi K3, 2026b. URL <https://forum.moonshot.ai/t/kimi-k3-is-here-our-most-capable-model/480>. Accessed September 26, 2026.

- Niels Mündler, Mark Niklas Mueller, Jingxuan He, and Martin Vechev. SWT-Bench: Testing and validating real-world bug-fixes with code agents. In *NeurIPS*, 2024. URL <https://openreview.net/forum?id=9Y8zU011EQ>.
- Magnus Müller and Gregor Žunič. Browser Use: Enable AI to control your browser, 2024. URL <https://github.com/browser-use/browser-use>.
- OpenAI. Introducing Codex, 2025. URL <https://openai.com/index/introducing-codex/>. Accessed September 26, 2026.
- OpenAI. Review GitHub pull requests with Codex. <https://developers.openai.com/codex/integrations/github>, 2026a. Accessed September 26, 2026.
- OpenAI. GPT-5.6, 2026b. URL <https://openai.com/index/gpt-5-6/>. Accessed September 26, 2026.
- OpenAI. GPT-6 Astra, 2026c. URL <https://openai.com/index/gpt-6-astra/>. Accessed September 26, 2026.
- Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with SWE-Gym. In *ICML*, 2025. URL <https://proceedings.mlr.press/v267/pan25g.html>.
- Qwen Team. Qwen3.8, 2026. URL <https://github.com/QwenLM/Qwen3.8/blob/main/README.md>. Accessed September 26, 2026.
- Mohit Raghavendra, Soham Dan, Miguel Romero Calvo, Yannis Yiming He, Johannes Baptist Mols, Gautam Anand, Cole McCollum, Edgar Arakelyan, Vijay Bharadwaj, Andrew Park, Jeff Da, MohammadHossein Rezaei, Bing Liu, Brad Kenstler, and Yunzhong He. SWE Atlas: Benchmarking coding agents beyond issue resolution. *arXiv preprint arXiv:2605.08366*, 2026a. URL <https://arxiv.org/abs/2605.08366>.
- Mohit Raghavendra, Anisha Gunjal, Aakash Sabharwal, and Yunzhong He. SWE-INTERACT: Reimagining SWE benchmarks as user-driven long-horizon coding sessions. *arXiv preprint arXiv:2606.30573*, 2026b. URL <https://arxiv.org/abs/2606.30573>.
- Muhammad Shihab Rashid, Christian Bock, Yuan Zhuang, Alexander Buchholz, Tim Esler, Simon Valentin, Luca Franceschi, Martin Wistuba, Prabhu Teja Sivaprasad, Woo Jung Kim, Anoop Deoras, Giovanni Zappella, and Laurent Callot. SWE-PolyBench: A multi-language benchmark for repository level evaluation of coding agents. *arXiv preprint arXiv:2504.08703*, 2025. URL <https://arxiv.org/abs/2504.08703>.
- Yuling Shi, Jinghan Xu, Kelin Fu, Wenhao Zeng, Shilin He, Lei Zhang, Yue Liu, Zelin Zhao, Terry Yue Zhuo, Jialun Cao, Siyu Ye, Tianyu Liu, Kai Cai, Shing-Chi Cheung, and Xiaodong Gu. SWE-Bench ProMax: Benchmarking agents on large-scale multilingual code refactoring. *arXiv preprint arXiv:2608.09802*, 2026. URL <https://arxiv.org/abs/2608.09802>. Conference version at COLM 2026.
- Zhao Tian, Pengfei Gao, Junjie Chen, and Chao Peng. Agent-based ensemble reasoning for repository-level issue resolution. In *ICSE*, 2026. URL <https://conf.researchr.org/details/icse-2026/icse-2026-research-track/233/Agent-Based-Ensemble-Reasoning-for-Repository-Level-Issue-Resolution>.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, et al. OpenHands: An open platform for AI software developers as generalist agents. In *ICLR*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/a4b6ad6b48850c0c331d1259fc66a69c-Abstract-Conference.html.
- xAI. Grok 4.6, 2026. URL <https://x.ai/news/grok-4-6>. Accessed September 26, 2026.
- Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Demystifying LLM-based software engineering agents. In *FSE*, 2025. URL <https://doi.org/10.1145/3715754>.

- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. OSWorld: Benchmarking multimodal agents for open-ended tasks in real computer environments. In *NeurIPS*, 2024. URL https://proceedings.nips.cc/paper_files/paper/2024/hash/5d413e48f84dc61244b6be550f1cd8f5-Abstract-Datasets_and_Benchmarks_Track.html.
- John Yang, Carlos Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *NeurIPS*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/5a7c947568c1b1328ccc5230172e1e7c-Abstract-Conference.html.
- John Yang, Carlos E Jimenez, Alex L Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press, Niklas Muenighoff, Gabriel Synnaeve, Karthik R Narasimhan, Diyi Yang, Sida I Wang, and Ofir Press. SWE-bench multimodal: Do AI systems generalize to visual software domains? In *ICLR*, 2025. URL <https://openreview.net/forum?id=riTiq3i21b>.
- John Yang, Kilian Lieret, Jeffrey Ma, Parth Thakkar, Dmitrii Pedchenko, Sten Sootla, Emily McMilin, Pengcheng Yin, Rui Hou, Gabriel Synnaeve, Diyi Yang, and Ofir Press. ProgramBench: Can language models rebuild programs from scratch? *arXiv preprint arXiv:2605.03546*, 2026. URL <https://arxiv.org/abs/2605.03546>.
- Naimeng Ye, Xiao Yu, Ruize Xu, Tianyi Peng, and Zhou Yu. AI agents for web testing: A case study in the wild. *arXiv preprint arXiv:2509.05197*, 2025.
- Z.ai. GLM-5.3, 2026. URL <https://z.ai/blog/glm-5.3>. Accessed September 26, 2026.
- Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Shulin Xin, Linhao Zhang, Qi Liu, Li Aoyan, Lu Chen, Xiaojian Zhong, et al. Multi-SWE-bench: A multilingual benchmark for issue resolving. In *NeurIPS*, 2025. URL https://papers.neurips.cc/paper_files/paper/2025/hash/5afa9cb1e917b898ad418216dc726fbd-Abstract-Datasets_and_Benchmarks_Track.html.
- Zhengran Zeng, Ruikai Shi, Keke Han, Yixin Li, Kaicheng Sun, Yidong Wang, Zhuohao Yu, Rui Xie, Wei Ye, and Shikun Zhang. SWR-Bench: Assessing LLM performance in real-world code review comment generation. In *FSE*, 2026. URL <https://conf.researchr.org/details/fse-2026/fse-2026-research-papers/78/SWR-Bench-Assessing-LLM-Performance-in-Real-World-Code-Review-Comment-Generation>.
- Daoan Zhang, Shuo Zhang, Zijian Jin, Jiebo Luo, Shengyu Fu, and Elsie Nallipogu. Sphinx: Benchmarking and modeling for LLM-driven pull request review. *arXiv preprint arXiv:2601.04252*, 2026. URL <https://arxiv.org/abs/2601.04252>.
- Linhao Zhang, Shilin He, Chaoyun Zhang, Yu Kang, Bowen Li, Chengxing Xie, Junhao Wang, Maoquan Wang, Yufan Huang, Shengyu Fu, Elsie Nallipogu, Qingwei Lin, Yingnong Dang, Saravan Rajmohan, and Dongmei Zhang. SWE-bench goes live! In *NeurIPS*, 2025. URL <https://www.microsoft.com/en-us/research/publication/swe-bench-goes-live/>.
- Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. AutoCodeRover: Autonomous program improvement. In *ISSTA*, 2024. URL <https://doi.org/10.1145/3650212.3680384>.
- Pujun Zheng, Zixin Shang, Shufan Jiang, Wenhui Tian, Dongsheng Zhu, Zerun Ma, Dingbo Yuan, and Qi Zhang. SWE-Bench Pro Verified: A reliable benchmark for software engineering agents. *arXiv preprint arXiv:2609.08149*, 2026. URL <https://arxiv.org/abs/2609.08149>.
- Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A realistic web environment for building autonomous agents. In *ICLR*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/4410c0711e9154a7a2d26f9b3816d1ef-Abstract-Conference.html.

Appendix

A Prompts

Every agent receives the prompt below, with `{{TASK_DESCRIPTION}}` replaced by the task’s open-ended request. In the ablation of Section 4.1, we remove the bullet that asks the agent to test in the browser, the bullet that mentions Chromium and Playwright, and the verification bullet about opening the product UI. For the detailed prompt of Section 4.2, the task description is replaced by a description of the target bug; the rest of the prompt is unchanged.

```
You are working in '/workspace/product'.

Please perform the following task:

{{TASK_DESCRIPTION}}

Important notes:
- Use only this checkout and these instructions. Do not search external repositories or known fixes.
- Do not fetch Git history, add remotes, or commit.
- This environment has no outbound internet for packages or source. npm, PyPI, GitHub, CDNs, and similar hosts are blocked and will return 403. Do not install, fetch, download, or rebuild anything from the network. Dependencies and tools are already on disk. If a command fails with a network or 403 error, treat that as final: do not retry, switch mirrors, or keep waiting. Continue with the local checkout and already-installed tools only.
- The product server is published on this machine at 127.0.0.1:13200.
- Use the browser to do end-to-end testing to discover and test issues. Drive the product UI like a real user in the tasked area and look for bugs that show up during normal use. Do not stop after the first issue you can patch; keep exploring other screens and controls in the same area.
- Chromium and Playwright are preinstalled.

How to verify your work:
- Product commands must run in the product container through '/workspace/task-shell'.
- Open the product UI at http://127.0.0.1:13200 with the preinstalled Playwright/Chromium in this environment. Do not launch the browser through '/workspace/task-shell'.
- Do not call 'yarn', 'pnpm', or 'npm' directly; use the 'product-*' wrappers in the container PATH.
- Find the right test or build scripts from package.json, README, or existing CI, then run them through '/workspace/task-shell'.

Deliverable:
- Modify product source in place and leave a clean, reviewable diff with at least one real code change.
- An empty diff is not acceptable. Do not commit.
```

For the Focalboard task, the two task descriptions are:

```
Open-ended: Find and fix issues in Focalboard’s board sharing and membership.

Detailed: Fix Focalboard board sharing so adding a member via Share also puts that board in the member’s default sidebar category. After a user shares a board with another member, that member should see the board under Boards in their left sidebar.
```

B Example Behavior Test

Table 6 lists the complete target behavior test for the Focalboard task in Figure 5, reproduced verbatim from the benchmark. The verifier follows these steps in the browser and returns a verdict from the final step. Setup problems, such as a failed login, are reported separately from product failures.

Step	Evaluation instruction
Target behavior test: After miraDunn shares a newly created board with kadeRue via Share, kadeRue’s left sidebar Boards category must list that board. On the buggy build kadeRue is a member but the board never appears in his sidebar.	
1	Open exactly <code>http://127.0.0.1:13200/login</code> and wait until the official Focalboard Log in page is visible: a styled rounded login card (not a bare unstyled text shell), the heading Log in, the username field with placeholder Enter username (id login-username), the password field with placeholder Enter password (id login-password), a filled Log in submit button, and the link or create an account if you don’t have one.
2	Type exactly miraDunn into the username field. Type exactly Slate6quay into the password field. Click the Log in button. Wait until the signed-in Focalboard chrome loads: the left sidebar shows the Focalboard logo wordmark and version, and + Add board. If a Welcome To Boards page appears, click No thanks, I’ll figure it out myself. If login fails or the sidebar never appears, stop and treat that as a setup problem, not a product pass/fail.
3	In the left sidebar click + Add board. Wait until the official Create a board template selector is visible (heading Create a board, templates on the left, preview on the right, buttons Use this template and Create empty board). Click Create empty board. Wait until a new board canvas opens with the title placeholder Untitled board and a top-right Share button.
4	Click the Untitled board title (the large board title, not the view name Board view). Type exactly Q3 Launch Desk and press Enter or click away so the title saves. Confirm the board title now reads Q3 Launch Desk and that the same name appears under the sidebar Boards category. If the empty board never opened or the title would not save, stop and treat that as a setup problem.
5	Click the top-right Share button (title Share board). Wait until the Share Board dialog is visible with the search placeholder Search for people and channels. Click that search field, type exactly kadeRue, and wait until a people result for @kadeRue appears. Click the @kadeRue result so kadeRue is added. Wait until the member list in the same dialog shows kadeRue (not only miraDunn / You). Close the dialog. If kadeRue cannot be found or added, stop and treat that as a setup problem, not the sidebar bug.
6	Open the Focalboard logo menu at the top of the left sidebar (the wordmark Focalboard). Click Log out. Wait until the Log in page is visible again.
7	Type exactly kadeRue into the username field. Type exactly Moss3vale into the password field. Click Log in. Wait until kadeRue’s signed-in Focalboard chrome loads (sidebar Focalboard logo and + Add board). If Welcome To Boards appears, click No thanks, I’ll figure it out myself. Use the browser reload/refresh once on this signed-in page so the sidebar refetches categories. If login fails, stop as a setup problem.
8	After that refresh, look only at kadeRue’s left sidebar list under the Boards category (the system category named Boards). Do not open Find Boards / the board switcher (Ctrl/Cmd+K) and do not paste a board URL. Judge only whether the sidebar category lists a board titled Q3 Launch Desk.
9	Final verdict — FAIL only if kadeRue is signed in, the left sidebar Boards category is visible, and Q3 Launch Desk is absent from that sidebar list (the shared board never landed in kadeRue’s default category). PASS only if that same sidebar Boards category lists Q3 Launch Desk. Do not PASS because miraDunn could share, or because kadeRue can log in. Do not FAIL for login/navigation/template-selector problems. Do not use fetch, XHR, DevTools Network, API JSON, or database reads for setup or pass/fail.

Table 6. Target behavior test for the Focalboard task. The verifier follows these instructions in the browser. Wording is reproduced verbatim; only typesetting and line wrapping are changed.