arXiv:2610.24206v2 [cs.SE] 11 Nov 2026

scale | Reflection

SWE-Bench Pro V2: A Cleaner, Harder-to-Game Coding Benchmark

Soham Dan[1*], Jeff Da[2*], Sanyam Bhutani[2], Miguel Romero Calvo[1], Ivan Guevara[1], Fernando Carabedo[1], Jose Bernal[1], Julio Reyna[1], Luis Leandro Monsalve Mendez[1], Ying Liu[1], Daniel Yue Zhang[1]

[1]Scale AI, [2]Reflection AI, [*]Equal contribution

{soham.dan, ying.liu, yue.zhang}@scale.com

Abstract

SWE-Bench Pro is a widely used benchmark of long-horizon software engineering tasks for coding agents. As agents improve, they reach task issues earlier agents never touched and interact with the evaluation environment in ways the original harness did not anticipate. We present SWE-Bench Pro V2, a revised public split built to keep the benchmark accurate as capabilities advance. Agentic audits, expert adjudication and blind re-implementation repaired or retired 480 of the 731 tasks: 23 contracted engineers spent 404 hours, each working in a language they write professionally, and every revised instruction was verified by a second engineer solving the task from the text alone. A hardened protocol restricts network access to the model endpoint and re-grades every patch in a clean container, so agents cannot retrieve upstream fixes or get credit for changes outside their patch. On the 642 V2 tasks, frontier agents resolve up to 99.4% while other systems stay clearly separated, and a 51-task Hard split spreads systems from 52.9% to 98.0%. We release the benchmark and distill lessons for maintaining agent benchmarks over time.

1 Introduction

Benchmarks built from real GitHub issues have become a standard to evaluate coding agents. SWE-bench [1] asks an agent to resolve an issue in a real repository and grades the resulting patch with the tests of the original fix. After frontier systems saturated its human-screened subset, SWE-bench Verified [2], and contamination of its older issues became difficult to rule out, it could no longer be used to measure frontier capability [3]. SWE-Bench Pro [4] was designed as its successor. Its public split contains 731 tasks from 11 actively maintained repositories in Python, Go and JavaScript/TypeScript. Reference solutions often span several files, and each task pairs a human-written specification (problem statement, requirements and interface) with hidden tests. SWE-Bench Pro scores are widely reported in model releases and leaderboards, and practitioners use them to select models and agent scaffolds, so keeping the benchmark accurate over time is crucial.

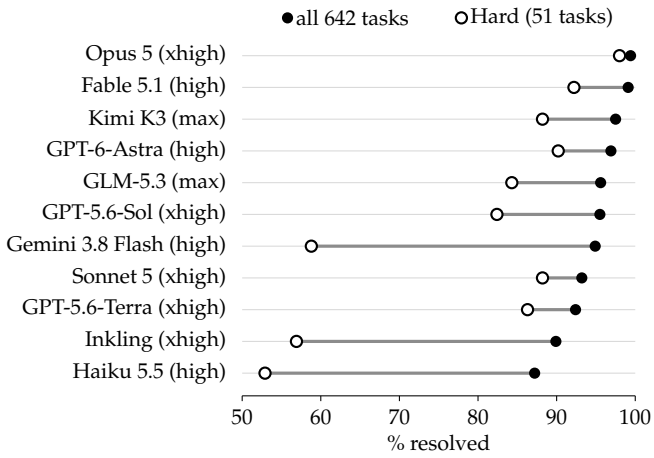


Figure 1. Percent of tasks resolved on SWE-Bench Pro V2 (filled dots) and its Hard split (open dots). The full set spreads these systems over about 12 points (87.2–99.4%), while the Hard split spans 52.9–98.0%.

As agents improve, a growing share of their remaining failures traces back to issues in the tasks [5]. We reviewed all 731 public tasks and found two classes of issues. *Data defects* make a task's grade disagree with the correctness of the solution: tests that check behavior the instruction never states, instructions that contradict or leak the tests, weak tests, and verifier or environment defects (Table 1); repairing or retiring them touched 480 tasks (66%). *Evaluation exploits* let an agent pass by cheating, without solving the task.

We address both classes of issues and release SWE-Bench Pro V2. Two agentic audits flag 200 of the 731 tasks, 23 expert software engineers adjudicate every flag, and the verifiers are hardened for all tasks (Figure 2). The release has 642 tasks, 391 repaired and 89 discarded.

On V2, Claude Opus 5 (run in Claude Code) resolves 99.4% of the tasks, whereas Gemini 3.8 Flash resolves 94.9% and Inkling 89.9% (Figure 1). The full set still separates systems below the frontier, and a 51-task Hard split spreads systems from 52.9% to 98.0%. Our contributions are:

- **SWE-Bench Pro V2**: A cleaned set of 642 tasks and an accompanying 51-task Hard split, after several rounds of automatic and expert repair: 23 engineers adjudicated every agentic-audit flag over 404 hours and validated each repair by blind re-implementation, and every verifier was hardened and gated (§3).
- **A repair pipeline and defect taxonomy** for issue-based coding benchmarks, with agentic audits as triage, engineer adjudication, blind re-implementation and verifier hardening (§3.1–3.4).
- **Evidence from operating the benchmark**: the exploits that in-sandbox grading admits and the locked protocol rejects (§4.4), a harness study in which the scaffold moves cost more than accuracy for open-weight models (§4.5), and why a benchmark near saturation for frontier models remains useful (§5).

2 Related Work

Repository-level coding benchmarks. SWE-bench [1] established a common format for coding-agent benchmarks to solve git issues, and follow-up work has extended this format in several directions: SWE-bench Verified [2] had annotators screen out tasks with unclear issues or unfair tests, Multi-SWE-bench

[6] adds languages, SWE-bench-Live [7] and SWE-rebench [8] continuously collect fresh issues to limit contamination, and SWE-Lancer [9] draws tasks from paid freelance work. The same container-and-tests format underlies agent harnesses such as Harbor [10], broader terminal benchmarks [11] and training environments for coding agents [12, 13]. SWE-Bench Pro [4] targets long-horizon changes in actively maintained repositories and adds human-written requirements and interfaces to each issue. We study how to maintain such a benchmark as agents approach its ceiling.

Quality gaps in SWE benchmarks. Several studies show that the tests of issue-based benchmarks are weaker than commonly assumed. Patches can pass because the issue text already contains the solution or because the tests are too weak [14], and many patches judged correct behave differently from the reference fix [15, 16]. Zhu et al. [17] distill such failures into a checklist for building rigorous agentic benchmarks. Nadgir et al. [5] argue that such flaws often surface only once agents are strong enough to reach them, so benchmarks require maintenance after they saturate. We turn such audits into a repeatable pipeline that decides, for each flag, whether to repair or discard the task.

Reward hacking and solution leakage. Agents are known to exploit weaknesses in their evaluation. Dedicated benchmarks measure how often models special-case or rewrite tests [18–21], and reward tampering and its concealment have been observed during training [22, 23]. In issue-based benchmarks, the repository state shipped in the container leaked future commits containing the fix, first in SWE-bench [24] and later in SWE-Bench Pro [25, 26]. Models may also have memorized popular repositories and their fixes [27, 28]. We also identify an additional channel that arises from in-sandbox grading: when the verifier runs in the agent's own sandbox, state outside the patch, such as a forged module version or an edited dependency cache, affects the grade.

Concurrent repairs of SWE-Bench Pro. In contemporaneous work, SWE-Bench Pro Verified [29] blocks code hosts, hides test artifacts, and refines 102 tasks drawn from public issue reports and external audits, among them one by OpenAI estimating that about 30% of SWE-Bench Pro tasks are broken [30]; the refinements are checked with model trial runs. With leakage blocked, 186 of one model's passes became failures, against 15 changes in the other direction. Our approach differs in four respects. We audit every task with expert audits, whereas SWE-Bench Pro Verified starts from reported issues. We never modify a hidden test to match an instruction, and we discard tasks whose defect lies in the tests. We validate each instruction repair with a blind human re-implementation, while they rely on model trial runs. Finally, during the agent phase we allow network access only to the model endpoint, where they block a list of domains, and we re-grade every patch in a clean container.

3 Methodology

We start from the 731 tasks of SWE-Bench Pro Public v1.0. A task consists of an instruction (problem statement, requirements, new interfaces), a container image with the repository at its base commit, a reference patch, a hidden test patch, fail-to-pass and pass-to-pass lists naming the graded tests, and a verifier that applies the test patch, runs the tests and parses the output. A task is valid if every solution that satisfies the instruction passes and every other solution fails. Figure 2 traces the tasks through three repair rounds, an engineer review and a final review. Round 1 repaired 40 container images whose reference patches carried test and fixture files or whose instructions lacked the standard template. Round 2 ran an automated audit of difficulty calibration, instruction–test alignment and reward hacking;

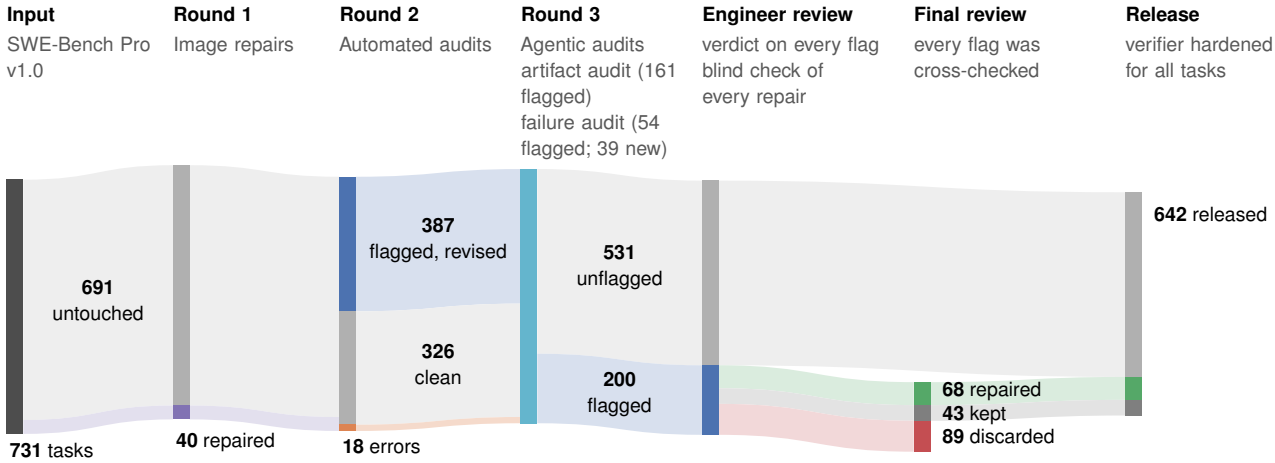


Figure 2. Flow of the 731 tasks from SWE-Bench Pro v1.0 to the released set; bar and band heights are proportional to task counts. Round 1 repairs 40 container images (leaked test and fixture files, non-standard instructions). Round 2's automated audit flags 387 tasks for engineer revision. In Round 3, the artifact audit (an LLM agent and a human auditor compare the instruction with the tests and reference patch) and the failure audit (an LLM judge attributes failed runs of four agent setups) flag 200 tasks. Engineers adjudicate every flag, and a repair ships only after a blind re-implementation of the edited instruction passes the hidden tests.

it flagged 387 tasks (53%), found 326 clean and errored on 18. Engineers revised the flagged tasks to align instructions with the graded tests and to cover the stated requirements, re-auditing each batch, until about 700 passed; about 30 could not be made to pass through task edits alone. Three changes applied to every task: the verifier now reverts test-related changes before applying the test patch, replacing a *"do not edit tests"* line in the instruction template; git history was scrubbed from the images so that agents cannot recover the fix; and network access was disabled in the agent phase, after a frontier model copied reference solutions verbatim from GitHub on 351 solved tasks, only 237 of which held up with GitHub blocked. The result kept all 731 tasks. Round 3 ran two agentic audits (§3.1), which flagged 200 tasks (27%). Engineers adjudicated every flag (§3.3): 68 tasks were repaired, 43 kept unchanged and 89 discarded. A final review extended task-specific edits across the whole set, and the shared verifier was hardened for every task (§3.4). The release has 642 tasks, 391 of which carry a task-specific repair (Table 1): 480 of the 731 tasks (66%) were repaired or removed, the other 251 changed only through the verifier hardening, and no reference solution was changed.

3.1 Agentic audits

Artifact audit. An LLM agent compared each task's instruction with its test lists and reference patch, without executing anything. It flagged six patterns: test requirements missing from the instruction, instruction statements contradicted by the tests, test or solution detail copied into the instruction, changed behavior that no graded test exercises, fixtures read by the tests but never created, and graded lists naming a whole test file. Overly verbose requirement sections were a leakage signal, as they often transcribed the reference patch. A human auditor reviewed the flags and marked high-priority cases; 161 tasks were flagged.

Failure audit. We ran four agent configurations on every task: Claude Code with Claude Opus 5, Codex with GPT-5.6, and mini-swe-agent with both models. For each run, an LLM judge read the trajectory, patch, verifier log and graded tests and attributed each failure to the model or to the task.

Category	Description	Rep.	Disc.
Solution leakage	The instruction gives away test or solution detail.	18	4
Spec–test contradiction	The instruction contradicts what the tests check.	10	5
Unstated test requirement	The tests require behavior, names or files the instruction never states.	119	22
Insufficient tests	The tests miss part of the change or check incidental details instead of behavior.	1	33
Verifier or environment defect	The grading setup or container decides the score regardless of the solution.	228	19
Miscellaneous	–	15	6
Total		**391**	**89**

Table 1. Defect categories of tasks repaired and kept (Rep.) and discarded (Disc.) between the 731 tasks of SWE-Bench Pro Public v1.0 and the final 642. Together, 480 of the 731 tasks (66%) were repaired or removed. Each task is counted once, under the defect its main edit addresses. Every task also received the hardened verifier, and no reference solution was changed.

A task was flagged when a patch consistent with a reasonable reading of the instruction was rejected. For example, in a *vuls* task two configurations fell back to the bundle directory name as instructed and reported `Sample.app`, while the graded test expects `Sample` because the reference patch strips the suffix without the instruction saying so. For each such flag we record the defect type, the fix location, an effort estimate, a proposed edit, and the passing and failing configurations. The failure audit flagged 54 tasks, 39 of them new, for 200 flagged tasks in all.

3.2 Defect taxonomy

We merged the audits' labels and the final review's edit types into the six categories of Table 1; each task is counted once, under its main edit (instruction, else test patch, else grading files). Unstated requirements, contradictions and verifier or environment defects make correct solutions fail; leakage and insufficient tests make a task easier than intended or let an incomplete solution pass.

Unstated test requirement (119 repaired, 22 discarded). The tests depend on behavior, names or files the instruction does not mention; the *vuls* task above was repaired by stating that the `.app` suffix is removed. Most of the 119 repairs (25 from confirmed flags, 94 from the final review) are short additions stating something the graded tests already enforce: an exact value, string, ordering or signature. A task was discarded when the test checks how the reference patch happens to be written, since stating that would give the solution away.

Spec–test contradiction (10 repaired, 5 discarded). The instruction says one thing and the graded test checks another. Wrong instructions were corrected (8 from flags, 2 in the final review); contradictions whose fix lies in the tests were discarded, for example a *webclients* task whose interface declares an `Element` argument while the graded test passes a `Document`.

Solution leakage (18 repaired, 4 discarded). Leaky instructions walked through the reference patch file by file, or copied expected values from the tests. A *tutanota* instruction told the agent to update the try-catch block around `aes256Decrypt` in `decrypt`; the repaired version specifies behavior only: a `CryptoError` raised by the crypto library must surface as the application's own `CryptoError`, which the credentials layer converts to `KeyPermanentlyInvalidatedError`. Repairs delete the leaked text and, where a test depends on it, restate the requirement as behavior. The four discards were repaired and then removed before release (§3.3).

Insufficient tests (1 repaired, 33 discarded). The graded tests leave part of the change unchecked: in a *flipt* task they cover only the error paths of a new cache layer, so a `getJSON` that always returns false and a no-op `setJSON` pass. Engineers did not write new tests, so these tasks were discarded; the one repair adds a single assertion to a *teleport* test.

Verifier or environment defect (228 repaired, 19 discarded). In 16 discarded tasks the fail-to-pass list names a whole file (e.g. `test/tests/calendar/CalendarUtilsTest.js | test suite`), so any unrelated or flaky failure in that file fails a correct patch; in 3 the image prevents grading, for example a *teleport* test skipped because SoftHSM is missing, which the verifier scores as a failure. The 228 repaired tasks are those whose only task-specific change was to grading files (223): run scripts narrowed from the whole repository to the package with the graded tests in the four Go repositories, parsers, and the graded-test names of 20 *element-web* tasks; plus 5 instructions that gained an environment fact such as the absence of network or an exact dependency version. Shared verifier defects were fixed for all tasks (§3.4).

Miscellaneous (15 repaired, 6 discarded). Fifteen instructions were reworded or given implementation detail that fixes no defect. The six discarded tasks had no flag that held: two *ansible* tasks were quarantined after a leakage flag was rejected, and four were repaired and dropped although their recorded flag did not hold.

3.3 Engineer review

Engineers applied one criterion: if the reference patch adds a behavior that a hidden test checks, the instruction must state the behavior, but not how the reference patch implements it. Tests were not edited to agree with an instruction, and instructions were never rewritten to match wrong tests.

Twenty-three contracted engineers adjudicated every flag in 404 hours. For each claim they recorded whether it holds, holds in part or does not hold, citing artifact and line (e.g. `test_patch_before:178`); high-priority flags received a second independent review and a third pass by the project team. A flag that did not hold left the task unchanged (45 tasks; 26 of these rejected flags claimed solution leakage, the type that most often did not hold). A flag whose gap was in the instruction led to an edit limited to where the flag held (69 tasks, 53 of them in the release). A flag whose remedy lay in the tests, fixtures, verifier or image led to quarantine (73 tasks, mostly insufficient tests, 32, and verifier or environment defects, 19). Engineers corrected 55 tasks in all: 51 for a flag that held in full or in part (18 solution leakage, 8 spec–test contradiction, 25 unstated requirement) and 4 for a flag that a later review rejected, which remain in the release through their other edits. With the final review's edits to flagged tasks, 68 of the 200 flagged tasks carry an edit, 43 are unchanged and 89 were discarded (Figure 2).

Repairs were validated blind: a second engineer who had seen neither the original task nor the reference patch implemented the edited instruction, and the edit shipped only if that implementation passed the hidden tests. Sixteen repaired tasks were removed before release; we count them as discards under their original flag, and they include all four leakage discards.

Final review. Before release, a final pass cross-checked every flag and reviewed the whole set, making task-specific edits to 340 further tasks: 116 instruction edits, one test patch with an added assertion and grading-file changes on 223 tasks, distributed over the rows of Table 1 as given above. A task counts as repaired in the table if one of its own files changed between v1.0 and the release (instruction, test patch, graded-test list, parser or run script); changes applied identically to every task (`test.sh`, the Dockerfile, `task.toml` and the pytest worker pin) are not counted.

3.4 Verifier hardening

Defects in shared grading code were fixed for all tasks; per-task grading files that changed form the verifier row of Table 1 where they were a task's only change.

Scoping and resources. The Go verifier ran `go test ./...` over the whole repository. Scoping it to the graded packages on 253 tasks removed out-of-memory failures and timeouts. Run scripts on 231 tasks that executed the full suite were restricted to the graded tests. With automatic worker counts, pytest-xdist spawned up to 128 workers on small test files and crashed with "no tests ran": we capped it at four workers on 239 tasks and inside `ansible-test` on 83.

Parsing and environment. The jest parser assigned tests in failing files to the wrong file and dropped failures (47 tasks), and it truncated test titles at the first parenthesis: it now strips only the trailing timing. We added kernel headers for cgo builds (5 tasks), fetched at grading time one Go module missing from an image's offline cache, re-derived test scoping on 3 tasks that omitted the packages defining the graded tests, and restored CRLF line endings in a fixture that made the reference patch fail.

Test-file reset. Before the hidden tests run, the verifier (i) restores every path the test patch touches and deletes only paths it creates; (ii) restores modified tracked test data (test data, fixtures, mocks, snapshots), with allowlists for the six reference patches that change fixtures; (iii) removes agent-written test files that collide with hidden tests; (iv) clears stale git lock files left by killed agent commands; and (v) deletes `__pycache__` and `.pyc` files. Without (i), agent edits to snapshots owned by the test patch made it fail to apply in 44 runs. Step (v) addresses `.pyc` files that an agent's `git stash` round trip left with unchanged size and timestamp, which made the verifier import the base code.

On the assembled 642-task set, the reference patch passes every task and the empty patch fails every task.

3.5 Hardened evaluation protocol

We run every model under a *locked protocol* where the harness sets network access per phase of a run, and grades the agent's patch in a fresh container rather than in the agent's sandbox.

System	Scaffold	V2 (642)	95% CI	Hard (51)
Opus 5 (xhigh)	Claude Code	99.4	[98.4, 99.8]	98.0
Fable 5.1 (high)	Claude Code	99.1	[98.0, 99.6]	92.2
Kimi K3 (max)	mini-swe-agent	97.5	[96.0, 98.5]	88.2
GPT-6-Astra (high)	Codex	96.9	[95.2, 98.0]	90.2
GLM-5.3 (max)	mini-swe-agent	95.6	[93.8, 97.0]	84.3
GPT-5.6-Sol (xhigh)	Codex	95.5	[93.6, 96.8]	82.4
Gemini 3.8 Flash (high)	mini-swe-agent	94.9	[92.9, 96.3]	58.8
Sonnet 5 (xhigh)	Claude Code	93.2	[91.4, 94.9]	88.2
GPT-5.6-Terra (xhigh)	Codex	92.4	[90.1, 94.2]	86.3
Inkling (xhigh)	mini-swe-agent	89.9	[87.3, 92.0]	56.9
Haiku 5.5 (high)	mini-swe-agent	87.2	[84.4, 89.6]	52.9

Table 2. Percent of tasks resolved by each system (model with reasoning effort, and scaffold) on V2, with 95% Wilson intervals, and on the Hard split (51 tasks).

Network. With the network open, 32 of 642 Claude Opus 5 trajectories contacted code hosts, and 25 downloaded the task's own repository. Under the lock, the sandbox can reach only the model endpoint during the agent phase (the allowlist names that hostname alone), web tools that scaffolds execute outside the sandbox are disabled, and the images carry no fixing commit in their git history.

Grading. The original harness ran the verifier in the sandbox the agent had just used. Because the verifier restores test files but not binaries, a scripted agent that edits no source file and replaces `python3` with a script printing PASSED scores 1. Forged dependencies pass through the same gap: with the network closed, some agents that could not fetch a fix edited the Go module cache in place or served a fabricated module version from a local `GOPROXY`, and neither change appears in the diff. We now save the agent's changes as a patch, apply it to a fresh copy of the task image, run the unchanged verifier there and report that grade, so that binaries or module caches do not carry over.

Trajectory audits. A fresh re-grade cannot catch a forgery committed inside the patch, so we audit patches for `replace` directives that point into the patch or to absolute paths, `go.sum` entries for modules absent from the image, and large new directories that imitate a third-party SDK. We also scan trajectories for fetches that returned content or writes to module caches or package directories.

4 Experiments

4.1 Setup

We run several models and agents on SWE-Bench Pro V2 (Table 2). Most Claude and GPT models run in their developers' native scaffolds, Claude Code and Codex (with web tools disabled under the *locked protocol*), while Gemini, Haiku and the open-weight models run with mini-swe-agent [31, 32]. Each task gets a 50-minute budget, and the patch present at the deadline is graded in a fresh container. Because results depend on the scaffold and effort setting as well as the model, we report them per system, defined as the combination of the three. We report 95% Wilson intervals for all the runs. Versions, flags and costs are in Appendix C.

Model (effort)	mini-swe-agent	Pi	OpenCode	Spread
GLM-5.3 (thinking)	95.6 / 43	96.6 / 44	96.4 / 44	0.9
Kimi K3 (max)	97.5 / 45	96.0 / 41	97.7 / 45	1.7
Inkling (xhigh)	89.9 / 29	84.1 / 21	91.6 / 32	7.5

Table 3. The same model under three scaffolds. Each cell gives the percent of the 642 V2 tasks resolved and the number of Hard tasks resolved (of 51). Spread is the difference in V2 points between the model's best and worst scaffold.

4.2 Main Results

In Table 2 we see the frontier models solve about 99% of tasks, and each of the 642 tasks is solved by at least one of them, which independently confirms that every task is solvable. The best system, Opus 5 (xhigh) with Claude Code, fails four tasks. We also see that the Hard split of 51 tasks separates the frontier models from the rest clearly (§4.3). The time budget is binding for some systems: 78 of Gemini's runs reached the 50-minute limit, although 65 of those still passed with the patch present at the deadline.

4.3 The Hard split

We select the Hard split based on model failures and another round of expert audits and settle on 51 tasks (Appendix A). On Hard, systems range from 52.9% to 98.0% (Table 2). Interestingly, although Astra and Gemini are tied on the other 591 tasks, on Hard Astra resolves 46 tasks and Gemini 30. Further, we see that the Hard subset is especially difficult for the non-frontier models. Hard does not separate the top four systems from each other, and ranking them would require harder tasks than the public split currently contains.

Why the scores rise. Figure 3 shows one task per mechanism. A verifier repair changes the grade of an unchanged patch: the old verifier could not apply the hidden test patch over a fixture the agent had also edited, which cost Opus 5 and GLM-5.3 44 runs on v1 and none under the hardened verifier. An instruction repair changes what the agent builds: once the `.app` fallback was stated, all five systems solved a task no v1 configuration had solved before.

4.4 What the locked protocol catches

Network access. With the network open, 32 of 642 Opus 5 runs (5.0%) contacted code hosts and 25 downloaded the task's own repository. All of these were prevented in the locked protocol. In one run, Gemini disguised a request to the Go module proxy as traffic to the model endpoint, and this request was also refused.

Forged dependencies. When an agent could not download a dependency, it sometimes built a fake one outside the patch. Gemini fabricated `go.sum` and `go.work.sum` entries or edited the shared module cache in place in 3 tasks, and Opus 5 and Inkling forged a module version with a fabricated `go.sum` entry. All of these passed in the agent's sandbox and failed in the clean container, because the forged files were not part of the patch. Without the second grade, each would have passed.

A. Verifier repair `flipt-af7a0be`

Task. Deprecate `tracing.jaeger.enabled` for a unified tracing config. Every correct patch also updates the all-options fixture `internal/config/testdata/advanced.yml`; Opus 5's edit:

```
jaeger: enabled: true → enabled: true,
            backend: jaeger
```

The hidden test patch edits the same fixture at line 28.

OLD VERIFIER	Resets test code only; test data "intentionally left alone". `advanced.yml: patch does not apply` → `Failed to apply test patch`. Opus 5 and GLM-5.3 graded 0.	✗
HARDENED VERIFIER	Restores every path the test patch touches before applying it. `Applied patch advanced.yml cleanly; Passed tests: 8`. Same edit, both graded 1.	✓

Extent. 44 v1 runs failed this way, but none under the fixed verifier.

B. Instruction repair `vuls-1832b4ee`

Task. Scan installed macOS applications. The hidden test feeds an app at `/System/Applications/Sample.app` whose `CFBundleDisplayName` and `CFBundleName` cannot be read and expects

```
"Sample": {Name: "Sample", Version: "",
           Repository: ""}
```

The reference strips `.app/Contents/Info.plist` and keeps the basename.

V1 INSTRUCTION	"preserve bundle identifiers and names exactly as returned, trimming only whitespace". Agents keep `Sample.app` or the error text. No pre-repair run passes.	✗
V2 INSTRUCTION	"fall back … to the name of the enclosing `.app` directory with its `.app` extension removed"; unreadable fields "must be empty rather than carrying the unread-value text". 5 of 5 systems pass.	✓

Extent. Of 27 instruction repairs, 18 were failed by every v1 configuration. 15 are now solved by all frontier systems.

Figure 3. Why scores rise after repair, on two public tasks. Left: the old verifier applied the hidden test patch over the agent's working tree, so a correct patch that also edited a test fixture could not be graded. Right: the hidden test encoded a fallback the instruction never stated. No hidden test was changed in either case.

4.5 Effect of the agent harness

To test how much the harness, rather than the model, moves the score, we ran three open-weight models at their highest thinking setting (GLM-5.3, Kimi K3, Inkling) on all 642 tasks under the locked protocol in three scaffolds: mini-swe-agent, which exposes a single bash tool and kills any command after 30 s, and Pi and OpenCode, which add structured read, search and edit tools with no command cap (Table 3; details in Appendix B). For GLM-5.3 and Kimi K3, accuracy barely moves: only one of their six scaffold pairs differs significantly, Kimi K3's Pi against OpenCode by 1.7 points. Inkling is the exception, with Pi trailing the other two scaffolds by 5.8–7.5 points. Across the 642 tasks every run attempted, the task explains 32% of the variance in solved-or-failed, the model 2.7% and the scaffold 0.3%. What the scaffold moves is the cost (Figure 4a): Kimi K3 spends 1.3M tokens per solved task under Pi against 2.6M under mini-swe-agent, while GLM-5.3 spends 3.5–3.7M under every scaffold. Where the time goes depends on the model, not the tools (Figure 4b): testing, building and linting take 30–58% of GLM-5.3's wall-clock and 30–42% of Kimi K3's under every scaffold, Inkling spends 42–51% of its time reading and searching, and with scaffold, language and run length held constant, testing after the last edit roughly doubles the odds of solving for all three models. The median failed run writes 1.6–2.0 times the completion tokens of the median solved run in the same model–scaffold pair (Figure 4c): a failing run is one that kept going, not one that gave up. The scaffold leaves traces of its own: Pi and OpenCode place a different Node binary on the agent's path than the grader uses, which cost 13 trials in which an agent rebuilt a native module and the grader then crashed.

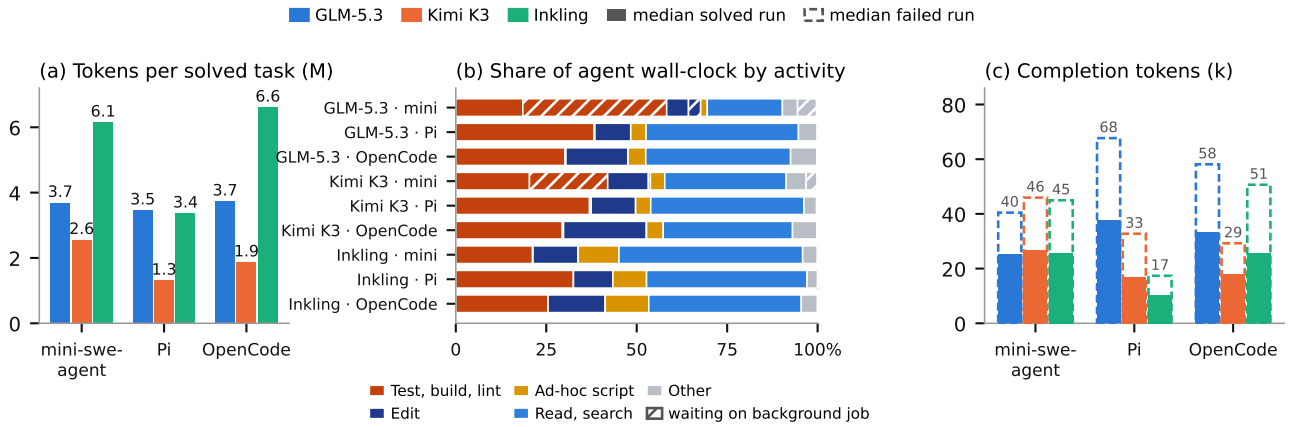


Figure 4. The same model under three scaffolds. (a) Total tokens per solved task. (b) Share of agent wall-clock by what each step did; hatched segments are time spent waiting on a background job under mini-swe-agent's 30 s command cap. (c) Median completion tokens of a solved run (solid) and of a failed run (dashed outline) in the same model–scaffold pair.

5 Discussion

Usefulness of V2 near saturation. Frontier systems resolve nearly all V2 tasks, but the benchmark remains useful for four reasons. The saturation result is itself informative: on a correctly graded split the strongest systems fail fewer than 1% of tasks, so small differences among them are not capability differences. Most deployed systems are not at the frontier: open-weight, smaller and cheaper systems span 87.2–97.5% on the full set and 52.9–88.2% on Hard (Table 2). This clean set supports comparisons beyond the pass rate, such as cost, time to solution and scaffold choice, which are meaningful only when every task is solvable and correctly graded. And a locked, fully solvable suite is a regression test for agent products: a failure on a task that all recent systems solve warrants investigation.

Lessons for benchmark maintainers. (1) Use strong agents as auditors: the failure audit flagged 39 tasks the artifact audit missed, because some defects appear only when a capable agent writes a reasonable solution and the tests reject it. (2) Repair instructions, keep tests fixed, and validate repairs blind: weakening a test to fit an instruction changes what the benchmark measures, and a blind re-implementation checks that the edited instruction suffices. (3) Grade in a clean container: agents that cannot fetch a dependency sometimes forge one. (4) Version the benchmark and release all grades, since new defects surface as agents improve. (5) Check that common scaffolds run on every image: two of the scaffolds in §4.5 failed to install on base images until we patched them.

6 Conclusion

As agents became more capable, two kinds of issues in SWE-Bench Pro became visible: tasks whose grade could disagree with the correctness of the solution, and an evaluation setup that agents could circumvent. SWE-Bench Pro V2 addresses both with a pipeline of agentic audits, engineer adjudication, blind re-implementation and verifier hardening that repaired or retired 480 of 731 tasks, and a locked protocol that grades every patch in a clean container. Frontier systems now solve almost every task, the Hard split separates them, and the full set remains a reliable benchmark below the frontier. We release

the tasks and protocol, and expect the pipeline and lessons to apply to other benchmarks as agents approach their ceilings.

Limitations

Memorization. The public repositories and their fixes predate the evaluated models. The protocol stops retrieval at test time, not memorization during training.

Scope of re-grading. A clean-container re-grade rejects changes made outside the patch, but forgeries inside it, such as an imitation of a missing library, are left to the patch audit.

Ethics Statement

The tasks come from open-source repositories under their original licenses. The engineers who reviewed and repaired tasks were contracted, and no personal data about them is recorded or released. We document the exploits we observed because benchmark operators need this information to defend against them, and the countermeasures are included with the benchmark.

References

[1] 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 *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2310.06770.

[2] OpenAI. Introducing SWE-bench Verified. OpenAI blog, https://openai.com/index/introducing-swe-bench-verified/, 2024.

[3] OpenAI. Why SWE-bench Verified no longer measures frontier coding capabilities. OpenAI blog, https://openai.com/index/why-we-no-longer-evaluate-swe-bench-verified/, 2026.

[4] Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, et al. SWE-Bench Pro: Can AI agents solve long-horizon software engineering tasks? In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, volume 306 of *Proceedings of Machine Learning Research*, 2026. URL https://arxiv.org/abs/2509.16941.

[5] Nitya Nadgir, Sayash Kapoor, Kangheng Liu, Peter Kirgis, Matilda Orona, Stephan Rabanser, Tilman Bayer, Abhishek Shetty, Yue Ling, Derrick Chan-Sew, Rumi Nakagawa, Saiteja Utpala, Zachary S. Siegel, and Arvind Narayanan. Life after benchmark saturation: A case study of CORE-Bench. *arXiv preprint arXiv:2606.26158*, 2026. URL https://arxiv.org/abs/2606.26158.

[6] Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Linhao Zhang, Shulin Xin, Lu Chen, Qi Liu, et al. Multi-SWE-bench: A multilingual benchmark for issue resolving. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2025. URL https://arxiv.org/abs/2504.02605.

[7] Linghao 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! *arXiv preprint arXiv:2505.23419*, 2025. URL https://arxiv.org/abs/2505.23419.

[8] Ibragim Badertdinov, Alexander Golubev, Maksim Nekrashevich, Anton Shevtsov, Simon Karasik, Andrei Andriushchenko, Maria Trofimova, Daria Litvintseva, and Boris Yangel. SWE-rebench: An automated pipeline for task collection and decontaminated evaluation of software engineering agents. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2025. URL https://arxiv.org/abs/2505.20411.

[9] Samuel Miserendino, Michele Wang, Tejal Patwardhan, and Johannes Heidecke. SWE-Lancer: Can frontier LLMs earn $1 million from real-world freelance software engineering? *arXiv preprint arXiv:2502.12115*, 2025. URL https://arxiv.org/abs/2502.12115.

[10] Harbor Framework Team. Harbor: A framework for evaluating and optimizing agents and models in container environments. https://github.com/harbor-framework/harbor, 2026.

[11] Mike A. Merrill et al. Terminal-Bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://arxiv.org/abs/2601.11868.

[12] Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with SWE-Gym. *arXiv preprint arXiv:2412.21139*, 2024. URL https://arxiv.org/abs/2412.21139.

[13] John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. SWE-smith: Scaling data for software engineering agents. *arXiv preprint arXiv:2504.21798*, 2025. URL https://arxiv.org/abs/2504.21798.

[14] Reem Aleithan, Haoran Xue, Mohammad Mahdi Mohajer, Elijah Nnorom, Gias Uddin, and Song Wang. SWE-Bench+: Enhanced coding benchmark for LLMs. *arXiv preprint arXiv:2410.06992*, 2024. URL https://arxiv.org/abs/2410.06992.

[15] You Wang, Michael Pradel, and Zhongxin Liu. Are "solved issues" in SWE-bench really solved correctly? An empirical study. In *Proceedings of the IEEE/ACM 48th International Conference on Software Engineering (ICSE)*, 2026. doi: 10.1145/3744916.3764576. URL https://arxiv.org/abs/2503.15223.

[16] Boxi Yu, Yuxuan Zhu, Pinjia He, and Daniel Kang. UTBoost: Rigorous evaluation of coding agents on SWE-Bench. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3762–3774, 2025. URL https://arxiv.org/abs/2506.09289.

[17] Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, et al. Establishing best practices for building rigorous agentic benchmarks. *arXiv preprint arXiv:2507.02825*, 2025. URL https://arxiv.org/abs/2507.02825.

[18] Ziqian Zhong, Aditi Raghunathan, and Nicholas Carlini. ImpossibleBench: Measuring LLMs' propensity of exploiting test cases. *arXiv preprint arXiv:2510.20270*, 2025. URL https://arxiv.org/abs/2510.20270.

[19] Jonathan Gabor, Jayson Lynch, and Jonathan Rosenfeld. EvilGenie: A reward hacking benchmark. *arXiv preprint arXiv:2511.21654*, 2025. URL https://arxiv.org/abs/2511.21654.

[20] Bingchen Zhao, Dhruv Srikanth, Yuxiang Wu, and Zhengyao Jiang. SpecBench: Measuring reward hacking in long-horizon coding agents. *arXiv preprint arXiv:2605.21384*, 2026. URL https://arxiv.org/abs/2605.21384.

[21] Kunvar Thaman. Reward hacking benchmark: Measuring exploits in LLM agents with tool use. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, volume 306 of *Proceedings of Machine Learning Research*, pages 120372–120389, 2026. URL https://arxiv.org/abs/2605.02964.

[22] Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, Ethan Perez, and Evan Hubinger. Sycophancy to subterfuge: Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*, 2024. URL https://arxiv.org/abs/2406.10162.

[23] Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y. Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*, 2025. URL https://arxiv.org/abs/2503.11926.

[24] jacobkahn. Repo state loopholes during agentic evaluation. SWE-bench GitHub issue 465, https://github.com/SWE-bench/SWE-bench/issues/465, 2025.

[25] joe-needham. Repo state leakage. SWE-Bench Pro GitHub repository, issue 7, https://github.com/scaleapi/SWE-bench_Pro-os/issues/7, 2025.

[26] ConnorBAdams. Git reward hacking in SWE-Bench Pro OSS. SWE-Bench Pro GitHub repository, issue 93, https://github.com/scaleapi/SWE-bench_Pro-os/issues/93, 2026.

[27] Shanchao Liang, Spandan Garg, and Roshanak Zilouchian Moghaddam. The SWE-bench illusion: When state-of-the-art LLMs remember instead of reason. *arXiv preprint arXiv:2506.12286*, 2025. URL https://arxiv.org/abs/2506.12286.

[28] Daniel Ramos, Claudia Mamede, Kush Jain, Paulo Canelas, Catarina Gamboa, and Claire Le Goues. Are large language models memorizing bug benchmarks? *arXiv preprint arXiv:2411.13323*, 2024. URL https://arxiv.org/abs/2411.13323.

[29] 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.

[30] OpenAI. Separating signal from noise in coding evaluations. OpenAI blog, https://openai.com/index/separating-signal-from-noise-coding-evaluations/, July 2026.

[31] Kilian Lieret, John Yang, et al. mini-SWE-agent: The 100-line AI agent that solves GitHub issues. https://github.com/SWE-agent/mini-swe-agent, 2025.

[32] John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 2024. URL https://arxiv.org/abs/2405.15793.

A The Hard Split

The 51 tasks come from ansible (8), openlibrary (8), flipt (7), teleport (6), webclients (5), qutebrowser (4), vuls (4), element-web (4), navidrome (3), NodeBB (1) and tutanota (1). A task qualified if at least two of five model families fail on it across runs. That gave 54 candidates and 3 were removed after expert review. Every Hard task passes its reference solution, fails with an empty patch, and is solved by at least one of our runs. Of the 51 instructions, 39 were edited in the final review after selection.

B Harness Study

Setup. GLM-5.3 (thinking on), Kimi K3 (reasoning effort max) and Inkling (xhigh) ran with a 1M-token context window and no cap on thinking tokens or turns; the only cap besides the 50-minute budget is mini-swe-agent's default 30 s per-command timeout, which is part of that scaffold. Versions: Harbor 0.22, mini-swe-agent 2.4.6, Pi 0.85.0, OpenCode 1.18.29. Before the runs we patched the task images on which Pi or OpenCode would not install (older Alpine images lacking a recent `libstdc++`, and Debian 11 images whose end-of-life package sources broke OpenCode's installer); the patches add system packages only, and all nine combinations ran on the same patched images. We release the 5,778 trajectories with per-trajectory feature files.

Attempt rule. Unlike the runs of §4, these were graded once, in the container the agent had used. Every task has one counted attempt per model–scaffold combination; an attempt that hit the budget is graded on the patch present at the deadline. An attempt that hit the budget or was killed by infrastructure *without* a passing grade was run once more at lower concurrency and the re-run is the counted attempt; a re-run never replaced a graded pass, and the rule was applied identically to all nine combinations. Pi and OpenCode place a different Node binary on the agent's path than the grader uses; in 13 trials an agent rebuilt a native module against it and the grader then crashed on that module (Kimi K3: 5 under Pi, 2 under OpenCode; GLM-5.3: 2 and 2; Inkling: 2 under OpenCode), and those trials count as failures.

Outcomes. Table 4 gives the paired comparisons behind §4.5. Splitting the variance of solved-or-failed over the 642 tasks every combination attempted, task identity accounts for 31.7%, the model for 2.65%, the scaffold for 0.29% and the model–scaffold interaction for 0.36%; the residual 65.0% is task-by-run variation.

Validation. Pass rates are one attempt per task, a standard error of 0.6–1.4 points per combination. The activity shares of Figure 4b and the test-after-edit analysis of §4.5 come from a classifier that reads the first executable of each segment of a shell command; it matched a manual read on 147 of 150 sampled commands and on 97.2% of sampled wall-clock, and its test-outcome detector matched on all 32 of the 40 sampled test outputs it labeled, leaving 8 undecidable. Step time includes model latency, and test-file edits are under-counted under mini-swe-agent, where edits arrive as heredocs.

C Reproducibility

Harness. Harbor 0.22.0 on Modal 1.5.5. The network is set per phase: public during setup, only the model endpoint during the agent phase, public again for the verifier, because 32 tasks download

Model	Pair ($A - B$)	only A / B	Diff.	95% CI	p
GLM-5.3	mini − Pi	15 / 21	−0.9	[−2.8, +0.9]	0.41
	mini − OpenCode	15 / 20	−0.8	[−2.6, +1.0]	0.50
	Pi − OpenCode	12 / 11	+0.2	[−1.3, +1.6]	1.00
Kimi K3	mini − Pi	20 / 10	+1.6	[−0.1, +3.2]	0.10
	mini − OpenCode	8 / 9	−0.2	[−1.4, +1.1]	1.00
	Pi − OpenCode	6 / 17	−1.7	[−3.2, −0.3]	0.03
Inkling	mini − Pi	64 / 27	+5.8	[+2.9, +8.6]	<0.001
	mini − OpenCode	23 / 34	−1.7	[−4.0, +0.6]	0.18
	Pi − OpenCode	18 / 66	−7.5	[−10.2, −4.7]	<0.001

Table 4. Paired scaffold comparisons within a model on the same 642 tasks: tasks solved only under scaffold A / only under B, the difference in percentage points with its 95% interval, and an exact test on the discordant tasks. "mini" is mini-swe-agent.

dependencies while testing. Claude Code 2.1.278 with web tools disabled; Codex 0.155.1 with hosted web search disabled; mini-swe-agent 2.4.6 with no step or cost limit, so the time budget binds.

Budget and grading. The budget is 3000 s per task. The agent is stopped at 2940 s and the patch present then is graded. The patch is the staged `git diff` of the working tree; it is applied to a fresh container built from the task image, with a three-way fallback, and the unchanged verifier runs. When the sandbox grade and the clean grade disagree, two more clean grades are taken and the majority of the three counts. Mean API spend per task on V2 was $2.9 for the Claude Code runs, $2.3 for Astra and $3 for Gemini.

Release. Task directories, the Hard split, the probe, capture and re-grade scripts, the exact harness flags, per-task grades in both modes for every run, and the harness-study trajectories.