SENIOR SWE-BENCH: EVALUATING CODING AGENTS AS SENIOR SOFTWARE ENGINEERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Coding agents are expected to perform like *senior engineers*: they are tasked with building features from underspecified requirements, investigating and fixing bugs from user-reported symptoms, and writing high-quality code that meets broader codebase standards. However, popular benchmarks evaluate agents like "junior" software engineers, with highly specified instructions and a focus on behavioral correctness alone. This gap is a consequence of standard grading approaches: pre-written verifiers require full interface specification while LLM-as-a-judge approaches give up the reliability of executing code. Neither can reliably grade senior-level work, which admits many valid solutions that differ in quality. We close the gap with two mechanisms: a verification agent that writes verifier tests adapted to each solution, and a taste judge that assesses code quality in a grounded manner. Using these techniques, we built Senior SWE-bench,[1] a benchmark of 100 tasks (50 public and 50 held-out to resist contamination) drawn from real senior engineering work. Tasks are "design-and-build" (e.g., new features addressing user stories) or "investigate-and-fix" (e.g., diagnosing and solving complex bugs), and instructions are naturally underspecified (on average, 26% of the median length of those in SWE-Bench Pro). Evaluating 46 frontier agent configurations with these grading mechanisms, we find that the best behavioral-correctness pass@1 is 62.1%, while the best *overall* pass@1 (which also considers taste) is only 33.3%. More broadly, we show that for frontier agents, correctness and taste rankings often diverge, and that Senior SWE-bench tasks can reveal trends in model behavior.

1 INTRODUCTION

Despite the enormous popularity of coding agents, there is a surprising gap between the way they are used and benchmarked. Agents are used to perform work at the level of a *senior engineer*—they are tasked with underspecified goals communicated in natural language, required to understand latent context and needs, and expected to exercise multiple skills beyond coding (such as system design and runtime investigations).[2] The growing use of background agents takes this further, as agents receive instructions from surfaces like Slack and begin work in a sandbox that may require additional environment setup. In contrast, popular benchmarks evaluate agents in narrow settings with highly specified and even overspecified instructions. Bug tasks have pointers to the root cause, and feature tasks spell out exact interfaces and behaviors. This is precisely how we would treat a *junior engineer*. Unfortunately, this gap makes existing benchmarks largely uninformative for modern coding agent usage.

Importantly, this gap is not merely a matter of writing more natural task instructions or increasing scope. Instead, it is inherent to the most common grading mechanisms used in agent evaluation: pre-written verifiers and LLM-as-a-judge. While pre-written verifiers in the form of code tests can reliably exercise solutions by running code, they cannot reliably adapt to multiple valid solutions admitted by an underspecified instruction. This is true even for verifiers that are designed to test system interfaces behaviorally. On the other extreme, LLM-as-a-judge approaches can adapt to different implementations but sacrifice the ability to execute code and are notorious for over-scoring

[1] `https://[anonymized]` (withheld for review).

[2] All em-dashes are human-authored.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

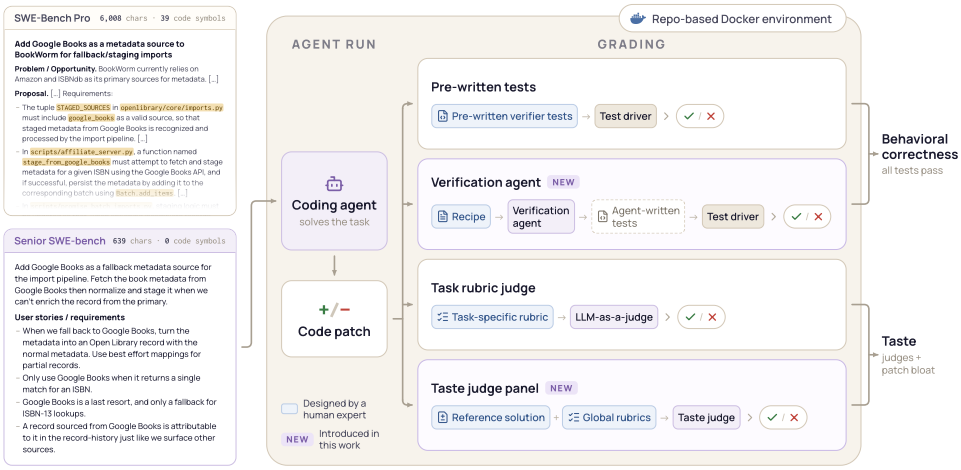

Figure 1: **Overview of Senior SWE-bench.** At left, the instruction for a Senior SWE-bench version of a SWE-Bench Pro (Deng et al., 2025) feature task (a Google Books fallback in Open Library) is shown. The SWE-Bench Pro version names files and symbols its tests call (highlighted), but the Senior SWE-bench version does not. At right, the grading mechanisms for Senior SWE-bench, including the newly introduced verification agent (which follows an expert-designed recipe to write tests adapted to the submitted solution) and the taste judge (which assesses code quality grounded by the codebase and reference solution).

solutions that look plausible. Neither of these typical approaches is sufficient for evaluating senior-level work.

To address these challenges, we introduce a verification agent that dynamically writes verifier tests adapted to the solution under test, following an expert-designed recipe. We demonstrate that this approach balances the reliability of pre-written verifiers and the flexibility of LLM-as-a-judge, while remaining cost-effective and aligned to human experts. We also introduce a taste judge that enables reliable scoring of code quality, going beyond behavioral correctness alone. The taste judge's assessments are grounded in the existing codebase and reference solutions, and similarly demonstrate its alignment to human code quality ratings.

Using these methods, we introduce a new coding benchmark, Senior SWE-bench, that better assesses the capabilities of agents when used to perform the work of senior engineers. It contains two complementary task families: *Design-and-Build*, which challenge agents to design and implement features or migrations from underspecified instructions; and *Investigate-and-Fix*, which require agents to diagnose and resolve complex bugs or performance problems from behavioral reports. We construct 100 tasks around these themes, drawn from real development work in 12 open-source repos across five programming languages. The median length of Senior SWE-bench's instructions is substantially shorter than that of earlier benchmarks (e.g., 26% of that of SWE-Bench Pro), reflecting the natural underspecification that real-world coding agents see.

We use Senior SWE-bench to study coding agents as they are used, evaluating 46 agent configurations across 18 models. We observe that the highest Behavioral Correctness pass@1 is 62.1%, whereas the highest *Overall* pass@1 (which also considers Taste) is only 33.3%. Several leading models have similar Overall scores despite substantial differences in Behavioral Correctness (e.g., ranging from 53.7% to 62.1%). These findings provide further evidence that behavioral correctness alone is insufficient to characterize progress toward autonomous senior-level engineering. Our contributions are:

- **Senior SWE-bench, a benchmark of senior-level engineering work.** It consists of 100 tasks that require independent design or investigation from naturally underspecified instructions, supported by reference implementations and a multi-stage, expert-led quality-control process.
- **Adaptive behavioral verification with the verification agent.** We introduce the verification agent, which translates an expert-designed recipe into executable tests adapted to each solution.
- **Code quality assessment with the taste judge.** An expert-calibrated taste judge evaluates implementation taste and codebase practice alignment, grounded in pre-existing code and reference solution; a separate per-task rubric judge also scores expert-defined criteria.

- **An empirical study targeting behavioral correctness and taste.** We characterize performance across task families and model behavior on senior-level tasks.

2 Related Work

Repository-level software engineering benchmarks. SWE-bench (Jimenez et al., 2024) and its descendants (Zan et al., 2025; Zhang et al., 2025; Badertdinov et al., 2025; Miserendino et al., 2025; Merrill et al., 2026; Deng et al., 2025; Zhou et al., 2026a; Du et al., 2025; He et al., 2025; Shetty et al., 2025; Huang et al., 2026b; Yang et al., 2026) score an agent with the tests of the pull request that fixed the issue—tests "written to validate a specific change, rather than to define an implementation-agnostic standard for solving the task" (OpenAI, 2026a). Audits of SWE-bench Verified flagged 61.1% of sampled issues for tests that may unfairly reject valid solutions and 38.3% for underspecified problem statements (Chowdhury et al., 2024; OpenAI, 2026b; Zheng et al., 2026). Senior SWE-bench instead states each task as an engineer would, leaving the design or the diagnosis to the agent (Section 3.1), and grades behavior rather than the merged implementation (Section 4).

Underspecification and interaction. Ambig-SWE, SWE-RPG, SWE-Interact and SWE-Together score the interaction itself (Vijayvargiya et al., 2026; Zhou et al., 2026b; Raghavendra et al., 2026b; Wu et al., 2026); SWE Atlas, Cursor and Meta grade solutions based on an open request (Raghavendra et al., 2026a; Cursor Research, 2026; Jha et al., 2026), as we do. Agents guess rather than ask (Ji et al., 2026), and no pre-written test can check what the instruction left out (Menet et al., 2026). Our verification agent follows an expert-designed recipe that carries that detail (Section 4.2), and guided variants measure how much difficulty the omission adds (Section 5).

Verifiers and judges for code. Pre-written verifiers overconstrain (reference-patch similarity (Wei et al., 2025) is the extreme) and underconstrain, since 7.8% of resolved Verified patches fail the full suite (Wang et al., 2026; Yu et al., 2025). LLM judges accept alternatives but reward style (Jain et al., 2025); rubric rewards (Gunjal et al., 2025; Viswanathan et al., 2025; Huang et al., 2026a) and agentic judges (Zhuge et al., 2025) help only partly. Tests generated from the artifact under test (Xin & Reiss, 2017; Ye et al., 2021; Mündler et al., 2024; Ahmed et al., 2025a; 2024) inherit its bugs (Zhao et al., 2026), and a fifth to a third of the patches they pass fail hidden tests (Ahmed et al., 2025b). Our verification agent (Section 4.2) instead generates tests from an expert-designed recipe of behavioral stories (North, 2006), only on tasks where the reference passes and a no-op fails (Raghavendra et al., 2026a; Li et al., 2026), and its verdicts are judge-screened and checked against expert labels (Starace et al., 2025; Patwardhan et al., 2025).

Beyond behavioral correctness: mergeability and taste. Benchmarks largely score behavioral correctness alone (Whitfill et al., 2026; Rein, 2025), though passing tests has never implied a mergeable patch (Qi et al., 2015; Smith et al., 2015; Mäntylä & Lassenius, 2009; Beller et al., 2014). FrontierCode grades quality with maintainer-designed rubrics (Lu et al., 2026). Others score instructions (Zhong et al., 2026), review-thread constraints (Yu et al., 2026), merge readiness (Xie et al., 2026), or uncalibrated rubrics (Vasilevski et al., 2026). We introduce an expert-calibrated "taste score" that measures each solution against codebase-wide best practices and the reference solution (Section 4.3, Appendix D).

3 Senior SWE-bench Design

Our goal is to craft realistic and diverse tasks, representative of the breadth of software development work with which we trust coding agents (Fig. 1). Like senior engineers, agents need to exercise skills like aligning to human intent, open-ended technical design, runtime debugging, and good taste.

- **Real senior engineering work.** Every task is built from PRs merged into one of 12 open-source repositories in five languages; the median source-PR author has 355 commits in that repository. Tasks are designed to reflect two categories of senior-level work: adding new features and resolving issues.
- **Multi-skill, long-horizon work.** Each task involves a range of senior-level "skills" drawn from a taxonomy of 26 (e.g. Data Model Design and Fault Localization), and tasks require a median of six. Tasks are also designed to be long-horizon, requiring multiple stages including setting up environments, running services locally, designing solutions, investigating issues, implementation, and verification.

- **Underspecified instructions.** Senior SWE-bench task instructions are phrased similarly to how a teammate might request a feature or a user might report an issue. They generally do not contain direct code references or interface specifications, in contrast to the unrealistic level of specification in popular benchmarks. As a result, tasks admit many valid implementations and most correct solutions diverge structurally in some way from the reference solution (Table A5).
- **Reliable grading for behavioral correctness and taste.** The verification agent and taste judge, introduced in this paper, combine with standard pre-written verifiers and LLM-as-a-judge to reliably grade both behavioral correctness and taste, even with underspecified instructions.

3.1 Task Format

Task sources and scope. Tasks are sourced from real-world pull requests (PRs) in open-source repositories, based on the approach introduced in SWE-bench (Jimenez et al., 2024). Source repositories were selected to represent a diversity of tech stacks, domains, and maturity. Half of the repos started in the last 5 years, capturing higher-velocity feature work and more modern technology stacks.

Table 1: **Sourcing high-quality PRs.** Twelve repositories (seven applications, two libraries, two tools, and one service; TypeScript in 7, Python 4, Elixir 3, Go 2, Rust 2). Full list in Table A1.

	Median	**Range**
Per repository ($n = 12$)		
Year started	2021	2015–2025
Stars	27.9k	2.5k–104k
Source lines	317k	148k–3.8M
Contributors	467	57–1,851
Merged PRs per week	57	15–554
Languages ($\geq$5% of code)	3	1–4
Per task ($n = 100$)		
PR author's commits in repo	355	1–4,633
Merger's commits in repo	652	19–4,633

Source PRs are filtered for recency (merged after February 2026, to reduce the risk of contamination) and the author's level of experience with the codebase (preferring maintainers with thousands of commits), then grouped with related PRs and LLM-scored for multi-skill complexity (Section 3.2 details the pipeline). Authoritative reference solutions, written by the frequent-committer PR authors, are retrieved from the source PR(s) for every task. When multiple PRs are used as a task source, their patches are combined using `git` operations (Table 1). Reference solution patches are then filtered to remove any hunks irrelevant to the task.

Tasks cover diverse scopes, from pinpoint bug fixes to multi-thousand-line features, and are long-horizon, requiring on the order of a hundred tool-calling steps per trial from even the strongest agents. See Fig. A1 for a comparison to other benchmarks.

Task segments. Tasks fall into two major segments, meant to reflect two common senior-level categories of work: Design-and-Build and Investigate-and-Fix (Table A2 in Appendix A contrasts the two). Design-and-Build tasks focus on the complexity of designing and implementing features and migrations without full specifications. Tasks are designed to have a wide design space, including crossing service boundaries or involving multiple languages when applicable. Investigate-and-Fix tasks focus on runtime diagnosis of tricky bugs and performance issues from behavioral reports. Agents generally need to first reproduce issues before diagnosing and resolving (Table A3 in Appendix A traces one such diagnosis).

Natural instructions. Instructions are written to reflect natural communication with background coding agents through surfaces like Slack. Design-and-build instructions are composed of an overall request and a set of behavioral requirements (or user stories) in line with what a product manager might provide. Investigate-and-fix instructions read like an issue report: describing the behavior observed so that the agent has to investigate instead of spelling out exactly how to make the fix. Instructions name user-visible concepts and observable symptoms rather than specific code pointers, since naming a file, class, or function turns an investigation into a directed code read (Fig. 1a; Appendix A gives full examples). Fig. A1 in Appendix A compares instruction length to other benchmarks; the median Senior SWE-bench instruction is 26% the length of the median SWE-Bench Pro instruction.

3.2 Task Construction and Quality Control

Fig. 2 depicts the task construction pipeline, which is designed to output challenging, realistic, and high-quality tasks. The pipeline stages are executed by a combination of human software engineering experts, deterministic programs, and specialized coding or evaluation agents.

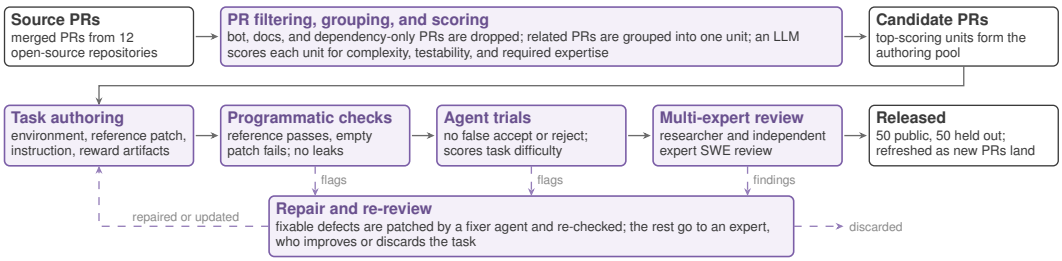

Figure 2: **Task construction and quality control.** Each purple box describes the conditions a task must meet to continue. Different stages combine tooling with human experts, including a final review by multiple human experts.

- **PR filtering, grouping, and scoring.** Sufficiently recent candidate PRs from the source repository first undergo filtering based on author, content, and other aspects. Related PRs are grouped into single units via batch LLM evaluation. An LLM then scores each unit for complexity, testability, and required skills on a 0–5 rubric. Units scoring 3 or more form the authoring pool (Section 3.1).
- **Task authoring.** From a candidate unit, authoring produces four artifacts: the task environment, the reference patch, the instruction (Section 3.1), and the reward artifacts of Section 4. These artifacts are expert-designed: drafted directly by the team or by a coding agent session following the team's procedures, then expert-reviewed. The environment is the repository at the pre-fix commit with only major system dependencies pre-installed (Appendix B). The reference patch is the maintainers' merged change, stripped of incidental changes.
- **Programmatic checks.** Before in-depth human review, a task must pass automated checks. Each task is executed multiple times with the reference solution and an empty patch to ensure that they are well formed and that non-deterministic elements (e.g. verification agent, taste judge) score consistently for solutions with known outcomes. A task is admitted only if the reference solution has pass^3 $= 1$ and the empty patch has pass@3 $= 0$. Dozens of high-precision LLM checks screen for instruction leakage (root cause, solution shape, or internal symbols), grading reliability risks (e.g. a test that rejects a valid alternative implementation), mismatches between task scope and grading, and benchmark leaks in agent-facing prose.
- **Agent trials.** Before a task is finalized, multiple agents of varying capability levels attempt it. Their solutions are analyzed to identify risks of false accepts (accepting solutions that are materially worse than the reference) or false rejects (rejecting solutions that are near-optimal with respect to the reference), and discarded or revised as a result. Agent trials are also used to score task difficulty.
- **Multi-expert review.** Contributing researchers review all tasks for design and implementation quality. Multiple independent software engineers review each task and example trajectories against an extensive rubric. A subset of these reviews also capture labels to calibrate the verification agent (Section 4.2) and LLM judges (Section 4.3).
- **Repair and re-review.** Identified task defects with clear fixes can be given to a fixer coding agent for repair; this agent can use programmatic checks as verification. Other defects identified by automated or human review go to a human software engineer, who either improves or discards the task. Repaired and updated tasks re-enter at authoring and repeat every check.

4 Task Grading

Senior SWE-bench uses a combination of standard and novel reward mechanisms to evaluate agent performance across two major grading components: behavioral correctness and taste (Fig. 1). Behavioral correctness is assessed by graders that exercise code with standard testing frameworks (e.g. pytest, go test). This consists of standard *pre-written verifier tests* and the *verification agent*, which follows an expert-designed recipe to write tests that adapt to the submitted solution (Section 4.2). Taste is assessed by LLM-as-a-judge approaches (Zheng et al., 2023) that review the code, including a standard *task rubric judge* that evaluates task-specific quality criteria and a novel *taste judge* that applies global code quality rubrics to every task (Section 4.3).

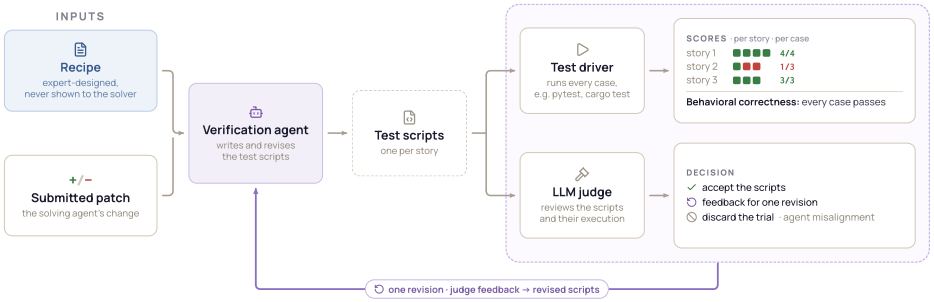

Figure 3: **The verification agent process.** The verification agent turns the expert-designed recipe into one test script per story. The driver runs every case, and a judge accepts the scripts, returns feedback for the one permitted revision (dashed), or discards the trial. Behavioral correctness requires every case to pass, in addition to any pre-written verifiers.

4.1 Writing Behavioral Tests

Behavioral tests, whether pre-written or written by the verification agent, are meant to be an implementation-agnostic standard for demonstrating the correct system behavior. As such, they are generally written against system-level interfaces like HTTP endpoints, CLI commands, or rendered UI components. Coding agents, like senior engineers, are expected to write code that aligns with both explicitly stated requirements and unstated requirements (such as implementing functionality within a system-wide framework). Senior SWE-bench tests both stated and unstated requirements in the following categories:

- **Behavioral contract.** Behavioral requirements stated in the task instruction. Example: new functionality added to API.
- **Established codebase practices.** Strong and consistent codebase patterns that impact functionality. Example: new API endpoint is registered with an existing auth framework.
- **Clear best practice.** Practice obvious to a senior engineer with no defensible alternatives. Example: API endpoint that lists an unbounded collection must be capped or paginated.

Investigate-and-fix tasks are designed to not modify interfaces, and so only use pre-written verifiers. Design-and-build tasks use a combination of the verification agent and pre-written verifiers (e.g. for regression tests). Verifiers and the verification agent can encode both pass-to-pass and fail-to-pass tests. Pass-to-pass tests succeed on the pre-patch codebase and act as regression guards, assuring that existing behavior continues to work. Fail-to-pass tests succeed only after a successful patch has been applied and discriminate between valid and invalid solutions.

4.2 Verification Agent

We introduce a verification agent which uses expert-designed recipes to write system-level, behavioral test scripts adapted to specific solutions (Fig. 3). Each design-and-build task has a human-designed *recipe*. The recipe consists of a set of behavioral user stories (or functional requirements), each with a procedure for writing its verifier test and several parameterized test cases (207 stories and 565 cases over the 50 tasks). The behavioral requirements, written in natural language, are injected from the recipe into the task instruction. No other details from the recipe (such as the procedure for testing these requirements) are shown to the solving agent. Once the verification agent has completed writing the test scripts, they are executed programmatically by the appropriate test driver (e.g. pytest, cargo test). An LLM judge then reviews the test execution and the test scripts, and decides whether to accept the scripts, write feedback for revision, or discard the trial due to poor verification agent performance or misalignment.

We found that mini-SWE-agent (Lieret & Jimenez, 2025) with Claude Sonnet 4.6 offered the best balance of performance and cost. The verification agent stage adds a median of 4.6 minutes and $0.72 per trial. With Claude Opus 5 at high effort as the solving agent, this represents 13% of the solving agent's time and 7.7% of its cost. Appendix C gives the recipe format, an example story, and the agent and judge protocol.

Validation. We perform a validation study to measure the verification agent's agreement with expert labels and its stability across repeated runs on 3 agent patches for each of 20 tasks, all labeled as correct or incorrect by experts at the story level. We compare it with three baselines: *static tests*, verifier tests pre-written against the reference solution; an *LLM-as-a-judge*, which reviews the recipe and patch without running any code; and a *weaker verification agent*, given the story descriptions alone, without the recipe's procedures and cases or the judge's revision loop. Fig. 4 plots each grader's error rates. The verification agent demonstrates the lowest false accept and false reject rates. The static tests are not able to adapt to multiple valid implementations, and the LLM-as-a-judge demonstrates the typical behavior of over-scoring plausible-looking code. Additionally, the expert-designed procedures ground the verification agent's behavior; without them, the weaker verification agent struggles to write appropriate tests. The verification agent is also the most stable model-based grader, returning a unanimous verdict on 96% of tasks across three runs.

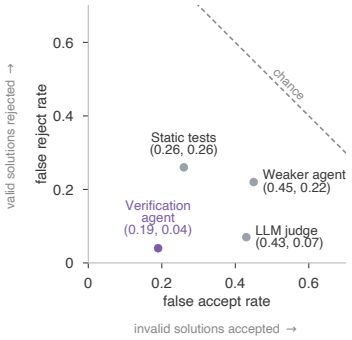

Figure 4: **Behavioral-correctness graders against the expert's labels.** Only the verification agent successfully balances the reliability-flexibility trade-off between static tests and LLM judges.

Using a Design-and-Build task's overall behavioral score (combining both pre-written verifiers and the verification agent) further decreases the false reject rate to 0.11, producing a balanced accuracy score of 0.92.

4.3 Taste Judge

We introduce a taste judge, which uses observational and comparative signals to grade the agent's solution on code quality. The judge uses two rubrics, each with five dimensions scored from 1 to 5 (Table A4): *relative code quality* against the reference implementation, and *codebase practice alignment* against the conventions the judge observes from the repository itself. Critically, the rubrics are designed to be technology- and practice-agnostic, grounding the judge's assessment in observed codebase practices and the reference solution. For example, a model is not penalized for not writing documentation in a codebase that doesn't broadly document its code (in fact, the judge rewards parsimony with documentation). To decrease variance and reduce risk of bias towards code from certain models, we use a panel of Claude Sonnet 4.6 and GPT-5.4 and average their scores per dimension and per criterion before rolling up.

In addition, we measure patch bloat, the ratio of source lines of code (SLOC) in the agent's solution to SLOC in the reference implementation. The taste judge score, patch bloat, and task-specific rubric score are combined to create the Taste Score (Section 4.4).

Validation. We validate the taste judge by showing that it behaves similarly to a human expert across both relative quality and codebase practice alignment. Two human software engineers assessed the taste rubric dimensions on 4 agent patches for each of 30 tasks. Code quality is subjective, so the ceiling is the experts' own agreement: two experts match only about half the time but land within one point of each other on 96% of patches. The judge matches the experts' consensus just as often and is within one point of it on 90% of patches for relative quality and 87% for practice alignment (Table 2). The scores are also stable: running the taste judge three times on a patch gives an axis score a standard deviation of 0.15 and 0.21 points (on a 1–5 scale). The panel of Claude Sonnet 4.6 and GPT-5.4 had a lower standard deviation than either judge alone.

Table 2: **Taste judge against the experts.** On the study's 119 patches, agreement on the 1–5 scale between every pair of experts (the ceiling) and between the panel and the experts' consensus.

	Pairs	Match	Within 1
Relative quality			
Expert vs. expert (ceiling)	714	54%	96%
Judge vs. experts	119	54%	90%
Practice alignment			
Expert vs. expert (ceiling)	714	51%	96%
Judge vs. experts	119	50%	87%

4.4 Scoring

Behavioral Correctness Score. $C = \frac{1}{N}\sum_t c_t$ is the fraction of tasks on which every pre-written test and every verification agent-authored test case passes.

Table 3: **Leaderboard.** Overall score, behavioral correctness and taste (%; Section 4.4). One row per model at its best Overall pass@1 effort tier, sorted by Overall pass@1. D&B / I&F show Overall pass@1 by segment. Steps: mean agent steps per trial. Out. $: mean output-token cost per task at list prices. Best value per column in bold. The three lowest-Overall models (GPT-5.4, Gemini 3.5 Flash, Gemini 3.1 Pro) are omitted.

		Overall			Behavioral Correctness			Taste	Overall (seg.)			
Model	**Effort**	pass@1	pass@3	pass^3	pass@1	pass@3	pass^3	pass@1	D&B	I&F	**Steps**	**Out. $**
Claude Fable 5.1	medium	**33.3**	**42.2**	20.4	57.9	65.3	43.2	57.6	38.5	28.5	77	1.89
Claude Fable 5	high	33.0	41.5	**21.9**	53.7	64.2	43.2	61.4	37.0	**29.2**	119	2.92
Claude Opus 5	high	32.4	42.1	18.8	**62.1**	**70.5**	**44.2**	52.2	**41.2**	24.1	141	1.77
GPT-5.6 Sol	xhigh	30.0	39.1	15.9	53.7	63.2	40.0	55.9	34.8	25.5	50	0.98
GPT-5.5	xhigh	28.0	34.4	13.2	52.6	61.1	36.8	53.3	31.0	25.3	85	1.08
Claude Opus 4.8	max	27.0	35.3	17.5	44.2	53.7	34.7	61.1	31.6	22.7	138	3.36
Claude Sonnet 5	max	26.3	33.6	14.7	50.5	59.1	34.4	52.0	34.7	18.7	318	2.71
Grok 4.6	high	25.5	36.7	15.7	51.6	65.3	35.8	50.5	30.6	20.8	125	0.14
GPT-5.6 Terra	xhigh	24.3	34.8	10.3	38.9	55.8	22.1	**62.4**	31.3	17.7	47	0.48
Grok 4.5	medium	24.2	35.4	11.9	50.5	64.2	31.6	47.9	31.7	17.2	72	0.13
MiniMax M3	adaptive	21.3	29.8	9.3	41.1	50.5	21.1	52.0	28.4	14.7	335	0.09
Claude Opus 4.7	max	20.6	30.8	12.6	46.8	56.4	31.9	43.9	28.7	12.8	285	2.52
Kimi K3	max	19.1	34.0	8.5	40.4	57.4	22.3	47.3	26.5	12.1	224	0.82
Gemini 3.7 Flash	high	18.9	29.3	10.5	44.2	61.1	31.6	42.8	25.5	12.7	254	0.18
GLM-5.2	max	17.4	27.4	8.1	35.8	49.5	23.2	48.5	18.7	16.1	194	0.09

Taste Score. $T = \sum_t c_t\, q_t \,/\, \sum_t c_t$ is the mean credit the judges give a correct solution, measured over solved tasks only. The credit q_t multiplies one logistic factor $\sigma(z) = 1/(1+e^{-z})$ per taste signal, each in the signal's margin m_{tk} above its bar: the rubric judge's fraction u_t (if any), the bloat ratio b_t, and the relative-quality and practice-alignment medians r_t and p_t of the taste judge,

$$q_t = \prod_k \sigma(m_{tk}/w), \qquad m_t = \big(u_t - 0.5,\ 1 - \log_2 b_t,\ r_t - 2,\ p_t - 2\big), \qquad w = 0.25. \tag{1}$$

Overall Score. The Overall score is the product $O = C\,T$ of Behavioral Correctness C and Taste T.

5 Results and Discussion

Experimental Setup. We run each agent three times and at all available effort levels for the top performing models. In total we score 46 configurations across 18 models, all using the mini-SWE-agent harness for standardization. With $s_{t,i} \in [0,1]$ the outcome of trial i on task t ($c_{t,i}$ for Behavioral Correctness, $c_{t,i}\, q_{t,i}$ for Overall), we report pass@1, the outcome of one trial per task. We also report pass@3 $= \max_i s_{t,i}$ (coverage) and pass^3 $= \min_i s_{t,i}$ (reliability) (Chen et al., 2021; Yao et al., 2024). Scores are reported over a set of 95 tasks.[3] For models with multiple effort levels, the effort level with the best Overall pass@1 is shown. See Table 3.

Profiles of top-performing models. Out of the top four models (Claude Fable 5.1, Claude Fable 5, Claude Opus 5, and GPT-5.6 Sol), different performance profiles can be seen. The Fable models have the highest Overall and Taste scores; Opus 5 has the highest Behavioral Correctness but the lowest Taste; and GPT-5.6 Sol is the most efficient (at around one-third of the output token cost of Fable 5).

Models are challenged by instructions that are naturally underspecified. To assess the impact of underspecification, additional guidance was appended to the instruction of 30 tasks. The additional guidance does not specify solutions at the level of other benchmarks' standard instructions, but instead provides recommendations for issue diagnosis or feature design, and was run with five diverse models. With the additional guidance, Behavioral Correctness increased an average of 24.6 points (95% interval 13.4–36.2) while Taste increased only by an average of 2.3 points. This implies that models are challenged specifically by underspecification, and Taste is more dependent on the model than on task specification. See Fig. A2, Appendix E.1.

Higher effort increases behavioral correctness, not taste. For models with multiple effort levels, going from a model's lowest to its highest tier adds 10.3 points of Behavioral Correctness on average. However, taste increases only by an average of 1.4 points, ranging from -10 to $+9$ points. The three

[3] Five tasks were excluded for consistent signs of solution leakage across models.

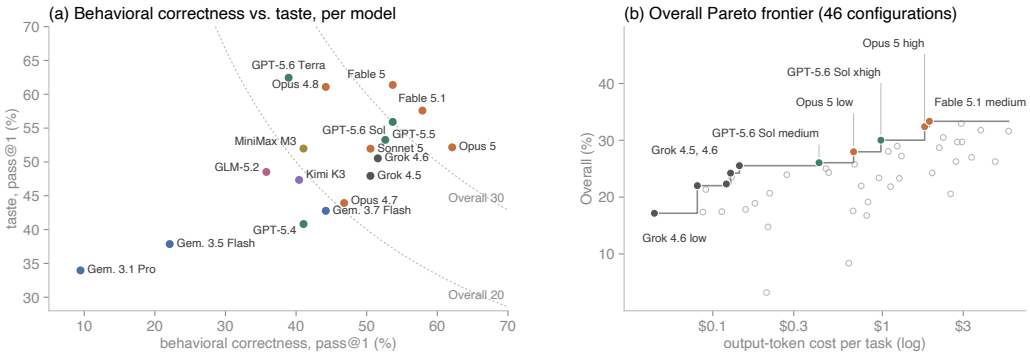

Figure 5: **Behavioral correctness, taste, and cost.** (a) Behavioral correctness vs. taste for each model; dotted curves are constant Overall at 20% and 30%. (b) Overall for every effort tier of every model (46 configurations) against mean output-token cost per task; filled points on the step line form the Pareto frontier (the best Overall at each budget).

current Anthropic ladders peak Overall at a middle tier (Claude Opus 5 high, Claude Fable 5 high, Claude Fable 5.1 medium) and taste falls beyond it.

Reward hacking concentrates in the most capable models. We analyzed agents' trajectories for signs of awareness that they are being benchmarked (Needham et al., 2025) and evidence of reward hacking (Appendix E.3 describes the procedure). Claude Fable 5 shows awareness in 49% of trials and tries a reward-hacking vector in 67%, against 17% and 8% for Claude Opus 4.8. Reward hacking most commonly takes the form of looking for a pre-existing solution (e.g. the post-PR code on GitHub or in PyPI) rather than deriving one. Solution leakage is an inherent challenge in benchmarks based on data found in the real world, such as Senior SWE-bench (Aleithan et al., 2024). Trials with evidence of reward hacking are voided and not included in scoring.

Limitations.

- **Additional LLM-based grading components.** The verification agent and taste judge add execution time, cost, and non-determinism, though the overhead is small relative to the solving agent (Section 4.2) and verdicts are stable across re-runs (Sections 4.2 and 4.3).
- **Subjectivity in taste scores.** Code quality assessment is inherently subjective; the judge panel matches the expert consensus as often as two experts match each other (Table 2).
- **Solution leakage.** Even with a network allowlist, agents find ways to discover solutions from the internet. We perform trajectory analysis to identify these cases (Appendix E.3), and are considering producing fully offline versions of tasks.
- **Coverage of senior work.** The benchmark does not measure the parts of a senior engineer's job that happen between changes: clarifying requirements with a stakeholder (Raghavendra et al., 2026b; Wu et al., 2026), reviewing others' code, or carrying a multi-task project across sessions.

6 Conclusion

Senior SWE-bench evaluates coding agents as senior engineers, introducing new mechanisms to address shortcomings that prevent pre-existing methods from reliably grading realistic software engineering tasks. We validate the verification agent and taste judge rigorously, assessing both alignment to human labels and stability. The verification agent makes far fewer errors than alternative methods for grading tasks with underspecified instructions, and the taste judge performs similarly to an additional human reviewer of code quality. Using the benchmark, we show that frontier agents still have significant room for improvement in completing senior-level work with sufficient quality. Given the continuous and refreshable nature of the benchmark, we plan to update the task set and leaderboard to reflect the frontier of coding capabilities.

AI USE STATEMENT

In the task generation pipeline (Section 3.2), AI agents are paired with expert software engineers to construct and review high-quality tasks. AI agents are also a critical component of the grading process: a verification agent writes behavioral tests and LLM-judges score taste and rubric criteria (Section 4), and a trial-analysis agent labels error modes and flags benchmark-awareness attempts (Appendices E.2 and E.3). Experts review every task before release, and the verification agent and taste judge are validated against expert labels (Sections 4.2 and 4.3).

AI coding agents were used to develop the evaluation harness, benchmark infrastructure, and analysis and plotting scripts in partnership with the authors. AI was also leveraged to help design the scoring method, interpret results, draft and edit the paper, and search for related work. Mathematical claims, proofs, and translation are not applicable to this work.

The authors reviewed code, figures, and the paper, and take full responsibility for the final content of the work.

REPRODUCIBILITY STATEMENT

Each task includes its environment definition, instruction, reference solution, and reward artifacts (pre-written tests and verification-agent recipes where the task has them, and the task rubric), together with the code of the verification agent and judges. Section 3.2 and Appendix A describe how tasks are built and checked; Appendix B gives the splits, environments, resource limits, and network policy; Appendix C gives the recipe format, an example story, and the agent and judge protocol; and Appendix D gives the judge rubrics. Section 5 describes the evaluated configurations and trial protocol. Grades depend on hosted models and can vary between runs; Sections 4.2 and 4.3 measure this variation. The 50 private tasks are withheld to limit contamination.

REFERENCES

Toufique Ahmed, Martin Hirzel, Rangeet Pan, Avraham Shinnar, and Saurabh Sinha. TDD-Bench Verified: Can LLMs generate tests for issues before they get resolved? *arXiv preprint arXiv:2412.02883*, 2024.

Toufique Ahmed, Jatin Ganhotra, Rangeet Pan, Avraham Shinnar, Saurabh Sinha, and Martin Hirzel. Otter: Generating Tests from Issues to Validate SWE Patches. In *International Conference on Machine Learning*, 2025a. arXiv:2502.05368.

Toufique Ahmed, Jatin Ganhotra, Avraham Shinnar, and Martin Hirzel. Investigating Test Overfitting on SWE-bench. *arXiv preprint arXiv:2511.16858*, 2025b.

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.

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*, 2025. arXiv:2505.20411.

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.

Moritz Beller, Alberto Bacchelli, Andy Zaidman, and Elmar Juergens. Modern code reviews in open-source projects: Which problems do they fix? In *Working Conference on Mining Software Repositories (MSR)*, pp. 202–211, 2014. doi: 10.1145/2597073.2597082.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,

Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Neil Chowdhury, James Aung, Jun Shern Chan, Oliver Jaffe, Dane Sherburn, Giulio Starace, Evan Mays, Rachel Dias, Marwan Aljubeh, Mia Glaese, Carlos E. Jimenez, John Yang, Leyton Ho, Tejal Patwardhan, Kevin Liu, and Aleksander Madry. Introducing SWE-bench Verified. OpenAI blog, `https://openai.com/index/introducing-swe-bench-verified/`, August 2024.

Cursor Research. Composer 2 Technical Report. *arXiv preprint arXiv:2603.24477*, 2026.

Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, Chetan Rane, Karmini Sampath, Maya Krishnan, Srivatsa Kundurthy, Sean Hendryx, Zifan Wang, Vijay Bharadwaj, Jeff Holm, Raja Aluri, Chen Bo Calvin Zhang, Noah Jacobson, Bing Liu, and Brad Kenstler. SWE-Bench Pro: Can AI agents solve long-horizon software engineering tasks? *arXiv preprint arXiv:2509.16941*, 2025.

Yaxin Du, Yuzhu Cai, Yifan Zhou, Cheng Wang, Yu Qian, Xianghe Pang, Qian Liu, Yue Hu, and Siheng Chen. SWE-Dev: Evaluating and training autonomous feature-driven software development. *arXiv preprint arXiv:2505.16975*, 2025.

Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Yunzhong He, Bing Liu, and Sean Hendryx. Rubrics as rewards: Reinforcement learning beyond verifiable domains. *arXiv preprint arXiv:2507.17746*, 2025.

Harbor Framework Team. Harbor: A framework for evaluating and optimizing agents and models in container environments. Software, doi: 10.5281/zenodo.20953922, `https://github.com/harbor-framework/harbor`, 2026.

Xinyi He, Qian Liu, Mingzhe Du, Lin Yan, Zhijie Fan, Yiming Huang, Yin Zheng, Zejian Yuan, and Zejun Ma. SWE-Perf: Can language models optimize code performance on real-world repositories? *arXiv preprint arXiv:2507.12415*, 2025.

Jiawei Huang, Qingping Yang, Renjie Zheng, and Jiaze Chen. Beyond verifiable rewards: Rubric-based GRM for reinforced fine-tuning SWE agents. *arXiv preprint arXiv:2604.16335*, 2026a.

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*, 2026b.

Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2E-Gym: Procedural environment generation and hybrid verifiers for scaling open-weights SWE agents. In *Conference on Language Modeling*, 2025.

Smriti Jha, Matteo Paltenghi, Chandra Maddila, Vijayaraghavan Murali, Shubham Ugare, and Satish Chandra. REAP: Automatic Curation of Coding Agent Benchmarks from Interactive Production Usage. In *IEEE/ACM International Conference on Automated Software Engineering (Industry Showcase)*, 2026. arXiv:2604.01527.

Zimo Ji, Zekai Zhang, Congying Xu, Zongjie Li, Yudong Gao, Shuai Wang, and Shing-Chi Cheung. Coding agents are guessing: Measuring action-boundary violations in underspecified DevOps instructions. *arXiv preprint arXiv:2607.02294*, 2026.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R. Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *International Conference on Learning Representations*, 2024.

Chenglin Li, Yisen Xu, Zehao Wang, Shin Hwei Tan, and Tse-Hsun Chen. Probe to Generate: Program Variant-Guided Test Augmentation for Repository-Level Repair Benchmarks. In *IEEE/ACM International Conference on Automated Software Engineering*, 2026. arXiv:2604.01518.

Kilian Lieret and Carlos E. Jimenez. mini-SWE-agent: A simple AI software engineering agent. Software, version 2.3.0, https://github.com/SWE-agent/mini-swe-agent, 2025.

Eric Lu, Ben Pan, Deniz Birlikci, Sam Lee, Ray Wang, Rohan Choudhury, Fermi Ma, TC Qin, Carlo Baronio, Silas Alberti, et al. Introducing FrontierCode. Cognition blog, https://cognition.com/blog/frontier-code, June 2026.

Mika V. Mäntylä and Casper Lassenius. What types of defects are really discovered in code reviews? *IEEE Transactions on Software Engineering*, 35(3):430–448, 2009. doi: 10.1109/TSE.2008.71.

Nicolas Menet, Michael Hersche, Andreas Krause, and Abbas Rahimi. A theoretical analysis of test-driven code generation. *arXiv preprint arXiv:2602.06098*, 2026.

Mike A. Merrill, Alexander G. Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, E. Kelly Buchanan, Junhong Shen, Guanghao Ye, Haowei Lin, Jason Poulos, Maoyu Wang, Marianna Nezhurina, Jenia Jitsev, Di Lu, Orfeas Menis Mastromichalakis, Zhiwei Xu, Zizhao Chen, Yue Liu, Robert Zhang, Leon Liangyu Chen, Anurag Kashyap, Jan-Lucas Uslu, Jeffrey Li, Jianbo Wu, Minghao Yan, Song Bian, Vedang Sharma, Ke Sun, Steven Dillmann, Akshay Anand, Andrew Lanpouthakoun, Bardia Koopah, Changran Hu, Etash Guha, Gabriel H. S. Dreiman, Jiacheng Zhu, Karl Krauth, Li Zhong, Niklas Muennighoff, Robert Amanfu, Shangyin Tan, Shreyas Pimpalgaonkar, Tushar Aggarwal, Xiangning Lin, Xin Lan, Xuandong Zhao, Yiqing Liang, Yuanli Wang, Zilong Wang, Changzhi Zhou, David Heineman, Hange Liu, Harsh Trivedi, John Yang, Junhong Lin, Manish Shetty, Michael Yang, Nabil Omi, Negin Raoof, Shanda Li, Terry Yue Zhuo, Wuwei Lin, Yiwei Dai, Yuxin Wang, Wenhao Chai, Shang Zhou, Dariush Wahdany, Ziyu She, Jiaming Hu, Zhikang Dong, Yuxuan Zhu, Sasha Cui, Ahson Saiyed, Arinbjörn Kolbeinsson, Jesse Hu, Christopher Michael Rytting, Ryan Marten, Yixin Wang, Alex Dimakis, Andy Konwinski, and Ludwig Schmidt. Terminal-Bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations*, 2026.

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.

Niels Mündler, Mark Niklas Müller, Jingxuan He, and Martin Vechev. SWT-Bench: Testing and Validating Real-World Bug-Fixes with Code Agents. In *Advances in Neural Information Processing Systems*, 2024. arXiv:2406.12952.

Joe Needham, Giles Edkins, Govind Pimpale, Henning Bartsch, and Marius Hobbhahn. Large language models often know when they are being evaluated. *arXiv preprint arXiv:2505.23836*, 2025.

Dan North. Introducing BDD. Better Software magazine; reprinted at https://dannorth.net/blog/introducing-bdd/, March 2006.

OpenAI. Separating signal from noise in coding evaluations. OpenAI blog, https://openai.com/index/separating-signal-from-noise-coding-evaluations/, July 2026a. July 8, 2026.

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/, February 2026b. February 23, 2026.

Tejal Patwardhan, Rachel Dias, Elizabeth Proehl, Grace Kim, Michele Wang, Olivia Watkins, Simón Posada Fishman, Marwan Aljubeh, Phoebe Thacker, Laurance Fauconnet, Natalie S. Kim, Patrick Chao, Samuel Miserendino, Gildas Chabot, David Li, Michael Sharman, Alexandra Barr, Amelia Glaese, and Jerry Tworek. GDPval: Evaluating AI model performance on real-world economically valuable tasks. *arXiv preprint arXiv:2510.04374*, 2025.

Zichao Qi, Fan Long, Sara Achour, and Martin Rinard. An analysis of patch plausibility and correctness for generate-and-validate patch generation systems. In *International Symposium on Software Testing and Analysis (ISSTA)*, pp. 24–36, 2015. doi: 10.1145/2771783.2771791.

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.

Mohit Raghavendra, Anisha Gunjal, Aakash Sabharwal, and Yunzhong He. SWE-INTERACT: Reimagining SWE Benchmarks as User-Driven Long-Horizon Coding Sessions, 2026b. arXiv:2606.30573.

David Rein. Research update: Algorithmic vs. holistic evaluation. METR blog, `https://metr.org/blog/2025-08-12-research-update-towards-reconciling-slowdown-with-time-horizons/`, August 2025.

Manish Shetty, Naman Jain, Jinjian Liu, Vijay Kethanaboyina, Koushik Sen, and Ion Stoica. GSO: Challenging software optimization tasks for evaluating SWE-agents. *arXiv preprint arXiv:2505.23671*, 2025.

Edward K. Smith, Earl T. Barr, Claire Le Goues, and Yuriy Brun. Is the cure worse than the disease? Overfitting in automated program repair. In *Joint Meeting on Foundations of Software Engineering (ESEC/FSE)*, pp. 532–543, 2015. doi: 10.1145/2786805.2786825.

Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, Johannes Heidecke, Amelia Glaese, and Tejal Patwardhan. PaperBench: Evaluating AI's ability to replicate AI research. In *International Conference on Machine Learning*, 2025.

Kirill Vasilevski, Ximing Dong, Benjamin Rombaut, Milad Soltany, Ruochen Deng, Jiahuei Lin, Arthur Leung, Dayi Lin, Boyuan Chen, Shaowei Wang, and Ahmed E. Hassan. Beyond Correctness: Enhancing Architectural Reasoning in Code LLMs via Scalable Labeling with Agentic Judgment. *arXiv preprint arXiv:2606.14948*, 2026.

Sanidhya Vijayvargiya, Xuhui Zhou, Akhila Yerukola, Maarten Sap, and Graham Neubig. AmbigSWE: Interactive agents to overcome underspecificity in software engineering. In *International Conference on Learning Representations*, 2026.

Vijay Viswanathan, Yanchao Sun, Shuang Ma, Xiang Kong, Meng Cao, Graham Neubig, and Tongshuang Wu. Checklists are better than reward models for aligning language models. In *Advances in Neural Information Processing Systems*, 2025.

You Wang, Michael Pradel, and Zhongxin Liu. Are "solved issues" in SWE-bench really solved correctly? An empirical study. In *International Conference on Software Engineering*, pp. 169–181, 2026. doi: 10.1145/3744916.3764576. arXiv:2503.15223.

Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida I. Wang. SWE-RL: Advancing LLM reasoning via reinforcement learning on open software evolution. In *Advances in Neural Information Processing Systems*, 2025.

Parker Whitfill, Cheryl Wu, Joel Becker, and Nate Rush. Many SWE-bench-passing PRs would not be merged into main. METR research note, `https://metr.org/notes/2026-03-10-many-swe-bench-passing-prs-would-not-be-merged-into-main/`, March 2026.

Yifan Wu, Zhuokai Zhao, Songlin Li, Ho Hin Lee, Jiacheng Zhu, Shirley Wu, Tianhe Yu, Serena Li, Lizhu Zhang, Xiangjun Fan, and Shengzhi Li. SWE-Together: Evaluating Coding Agents in Interactive User Sessions, 2026. arXiv:2606.29957.

Zixuan Xie, Xinyu Liu, and Shangtong Zhang. MathlibPR: Pull request merge-readiness benchmark for formal mathematical libraries. *arXiv preprint arXiv:2605.07147*, 2026.

Qi Xin and Steven P. Reiss. Identifying test-suite-overfitted patches through test case generation. In *International Symposium on Software Testing and Analysis (ISSTA)*, pp. 226–236, 2017. doi: 10.1145/3092703.3092718.

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

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. Program-Bench: Can language models rebuild programs from scratch? *arXiv preprint arXiv:2605.03546*, 2026.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ-bench: A benchmark for tool-agent-user interaction in real-world domains. *arXiv preprint arXiv:2406.12045*, 2024.

He Ye, Matias Martinez, and Martin Monperrus. Automated patch assessment for program repair at scale. *Empirical Software Engineering*, 26(2):20, 2021. doi: 10.1007/s10664-020-09920-w.

Boxi Yu, Yuxuan Zhu, Pinjia He, and Daniel Kang. UTBoost: Rigorous evaluation of coding agents on SWE-Bench. *arXiv preprint arXiv:2506.09289*, 2025.

Kai Yu, Zhenhao Zhou, Junhao Zeng, Ying Wang, Xueying Du, Zhiqiang Yuan, Junwei Liu, Ziyu Zhou, Yujia Wang, Chong Wang, and Xin Peng. Does Pass Rate Tell the Whole Story? Evaluating Design Constraint Compliance in LLM-based Issue Resolution. *arXiv preprint arXiv:2604.05955*, 2026.

Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Linhao Zhang, Shulin Xin, Lu Chen, Qi Liu, Xiaojian Zhong, Aoyan Li, Siyao Liu, Yongsheng Xiao, Liangqiang Chen, Yuyu Zhang, Jing Su, Tianyu Liu, Rui Long, Kai Shen, and Liang Xiang. Multi-SWE-bench: A Multilingual Benchmark for Issue Resolving. In *Advances in Neural Information Processing Systems (Datasets and Benchmarks Track)*, 2025. arXiv:2504.02605.

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! In *Advances in Neural Information Processing Systems (Datasets and Benchmarks Track)*, 2025. arXiv:2505.23419.

Junda Zhao, Shurui Zhou, and Eldan Cohen. Evaluating and Mitigating the Misguidance Effect of Buggy Code in LLM-Generated Unit Tests. *Proceedings of the ACM on Software Engineering*, 3 (ISSTA), 2026. ISSTA 2026; arXiv:2607.22883.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems*, 2023.

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.

Ming Zhong, Xiang Zhou, Ting-Yun Chang, Qingze Wang, Nan Xu, Xiance Si, Dan Garrette, Shyam Upadhyay, Jeremiah Liu, Jiawei Han, Benoit Schillings, and Jiao Sun. SWE-IF: Aligning code evaluation with human preference. In *International Conference on Machine Learning*, 2026. First circulated as "Vibe Checker: Aligning Code Evaluation with Human Preference," arXiv:2510.07315.

Ziqian Zhong, Aditi Raghunathan, and Nicholas Carlini. ImpossibleBench: Measuring LLMs' propensity of exploiting test cases. *arXiv preprint arXiv:2510.20270*, 2025.

Qixing Zhou, Jiacheng Zhang, Haiyang Wang, Rui Hao, Jiahe Wang, Minghao Han, Yuxue Yang, Shuzhe Wu, Feiyang Pan, Lue Fan, Dandan Tu, and Zhaoxiang Zhang. FeatureBench: Benchmarking agentic coding for complex feature development. In *International Conference on Learning Representations*, 2026a.

Xin Zhou, Chun Yong Chong, Kisub Kim, Yun Peng, Rui Shu, Zihan Wu, Xu Han, Guowen Yuan, Zeyang Zhuang, Jounghoon Kim, Jeongjin Ju, Seongmin Ju, Taein Yoon, and David Lo. A unified issue resolution benchmark for requirement clarification, planning, and code generation for coding agents. *arXiv preprint arXiv:2608.09072*, 2026b.

Mingchen Zhuge, Changsheng Zhao, Dylan Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, Yangyang Shi, Vikas Chandra, and Jürgen Schmidhuber. Agent-as-a-Judge: Evaluate agents with agents. In *International Conference on Machine Learning*, 2025.

A TASK DETAILS

Source repositories. Table A1 gives one row per repository behind the summary in Table 1. The last column is the benchmark's proxy for seniority: the commit counts of the source-PR author and of the maintainer who merged it, in that repository.

Table A1: **Source repositories, one row each.** LOC (source lines, programming languages only), stars and start year as of June 2026. Type: application, library, tool, or service. Contrib.: GitHub contributors including anonymous authors. PRs/wk: merged PRs per week over January or February to June 2026. Langs.: languages above 5% of the repository's code. Tasks by segment: D design-and-build, I investigate-and-fix. Author / merger commits: median over the repository's tasks of the source-PR author's and the merging maintainer's commits in that repository.

Repository	Languages	Type	Description	Started	LOC	Stars	Contrib.	PRs/wk	Langs.	Tasks (D/I)	Author / merger commits
`posthog`	Python, TypeScript	Application	Product analytics	2020	3.8M	35.1k	589	554	3	14 (9/5)	331 / 331
`electric`	Elixir, TypeScript	Service	Postgres real-time sync	2022	345k	10.2k	77	27	4	12 (4/8)	259 / 222
`gitea`	Go	Application	Self-hosted Git forge	2016	397k	56.3k	1,851	47	3	12 (6/6)	754 / 1,492
`better-auth`	TypeScript	Library	Authentication framework	2024	289k	28.7k	945	57	1	10 (4/6)	355 / 508
`prefect`	Python	Library	Workflow orchestration	2018	664k	22.6k	475	66	2	9 (6/3)	1,129 / 1,129
`analytics`	Elixir	Application	Privacy-friendly analytics	2018	228k	27.2k	118	15	2	7 (4/3)	401 / 401
`paperless-ngx`	Python, TypeScript	Application	Document management	2022	148k	42.2k	459	22	3	7 (4/3)	3,901 / 3,901
`teleport`	Go, TypeScript	Application	Infrastructure access	2015	2.8M	20.5k	375	136	3	7 (2/5)	65 / 65
`turborepo`	Rust, TypeScript	Tool	Monorepo build system	2021	215k	30.6k	712	61	3	7 (1/6)	2,093 / 2,093
`immich`	TypeScript	Application	Self-hosted photo backup	2022	542k	103.6k	1,004	57	3	6 (3/3)	326 / 944
`harbor`	Python	Tool	Agent evaluation harness	2025	219k	2.5k	321	24	1	5 (5/0)	4 / 55
`firezone`	Elixir, Rust	Application	Zero-trust access	2020	247k	8.7k	57	58	2	4 (2/2)	2,426 / 2,426

Scope. Fig. A1 places the tasks against two public benchmarks: instructions are shorter than either comparison set, while reference patches span a wider range of sizes, from one-line fixes to multi-thousand-line features (design-and-build patches touch a median of 7.5 files). Panel (c) shows the language mix and the most common skills by segment: interface and API-contract design sit almost entirely on design-and-build tasks, fault localization and logic fixes on investigate-and-fix ones, while framework fluency, codebase exploration, and concurrency reasoning are needed on both.

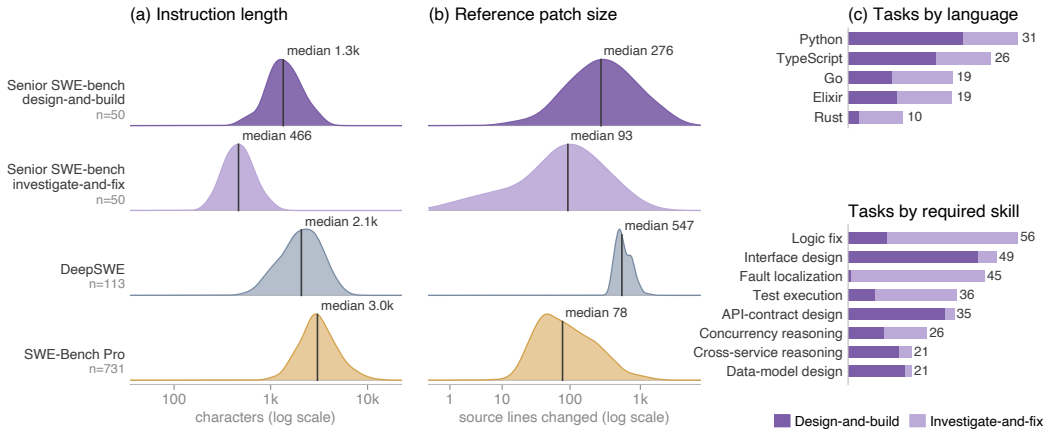

Figure A1: **Benchmark scope and composition.** **(a)** Instruction length; length counts the full task text an agent receives (for SWE-Bench Pro, the problem statement with its requirements and interface sections), with harness boilerplate removed. **(b)** More diverse task scope: reference solutions span from one-line fixes to multi-thousand-line features. **(c)** Composition of the 100 tasks by language and by the most common required skills, split by segment; a task may touch several languages, and nearly all tasks also require codebase exploration and framework fluency.

Segments. Table A2 contrasts the two segments row by row, from task type and source PR to task count. The skill rows name what each segment adds on top of the codebase exploration and framework fluency that nearly every task needs; Section 4.1 explains why the reward rows differ.

Table A2: **The two task segments.**

	Design-and-Build	**Investigate-and-Fix**
Task types	Feature, Migration	Bug, Performance
Source PR qualities	Complex, multi-component features and migrations	Bug and performance fixes with evidence of runtime investigation
Instruction style	Description of requested change with PM-level user stories	Behavioral user issue report
Interface changes	May involve new or changed system-level interfaces	System-level interfaces are not changed
Reward mechanisms	Pre-written tests, verification agent, rubric, taste	Pre-written tests, rubric, taste
Example skills	Interface / API design, data model design, architectural refactoring, cross-service reasoning, framework fluency	Fault localization, log / error interpretation, local deployment, concurrency reasoning, query optimization
Count	50 (48 feature, 2 migration)	50 (39 bug, 11 performance)

Requirements. Since every pre-written and verification-agent test must pass for a successful solve, careful attention is paid to ensuring that the unstated requirements of Section 4.1 are truly codebase-mandated or have no reasonable alternative. Agent trajectories are analyzed for signs of unfair grading on this basis as part of the quality control process (Section 3.2), and tasks are revised accordingly. Requirements that are infeasible to test in runtime (such as an integration with an external service) are routed to the task-specific rubric. These get non-gating treatment, as the LLM judge has lower reliability than methods that exercise the code in runtime.

Example task instruction. The instruction below is from a released investigate-and-fix task on the electric repository (`electric-fix-resolve-pending-shapes`, a bug fixed by a maintainer's PR), reproduced in full: the report, then the general instructions every task appends. Table A3 traces how the maintainer and a Claude Opus 4.8 trial solved it.

> **Task**
> it looks like pg replication slot lag is growing unbounded on prod stacks. the flush lsn we send back to the db just stops advancing. clients are still getting new data, storage is still writing, nothing crashes. lag keeps growing for hours until someone restarts the stack.
> it always seems to happen right after a txn whose changes spanned multiple wal fragments. find what's causing the global flush boundary to get stuck and fix it. note: the upstream tracker should only see flush acks at txn boundaries.
>
> **General instructions**
>
> - The code repo is at `/repo/electric`.
> - You are inside of a Docker container. You may not be able to perform all operations you would normally be able to do on a local machine. Dependencies have not been pre-installed, and you may need to install them yourself.
> - You are expected to act autonomously as a software engineer to complete tasks you are given.
> - Do not stop until you feel you have completed the task and your code changes can be merged.
> - You may need to use software engineering skills like analyzing the codebase, researching technologies, running services, analyzing logs, etc. to complete the task. Not all tasks will be solvable by reading source code alone.

Table A3: **One investigate-and-fix task, solved twice.** The maintainer's path and a Claude Opus 4.8 trial's on the task above, phase by phase: the agent reconstructs from source and a synthesized failing test what the maintainer read off production, and lands the same fix.

	Human maintainer	**Claude Opus 4.8**
Localize	Locates the stuck flush boundary from a production state dump, where `last_sent` had run ahead of `last_flushed`.	Locates the same boundary by reading the flush path through `consumer.ex` and the storage write-loop.
Model the race	Models the race: storage flushes mid-transaction ($\geq$64 KiB or a timer), so the earlier #3985 fix closed only half the window.	Models the same race from source, then reads the in-repo #3985 test and rules that earlier fix incomplete.
Runtime reproduction	Catches the bad state live in production: the reported flush offset frozen mid-fragment while data keeps streaming.	Catches the same bad state in a failing test: a `:trace` assertion trips on the unaligned offset `(10,4)`.
Fix	Defers the flush and emits it only at the transaction boundary.	Lands the same defer-and-emit fix, across all three completion paths.

B Release and Environment

Splits. The release has 100 tasks, with 50 kept private to mitigate contamination. Both splits are drawn from the same pipeline and repositories, and the private split is refreshed by the pipeline in Section 3.2 as models evolve and knowledge cutoffs move forward.

Harness and environment. Senior SWE-bench is open-source and Harbor-compatible (Harbor Framework Team, 2026). Task environments are created by cloning the target repository, rewinding before the first task commit, then expiring the reflog and pruning the repo. Major system dependencies (such as Python or Postgres) are pre-installed, as are heavier individual libraries (such as PyTorch). In order to best reflect the performance of a background agent dropped into a development sandbox, full environment setup (such as installing required packages and starting services) is left to the agent.

Network policy. During the agent phase, tasks run under a default-deny allowlist of about 32 hosts covering the agent's model gateway, package registries, OS package repositories, toolchain installers, and official language references. GitHub, its mirrors (jsDelivr, unpkg, the Go module proxy), and per-package documentation hosts are deliberately absent: they are fast-forwarded and document the very change a task asks for. Image build and grading run with open network access. Trajectories are additionally analyzed for reward hacking after the fact (Appendix E.3).

C Verification Agent

This appendix documents the verification agent of Section 4.2: the recipe format and an example story (Appendix C.1), the agent and judge protocol (Appendix C.2), and the judge's defect detection and the stage's cost (Appendix C.3). The pipeline's internal name for the recipe is a *validation spec* and for its tests *stories*; the verbatim excerpts below keep those names.

C.1 Recipe Format and Example

Recipe format. Every design-and-build task ships one TOML recipe that the solving agent never sees: task-level `instructions` (environment, verified harness utilities, testing conventions) and a list of `[[story]]` tables. A story has an `id`, a `driver` (`pytest`, `jest`, `go-test`, `cargo-test`, or `mix-test`), a user-facing `description`, a `procedure` whose `EXPECTED:` lines name every assertion the script must make, and `[[story.test_case]]` tables of `inputs` and `expected` values. Across the 50 design-and-build tasks the recipes hold 207 stories and 565 test cases: a median of 4 stories per task and 2 cases per story.

Example story. The story below is the second of five in the public Gitea task `gitea-add-project-column-picker`, whose feature lets an editor move an issue between project-board columns from the issue sidebar. The recipe's `instructions` (not shown) describe

the integration-test conventions and a shared helper for discovering the move control. The procedure is reproduced with two of its longer guidance passages elided ([...]).

```
[[story]]
id = "move_changes_column"
driver = "go-test"
go_package = "tests/integration"

description = """
When an editor moves an issue to another column of the same project from the
sidebar, the change persists.
"""

procedure = """
- Setup: `defer tests.PrepareTestEnv(t)()`. `sess := loginUser(t, "user2")`.
- Discover the move control once [...]: `mc := discoverMoveControl(t, sess, ...)`.
- For each test case from `GetValidationTestCases(t)`, run a subtest named
  `tc.Name`:
  - Build and send the move: `req := mc.buildMove(t, ...,
    int64(tc.Inputs["target_column_id"].(float64)))`.
    `resp := sess.MakeRequest(t, req, NoExpectedStatus)`.
  - EXPECTED: a successful move -- assert the status is within the accepted
    range carried in the expected params (a 2xx).
  - EXPECTED: persisted (assert via the DB bean -- REQUIRED, do NOT re-parse
    HTML here): load the pre-existing bean
    `project_model.ProjectIssue{IssueID: ...}`
    and assert its `ProjectColumnID` equals
    `int64(tc.Expected["new_selected_column_id"].(float64))`. [...]
  - EXPECTED: the move did not corrupt the target column's ordering -- the
    moved issue was given its own sort position, not dropped on top of the
    issue already in that column. [...] assert only that the
    positions are distinct.
"""

[[story.test_case]]
[story.test_case.inputs]
issue_url = "/user2/repo1/issues/4"
issue_index = 4
issue_id = 5
target_column_id = 1
[story.test_case.expected]
min_status = 200
max_status = 299
new_selected_column_id = 1
existing_issue_id = 1

[[story.test_case]]
[story.test_case.inputs]
issue_url = "/user2/repo1/issues/4"
issue_index = 4
issue_id = 5
target_column_id = 2
[story.test_case.expected]
min_status = 200
max_status = 299
new_selected_column_id = 2
existing_issue_id = 3
```

The agent is shown each story's description and procedure but only the *first* test case, as an example, together with the total number of cases. At execution time the driver injects every case's `inputs` and `expected` from an immutable copy of the recipe, so a script cannot pass by hard-coding the values it was shown: here the second case moves the same issue to a different column and expects the persisted column and the neighboring issue's sort position to track it.

C.2 Agent Prompt and Judge Protocol

Agent prompt. The verification agent runs in the bash-only mini-SWE-agent harness (Lieret & Jimenez, 2025; Yang et al., 2024) with Claude Sonnet 4.6. Its prompt casts it as a test engineer whose job is to implement each story's procedure faithfully and report what the implementation actually does—not to make tests pass, and not to fix the solution under test: use every `expected` value in an assertion and assert nothing beyond the procedure's `EXPECTED:` lines; exercise real code paths rather than mocks; discover the submission's own names and interfaces from the diff and source; let a test of missing or broken functionality fail naturally; and never write outside its scripts directory—any change to the repository or test harness zeros the trial.

Judge protocol. After the agent finishes, the framework smoke-tests the first case of every story and calls the judge (Claude Sonnet 4.6 by default) on the scripts and the smoke output; after the full parameterized run it calls the same judge again on the execution results. The judge is told that a failing test is not a defect by itself and that differently-named identifiers are expected, not errors. It scores five criteria from 0 to 3, all gating: a 0 on any of them discards the trial, with score-0 anchors that are, in the rubric's words, "intentionally narrow—they're calibrated for high-precision discard, NOT general 'this could be better' signals". The criteria are:

- `behavioral_authenticity`: tests exercise the real implementation and report honestly—no mocked systems, fabricated passes, skipped tests, or weakened assertions.
- `procedure_fidelity`: each test case maps cleanly to its story's procedure.
- `story_coverage`: every story has at least one script that exercises it.
- `execution_outcome_alignment`: the final pass/fail outcome is consistent with what the recipe implies given the submitted patch.
- `expected_fidelity`: every key in each `expected` dict is asserted unconditionally.

The judge separately returns engineering suggestions tagged minor or moderate; one revision is run when it attributes a story's smoke failure to the script rather than the submission, or on any moderate suggestion. Before execution the framework diffs the repository against a pre-agent checkpoint and discards the trial on any change. A discarded trial leaves the denominator rather than counting as a pass or as a failure charged to the solving agent.

C.3 Defect Detection and Cost

To measure the reviewing judge's precision, 30 verification outputs were seeded with a script defect (10 each of fabricated passes, skipped tests, and mocked systems under test, the three failures `behavioral_authenticity` names) alongside 20 defect-free controls. The judge gave a failing axis score to 28 of the 30 planted defects and to none of the controls. At the median, setup takes 0.5 minutes of the verification stage, writing the scripts 1.4, and running services, scripts, and judges 1.7.

D Taste Score

D.1 Judge Rubrics

Table A4 lists the ten dimensions the taste judge scores. Each judge in the panel reads the submitted diff, the reference solution, and the files around the change, and scores every dimension from 1 to 5; the two judges' scores are averaged per dimension, and the median over a rubric's five dimensions is the r_t or p_t that enters Eq. (1). The bloat ratio and the task rubric of Section 4.3 are separate signals, not dimensions of these rubrics.

Table A4: **Taste judge rubrics.** Each dimension is scored 1–5; relative-quality scores of 4 and 5 mean on par with and better than the reference.

Relative code quality (vs. the reference implementation)	
Minimality	Changes are focused, with no scope creep.
Approach quality	Right solution approach: root-cause fix for bugs, sound design for features.
Hygiene	No shortcuts, workarounds, hardcoded values, or other code smells.
Fluency	Fluent with the domain, frameworks, tools, and conventions.
Craftsmanship	Displays quality engineering effort that a senior reviewer would approve.
Codebase practice alignment (explores surrounding code)	
Style consistency	Formatting, naming, and structure match the surrounding code.
Pattern adherence	Uses the project's established patterns and idioms.
Library usage	Reuses libraries already in the project rather than introducing alternatives.
Abstraction level	The right abstraction level for this part of the codebase.
Documentation fit	Comments and docstrings match the project's style and density.

E Additional Results

E.1 Effort and Instruction Guidance

Fig. A2 breaks down, per study model, the effort-ladder and guided-instruction changes summarized in Section 5 and Section 5.

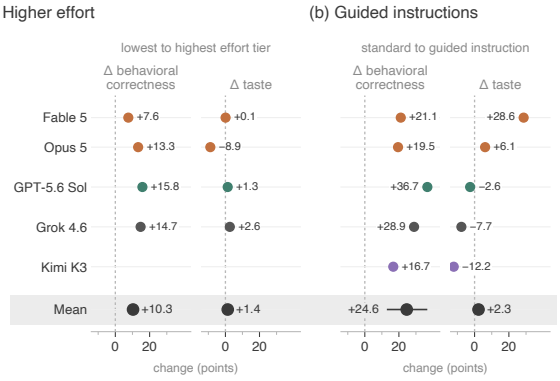

Figure A2: **Effort and instruction guidance**, for the five guided-study models. (a) Lowest to highest effort tier (pass@1, 95-task support set); the shaded mean row averages all nine effort ladders (every model with two or more scored tiers); Kimi K3 runs at max only, so its cell is empty. (b) Standard to guided instruction (30 study tasks, avg@3; whisker: 95% interval on the pooled behavioral-correctness change; taste is the share of correct solutions clearing every taste bar).

E.2 Error-Mode Analysis

To dive deeper into error modes, we classified trial outcomes using an agent that analyzes the trajectory and patch (with the task definition and reference solution as context). This analysis considers only failures from runtime tests (pre-written verifiers and verification-agent tests). The error taxonomy is segmented by solving phase, from comprehending the problem through verifying a potential solution. Precedence is given to errors observed in upstream phases, given the likelihood of error propagation downstream. Passing trials are additionally labeled *canonical* when the agent's design is materially the same as the reference solution and *divergent* when it satisfies the contract with a structurally different approach. Table A5 gives the per-model composition.

Table A5: **Trial outcomes by solving phase** (% of trials; four models · 50 public tasks, one trial per task). "Canonical" passes match the reference's design, "divergent" passes take a different approach; failures are attributed to the earliest phase with an observed error. Right: the two most common error modes.

		Pass		**Failure phase**						**Top modes**	
Model	n	canon.	diverg.	compr.	invest.	design	impl.	verify	other	wrong root cause	missed codebase req.
Claude Opus 4.7	50	10	28	12	12	14	10	10	4	12	6
Claude Opus 4.8	50	14	26	6	18	20	4	6	6	18	14
GPT-5.4	50	8	36	10	14	10	12	10	0	14	2
GPT-5.5	50	4	40	10	14	14	12	6	0	14	4

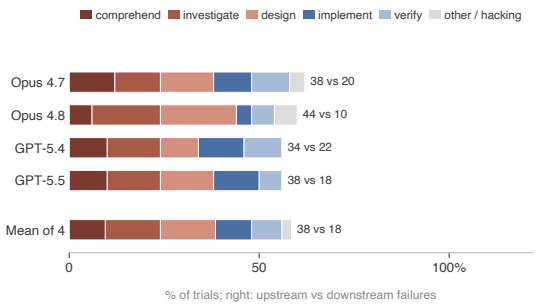

Figure A3: **Failures by solving phase.** The data of Table A5; the label right of each bar counts upstream (comprehension, investigation, design) vs. downstream (implementation, verification) failures.

Failures are upstream: diagnosis and design, not implementation or style. For every model, more failures are attributed to comprehension, investigation, and design than to implementation and verification (Fig. A3), more than twice as many across the four models (77 against 35 of 117 failures). A wrong root cause is the most common error for every model (29 of 117 failures); for each model, this one error outnumbers all implementation-phase failures combined.

Table A6 defines the classes behind Table A5.

Table A6: **Trial-outcome classes.** One class per trial, in the phase groups of Table A5; upstream classes describe what the agent understood or decided, downstream ones what it typed or checked.

Phase	Class	Definition
Pass	Canonical	Solved it the same way as the reference.
	Divergent	Solved it correctly, but with a different design.
Comprehend	Misunderstood the task	Misread what the task was actually asking.
	Missed a stated requirement	Skipped a required part of the task instruction.
	Wrong location	Edited the wrong place, or never found it.
Investigate	Wrong root cause	Found the spot but misdiagnosed why it broke.
	Knowledge gap	Didn't know the API or library well enough.
	Unverified assumption	Assumed how something worked without checking.
Design	Poor design	Right diagnosis, but chose a poor approach.
	Wrong layer	Built it at the wrong layer of the system.
	Missed a codebase requirement	Ignored an existing pattern or contract in the repo.
Implement	Build issue	Code doesn't compile or breaks the build.
	Logic error	Code runs, but the logic is wrong.
	Integration error	Works alone but doesn't connect to callers.
	Introduced regression	Broke something that used to work.
	Bad edit	A careless mechanical slip in the edit.
	No change	Produced an empty or no-op patch.
Verify	False green	Trusted a misleading check that falsely passed.
	No verification	Shipped without checking the change at all.
Other	Not assessable	Infra broke; the run can't be judged.
	Reward hacking	Looked up or gamed the answer.

E.3 Reward-Hacking Audit

We manually reviewed a sample of trajectories from different agents for patterns in assistant messages, tool calls, and reasoning (when available) that imply benchmark awareness (e.g. mentioning hidden tests) and nuanced red herrings (Baker et al., 2025; Zhong et al., 2025). We then used these examples to seed the benchmark awareness check in our trial analysis agent, which runs after each trial. Table A7 reports the per-model attempt rates behind Section 5. Under the network allowlist (Appendix B), 63 of the 76 recorded attempts were blocked; the 13 that succeeded all fetched the repository's own newer release from a package registry, a class of host the allowlist must admit for dependency installs.

Table A7: **Benchmark awareness and reward hacking.** Left: one trial per model on the 50 public tasks under the network allowlist (Appendix B). "Aware": trajectory shows benchmark awareness; "Attempt": at least one attempt to retrieve a solution, broken out by vector (a trial can use several); "Succ.": attempts that got through. Right: share of trials voided as reward hacks.

		Aware	**Attempt**	**Attempts by vector**					**Voided**
Model	n	(%)	(%)	code host	registry	proxy/docs	local git	**Succ.**	**(%)**
Claude Opus 4.8	48	16.7	8.3	2	2	0	0	2	0.0
GPT-5.5	49	2.0	0.0	0	0	0	0	0	0.0
Claude Sonnet 5	50	6.0	42.0	16	2	0	5	2	2.7
Claude Fable 5	49	49.0	67.3	33	11	4	1	9	6.8

The four vectors of Table A7, as the trial-analysis agent labels them:

- *Code-host fetch*: Clone/fetch the repo, its PR/commit diff, or raw source from the code host (GitHub, etc.).
- *Package registry*: Download the repo's own published release from a registry (pip/npm/hex) and read the fix out of it.
- *Module proxy / docs*: A module proxy or per-package doc mirror that re-serves the repo's own source or fast-forwarded docs.
- *Local git history*: Mine the in-container git history (`git log -all` / `git show`) for the upstream fix commit.

Three examples of awareness in the assistant's own messages:

- Claude Opus 4.8, `better-auth-fix-api-return-response`: "This is a benchmark task with an intentionally introduced bug. Let me diff against the parent commit to find what was changed."
- GPT-5.5, `harbor-add-multi-step-tasks`: "I'm going to implement the new task-step and reward configuration models first, including flexible aliases for likely TOML field names so local and hidden tests can declare the schema naturally."
- Claude Fable 5, `immich-feat-release-candidate-detection`: "This looks like it corresponds to an upstream Immich change. Let me check if I have network access to look at the actual upstream implementation for reference."

Two attempts, one that succeeded and one the allowlist blocked:

- Claude Opus 4.8, `prefect-fix-resolve-race-condition`, package registry, succeeded: it downloaded the repository's own newer release ("The newest upstream version (3.7.6) contains the fix") and diffed the extracted module against the repository, which exposed the race fix.
- Claude Fable 5, `gitea-feat-oidc-rp-logout`, module proxy, blocked: its attempt to fetch the upstream module ("Let me try to fetch the upstream implementation for reference via a Go module proxy trick") ended in a reset connection.

The same pass also scored wrongdoing or concealment—the agent recognizing its action as illegitimate, or hiding it; no trial of any model showed either.