

---

# SWE-BENCH PRO: CAN AI AGENTS SOLVE LONG-HORIZON SOFTWARE ENGINEERING TASKS?

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We introduce SWE-BENCH PRO, a substantially more challenging benchmark that builds upon the best practices of SWE-Bench Zhang et al. (2025), but is explicitly designed to capture realistic, complex, enterprise-level problems beyond the scope of SWE-Bench. SWE-BENCH PRO contains 1,865 problems sourced from a diverse set of 41 actively maintained repositories spanning business applications, B2B services, and developer tools. The benchmark is partitioned into a *public* set with open access to problems sourced from 11 repositories, a *held-out* set of 12 repositories and a *commercial* set of 18 proprietary repositories where we have formal partnership agreements with early-stage startups. Problems in the held-out and the commercial set are not publicly accessible, but we release results on the commercial set. Our benchmark features long-horizon tasks that may require hours to days for a professional software engineer to complete, often involving patches across multiple files and substantial code modifications. All tasks are human-verified and augmented with sufficient context to ensure resolvability. In our evaluation of widely used coding models, under a unified scaffold, we observe that their performance on SWE-BENCH PRO remains below 25% (Pass@1), with GPT-5 achieving the highest score to date at 25.9%. To better understand these limitations, we cluster the failure modes observed in the collected agent trajectories for a clearer characterization of the error patterns exhibited by current models. Overall, SWE-BENCH PRO provides a contamination-resistant testbed that more faithfully captures the complexity and diversity of real-world software development, advancing the pursuit of truly autonomous software engineering agents at a professional level.

## 1 INTRODUCTION

Large Language Model (LLM) agents have been widely adopted in modern software development workflows. SWE-bench Jimenez et al. (2024) and related works Yang et al. (2024b); Zan et al. (2024); Yang et al. (2024a); Zhang et al. (2025); OpenAI (2024) establish the task of issue resolution as a de-facto standard for assessing their capability and usefulness. In this setting, an agent is given an entire codebase, a task description (e.g., a bug report or feature request) in natural language and is instructed to produce a code patch that resolves the issue and passes the repository’s test suite. These benchmarks have been instrumental in demonstrating both the substantial potential and the persistent limitations of current models as SWE agents.

Current coding benchmarks face several limitations. First, many benchmarks are susceptible to *contamination* Xu et al. (2024); Deng et al. (2024); Zhang et al. (2024a); White et al. (2024), as exemplified by recent works Xu et al. (2024); Deng et al. (2024); Cheng et al. (2025) and social media posts Aleithan et al. (2024); Zhang et al. (2025). This risk arises because widely used open-source repositories—particularly those distributed under permissive licenses (e.g., MIT, Apache 2.0, BSD)—are prime candidates for inclusion in the large-scale web-crawled corpora used to pre-train LLMs Brown et al. (2020). As a result, constructing benchmarks from public GitHub repositories is inherently difficult, since many are already accessible as training data. Second, existing tasks may *not* adequately capture the complexity of real-world software engineering. For example, SWE-Bench Verified Jimenez et al. (2024) includes a substantial proportion of relatively trivial problems (161 out of 500) that require only one- to two-line modifications. In contrast, industrial software engineering, particularly in enterprise settings, often demands multi-file modifications spanning hundreds of lines Hassan (2009); Steidl et al. (2017). This discrepancy raises concerns about whether current benchmarks truly reflect the challenges faced in practical development scenarios.

---

054 Our first contribution in SWE-BENCH PRO is a novel data collection strategy designed to **mitigate**  
055 **data contamination**. Specifically, our approach involves two complementary measures: (1) exclu-  
056 sively selecting repositories distributed under strong *copyleft* licenses (GPL) to construct a public set  
057 (11 repositories) and a held-out set (12 repositories), and (2) *acquiring commercial codebases* from  
058 real startups to capture enterprise-grade problems in a commercial set (18 repositories). In doing  
059 so, we reduce contamination risks by leveraging both legal protections and restricted data access.  
060 While analogous efforts may have been undertaken in industry using proprietary codebases, to the  
061 best of our knowledge, this work is the first to systematically apply such a methodology for curat-  
062 ing a benchmark in the research community. The three subsets are made available under different  
063 access policies. The public set provides both problems and evaluation results openly. The held-out  
064 set remains private, preserving it for future overfitting checks against the public set. Finally, for the  
065 commercial set, we release evaluation results while keeping the underlying codebases private.

066 The second contribution of SWE-BENCH PRO is its emphasis on **challenging, diverse, and indus-**  
067 **trially relevant** tasks. To ensure task complexity, we exclude trivial edits (1–10 lines of code) and  
068 retain only problems requiring substantial, multi-file modifications. On average, the reference solu-  
069 tions span 107.4 lines of code across 4.1 files. Every problem involves at least 10 lines of change,  
070 and over 100 tasks demand more than 100 lines of modification. In addition to complexity, we prior-  
071 itize diversity and representativeness. The curated repositories are all actively maintained and span a  
072 range of domains, including consumer applications, B2B services, and developer tooling platforms.  
073 Each repository contributes between 50 and 100 instances, with a strict cap of 100 instances, thereby  
074 reducing the risk of overfitting to any single repository.

075 The third contribution of SWE-BENCH PRO is to demonstrate a **human-centered augmentation**  
076 **and verification workflow** to ensure task resolvability. We design a novel three-stage human-in-  
077 the-loop process that serves dual purposes: (1) clarifying ambiguity and adding missing context to  
078 preserve core technical challenges, and (2) recovering unit tests as robust verifiers by constraining  
079 solution spaces to avoid false negatives while maintaining implementation flexibility.

080 Taken together, SWE-BENCH PRO aims to serve the community by providing a contamination-  
081 resistant and industrially realistic benchmark, supported by a transparent curation process and fine-  
082 grained diagnostic analyses. We release both the problems and evaluation results for the public  
083 set, retain the held-out set to monitor potential overfitting, and report results on the commercial set  
084 while preserving the privacy of its underlying codebases. Combined with standardized evaluation  
085 protocols and trajectory-level failure analyses, SWE-BENCH PRO offers a rigorous foundation for  
086 measuring progress beyond the saturation of SWE-Bench Verified, establishing a common yardstick  
087 for researchers and practitioners developing next-generation coding agents.

## 089 2 RELATED WORK

### 092 2.1 CODE AND SOFTWARE ENGINEERING BENCHMARKS

094 The evaluation of code generation capabilities has evolved from simple function-level tasks to com-  
095 plex repository-level challenges. Chen et al. (2021) introduced HumanEval, a foundational bench-  
096 mark of 164 handwritten programming problems that established the standard for measuring func-  
097 tional correctness in generated code. This was complemented by MBPP (Austin et al., 2021), which  
098 provided approximately 1,000 crowd-sourced Python problems designed for entry-level program-  
099 mers. For more challenging algorithmic tasks, APPS (Hendrycks et al., 2021) introduced 10,000  
100 programming problems spanning from simple to complex algorithmic challenges.

101 The field has since recognized the limitations of function-level evaluation. Jimenez et al. (2024)  
102 pioneered repository-level evaluation with SWE-bench, presenting 2,294 real GitHub issues from  
103 12 Python repositories that require understanding entire codebases to resolve. This revealed a sig-  
104 nificant performance gap, with state-of-the-art models resolving only the simplest issues. Building  
105 on this foundation, Zan et al. (2024) extended the approach to multiple programming languages with  
106 Multi-SWE-bench, covering Java, TypeScript, JavaScript, Go, Rust, C, and C++ with 1,632 expert-  
107 curated instances. Da et al. (2025) shows that these instances can be used for RL training as well as  
evaluation.

---

## 2.2 SOFTWARE ENGINEERING AGENTS

The development of autonomous agents capable of resolving real-world software engineering tasks has seen rapid progress. Yang et al. (2024a) introduced SWE-agent, emphasizing the critical importance of agent-computer interfaces (ACIs) in enabling effective code manipulation, achieving 12.5% resolution rate on SWE-bench. This work highlighted how interface design can be as important as model capabilities for agent performance. Zhang et al. (2024b) developed AutoCodeRover, which combines LLMs with sophisticated AST-based code search capabilities, achieving 19% on SWE-bench-lite while maintaining low operational costs.

## 3 DATASET OVERVIEW

### 3.1 CHARACTERISTICS OF SWE-BENCH PRO

**Industrially-Relevant, Diverse, and Challenging Tasks.** First, all repositories selected in SWE-BENCH PRO are actively maintained professional projects with substantial user bases, comprehensive documentation, and established development practices. In addition, we source commercial repositories. These repositories are private and sourced from startups, where we contacted the company and purchased their engineering repos. We sample repositories from a diverse range of topics, including consumer applications with complex UI logic, B2B platforms with intricate business rules, and developer tools with sophisticated APIs. Second, we limit each repository to contribute 50-100+ instances. This avoids the situation where models get an advantage by being especially good at a single repository, rewarding models that can truly generalize. Finally, we require edits to span multiple files and contain a substantial code change, similar to real software engineering tasks. Subsequently, SWE-BENCH PRO problems are naturally challenging – the best model performance is around 25%.

**Verified and Human-Augmented.** Similar to SWE-Bench Verified, each problem in SWE-BENCH PRO goes through a human augmentation and verification process. This ensures that task descriptions are not missing critical information, tests are well specified to validate the generated solution, and problems are representative of real-world software engineering tasks. In particular, we augment each issue with a list of human-written requirements – simulating the standard engineering practice of resolving issues follow problem specification and provide additional guarantee that the problems are self-contained. Note that real software engineering tasks can be under-specified (for example, may require exploration before solving), and that the setting *without requirements* is potentially interesting.

**Contamination-Resistant by Design.** By exclusively using repositories with GPL and other copy-left licenses, we ensure benchmark content is unlikely to appear in proprietary model training sets, as the nature of these licenses creates legal barriers to their inclusion in commercial training corpora. In addition, we use commercial repositories purchased from startups, which are private.

### 3.2 TASK SPECIFICATION

Each task instance in SWE-BENCH PRO is complete with human-augmented problem statement, requirements and interface as the task description for the model. The model must generate a patch file to resolve the issue and pass a suite of human-reviewed tests as validation.

**Problem Statement.** Similar to SWE-Bench, we provide a problem statement describing the issue to solve. We use content from the original commits, PR and issue, then rewrite it in the style of issues and add in missing information when necessary. Agents should be able to solve the task using only the problem statement.

**Requirements.** Problems in SWE-BENCH PRO can be more complex than previous iterations of SWE-Bench, and thus, we introduce requirements to resolve any potential ambiguity issues. For each problem, we list out a set of requirements that give additional detail on what is needed to solve the task. These requirements are grounded on the unit tests that are used for validation. For example, a requirement might specify the route names and functionality expected for an API.

**Interface.** SWE-BENCH PRO tasks include feature additions and code enhancements, such as refactoring, that involve creating or altering classes and methods. For these tasks, a common false negative in unit-test verifiers is when a model submits a valid solution with different interfaces than

162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215

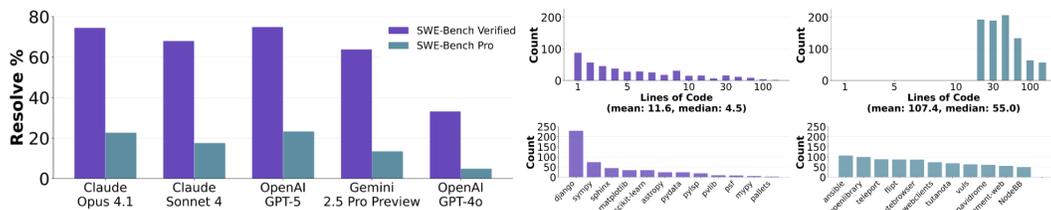


Figure 1: Frontier models, such as GPT-5 and Claude Opus 4.1, score 70% of SWE-Bench Verified but about 25% on SWE-BENCH PRO. Patches are larger, more challenging, and require expert knowledge about a variety of topics. Patches are generated with SWE-Agent Yang et al. (2024a) and evaluated on the public subset of SWE-BENCH PRO.

what the unit test is expecting. Here, we explicitly define the class and function names expected by the tests to avoid this failure mode when relevant.

**Environments.** Each task is evaluated in a containerized, language-specific environment with full dependency resolution. Python tasks use isolated virtual environments, JavaScript/TypeScript tasks use Node.js with npm/yarn, and Go tasks use module-aware environments with proper GOPATH configuration. All environments will be released as pre-built docker images to ensure that they are fully reproducible.

**Tests.** Every task includes human-reviewed test suites with `fail2pass` tests that verify issue resolution and `pass2pass` tests that ensure existing functionalities remain intact.

### 3.3 PUBLIC, COMMERCIAL, AND HELD-OUT SWE-BENCH PRO

SWE-BENCH PRO consists of a total of 1865 human-verified and augmented problems, divided as three subsets: public, commercial, and held-out.

- **Public.** We release 731 instances openly on HuggingFace and report the relevant statistics and model performances in this paper. These are sourced from public repositories with copy-left license.
- **Commercial.** For the commercial set of 276 problems sourced from startup repositories, we keep it private but report results publicly in this paper and will update in the leaderboard. This is the only set containing proprietary repositories from startups, which we cannot release for legal reasons.
- **Held-Out.** We hold out a set of 858 problems mirroring the public set but use a separate set of repositories. We keep this set private to test for overfitting in the future.

## 4 DATASET CREATION

Each problem from SWE-BENCH PRO consists of three components: a task description that prompts a SWE agent to resolve an issue, a set of relevant tests that verifies whether the issue has been resolved, and a working environment to run the codebase. To ensure a faithful and reliable evaluation, we manually verified and cleaned the test suite, and conduct human augmentation of the task description to include problem statement, requirements and interface that specify all the details necessary to pass the test suite.

### 4.1 SOURCING PROBLEMS

To collect the problems, we leverage the evolution of a codebase through its commit history. Specifically, we identify pairs of consecutive commits that together capture the resolution of an issue. In each pair, we refer to the older commit as the base and the newer commit as the instance. We define the test patch as the diff of test related files between the two commits. In other words, it consists of the new or modified tests introduced in the instance commit but absent in the base commit. The remaining diff, excluding the test patch, is referred to as the gold patch.

216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269

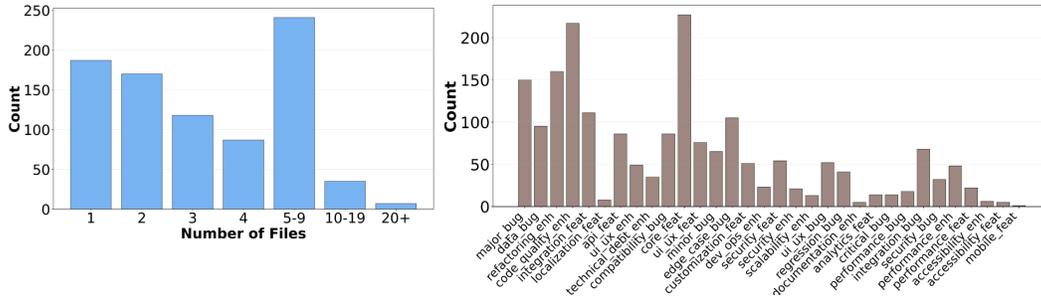


Figure 2: Distributions in the public set of SWE-BENCH PRO. SWE-BENCH PRO contains complex, long-horizon tasks involving several files and across a variety of task types. We include a diverse selection of feature requests as well as bug fixes, across optimization, security, UI/UX, and backend changes.

#### 4.2 CREATING TASK DESCRIPTIONS

SWE-BENCH PRO leverages human-driven augmentation, which makes it possible to construct problems beyond existing issues or PRs on Github. The goal of augmentation is to equip the SWE agent with sufficient context to resolve the issue without failing due to an underspecified task description. Although metadata are collected during commit scraping, commit messages are often unstructured, incomplete, or entirely missing. In practice, issue reproduction and problem solving typically requires extended communication among users, contributors, and codebase maintainers, often including screenshots, links, or other media. To address this gap, we collect and organize the available information from original sources, such as issue discussions, commit messages, or pull requests, and produce the final task description with two artifacts: (1) a problem statement, which captures the motivation for the change without extending beyond sources, and (2) a list of requirements and optionally interface, which provides the necessary details to fully understand and resolve the issue, grounded in the gold patch and test expectations when applicable. Importantly, the requirements specify the expected behavior but does not prescribe how the solution should be implemented.

#### 4.3 CREATING ENVIRONMENTS

We create environments through 3 steps: First, we construct environments manually with software engineering experts. Second, we use an in-house pipeline to validate that test are not flaky and that golden tests can pass the test suite successfully. Finally, we have a human-verification of all tests in the `fail2pass` test list, in which irrelevant tests are dropped.

**Environment construction.** We leveraged professional software engineers to create Docker-based environments. The engineers systematically incorporated system packages, repository documentation, build tools, and dependencies from each codebase into customized Dockerfiles and refined them until the resulting Docker images could successfully run the codebase and its tests. This process ensures that any agent can access the codebase and execute the tests out of the box.

**Environment verification.** We use automatic verification to ensure that the environment is working as expected. For each environment, we run the gold tests several times and ensure that they pass consistently. This ensures that the environment can be used properly, and also that there are not any flaky tests that may change run by run. We drop any problems that do not pass this criteria.

**Test verification.** We additionally send all tests through a human verification pipeline, where each tests is checked if it is relevant to the task description, and if it is not too broad. In either case, we drop tests that fall into either category: a) it is irrelevant to the task description, and b) it is too broad. In the case that all tests are too broad or not relevant, we drop the problem.

MODEL	RESOLVE (%)
OPENAI GPT-5 (HIGH)	25.9
OPENAI GPT-5 (MEDIUM)	23.3
CLAUDE OPUS 4.1	22.7
CLAUDE SONNET 4	17.6
OPENAI GPT-OSS 20B	16.2
GEMINI 2.5 PRO PREVIEW	13.5
SWE-SMITH-32B	6.8
OPENAI GPT-4O	4.9
QWEN-3 32B	3.4

Table 1: Model performance on the public set of SWE-BENCH PRO (N=731). Models are evaluated using SWE-Agent Yang et al. (2024a), without any ambiguity (e.g. we provide the augmented problem statement, requirements, interface).

MODEL	RESOLVE (%)
CLAUDE OPUS 4.1	17.8
OPENAI GPT-5 (HIGH)	15.7
OPENAI GPT-5 (MEDIUM)	14.9
GEMINI 2.5 PRO PREVIEW	10.1
CLAUDE SONNET 4	9.1
OPENAI GPT-4O	3.6

Table 2: Model performance on the commercial set of SWE-BENCH PRO (N=276). Commercial problems are sourced from startup repositories, where each problem is augmented with an environment and relevant information.

## 5 RESULTS

We present the results on SWE-BENCH PRO. Below, we detail the evaluation criteria, scaffold, and settings for reproducibility. We evaluate a suite of models, including frontier models, open-weight models, and models fine-tuned on SWE-bench-like trajectories (e.g. SWE-Smith).

**Scaffold.** We use the SWE-Agent Yang et al. (2024a) scaffold. We also explore another popular scaffold, Agentless Xia et al. (2024). However, we find that Agentless has difficulty in multi-file editing, thus, produces low evaluation scores. We focus on SWE-Agent for our results.

**Evaluation settings.** All models use the latest versions as of September 18th, 2025. For open-source LLMs, we use vllm to host each model. Models are hosted on a single node, with 8 H100 Nvidia GPUs. We enable tool-use when possible, for open-weight models, we use syntax parsing to enable tool-use. Models have a maximum of 50 turns. We use the same prompt for all models, which is the default prompt from Yang et al. (2024a). We use Opus 4.1 without extended thinking as reported by Anthropic in their own evaluation on SWE-Bench Verified.

**Issue Ambiguity.** Models are evaluated in the setting without any ambiguity – that is, we include the problem statement, requirements and interface specification in the agent prompt. Here, models are evaluated on their ability to implement a given repair or patch after being given significant details (rather than their ability to resolve ambiguity).

**Evaluation sets.** Evaluations are done on the public set and commercial set. For all analysis, we use the public set to avoid potential leakage with the commercial set. Finally, we keep the private set held-out for future analysis.

**Results.** Table 1 shows the results of various models on SWE-BENCH PRO. We report Pass@1 as the resolve rate. OPENAI GPT-5 and CLAUDE OPUS 4.1 achieve the highest resolve rates at 25.9% and 22.7% respectively, substantially outperforming smaller models. CLAUDE 4 SONNET also achieves a 16.3% resolve rate, while earlier generation models like DeepSeek Qwen-3 32B and OpenAI GPT-4o show considerably lower performance at 3.4% and 3.9% respectively. There is also a significant performance gap between the public and commercial set, where the best models score less than 20% in the commercial set, highlighting the difficulty of navigating enterprise codebases.

## 6 ANALYSIS

In this section, we provide additional analysis for model performance on SWE-BENCH PRO. We include analysis of performance on different types of issues, failure modes of agent trajectories for different models, and the efficacy of human augmentations for improving task resolvability.

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

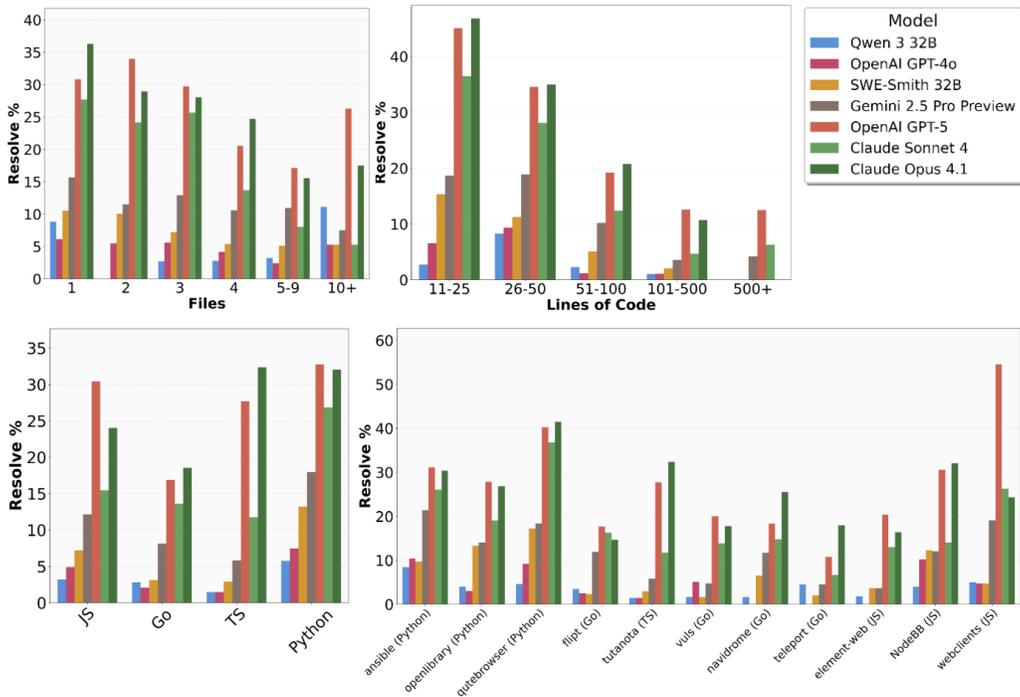


Figure 3: Model performance varies across languages, and models current perform better at Python. Resolve rates across different repos in the public set of SWE-BENCH PRO. SWE-BENCH PRO includes a variety of repos across different languages, with a similar number of problems per repo. Note that some categories, especially 10+ files and 500+ LOC, contain about 20-30 examples and thus have higher variance.

MODEL	PROBLEM STATEMENT, REQUIREMENTS, INTERFACE	PROBLEM STATEMENT ONLY
OPENAI GPT-5 (HIGH)	25.9%	8.40%
CLAUDE OPUS 4.1	22.7%	8.20%

Table 3: Comparison of model performance with and without human augmentations. The setting PROBLEM STATEMENT ONLY is without human augmentation. Without these augmentations, unit test verifiers are susceptible to false negatives.

## 6.1 ANALYSIS ON MODEL PERFORMANCE ON SWE-BENCH PRO

**Difficulty varies across programming languages.** As shown in Figure 3 (left), resolve rates differ markedly across programming languages. Go and Python generally show higher resolve rates across most models, with some models achieving resolve rates above 30% in these languages. JavaScript (JS) and TypeScript (TS) present more variable performance, with resolve rates ranging from near 0% to over 30% depending on the model.

**Resolve rate varies across repositories.** Figure 3 (right) demonstrates that resolve rates also vary considerably among different repositories in SWE-BENCH PRO. Some repositories show consistently low resolve rates across all models (below 10%), while others allow certain models to achieve resolve rates exceeding 50%. This suggests that repository-specific factors such as codebase complexity, documentation quality, or problem types significantly impact model performance.

**Model behavior correlates with task complexity.** The relationship between file count and resolve rate (Figure 3, top left) reveals distinct performance degradation patterns. Models maintain relatively stable performance for single-file problems but exhibit sharp declines as file count increases. No-

---

ably, the performance gap between frontier and smaller models widens dramatically beyond 3 files, with Claude Opus 4.1 and OpenAI GPT-5 maintaining above 10% resolve rates even for problems involving 10+ files, while open-source alternatives approach near-zero performance. This suggests that handling multi-file contexts requires capacity that current open-source models lack.

**Frontier models show more consistent cross-domain performance.** Claude Opus 4.1 and OpenAI GPT-5 maintain relatively high performance across most repositories and languages compared to smaller models, which show more erratic performance patterns that yield near-zero resolve rates on certain repositories.

## 6.2 ABLATION: REMOVING HUMAN AUGMENTATIONS

We perform an ablation in the importance of augmentations in SWE-BENCH PRO, namely the requirements and interface. These augmentations provide the agent with enough information to mitigate false negatives from unit tests that expect specific APIs, signatures, or functional behavior. Table 3 shows results in two settings: our default setting (where we include the problem statement, requirements, and interface), and another setting where we include only the problem statement. Without the requirements and interface, both models tested (GPT-5 and CLAUDE OPUS 4.1) show significantly degraded performance. In this setting, agents are less constrained and can submit more diverse solutions (particularly for feature additions). Since unit tests expect a narrow set of solutions, verifiers are prone to false negatives, resulting in lower pass rates. Our human augmentations mitigate these false negatives by constraining the agent’s solution space such that verifiers are robust, providing appropriate context for measuring task resolution.

## 6.3 TRAJECTORY FAILURE MODES

We conduct an LLM-as-a-judge analysis for failure modes of different models, utilizing GPT-5 as the judge. Our work follows Yang et al. (2024a), who demonstrate 87% alignment of automated judgments with human categorization of failure modes.

**Method.** We begin by hand-curating buckets for common failure patterns of agents in software engineering tasks, as determined by heuristics and a random sample of agent trajectories. These buckets are shown in Table 4. For each of the models in Table 4, we programmatically filter to only unresolved instances of SWE-BENCH PRO and collect the last 20 turns of each rollout. We determined 20 turns to have the highest correspondence with human validations of failure mode compared to 10 turns and 40 turns. With a system prompt providing strict descriptions of the failure buckets and overall SWE-Agent format, we feed the trajectory input and prompt the GPT-5 judge to first produce a 1-paragraph reasoning and then an ultimate selection of one failure mode per instance.

**Categories.** Here, we detail several of the common categories of failure modes. For the full category descriptions, view the Appendix. **Wrong solution.** The agent produces a syntactically valid patch that is functionally incorrect, incomplete, or fails to address the core problem. **Tool-Use.** Failure is attributed to the agent’s incorrect use of its available tools. This misuse prevents the agent from gathering necessary information or applying changes correctly. **Syntax error.** The agent successfully modifies the target files but introduces syntactic errors that render the codebase uncompileable or unrunnable. **Incorrect file.** This failure occurs when the agent correctly understands the high-level goal but fails to locate the correct source file or function for modification.

**Results.** Table 4 shows the results. Frontier models fail on SWE-BENCH PRO for several reasons. OPUS 4.1 primarily fails on semantic understanding, with wrong solutions accounting for 35.9% of failures and syntax errors at 24.2%, suggesting strong technical execution but challenges in problem comprehension and algorithmic correctness. GPT-5 indicates potential differences in effective-tool-use, but fewer wrong solutions. Other models reveal distinct operational challenges. SONNET 4 has context overflow as its primary failure mode (35.6%) and substantial endless file reading behaviors (17.0%), suggesting limitations in context management and file navigation strategies. GEMINI 2.5 demonstrates more balanced failures across tool errors (38.8%), syntax errors (30.5%), and wrong solutions (18.0%), maintaining competence across multiple dimensions. QWEN3 32B, as an open-source model, exhibits the highest tool error rate (42.0%) which highlights the importance of integrated tool-use for effective agents.

Model	Overall		Submitted						Not-Submitted		
	Submitted	Not-Submitted	Wrong Solution	Syntax Error	Incorrect File	Instruction Following	Edge Case	Other	Tool-Use	Long-Context	Stuck in Loop
	CLAUDE OPUS 4.1	74.2% (511)	25.8% (178)	50.3% (257)	31.3% (160)	4.9% (25)	2.7% (14)	0.8% (4)	10.0% (51)	68.0% (121)	28.7% (51)
GPT-5 (HIGH)	27.2% (147)	72.8% (394)	39.5% (58)	29.3% (43)	8.8% (13)	4.8% (7)	0.0% (0)	17.7% (26)	96.4% (380)	2.5% (10)	1.0% (4)
CLAUDE SONNET 4	44.1% (235)	55.9% (298)	25.5% (60)	7.7% (18)	2.1% (5)	2.1% (5)	0.0% (0)	62.6% (147)	8.7% (26)	57.4% (171)	33.9% (101)
GEMINI 2.5 PRO PREVIEW	52.9% (377)	47.1% (335)	35.0% (132)	56.5% (213)	4.0% (15)	1.9% (7)	0.0% (0)	2.7% (10)	83.6% (280)	14.3% (48)	2.1% (7)
GPT-4o	72.1% (569)	27.9% (220)	45.2% (257)	36.7% (209)	11.2% (64)	6.2% (35)	0.0% (0)	0.7% (4)	100.0% (220)	0.0% (0)	0.0% (0)
QWEN3 32B	47.3% (304)	52.7% (339)	25.0% (76)	48.7% (148)	19.7% (60)	1.6% (5)	0.7% (2)	4.3% (13)	78.8% (267)	1.5% (5)	19.8% (67)

Table 4: Failure mode analysis for models on SWE-BENCH PRO public set. We use LLM-as-a-judge to classify failing trajectories into buckets. Top LLMs, such as Opus 4.1 and GPT-5, are strong agents but struggle to produce solutions on high-complexity tasks. Weaker models, such as smaller open-source models, struggle with syntax, formatting, and tool-use.

## 7 LIMITATIONS AND FUTURE WORK

### 7.1 LIMITATIONS

**Limited Language Coverage.** Although SWE-BENCH PRO includes multiple programming languages (Python, JavaScript, TypeScript, Go), the distribution is not uniform, and some widely-used languages like Java, C++, and Rust are underrepresented. This may limit the benchmark’s ability to assess agent performance across the full spectrum of modern software development.

**Dependency on Test Suite.** We rely on a test suite of `fail2pass` and `pass2pass` to verify problem solutions. However, real software engineering tasks may have a variety of correct solutions, even if they do not pass the original tests outlined in the task. Ideally, we might have a set of verifiers which can verify any valid solution.

### 7.2 FUTURE WORK

**Alternative Evaluation Metrics.** Developing evaluation approaches beyond test-based verification, such as rubrics, code quality assessment, security analysis, performance optimization, and adherence to software engineering best practices. This could include human evaluation of code maintainability, readability, and architectural soundness.

**Collaborative Development Scenarios.** Introducing problems that require coordination between multiple agents or human-agent collaboration, reflecting modern team-based software development practices. This could include scenarios involving code reviews, merge conflict resolution, and distributed development workflows.

## 8 CONCLUSION

In conclusion, our introduction of SWE-BENCH PRO marks a significant step forward in the rigorous and realistic evaluation of AI coding agents. By adhering to three core principles—diverse, real-world task selection; challenging, multi-file code changes; and strict contamination prevention—we have created a benchmark that more accurately reflects the complexity of professional software engineering. Our findings, which show top-tier models like Opus 4.1 and GPT-5 achieving a 23% success rate on SWE-BENCH PRO compared to over 70% on benchmarks like SWE-Bench Verified, highlight a critical gap between current agent capabilities and the demands of real-world development. This new baseline not only provides a more accurate measure of progress but also offers crucial insights into the specific limitations that must be addressed to advance the field. SWE-BENCH PRO serves as a robust, contamination-resistant testbed that can help guide future research toward developing truly autonomous and capable software engineering agents.

486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

---

## REFERENCES

- Reem Aleithan et al. Swe-bench+: Enhanced coding benchmark for llms. *arXiv preprint arXiv:2410.06992*, 2024.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde, Jared Kaplan, Harrison Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Yuxing Cheng, Zhenghao Li, and Yanzhao Zhou. A survey on data contamination for large language models. *arXiv preprint arXiv:2502.14425*, 2025.
- Jeff Da, Clinton J. Wang, Xiang Deng, Yuntao Ma, Nikhil Barhate, and Sean M. Hendryx. Agent-rlvr: Training software engineering agents via guidance and environment rewards. *ArXiv*, abs/2506.11425, 2025. URL <https://api.semanticscholar.org/CorpusID:279391657>.
- Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein, and Arman Cohan. Investigating data contamination in modern benchmarks for large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 8706–8719, Mexico City, Mexico, 2024. Association for Computational Linguistics.
- Ahmed E Hassan. Predicting faults using the complexity of code changes. In *2009 IEEE 31st International Conference on Software Engineering*, pp. 78–88. IEEE, 2009.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. Measuring coding challenge competence with apps. In *Neural Information Processing Systems*, 2021.
- 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 *The Twelfth International Conference on Learning Representations*, 2024.
- OpenAI, 2024. URL <https://openai.com/index/introducing-swe-bench-verified/>.
- Daniela Steidl, Benjamin Hummel, and Elmar Jürgens. Evaluating code complexity triggers, use of complexity measures and the influence of code complexity on maintenance time. *Empirical Software Engineering*, 22(2):971–1015, 2017.
- Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddhartha Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, Micah Goldblum, Abacus.AI, Nyu, and Nvidia. Livebench: A challenging, contamination-free llm benchmark. *ArXiv*, abs/2406.19314, 2024. URL <https://api.semanticscholar.org/CorpusID:270556394>.
- Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Agentless: Demystifying llm-based software engineering agents. *arXiv preprint arXiv:2407.01489*, 2024.
- Cheng Xu, Jiu Hai Guan, Xu Zhao, Chenyi Fu, Qiushi Xin, Zihan Wang, Libo Li, Jin Fu, Hao Wang, and Jun Liu. Benchmark data contamination of large language models: A survey. *arXiv preprint arXiv:2406.04244*, 2024.
- 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 *Neural Information Processing Systems*, 2024a.

---

540 John Yang, Carlos E Jimenez, Alexander Wettig, Karthik Narasimhan, and Ofir Press. Swe-  
541 bench multimodal: Do ai systems generalize to visual software domains? *arXiv preprint*  
542 *arXiv:2410.03859*, 2024b.

543  
544 Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Linhao Zhang, Shulin Xin, Lu Chen, Qi Liu,  
545 Xiaojian Zhong, Aoyan Li, et al. Multi-swe-bench: A multilingual benchmark for issue resolving.  
546 *arXiv preprint arXiv:2404.02605*, 2024.

547  
548 Chaoyun Zhang et al. Swe-bench goes live! *arXiv preprint arXiv:2505.23419*, 2025.

549  
550 Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao,  
551 Pranav Raja, Dylan Slack, Qin Lyu, Sean M. Hendryx, Russell Kaplan, Michele Lunati, and  
552 Summer Yue. A careful examination of large language model performance on grade school  
553 arithmetic. *ArXiv*, abs/2405.00332, 2024a. URL [https://api.semanticscholar.org/  
CorpusID:269484687](https://api.semanticscholar.org/CorpusID:269484687).

554  
555 Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. Autocoderover: Autonomous  
556 program improvement. In *ACM SIGSOFT International Symposium on Software Testing and*  
557 *Analysis*, 2024b.

558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593

594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647

---

## APPENDIX

In the appendix, we include more details regarding example instances of the dataset.

### A FAILURE MODE CATEGORY DESCRIPTIONS

- **Wrong solution.** The agent produces a syntactically valid patch that is functionally incorrect, incomplete, or fails to address the core problem.
- **Tool-Use.** Failure is attributed to the agent’s incorrect use of its available tools. This misuse prevents the agent from gathering necessary information or applying changes correctly.
- **Syntax error.** The agent successfully modifies the target files but introduces syntactic errors that render the codebase uncompileable or unrunnable.
- **Incorrect file.** This failure occurs when the agent correctly understands the high-level goal but fails to locate the correct source file or function for modification.

**Endless File Reading:** The agent enters a non-productive loop of exploratory actions, such as repeatedly reading the same files, searching for keywords, and viewing code snippets, without ever progressing to an implementation phase. It successfully gathers information but fails to synthesize it into a concrete code modification, eventually timing out or failing due to inaction.

**Misunderstood Problem Statement:** This category describes failures where the agent fundamentally misinterprets the task’s objective. Instead of implementing the required code changes, it pursues a tangential or incorrect goal, such as focusing on creating a complex runtime reproduction environment for a simple refactoring task. The agent’s actions are coherent with its flawed understanding but do not address the actual issue.

**Other:** A catch-all category for failures that do not fit into the more specific classifications above. This often includes trajectories where the agent exceeds computational or time limits (e.g., cost limits) due to an inefficient workflow, such as running an excessive number of verbose tests, rather than a single, clear technical mistake. It can also encompass a combination of minor, compounding issues.

### B EXAMPLE TASK INSTANCE

This section includes an example instance of SWE-BENCH PRO with descriptions of each key field.

#### B.1 PROBLEM STATEMENT

The problem statement describes the task that the agent needs to complete in the codebase. The structure of the problem statement is similar to a Github Issue, and includes the same markdown formatting and conventions found in common open-source repositories.

When creating problem statements, effort is made to keep the problem statements as close as possible to the real-world distribution, such as ensuring every problem statement uses the same default issue templates that are used in the repository for a specific task.

Problem statements are curated from existing commits, issues, and PRs in codebases, and are rewritten to be well-specified, as shown in Table 5

##### B.1.1 EXAMPLE

This example is a feature request for Open Library, an open source non-profit project run by the Internet Archive with the goal of creating a web page for every book published. As a real-world full-stack web application, Open Library is representative of the kind of repositories SWE-BENCH PRO includes to maximize environment realism.

```
### Add Google Books as a metadata source to BookWorm for fallback/staging imports
```

```
### Problem / Opportunity
```

---

648 BookWorm currently relies on Amazon and ISBNdb as its primary sources for metadata. This  
649 presents a problem when metadata is missing, malformed, or incomplete particularly for books with  
650 only ISBN-13s. As a result, incomplete records submitted via promise items or `'/api/import'` may  
651 fail to be enriched, leaving poor-quality entries in Open Library. This limitation impacts data quality  
652 and the success rate of imports for users, especially for less common or international titles.

653 **### Justify: Why should we work on this and what is the measurable impact?**

654 Integrating Google Books as a fallback metadata source increases Open Library's ability to sup-  
655 plement and stage richer edition data. This improves the completeness of imported books, reduces  
656 failed imports due to sparse metadata, and enhances user trust in the import experience. The impact  
657 is measurable through increased import success rates and reduced frequency of placeholder entries  
658 like "Book 978...".

659 **### Define Success: How will we know when the problem is solved?**

660 - BookWorm is able to fetch and stage metadata from Google Books using ISBN-13. - Automated  
661 tests confirm accurate parsing of varied Google Books responses, including: - Correct mapping of  
662 available fields (title, subtitle, authors, publisher, page count, description, publish date). - Proper  
663 handling of missing or incomplete fields (e.g., no authors, no ISBN-13). - Returning no result when  
664 Google Books returns zero or multiple matches.

665 **### Proposal**

666 Introduce support for Google Books as a fallback metadata provider in BookWorm. When an Ama-  
667 zon lookup fails or only an ISBN-13 is available, BookWorm should attempt to fetch metadata from  
668 the Google Books API and stage it for import. This includes updating source logic, metadata pars-  
669 ing, and ensuring records from `'google_books'` are correctly processed.

## 673 B.2 REQUIREMENTS

674 The requirements section includes a list of human-authored requirements that provide additional in-  
675 formation that the agent needs in order to create a valid solution that is verifiable by the unit tests.  
676 Requirements often specify expected behavior by the implemented solution that will be explicitly  
677 tested for. For example, if a unit test asserts for the presence of a specific error log string, a re-  
678 quirement is written to specify that the solution should produce the exact same error log string.  
679 Requirements never include specific code implementation and don't leak solutions.

### 682 B.2.1 EXAMPLE

683 This example includes the requirements that the agent must consider when implementing the feature  
684 addition to Open Library. It includes requirements for the expected behavior of the implemented  
685 solution, as well as specific details that the agent wouldn't otherwise have knowledge of (such as the  
686 URL to stage bookworm data). - The tuple `'STAGED_SOURCES'` in `'openlibrary/core/imports.py'`  
687 must include `"google_books"` as a valid source, so that staged metadata from Google Books is  
688 recognized and processed by the import pipeline.

689 - The URL to stage bookworm metadata is `"http://{affiliate_server_url}/isbn/{identifier}?high_priority=true&stage_import=true"`  
690 where the `affiliate_server_url` is the one from the `openlibrary/core/vendors.py`, and the param `identi-`  
691 `fier` can be either ISBN 10, ISBN 13, or B\*ASIN.

692 - When supplementing a record in `'openlibrary/plugins/importapi/code.py'` using `'supple-`  
693 `ment_rec_with_import_item_metadata'`, if the `'source_records'` field exists, new identifiers must be  
694 added (extended) rather than replacing existing values.

695 - In `'scripts/affiliate_server.py'`, a function named `'stage_from_google_books'` must attempt to fetch  
696 and stage metadata for a given ISBN using the Google Books API, and if successful, persist the  
697 metadata by adding it to the corresponding batch using `'Batch.add_items'`.

698 - The affiliate server handler in `'scripts/affiliate_server.py'` must fall back to Google Books for  
699 ISBN-13 identifiers that return no result from Amazon, but only if both the query parameters  
700 `'high_priority=true'` and `'stage_import=true'` are set in the request.  
701

- 
- If Google Books returns more than one result for a single ISBN query, the logic must log a warning message and skip staging the metadata to avoid introducing unreliable data.
  - The metadata fields parsed and staged from a Google Books response must include at minimum: 'isbn\_10', 'isbn\_13', 'title', 'subtitle', 'authors', 'source\_records', 'publishers', 'publish\_date', 'number\_of\_pages', and 'description', and must match the data structure expected by Open Library import system.
  - In 'scripts/promise\_batch\_imports.py', staging logic must be updated so that, when enriching incomplete records, 'stage\_bookworm\_metadata' is used instead of any previous direct Amazon-only logic.

### B.3 INTERFACE

The interface is an optional field that is only used when the task solution requires modifying or creating new public interfaces. It includes the interfaces for all classes and functions that have been modified or created, including their signatures, and their file path.

The interface plays an important role in mitigating false negatives for unit test verification. This is particularly relevant for code changes related to feature additions. When a new feature is added, the associated unit tests are written to a specific set of interfaces that the newly added classes and functions expose. Since SWE-BENCH PRO uses unit tests without modification, the interface helps the agent avoid the failure mode where it implements a viable solution, but uses a class name or module path that the unit test is not expecting.

#### B.3.1 EXAMPLE

This example includes all the public interfaces that were modified or created in the golden patch that added the new feature in Open Library. These interfaces are coupled to the associated unit tests implemented in the test patch for this commit. Function: `fetch_google_book` Location: `scripts/affiliate_server.py` Inputs: `isbn (str)` ISBN-13 Outputs: dict containing raw JSON response from Google Books API if HTTP 200, otherwise None Description: Fetches metadata from the Google Books API for the given ISBN.

Function: `process_google_book` Location: `scripts/affiliate_server.py` Inputs: `google_book_data (dict)` JSON data returned from Google Books Outputs: dict with normalized Open Library edition fields if successful, otherwise None Description: Processes Google Books API data into a normalized Open Library edition record.

Function: `stage_from_google_books` Location: `scripts/affiliate_server.py` Inputs: `isbn (str)` ISBN-10 or ISBN-13 Outputs: bool True if metadata was successfully staged, otherwise False Description: Fetches and stages metadata from Google Books for the given ISBN and adds it to the import batch if found.

Function: `get_current_batch` Location: `scripts/affiliate_server.py` Inputs: `name (str)` batch name such as "amz" or "google" Outputs: Batch instance corresponding to the provided name Description: Retrieves or creates a batch object for staging import items.

Class: `BaseLookupWorker` Location: `scripts/affiliate_server.py` Description: Base threading class for API lookup workers. Processes items from a queue using a provided function. Method: `BaseLookupWorker.run(self)` Location: `scripts/affiliate_server.py` Description: Public method to process items from the queue in a loop, invoking the `process_item` callable for each item retrieved.

Class: `AmazonLookupWorker` Location: `scripts/affiliate_server.py` Description: Threaded worker that batches and processes Amazon API lookups, extending `BaseLookupWorker`. Method: `AmazonLookupWorker.run(self)` Location: `scripts/affiliate_server.py` Description: Public method override that batches up to 10 Amazon identifiers from the queue, processes them together using the Amazon batch handler, and manages timing according to API constraints.

756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

Table 5: Problem Statement Comparison: Original vs. Rewritten

Original Commit Message	Human Authored Issue
<p>enable vCard v4.0 contact import (close #1328)</p> <p>No description provided.</p>	<p><b>Title:</b> Unable to import contacts encoded as vCard 4.0</p> <p><b>Description:</b> The application’s contact importer recognises vCard 2.1 and 3.0, but any file that starts with <code>VERSION:4.0</code> is treated as an unsupported format. The import either fails outright (returns <code>null</code>) or produces an empty contact, preventing users from migrating address books exported by modern clients that default to vCard 4.0.</p> <p><b>Impact:</b></p> <ul style="list-style-type: none"> <li>• Users cannot migrate their contact lists from current ecosystems (e.g. iOS, macOS, Google Contacts).</li> <li>• Manual conversion or data loss is required, undermining interoperability.</li> <li>• Breaks the expectation that the app can import the latest vCard standard.</li> </ul> <p><b>Steps to Reproduce:</b></p> <ol style="list-style-type: none"> <li>1. Export a contact as a vCard 4.0 file from a standards-compliant source (e.g. iOS Contacts).</li> <li>2. In the application UI, choose <b>Import contacts</b> and select the <code>.vcf</code> file.</li> <li>3. Observe that no contact is created or that the importer reports an error.</li> </ol> <p><b>Expected Behaviour:</b></p> <ul style="list-style-type: none"> <li>• The importer should recognise the <code>VERSION:4.0</code> header and process the file.</li> <li>• Standard fields present in earlier versions (FN, N, TEL, EMAIL, ADR, NOTE, etc.) must be mapped to the internal contact model as they are for vCard 2.1/3.0.</li> <li>• Unsupported or unknown properties must be ignored gracefully without aborting the import.</li> </ul> <p><b>Additional Context:</b></p> <ul style="list-style-type: none"> <li>• Specification: RFC 6350 vCard 4.0</li> <li>• Minimal sample input that currently fails:</li> </ul>

---

## 810 C TRAJECTORY FAILURE MODE ANALYSIS

811

812

### 813 C.1 LLM-AS-A-JUDGE PROMPT

814

815

816 You are an expert software engineer analyzing why a software engineering agent failed to resolve  
817 an issue.

818 INSTANCE ID: {instance\_id} {exit\_status\_desc}

819

820 AVAILABLE AGENT ACTIONS:

821

822 — BEGIN FUNCTION #1: bash — Description: Execute a bash command in the terminal. \* Can  
823 generate very large outputs when listing files (ls, find, grep) \* Output contributes directly to context  
824 window usage \* Commands like 'find /repo -name "\*.py"' can list thousands of files \* Large outputs  
can quickly fill the context window

825

826 Parameters: (1) command (string, required): The bash command to execute. Can be empty to view  
827 additional logs when previous exit code is '-1'. Can be 'ctrl+c' to interrupt the currently running  
process. — END FUNCTION #1 —

828

829 — BEGIN FUNCTION #2: submit — Description: Finish the interaction when the task is com-  
830 plete OR if the assistant cannot proceed further with the task. \* Used when agent thinks task is done  
(may be correct or incorrect solution) \* Also used when agent is stuck and cannot make progress \*  
831 No parameters are required for this function. — END FUNCTION #2 —

832

833 — BEGIN FUNCTION #3: str\_replace\_editor — Description: Custom editing tool for viewing,  
834 creating and editing files \* State is persistent across command calls and discussions with the user \*  
835 If 'path' is a file, 'view' displays the result of applying 'cat -n'. If 'path' is a directory, 'view' lists  
836 non-hidden files and directories up to 2 levels deep \* Directory views can generate large outputs  
837 contributing to context usage \* The 'create' command cannot be used if the specified 'path' already  
838 exists as a file \* If a 'command' generates a long output, it will be truncated and marked with  
'response clipped' \* The 'undo\_edit' command will revert the last edit made to the file at 'path'

839

840 Notes for using the 'str\_replace' command: \* The 'old\_str' parameter should match EXACTLY  
841 one or more consecutive lines from the original file. Be mindful of whitespaces! \* If the 'old\_str'  
842 parameter is not unique in the file, the replacement will not be performed. Make sure to include  
843 enough context in 'old\_str' to make it unique \* The 'new\_str' parameter should contain the edited  
lines that should replace the 'old\_str'

844

845 Parameters: (1) command (string, required): The commands to run. Allowed options are: 'view',  
846 'create', 'str\_replace', 'insert', 'undo\_edit'. (2) path (string, required): Absolute path to file or di-  
847 rectory, e.g. '/repo/file.py' or '/repo'. (3) file\_text (string, optional): Required parameter of 'create'  
848 command, with the content of the file to be created. (4) old\_str (string, optional): Required parame-  
849 ter of 'str\_replace' command containing the string in 'path' to replace. (5) new\_str (string, optional):  
850 Optional parameter of 'str\_replace' command containing the new string (if not given, no string will  
851 be added). Required parameter of 'insert' command containing the string to insert. (6) insert\_line  
852 (integer, optional): Required parameter of 'insert' command. The 'new\_str' will be inserted AFTER  
853 the line 'insert\_line' of 'path'. (7) view\_range (array, optional): Optional parameter of 'view' com-  
854 mand when 'path' points to a file. If none is given, the full file is shown. If provided, the file will  
855 be shown in the indicated line number range, e.g. [11, 12] will show lines 11 and 12. Indexing at  
856 1 to start. Setting '[start\_line, -1]' shows all lines from 'start\_line' to the end of the file. — END  
FUNCTION #3 —

857

858 — BEGIN FUNCTION #4: file\_viewer — Description: Interactive file viewer for opening and  
859 navigating files in the editor. \* open ;path; [;line\_number;]: Opens the file at path. If line\_number  
860 is provided, the view moves to include that line. \* goto ;line\_number;: Moves the window to show  
861 the specified line number. \* scroll\_down: Moves the window down 100 lines. \* scroll\_up: Moves  
the window up 100 lines.

862

863 Parameters: (1) command (string, required): One of 'open', 'goto', 'scroll\_down', 'scroll\_up'. (2)  
path\_or\_line (string/int, optional): For 'open', a path (and optional line). For 'goto', a line number.  
— END FUNCTION #4 —

---

864 — BEGIN FUNCTION #5: search\_tools — Description: Searching utilities for locating text or  
865 files within the workspace. \* search\_file ;search\_term; [;file;]: Searches for search\_term in file. If  
866 file is not provided, searches the current open file. \* search\_dir ;search\_term; [;dir;]: Searches for  
867 search\_term in all files in dir. If dir is not provided, searches in the current directory. \* find\_file  
868 ;file\_name; [;dir;]: Finds all files with the given name in dir. If dir is not provided, searches in the  
869 current directory.

870 Parameters: (1) subcommand (string, required): One of ‘search\_file’, ‘search\_dir’, ‘find\_file’. (2)  
871 arg1 (string, required): The search term or file name, depending on subcommand. (3) arg2 (string,  
872 optional): Target file (for search\_file) or directory (for search\_dir/find\_file). — END FUNCTION  
873 #5 —

874 — BEGIN FUNCTION #6: edit\_block — Description: Block editor for replacing ranges in the  
875 current open file and finalizing edits. \* edit ;n; ;m; ;replacement\_text;: Replaces lines n through m  
876 (inclusive) with the given text in the open file. Ensure indentation is correct. \* end\_of\_edit: Applies  
877 the pending changes. Python files are syntax-checked after the edit; if an error is found, the edit is  
878 rejected.

879 Parameters: (1) command (string, required): ‘edit’ or ‘end\_of\_edit’. (2) range\_and\_text (varies): For  
880 ‘edit’, a line range ‘n:m’ and the replacement text. — END FUNCTION #6 —

881 — BEGIN FUNCTION #7: create\_file — Description: Creates and opens a new file with the given  
882 name.

883 Parameters: (1) filename (string, required): Absolute or workspace-relative path to create. The file  
884 must not already exist. — END FUNCTION #7 —

885 PROBLEM STATEMENT: {problem\_statement}

886 FINAL ACTIONS TAKEN (Last {NUM\_PAST\_ACTIONS}): {chr(10).join(final\_actions[-  
887 NUM\_PAST\_ACTIONS:]) if final\_actions else "No actions recorded"}

888 FINAL OBSERVATIONS (Last {NUM\_PAST\_ACTIONS}): {chr(10).join(final\_observations[-  
889 NUM\_PAST\_ACTIONS:]) if final\_observations else "No observations recorded"}

890 TRAJECTORY SUMMARY: - Total steps: {len(trajectory\_steps)} - Final state: Failed (no success-  
891 ful patch generated)

892 ANALYSIS INSTRUCTIONS: The exit status indicates WHY the agent terminated. Consider how  
893 the final actions contributed to this specific exit condition.

894 Based on the information above, provide an error analysis in two parts: First, an explanation of the  
895 issue and why the trajectory failed. Second, a category for the error.

896 Wrap your explanation in ;description;/description tags.

897 For the category, choose EXACTLY one from the following set: identified\_incorrect\_file: The agent  
898 incorrectly identified the file that needed to be fixed., missed\_edge\_case: The agent missed an edge  
899 case in one of the test cases., misunderstood\_problem\_statement: The agent misunderstood the prob-  
900 lem statement., wrong\_solution: The agent generated a wrong solution., tool\_error: The agent en-  
901 countered an error while using a tool (e.g. by calling it incorrectly)., infinite\_loop: The agent entered  
902 an infinite loop (e.g. repeating the same sequence of steps)., endless\_file\_reading: The agent read the  
903 same file multiple times without making any changes., context\_overflow\_from\_listing: The agent’s  
904 file listing operations (ls, find, etc.) caused context overflow., syntax\_error: The agent generated  
905 syntactically incorrect code., other: The agent failed to resolve the issue for other reasons. Do NOT  
906 invent or propose new categories. If none fits, use "other".

907 Place the category at the end, separated by two newlines. Category must be all lowercase and only  
908 list the category name.

909 Remember to write two new lines before the category.

910

911

912

913

914

915

916

917