

SWE-smith: Scaling Data for Software Engineering Agents

John Yang¹, Kilian Lieret², Carlos E. Jimenez², Alexander Wettig², Kabir Khandpur³,
Yanzhe Zhang¹, Binyuan Hui⁴, Ofir Press², Ludwig Schmidt¹, Diyi Yang¹

¹Stanford University ²Princeton University ³Indepedent ⁴Alibaba Qwen

Abstract

Despite recent progress in Language Models (LMs) for software engineering, collecting training data remains a significant pain point. Existing datasets are small, with at most 1,000s of training instances from 11 or fewer GitHub repositories. The procedures to curate such datasets are often complex, necessitating hundreds of hours of human labor; companion execution environments also take up several terabytes of storage, severely limiting their scalability and usability. To address this pain point, we introduce SWE-smith, a novel pipeline for generating software engineering training data at scale. Given any Python codebase, SWE-smith constructs a corresponding execution environment, then automatically synthesizes 100s to 1,000s of task instances that break existing test(s) in the codebase. Using SWE-smith, we create a dataset of 50k instances sourced from 128 GitHub repositories, an order of magnitude larger than all previous works. We train SWE-agent-LM-32B, achieving 40.2% Pass@1 resolve rate on the SWE-bench Verified benchmark, state of the art among open source models. We open source SWE-smith (collection procedure, task instances, trajectories, models) to lower the barrier of entry for research in LM systems for automated software engineering. All assets available at <https://swesmith.com>.

1 Introduction

Language Model (LM) agents, such as SWE-agent (Yang et al., 2024a) or OpenHands (Wang et al., 2024), have made remarkable progress towards automating software engineering (SE) tasks, as tracked by benchmarks such as SWE-bench (Jimenez et al., 2024b). However, the most effective agents still rely on proprietary LMs. On the other hand, building open source LMs for SE remains bottlenecked by the lack of large-scale, high-quality training data. To ensure that open research remains relevant in this field, it is critical to develop infrastructure for collecting software engineering training data at scale.




Figure 1: **Scaling task instances** (left) and **performance** (right) for SWE-agent's with SWE-smith. Using SWE-smith, we can create 100s to 1000s of instances for any Python codebase, enabling us to train SWE-agent-LM-32B which achieves 40.2% on SWE-bench Verified.

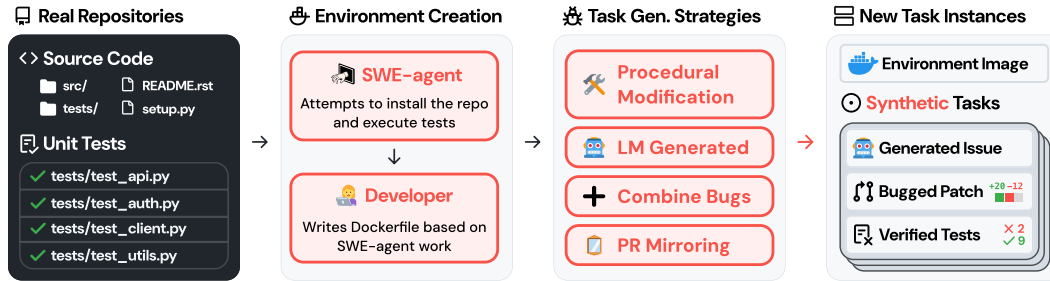


Figure 2: SWE-smith creates training data for software engineering agents by crafting bugs into real codebases. Given a codebase, we employ several strategies to create task instances that break existing tests. Using SWE-smith, we create 50k+ task instances with execution environments from 128 real world repositories.

The current open-source software ecosystem offers two kinds of data sources to train LMs on SE tasks. One simple approach is to crawl pull requests (PRs) from GitHub repositories. However, without execution environments or tests, these instances offer no reliable way of validating generated solutions, and LMs are limited to learning from the surface form of code (Xie et al., 2025a) or via rewards based on superficial string similarity (Wei et al., 2025).

In contrast, SWE-bench provides reliable validation by running unit tests against proposed solutions. Another line of work has simply extended the SWE-bench collection strategy to a new set of repositories for training purposes (Pan et al., 2024). This produces flexible environments for training and distilling LM agents, since we can generate agent trajectories and filter them based on the unit test results. However, the scalability of this approach is severely limited by the challenges associated with SWE-bench’s collection strategy. SWE-bench’s filtering process leaves only a small number of PRs that not only resolve a Github issue, but also make meaningful changes to unit tests. Also, setting up execution environments for each instance requires a substantial amount of human intervention.

In this paper, we introduce the SWE-smith toolkit, which marries the flexible execution environments of SWE-bench with scalable instance collection (Figure 1). SWE-smith features several techniques to automatically synthesize bugs in existing GitHub repositories, such as (1) generating errant rewrites of functions with an LM, (2) procedurally modifying the abstract syntax tree (AST) of functions, (3) undoing PRs, and (4) combining bugs. Our key insight is that execution-based validation can not only validate proposed solutions, but also identify bug candidates which cause substantial software regression (i.e., break tests).

In a nutshell, SWE-smith puts forth the following task creation workflow, as shown in Figure 2. Given a codebase, we automatically set up a corresponding environment using SWE-agent (Yang et al., 2024a). Within this environment, we then use the aforementioned techniques to synthesize 100s to 1,000s of task instances. Finally, we craft realistic issue descriptions automatically with LMs. SWE-smith’s design significantly reduces the amount of human labor and storage required for constructing execution environments. Using SWE-smith, we create a dataset of 50k task instances across 128 real-world GitHub repositories.

We train effective LMs with SWE-smith. We run SWE-agent with Claude 3.7 Sonnet on 20k task instances, collecting 5,016 expert trajectories. We fine-tune the Qwen 2.5 Coder Instruct 32B model on these trajectories, resulting in SWE-agent-LM-32B which achieves 40.2% (+33.4%) on SWE-bench Verified in a single attempt, without inference-time scaling. At the 32B model size, SWE-agent-LM-32B achieves a state of the art result.

The scale and diversity of the SWE-smith dataset enables us to begin establishing truths and investigate interesting phenomena about developing SWE-agents. Training on more instances, bug types, and repositories helps. LM generated issue text approximates real ones effectively. Using SWE-smith, we find that it’s possible to optimize LMs to perform well for specific repositories while only suffering minor generalization loss.

We release SWE-smith as an open-source toolkit — including instances, environments, and trajectories — to catalyze the development of stronger open-source LM agents.

2 SWE-smith: Software Task Generation at Scale

The core principle of SWE-smith’s collection strategy is to define an execution environment first, and then synthesize task instances within the environment. Conceptually, this is a simple inversion of SWE-bench’s approach, which instead prioritizes identifying task instances, and then attempts to build an environment for each. In this section, we describe the procedure in detail and show how, in practice, SWE-smith scales significantly better in terms of repositories, task instances, and storage.

2.1 Collection

Building execution environments for repositories with passing tests. Given a repository, we run SWE-agent (Yang et al., 2024a) on the latest commit for at most 100 steps, instructing it to install the codebase and run the test suite. We then manually verify the installation and testing instructions, check if more than 80% of existing tests pass, and finally create a Docker image for the repository. We target repositories for the 5,000 most downloaded packages listed in the Python Package Index (PyPI) as of November 18, 2024, sort the PyPI packages by GitHub stars, and then remove any PyPI package with less than 1,000 stars, as well as all 12 SWE-bench test repositories from consideration. More in §A.2.

Creating task instance candidates. Per repository, we employ four different strategies to create candidates. As shown in Figure 2, each strategy takes in a repository as input, then produces task instance candidates represented as `.diff` files. Extensive details in §B.

- **LM Generation:** Per repository, we identify all programmatic entities (functions, classes), then take two approaches: (1) provide an LM with the function and prompt it to introduce errant *modifications* (henceforth referred to as “LM Modify”), and (2) given only the function header and docstring, ask the LM to *rewrite* it (“LM Rewrite”). More in §B.1.
- **Procedural Modification:** Per function, we acquire an abstract syntax tree (AST) representation of the code, then randomly perform one or more transformations (e.g., remove a conditional/loop, change an operator, +11 more. See Table 8). More in §B.2.
- **Combine Bugs:** LM generation and Procedural Modification task instances exclusively edit one function or class. To create more complex tasks that require editing multiple portions of the codebase, we devise a “Patch Combination” strategy that creates a task instance by aggregating candidates from the same file(s) or module(s). More in §B.3.
- **Invert PRs** (or “PR Mirror”): Per repository, we collect all PRs that modify Python files. Per PR, we attempt to *undo* its revisions in the current version of the repository. To achieve this, we provide an LM with the PR’s code changes (a `.diff` plaintext) and prompt it to rewrite each affected file such that the PR edits are reverted. Unlike SWE-bench, we do *not* check out the PR’s base commit, as the install specifications determined in the previous step may not be compatible with older versions of the repo. More in §B.4.

Execution-based validation of candidates. We apply each candidate patch to the corresponding repository, run the test suite, and only keep patches that break one or more existing, passing tests (referred to as *Fail-to-Pass* or *F2P* test(s)). For efficiency purposes, we also limit testing runtime to two minutes; bug candidates that cause test runtimes in excess of this time limit are discarded. Minor additional details in §A.3.

Generating problem statements. The issue text associated with a bug can significantly alter the difficulty and feasibility of the task instance. Detailed descriptions of “expected” vs. “observed” behavior or bug-reproduction code in issue text heavily affect an agent’s capacity to localize bugs or iterate on proposed solutions. We explore several techniques covered fully in §D, and ultimately settle on a simple strategy. Per task instance, we provide an LM with the `.diff` patch, source code of a random F2P test, and execution output from running the repository’s test suite with the bug patch applied. We prompt the LM for GitHub issue-style text that includes reproduction code based on the F2P test.

What human labor remains? The steps requiring manual effort are (1) parsing the correct installation setup procedures from the agent trajectory (~ 7 min per repository), and (2)




Figure 3: Distribution of instances per repo for 128 repos grouped into 6 categories.

Bug Type	Yield %	# Insts	Cost	F2P	Lines
Combine	96.9%	10,092	0.00¢	15	11
LM Modify	56.0%	17,887	0.38¢	4	3
LM Rewrite	35.0%	4,173	3.93¢	4	24
PR Mirror	33.8%	2,344	5.53¢	3	14
Procedural	40.2%	15,641	0.00¢	7	5
Total	50.1	50,137	2.32¢	6	5

Table 1: Summary of SWE-smith statistics. “Yield %” is the % of candidates generated by a strategy that break 1+ tests. “Cost” is the average cost to generate one candidate. “F2P” (Fail to Pass tests), “Lines [Edited]” are median values.

implementing the parser for test outputs (~ 1 min per repository). Step two requires very little time because parsers can be reused for repositories with the same testing infrastructure (e.g., pytest). SWE-smith removes the need for manual efforts to determine installation specifications for multiple versions of a codebase across time, the most costly step of SWE-bench collection. Creating SWE-smith took one author ~ 20h of human labor.

2.2 Features

We apply SWE-smith to 128 Python repositories, generating a total of 50k instances. Table 1 captures the key statistics. On average, we generate 381 task instances per repository, with as many as 2277 for pandas-dev/pandas. We summarize the distribution of task instances per repository in Figure 3, where repositories are grouped into one of six general categories.

Bug generation strategies vary in cost and yield rate. Of methods relying on LMs, PR Mirrors are more expensive because the task entails rewriting entire files, as opposed to individual functions for LM Modify and LM Rewrite. Yield rates are limited by either lack of test coverage for the change or because the bug candidate did not actually introduce relevant issues. For example, for LM Rewrite, the LM is asked to re-implement the function; it is *not* explicitly asked for bugs. When requested outright (LM Modify), the yield is higher.

How difficult are SWE-smith task instances? As part of the SWE-bench Verified curation effort, Chowdhury et al. (2024) performed human annotations of 1699 SWE-bench task instances with four levels of difficulty: ≤ 15 min, 15m - 1hr, 1-4 hrs, and 4+ hrs. We map these ratings to three labels: easy (≤ 15 min), medium (15m - 1hr), and hard (1+ hrs). We then LoRA fine-tune (Hu et al., 2021) a Qwen 2.5 32B Instruct model to assign a difficulty label given the solution/bug patch and issue text. The 1699 instances are split into an 80/20 train/test split, and the SFT’ed model achieves an accuracy of 68.24% on the test set. To quantify difficulty, we assign scores of 1/5/9 to easy/medium/hard. Using this model, we can estimate difficulty for SWE-smith and entire datasets, as shown in Figure 4. More in §E.




Figure 4: Distribution of task instance difficulty (easy/medium/hard) for existing SWE-bench style datasets (left 5 bars) and SWE-smith (right 5 bars), assessed by our difficulty rating model. The average difficulty score for each dataset is listed above each bar. For SWE-smith, per bug strategy, we sample 1000 task instances with LM generated issue text.

Dataset	# Tasks	# Repos	Exec?	Source	Env. Size
R2E (Jain et al., 2024)	0.25k	137	✓	Synth	270 GBs
R2E-gym (Subset) (Jain et al., 2025)	4.6k	10	✓	Synth	4 TBs
SWE-bench-extra (Badertdinov et al., 2024)	6.38k	2k	✗	Real	-
SWE-bench-train (Jimenez et al., 2024b)	19k	37	✗	Real	-
SWE-fixer (Xie et al., 2025a)	115k	856	✗	Real	-
SWE-gym (Pan et al., 2024)	2.4k	11	✓	Real	6 TBs
SWE-smith	50k	128	✓	Both	295 GBs

Table 2: Comparison of open source training datasets for software engineering tasks. Relative to existing datasets, SWE-smith has multiple times the number of task instances, repositories, and environments at a fraction of prior storage costs. SWE-fixer and SWE-bench-train task instances do not have execution environments, so “Env. Size” is blank.

Scaling execution environments. Unlike SWE-bench which creates a Docker image per task instance, SWE-smith leverages a simpler design where tasks from the same repository share the same environment, reducing storage overhead significantly, as shown in Table 2. This approach not only makes scaling task instances more affordable, but also renders SWE-smith more accessible and maintainable than existing datasets. We estimate that creating a similar quantity of task instances (50k) using SWE-bench would require 50 to 150 TBs of storage for environments, a 500x difference. Extended discussion in §C.1.

Cost rundown and estimations. SWE-smith took about \$1360 to create (\$1000 on LM Modify/LM Rewrite/PR Mirror bugs, \$160 for automatic repository installation with SWE-agent, \$200 to generate issues for 10K bugs). Generating an issue costs 2.54¢ on average. Therefore, creating 1000 LM Rewrite bugs with SWE-smith costs an estimated \$65. Recreating SWE-bench with SWE-smith would take $2294 \times 0.053 = \$126.86$.

3 Experiments

To explore the utility of SWE-smith for training software engineering agents, we use rejection sampling fine-tuning (Yuan et al., 2023) as the primary procedure for improving a base LM with SWE-smith. Our experiment workflow is as follows. First, we curate a subset of SWE-smith task instances. Next, we run an agent system with an expert model on this subset. At this step, the trajectory corresponding to each run is recorded. Then, we fine-tune the base (or “student”) model on the trajectories corresponding to resolved instances. Finally, we evaluate the agent system run with the student model on a separate, test split.

Models. For expert models, we use `claude-3-7-sonnet-20250219` (Anthropic, 2025). For fair comparisons with prior works (Pan et al., 2024), we also use `claude-3-5-sonnet-20240620` and `gpt-4o-2024-08-06`. We use the `Qwen-2.5-Coder-Instruct` (Hui et al., 2024) 7B and 32B series as the base models. Training and hyperparameter details are in §F.1.

Agent system. We use SWE-agent (Yang et al., 2024a), an agent system for solving GitHub issues. SWE-agent provides a base LM with an Agent Computer Interface (ACI) that enables more effective interactions with a codebase. At each turn, SWE-agent prompts an LM to generate a ReAct (Yao et al., 2023b) style (thought, action) pair, where the action either edits a file or executes a shell command. We choose SWE-agent because, at the time of writing, SWE-agent with Claude 3.7 Sonnet is the top open source solution on SWE-bench. When generating trajectories with expert models, we run SWE-agent for at most 75 steps and \$2.00 cost limit. For inference of student models, we impose the same 75 step maximum and fix temperature at 0.0. Full configuration details are in §F.1.

Evaluation metrics. We evaluate on the SWE-bench Lite and Verified (Chowdhury et al., 2024) subsets. SWE-bench evaluates AI systems on their ability to solve software issues from 12 real world GitHub repositories. The Lite split is a subset of 300 instances, curated to be an easier evaluation set that’s less costly to run. The Verified split is a human-curated subset of 500 instances, selected for clearer problem statements and more reliable evaluation. We report the % **resolved** metric, the proportion of successfully resolved instances.

4 Results

In Table 3, we show the performance of the Qwen 2.5 Coder Instruct model at 7B and 32B, fine-tuned on 5,016 trajectories generated from SWE-smith task instances. We refer to the models as SWE-agent-LM-7B and SWE-agent-LM-32B. At the 32B model size, SWE-agent-LM-32B achieves state-of-the-art performance.

Model	System	Train Size	Lite	Verified
<i>Closed Weight Models</i>				
GPT-4o (OpenAI, 2024a)	Agentless	-	32.0	38.8
	OpenHands	-	22.0	-
	SWE-agent	-	18.3	23.0
Claude 3.5 Sonnet (Anthropic, 2024)	Agentless	-	40.7	50.8
	AutoCodeRover	-	-	46.2
	OpenHands	-	41.7	53.0
	SWE-agent	-	23.0	33.6
Claude 3.7 Sonnet (Anthropic, 2025)	SWE-agent	-	48.0	58.2
Llama3-SWE-RL-70B (Wei et al., 2025)	Agentless	11M	-	41.0
<i>Open Weight Models</i>				
Lingma-SWE-GPT-72B (Ma et al., 2024)	SWE-SynInfer	-	-	28.8
Qwen3-235B-A22B (Qwen et al., 2025)	OpenHands	-	-	34.4
R2E-Gym-32B (Jain et al., 2025)	OpenHands	3.3k	-	34.4
SWE-fixer-72B (Xie et al., 2025a)	SWE-Fixer	110k	24.7	32.8
SWE-gym-32B (Pan et al., 2024)	OpenHands	491	15.3	20.6
SWE-agent-LM-7B	SWE-agent	2k	11.7	15.2
SWE-agent-LM-32B	SWE-agent	5k	30.7	40.2

Table 3: Resolve rates for existing solutions on SWE-bench Lite and Verified, collected from Jimenez et al. (2024a), compared to models fine-tuned on SWE-smith. All performance numbers are pass@1. For units under “Train Size”, k refers to 1000 and M to a million. We report scores for systems that do *not* use verifiers or multiple attempts at test time.

The final dataset of 5,016 training points was curated as follows. We start by collecting a large pool of expert trajectories. First, we carried out each of the ablations in Section 4.1, giving us an initial set of 5,105 trajectories. Next, based on our observation that PR Mirror and LM Rewrite task instances yield the most effective expert trajectories (discussed below), we run the expert model on all task instances of these types, bumping up the total number to 6,457 task instances. Ultimately, we attempt to generate expert trajectories for 8,686 unique task instances, or 17.3% of the SWE-smith dataset. Reinforcing the difficulty rating findings from Section 2.2, we observe that SWE-smith task instances are non-trivial for the top agent systems today. The final pool of 6,457 represents a 36% resolve rate of all 17,906 attempts to solve one of the 8,686 task instances.

Next, we perform minor filtering of this collection. As reported in Pan et al. (2024), we also observe that “easier” trajectories – task instances that are repeatedly solved across multiple runs — degrades model performance. Therefore, we limit the number of times any SWE-smith task instance is represented in the training set to 3 trajectories. This leads to the final 5,016 training set. A full breakdown of trajectory collection is provided in §F.2.

Performance improves with more data points. Extending similar graphs in Pan et al. (2024) and Jain et al. (2025), Figure 1 shows increasing performance with more trajectories.

Comparison at the same training set size. To compare with SWE-gym-32B (Pan et al., 2024) and R2E-Gym-32B (Jain et al., 2025), we run expert trajectory generation on 1000 random SWE-smith task instances with SWE-agent + Claude 3.5 Sonnet (800) or GPT-4o (200). We then fine-tune the 32B model on 500 successful trajectories, a training set size both works report on. Our model achieves a 28.2% resolve rate on SWE-bench Verified, a relative difference of +8.2% with Pan et al. (2024) and +0.7% with Jain et al. (2025).

Strategy	# Trajs.	% Resolved	Issue	# Trajs.	% Resolved
LM Modify	802	5.7 (± 1.5)	Fixed	259	6.4 (± 1.5)
LM Rewrite	507	8.8 (± 1.7)	F2P Test	390	7.3 (± 1.9)
Procedural	745	8.6 (± 1.8)	LM	328	7.7 (± 1.5)
PR Mirror	557	9.2 (± 1.7)	Original	319	7.8 (± 1.8)

Table 4: Comparison of training on 1000 SWE-smith instances created with different strategies.

Table 5: Comparing training on 600 PR Mirror instances with varied issue text.

4.1 Ablations of SWE-smith

We perform several ablations of how SWE-smith’s bug and problem statement generation strategies impact the quality of training data. Unless otherwise specified, for all ablations we use Claude 3.7 Sonnet as the expert model, Qwen 2.5 7B Coder Instruct as the student model, and SWE-bench Verified for evaluation.

LM Rewrite and Procedural bugs are comparable to PR mirrors. We randomly sample 1000 instances per bug generation strategy (LM Modify, LM Rewrite, Procedural Modifications, PR Mirrors). Per set, we generate issue text with an LM and run expert trajectory generation. We then fine-tune separate student models per strategy, capping training points to the minimum number of successful trajectories from any strategy (507) for fair comparison.

Table 4 summarizes the results. Trajectories generated from PR mirrors are empirically the most effective training data — this is expected, since they are most reflective of SWE-bench. What’s noteworthy is that trajectories from Procedural Modification and LM Rewrite instances lead to competitive models. There is a steep drop off with LM Modify bugs.

LM generated issues are comparable to real issues. We randomly sample 600 PR Mirror task instances. We compare LM generated issues with three alternatives — fixed issue templates, the source code + testing logs of a random Fail-to-Pass test, and the original issue text associated with the PR. We again cap the number of training points to the minimum number of successful trajectories (259) for fairness.

As shown in Table 5, training on task instances with LM generated issues is empirically comparable to using the original issue text. Using fixed issue templates not only leads to the fewest successful trajectories, but also results in relatively homogeneous problem solving sequences. The expert trajectories from fixed issue templates have 31% fewer unique actions compared to LM generated text (379 vs. 550).

While providing a Fail-to-Pass test leads to more successful expert trajectories, leaking the evaluation criteria causes the model to skip over writing a reproduction script, which accounts for the performance drop. Of 500 SWE-bench Verified instances, the student model trained on LM-generated issues attempts to reproduce the bug for 379 of the runs. The model trained on test-based issues only does so for 127 cases, a 66% decrease.

Task difficulty correlates with solvability but not with effectiveness as training data. First, we run our difficulty rating model on 10k randomly selected SWE-smith task instances. From this pool, we curate subsets of 1000 instances corresponding to the three difficulty levels, then run expert trajectory generation per subset 3 times. For the easy/medium/hard subsets, the resolve rate by the expert model are 58.6%, 41.0%, and 17.0% respectively.

Next, from all successful trajectories, we create four fine-tuning datasets of 500 trajectories each corresponding to difficulty scores of 2, 4, 6, and 8. As mentioned in Section 2.2, the corresponding scores for easy/medium/hard are 1/5/9. Therefore, the SFT dataset for score 2 is made up of trajectories corresponding to 375 easy and 125 medium instances, and so on. Somewhat surprisingly, we do not observe strong correlation between increased difficulty and downstream performance. For the student models trained on the 2/4/6/8 difficulty SFT datasets, we get pass@1 scores of 12.4%, 10.8%, 13.6%, and 12.2% on SWE-bench Verified.

Training on more repositories improves general performance. We train models in four settings by sampling 700 expert trajectories on Procedural Modification tasks from pools of 4, 25, 50, and 100 repositories. Echoing similar findings for code generation tasks (Xie




Figure 5: We fine-tune a 7B base model and SWE-agent-LM-32B on 700 trajectories for SymPy. Specialization greatly improves SymPy performance, with a slight drop in generalization.




Figure 6: Keeping the number of training samples fixed at 700, we observe performance increases logarithmically with the number of repositories.

et al., 2025b), we find that increasing repositories represented in the training set improves performance, as shown in Figure 6. We observe an approximately logarithmic relation between model performance and the number of repos sampled from during training.

Repository-specialized models excel on the specialized repository with minor generalization loss. We experiment with training models to be specialists on one particular repository. To assess performance, we evaluate models on a subset of SWE-bench Verified tasks that are (1) from SymPy, and (2) created after January 1st, 2022, a total of 22 instances. We create SymPy specific training data as follows. We first select a base commit of SymPy just before the cutoff date. Next, we create 1276 Procedural Modification task instances, then generate 700 expert trajectories. We evaluate specialization in two settings: (1) single-repository fine-tuning, and (2) specialist stage fine-tuning, both shown in Figure 5. For single-repository tuning, we compare a model initialized with Qwen-2.5-Coder-Instruct 7B and trained on 700 instances sampled from 100 repositories, to the same Qwen base model but fine-tuned on the 700 SymPy instances only. For specialist stage fine-tuning, we simply compare SWE-agent-LM-32B to the same model further fine-tuned on the 700 SymPy instances.

Specializing on a single repository significantly boosts performance for the target repository with only slight drops in general performance in both the single-repository fine-tuning (21.2% vs. 13.6%) and specialist stage fine-tuning (42.4% vs. 33.3%) settings, compared to baselines trained across 128 repositories.

4.2 Analysis of Agent Behavior

This section analyzes the behavior, failure modes, and efficiency of SWE-agent when run with SWE-agent-LM-32B or Claude 3.7 on SWE-bench verified.

SWE-agent-LM-32B can solve tasks efficiently. SWE-agent-LM-32B resolves tasks in fewer steps on average (24.9) than Claude 3.7 (29.1), though the difference becomes marginal when accounting for different average difficulties of the resolved tasks: On the overlap of tasks that are resolved by both LMs, SWE-agent-LM-32B uses 24.8 steps compared to 25.6 used by Claude 3.7 (see Fig. 7). While shorter trajectories are not always preferred (additional actions can be used for additional validation purposes, for example), this shows that SWE-agent-LM-32B solves tasks very efficiently. At the same time SWE-agent-LM-32B also demonstrates that it can remain focused throughout long trajectories, with 31 instances being resolved after 40 steps or more. We further highlight that the accuracy of naturally terminating¹ agent submissions with SWE-agent-LM-32B achieve an accuracy nearly matching that of Claude 3.7 (60% vs 63%), showing that SWE-agent-LM-32B is adept at determining whether an instance has been resolved. As the overall cost and turn count averages scale strongly with the cost and turn limits, we reserve a more thorough analysis for §F.3.1.

¹i.e., excluding agent runs that are terminated due to errors or cost/step count limits. Note that SWE-agent still extracts and submits any changes performed by the agent in these cases and some




Figure 7: For instances solved by both models, SWE-agent-LM-32B is faster than Claude 3.7 Sonnet (solved in fewer steps).




Figure 8: For unsuccessfully resolved task instances, a frequent failure mode is that SWE-agent-LM-32B will repeat actions.

Repetitive actions are a key problem. We observe a tendency for SWE-agent-LM-32B to get stuck in long sequences of repetitive actions, in particular long sequences of calls that display different portions of a file instead of using search commands.² More than 25% of SWE-agent-LM-32B trajectories have a repetitive sequence of at least length 10, compared to less than 4% for Claude 3.7 (see Figure 8). The occurrence of long repetitive sequences correlates strongly with the agent’s ability to solve the corresponding task instance, largely because the LM continues issuing similar commands until either the agent cost or turn limit is reached, at which point the run is terminated. For example, repetitive sequences of length 10 correspond to an 89% failure probability. On the other hand, when only comparing successful instances, the distributions of longest repetitive sequence length are similar between SWE-agent-LM-32B and Claude 3.7 (see Figure 8).

Localization is the dominant failure mode. Guided by a short plan in the system prompt, SWE-agent typically starts by *localizing* (search and read actions), *reproducing* (test file creation and execution), before modifying source files and validating the fixes. If the agent gets stuck at any of these stages or keeps on iterating, the agent loop is eventually interrupted by runtime limits (cost, number of LM calls, runtime). While this rarely happens with Claude 3.7 Sonnet, 53% of SWE-agent-LM-32b’s failures are associated with such limits (Figure 9). The agent often already gets stuck during localization or initial efforts to reproduce a bug, with endlessly repeated actions being a persistent issue. More failure modes in §F.3.

5 Related Work

LMs for Software Engineering. As contemporary LMs have saturated traditional code generation tasks (Austin et al., 2021; Chen et al., 2021), software engineering benchmarks (Jimenez et al., 2024b; Yang et al., 2024b; Jain et al., 2024; Zhao et al., 2024; Zan et al., 2025), notably SWE-bench, have become a new de facto evaluation setting due to their diverse, complex, real-world programming challenges. The most significant source of open source progress on SWE-bench has been the development of LM-based workflows (Orwall, 2024; Xia et al., 2024; Zhang et al., 2024b) and agents (Antoniades et al., 2024; Wang et al., 2024; Yang et al., 2024a; Zhang et al., 2024a). Workflow based systems are typically human engineered decompositions of a task into a sequence of sub-goals. Yang et al. (2024b) suggests such pipelines may not generalize effectively to non-Python repositories, requiring additional human intervention to re-adapt. We therefore elect to focus on generating trajec-

of them can be successful (for example if the agent is terminated due to cost while testing already performed edits).

²In fact, these `str_replace_editor` view commands make up 73% of the longest repetitive sequences. For this analysis, we look at repetitions of the base command, i.e., without any flags and arguments. For more information on this analysis, see §F.3.




Figure 9: More than half of the unresolved instances of SWE-agent-LM-32B correspond to runs terminated by cost/step limits, and these limits are frequently reached before source code has been modified. See §F.3 for details on how failures are attributed.

tories with and for LM agent systems (Yao et al., 2023b; Yang et al., 2023; Sumers et al., 2024). Because no workflow is imposed, agent systems inherently rely more on the LM to plan and refine its actions, putting more focus on an LM’s capabilities, not inference scaffolds.

Training Datasets for Coding. Prior work around training data has mostly focused on instruction following (Luo et al., 2023; Muennighoff et al., 2024; Shypula et al., 2024; Wei et al., 2024a;b; Yu et al., 2024) and preference learning (Liu et al., 2024a;b) for code completion tasks. Several recent works introduce training sets for retrieval augmented generation (Jimenez et al., 2024b; Xie et al., 2025a), workflows (Wei et al., 2025), and agent (Badertdinov et al., 2024; Ma et al., 2024; Pan et al., 2024; Jain et al., 2025) approaches to SWE-bench. Our work can be thought of as applying Haluptzok et al. (2023) to a repository — by having an LM break a codebase, we drastically reduce the amount of human labor needed to identify a task instance and construct an environment for it. Concurrent to our work, Xie et al. (2025b) (RePOST) also constructs execution environments for repository functions, but differs significantly in methodology and evaluation. RePOST sandboxes a function and its dependencies to a separate script, then generates tests with an LM, removing the original codebase as context. The tasks’ source is repository-level; the environments and tasks are not. RePOST evaluates solely on code generation (e.g., HumanEval (Chen et al., 2021)). Jain et al. (2025) (R2E-Gym) improves open source LMs’ performance on SWE-bench with inference time scaling and verifiers. R2E-gym’s 51% resolve rate is not comparable to Table 3 results, as each instance is attempted 26 times. R2E-gym’s 4.6k training instances are collected using SWE-bench’s pipeline, with some augmentations around using LMs to synthesize issue text and tests. The fundamental limitations around constructing task instances with SWE-bench that we address are not the focus of any prior works.

6 Discussion

Limitations and future directions. First, SWE-smith’s collection pipeline is Python-centric. The mechanisms to identify programmatic objects (e.g. functions, classes) and perform transformations rely heavily on the Python specific ast library. That said, SWE-smith’s collection strategy is transferable to other languages. Second, due to both compute/budget constraints and our work’s primary stance as a dataset contribution, we only include fine-tuning as a demonstration of SWE-smith’s effectiveness. We do not explore other training techniques such as reasoning capabilities elicited via reinforcement learning.

Conclusion. We introduce SWE-smith, a dataset of 50k software engineering task instances from across 128 real world GitHub repositories. SWE-smith collection pipeline allows us to scale up task instances, environments, and trajectories at a fraction of prior costs without sacrificing faithfulness to open source software development practices. Using SWE-smith, we train SWE-agent-LM-32B, achieving a state-of-the-art 40.2% on SWE-bench Verified. Our experiments show how SWE-smith can be used to identify fundamental trends about developing SWE-agents. We believe SWE-smith provides the foundational data and infrastructure needed to train software engineering agents in a truly scalable manner.

Acknowledgments

We thank Princeton Language & Intelligence (PLI) for providing credits for running closed-source API models. Thanks to Samuel Ainsworth for his constant support of bitbop.io (<https://bitbop.io/>), the compute service for which the majority of the project was carried out with. We'd also like to thank Akshat Bubna, Howard Halim, Andrew Liu, Peyton Walters, and the great team at Modal (<https://modal.com/>) for providing credits that made fine-tuning and model serving efforts extremely easy for this project. This work is partially supported by ONR grant N000142412532 and NSF grant IIS-2247357. We also thank Open Philanthropy and Andreessen Horowitz for providing funding for this work. Finally, thanks to Tianyu Gao, William Held, Niklas Muennighoff, Rafael Rafailov, Yijia Shao, Chenglei Si, Anikait Singh, Tianyi Zhang, Kexin Pei, and Karthik Narasimhan for constructive discussions and support throughout this project.

References

- Anthropic. Introducing claude 3.5 sonnet, 2024. URL <https://www.anthropic.com/news/claude-3-5-sonnet>.
- Anthropic. Introducing claude 3.7 sonnet, 2025. URL <https://www.anthropic.com/news/claude-3-7-sonnet>.
- Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, and William Wang. SWE-Search: Enhancing Software Agents with Monte Carlo Tree Search and Iterative Refinement, December 2024. URL <http://arxiv.org/abs/2410.20285>. arXiv:2410.20285 [cs].
- 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, August 2021. URL <http://arxiv.org/abs/2108.07732>. arXiv:2108.07732 [cs].
- Ibragim Badertdinov, Maria Trofimova, Yuri Anapolskiy, Sergey Abramov, Karina Zainullina, Alexander Golubev, Sergey Polezhaev, Daria Litvintseva, Simon Karasik, Filipp Fisin, Sergey Skvortsov, Maxim Nekrashevich, Anton Shevtsov, and Boris Yangel. Scaling data collection for training software engineering agents. *Nebius blog*, 2024.
- Ben Bogin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish Sabharwal, and Tushar Khot. Super: Evaluating agents on setting up and executing tasks from research repositories, 2024. URL <https://arxiv.org/abs/2409.07440>.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. FireAct: Toward Language Agent Fine-tuning, October 2023. URL <http://arxiv.org/abs/2310.05915>. arXiv:2310.05915 [cs].
- 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, July 2021. URL <http://arxiv.org/abs/2107.03374>. arXiv:2107.03374 [cs].
- Neil Chowdhury, James Aung, Chan Jun Shern, Oliver Jaffe, Dane Sherburn, Giulio Starace, Evan Mays, Rachel Dias, Marwan Aljubeih, Mia Glaese, et al. Introducing swe-bench verified, 2024. URL <https://openai.com/index/introducing-swe-bench-verified>, 2024.

- Michael Han Daniel Han and Unsloth team. Unsloth, 2023. URL <http://github.com/unslothai/unsloth>.
- Aleksandra Eliseeva, Alexander Kovrigin, Ilia Kholkin, Egor Bogomolov, and Yaroslav Zharov. Envbench: A benchmark for automated environment setup, 2025. URL <https://arxiv.org/abs/2503.14443>.
- Patrick Haluptzok, Matthew Bowers, and Adam Tauman Kalai. Language models can teach themselves to program better, 2023. URL <https://arxiv.org/abs/2207.14502>.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. Qwen2.5-coder technical report, 2024. URL <https://arxiv.org/abs/2409.12186>.
- Naman Jain, Manish Shetty, Tianjun Zhang, King Han, Koushik Sen, and Ion Stoica. R2e: Turning any github repository into a programming agent environment. In *ICML 2024*, 2024.
- Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2e-gym: Procedural environments and hybrid verifiers for scaling open-weights swe agents, 2025. URL <https://arxiv.org/abs/2504.07164>.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench leaderboard, 2024a. URL <https://swe-bench.github.io/>.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can Language Models Resolve Real-World GitHub Issues?, November 2024b. URL <http://arxiv.org/abs/2310.06770>. arXiv:2310.06770 [cs].
- Jiawei Liu, Thanh Nguyen, Mingyue Shang, Hantian Ding, Xiaopeng Li, Yu Yu, Varun Kumar, and Zijian Wang. Learning code preference via synthetic evolution, 2024a. URL <https://arxiv.org/abs/2410.03837>.
- Zhihan Liu, Sheno Zhang, Yongfei Liu, Boyi Liu, Yingxiang Yang, and Zhaoran Wang. DSTC: Direct Preference Learning with Only Self-Generated Tests and Code to Improve Code LMs, December 2024b. URL <http://arxiv.org/abs/2411.13611>. arXiv:2411.13611 [cs].
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct, 2023. URL <https://arxiv.org/abs/2306.08568>.
- Yingwei Ma, Rongyu Cao, Yongchang Cao, Yue Zhang, Jue Chen, Yibo Liu, Yuchen Liu, Binhua Li, Fei Huang, and Yongbin Li. Lingma swe-gpt: An open development-process-centric language model for automated software improvement, 2024. URL <https://arxiv.org/abs/2411.00622>.
- Modal. Modal: High-performance ai infrastructure, 2025. URL <https://modal.com/>.
- Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. OctoPack: Instruction Tuning Code Large Language Models, February 2024. URL <http://arxiv.org/abs/2308.07124>. arXiv:2308.07124 [cs].

- Shikhar Murty, Dzmitry Bahdanau, and Christopher D. Manning. NNetscape Navigator: Complex Demonstrations for Web Agents Without a Demonstrator, October 2024. URL <http://arxiv.org/abs/2410.02907>. arXiv:2410.02907 [cs].
- Niels Mündler, Mark Niklas Müller, Jingxuan He, and Martin Vechev. Swt-bench: Testing and validating real-world bug-fixes with code agents, 2025. URL <https://arxiv.org/abs/2406.12952>.
- OpenAI. Gpt-4o system card, 2024a. URL <https://arxiv.org/abs/2410.21276>.
- OpenAI. Openai o3-mini system card, 2024b. URL <https://cdn.openai.com/o3-mini-system-card-feb10.pdf>.
- Albert Orwall. Moatless tools, 2024. URL <https://github.com/aorwall/moatless-tools>.
- Tianyue Ou, Frank F. Xu, Aman Madaan, Jiarui Liu, Robert Lo, Abishek Sridhar, Sudipta Sengupta, Dan Roth, Graham Neubig, and Shuyan Zhou. Synatra: Turning Indirect Knowledge into Direct Demonstrations for Digital Agents at Scale, September 2024. URL <http://arxiv.org/abs/2409.15637>. arXiv:2409.15637 [cs].
- Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training Software Engineering Agents and Verifiers with SWE-Gym, December 2024. URL <http://arxiv.org/abs/2412.21139>. arXiv:2412.21139 [cs].
- PyTorch. torchtune: Pytorch’s finetuning library, April 2024. URL <https://github.com/pytorch/torchtune>.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 Technical Report, January 2025. URL <http://arxiv.org/abs/2412.15115>. arXiv:2412.15115 [cs].
- Junhong Shen, Atishay Jain, Zedian Xiao, Ishan Amlekar, Mouad Hadji, Aaron Podolny, and Ameet Talwalkar. Scribeagent: Towards specialized web agents using production-scale workflow data, 2024. URL <https://arxiv.org/abs/2411.15004>.
- Alexander Shypula, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob Gardner, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, and Amir Yazdanbakhsh. Learning performance-improving code edits, 2024. URL <https://arxiv.org/abs/2302.07867>.
- Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. Cognitive architectures for language agents, 2024. URL <https://arxiv.org/abs/2309.02427>.
- Konstantinos Vergopoulos, Mark Niklas Müller, and Martin Vechev. Automated benchmark generation for repository-level coding tasks, 2025. URL <https://arxiv.org/abs/2503.07701>.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. OpenHands: An Open Platform for AI Software Developers as Generalist Agents, October 2024. URL <http://arxiv.org/abs/2407.16741>. arXiv:2407.16741 [cs].
- Yuxiang Wei, Federico Cassano, Jiawei Liu, Yifeng Ding, Naman Jain, Zachary Mueller, Harm de Vries, Leandro von Werra, Arjun Guha, and Lingming Zhang. SelfCodeAlign: Self-Alignment for Code Generation, November 2024a. URL <http://arxiv.org/abs/2410.24198>. arXiv:2410.24198 [cs].

- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. Magicoder: Empowering code generation with oss-instruct, 2024b. URL <https://arxiv.org/abs/2312.02120>.
- 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, 2025. URL <https://arxiv.org/abs/2502.18449>.
- Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Agentless: Demystifying LLM-based Software Engineering Agents, October 2024. URL <http://arxiv.org/abs/2407.01489>. arXiv:2407.01489 [cs].
- Jiannan Xiang, Tianhua Tao, Yi Gu, Tianmin Shu, Zirui Wang, Zichao Yang, and Zhiting Hu. Language Models Meet World Models: Embodied Experiences Enhance Language Models, October 2023. URL <http://arxiv.org/abs/2305.10626>. arXiv:2305.10626 [cs].
- Chengxing Xie, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou, and Kai Chen. Swe-fixer: Training open-source llms for effective and efficient github issue resolution, 2025a. URL <https://arxiv.org/abs/2501.05040>.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments, May 2024. URL <http://arxiv.org/abs/2404.07972>. arXiv:2404.07972 [cs].
- Yiqing Xie, Alex Xie, Divyanshu Sheth, Pengfei Liu, Daniel Fried, and Carolyn Rose. Repost: Scalable repository-level coding environment construction with sandbox testing, 2025b. URL <https://arxiv.org/abs/2503.07358>.
- Yiheng Xu, Dunjie Lu, Zhennan Shen, Junli Wang, Zekun Wang, Yuchen Mao, Caiming Xiong, and Tao Yu. AgentTrek: Agent Trajectory Synthesis via Guiding Replay with Web Tutorials, December 2024. URL <http://arxiv.org/abs/2412.09605>. arXiv:2412.09605 [cs].
- John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. InterCode: Standardizing and Benchmarking Interactive Coding with Execution Feedback, October 2023. URL <http://arxiv.org/abs/2306.14898>. arXiv:2306.14898 [cs].
- 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, November 2024a. URL <http://arxiv.org/abs/2405.15793>. arXiv:2405.15793 [cs].
- John Yang, Carlos E. Jimenez, Alex L. Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press, Niklas Muennighoff, Gabriel Synnaeve, Karthik R. Narasimhan, Diyi Yang, Sida I. Wang, and Ofir Press. SWE-bench Multimodal: Do AI Systems Generalize to Visual Software Domains?, October 2024b. URL <http://arxiv.org/abs/2410.03859>. arXiv:2410.03859 [cs].
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. WebShop: Towards Scalable Real-World Web Interaction with Grounded Language Agents, February 2023a. URL <http://arxiv.org/abs/2207.01206>. arXiv:2207.01206 [cs].
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing Reasoning and Acting in Language Models, March 2023b. URL <http://arxiv.org/abs/2210.03629>. arXiv:2210.03629 [cs].
- Zhaojian Yu, Xin Zhang, Ning Shang, Yangyu Huang, Can Xu, Yishujie Zhao, Wenxiang Hu, and Qiufeng Yin. Wavocoder: Widespread and versatile enhancement for code large language models by instruction tuning, 2024. URL <https://arxiv.org/abs/2312.14187>.

- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and Jingren Zhou. Scaling relationship on learning mathematical reasoning with large language models, 2023. URL <https://arxiv.org/abs/2308.01825>.
- 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, 2025. URL <https://arxiv.org/abs/2504.02605>.
- Kexun Zhang, Weiran Yao, Zuxin Liu, Yihao Feng, Zhiwei Liu, Rithesh Murthy, Tian Lan, Lei Li, Renze Lou, Jiacheng Xu, Bo Pang, Yingbo Zhou, Shelby Heinecke, Silvio Savarese, Huan Wang, and Caiming Xiong. Diversity Empowers Intelligence: Integrating Expertise of Software Engineering Agents, August 2024a. URL <http://arxiv.org/abs/2408.07060>. arXiv:2408.07060 [cs].
- Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. AutoCodeRover: Autonomous Program Improvement, July 2024b. URL <http://arxiv.org/abs/2404.05427>. arXiv:2404.05427 [cs].
- Wenting Zhao, Nan Jiang, Celine Lee, Justin T. Chiu, Claire Cardie, Matthias Gallé, and Alexander M. Rush. Commit0: Library Generation from Scratch, December 2024. URL <http://arxiv.org/abs/2412.01769>. arXiv:2412.01769 [cs].
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A Realistic Web Environment for Building Autonomous Agents, April 2024. URL <http://arxiv.org/abs/2307.13854>. arXiv:2307.13854 [cs].

Appendix


The appendix is generally structured as follows. In Sections A to D, we review details about SWE-smith’s infrastructure and collection strategies for curating the SWE-smith task instances and execution environments, providing comparisons to existing datasets such as SWE-bench and SWE-gym along the way. In Sections E and onward, we discuss more about how we created the trajectories dataset, then provide additional ablations and results showcasing the effectiveness of SWE-smith as a dataset.

SWE-smith Overview

1. Given a GitHub repository, turn it into an execution environment



2. Given an execution environment, synthesize task instances



3. Given an execution environment + task instances, train SWE-agents!




Figure 10: An overview of pipelines in SWE-smith. Scripts/functions and manual steps are highlighted in blue. Artifacts that are also the inputs and outputs of these scripts are in orange. SWE-smith fits in seamlessly with the SWE-bench and SWE-agent ecosystem. Use SWE-smith to construct execution environments and generate task instances. Use SWE-agent to generate expert trajectories on SWE-smith task instances and run inference with models trained on these trajectories. Use SWE-bench to evaluate how good your models are at resolving GitHub issues and performing software engineering tasks.

A Infrastructure

We cover additional details about how SWE-smith works, specifically

- The form factor of a SWE-smith task instance.
- How we identify repositories and the SWE-agent configuration we use to automatically install them.
- How the task validation and evaluation harnesses work.

A.1 SWE-smith Task Instance

We briefly review the format of a SWE-smith task instance, highlight how it is different from a SWE-bench task instance, and discuss why SWE-smith’s relatively simple infrastructure compared to SWE-bench allows us scale task collection much more efficiently.

A SWE-smith task instance is very similar to the form factor of a SWE-bench task instance, with several minor differences. A SWE-smith task instance includes the following fields:

- `repo`: The repository the task instance is from.
- `instance_id`: A unique identifier (usually `(repo).(bug_type).(hash)`)
- `base_commit`: Hash of the GitHub branch that points to the repository with the bug patch applied.
- `patch`: The diff that causes the bug. It is applied to the original codebase to create the bug. Reverting this patch is effectively the solution.
- `problem_statement`: The generated issue text that conveys the bug. It is provided to a model or system before it begins attempting a fix.
- `created_at`: A timestamp matching when the bug was successfully validated and pushed to the mirror repository as a branch.
- `FAIL_TO_PASS`: The unit tests that break when the test suite is run with the bug patch applied.
- `PASS_TO_PASS`: The unit tests that do not break. These correspond to the set of all tests minus the `FAIL_TO_PASS` tests.

We summarize the key distinctions between a SWE-smith and SWE-bench task instance:

- SWE-smith task instances do not include the `version` or `environment_setup_commit` fields, which SWE-bench requires as additional identifiers for specifying repository-specific installation instructions across time. In SWE-smith, unique installation instructions are specified for each (repository, commit).
- The `hints_text` field is not included. In SWE-bench, this refers to the issue and PR thread comments written after the first commit of the corresponding PR.
- The `created_at` field is assigned the timestamp reflecting when the bug was successfully validated. Originally, `created_at` refers to when a PR was created.
- There is no `test_patch` field, as the SWE-smith collection pipeline does not create or synthesize any hidden tests. All `FAIL_TO_PASS` bugs are visible and runnable in the repository at inference time.

A.2 Repository Selection

In addition to the criteria discussed in Section 2.1, we also ensure that a repository has a license that allows non-proprietary use. The majority of software licenses are permissive (BSD, MIT, Apache), while the remainder are largely protective licenses (GPL) that still allow for non-commercial use. We inspected the repositories with custom licenses and confirmed they allowed for the use cases exercised in our work. The licenses for each repository are fully listed in Table 6.

We deliberately limit the search scope for repositories to those predominantly written in Python. Following precedents, focusing on Python repositories allowed us to form assumptions about installation and testing procedures (e.g. repository is organized as a PyPI package, pytest is the testing framework) that made scaling up automatic repository setup with SWE-agent more tractable. A worthwhile direction to consider for future work is expanding the coverage of repositories to be more comprehensive of codebases written in different programming languages, as Yang et al. (2024b) does, extending SWE-bench style evaluation to JavaScript repositories with multimodal inputs.

Apache License 2.0	Project-MONAI/MONAI; alanjds/drf-nested-routers; arrow-py/arrow; buriy/python-readability; facebookresearch/fvcore; getmoto/moto; google/textfsm; iterative/dvc; jax-ml/jax; jd/tenacity; kayak/pypika; modin-project/modin; pyca/pyopenssl; spulec/freezegun; tkrajina/gpxpy; tornadoweb/tornado; weaveworks/grafanalib
BSD 2-Clause "Simplified" License	madzak/python-json-logger; pyasn1/pyasn1; pygments/pygments; sunpy/sunpy
BSD 3-Clause "New" or "Revised" License	Suor/funcy; alecthomas/voluptuous; andialbrecht/sqlparse; cookiecutter/cookiecutter; dask/dask; django/channels; django/daphne; encode/starlette; gawel/pyquery; gweis/isodate; john-kurkowski/tldextract; lepture/mistune; oauthlib/oauthlib; pallets/click; pallets/flask; pallets/jinja; pallets/markupsafe; pandas-dev/pandas; scrapy/scrapy; theskumar/python-dotenv
GNU General Public License v3.0	Cog-Creators/Red-DiscordBot; adrierverge/yamllint
GNU Lesser General Public License v2.1	chardet/chardet; paramiko/paramiko; pylint-dev/astroid
GNU Lesser General Public License v3.0	Knio/dominate
ISC License	kennethreitz/records
MIT License	amueller/word_cloud; borntyping/python-colorlog; bottlepy/bottle; cantools/cantools; cdgriffith/Box; cknd/stackprinter; conan-io/conan; cool-RR/PySnooper; datamade/usaddress; dbader/schedule; erikrose/parsimonious; facebookresearch/hydra; facelessuser/soupsieve; getnikola/nikola; graphql-python/graphene; hukkin/tomli; jaraco/infect; jawah/charset-normalizer; joke2k/faker; keleshev/schema; life4/textdistance; luozhouyang/python-string-similarity; marshmallow-code/apispec; marshmallow-code/marshmallow; marshmallow-code/webargs; martinblech/xmltodict; matthewwithanm/python-markdownify; mewwts/addict; mido/mido; mozillaz/python-pinyin; msiemens/tinydb; pdfminer/pdfminer; pndurette/gTTS; pudo/dataset; pydantic/pydantic; pyparsing/pyparsing; pytest-dev/iniconfig; python-hyper/h11; python-jsonschema/jsonschema; python-xml/Python-xml; pyupio/safety; pyvista/pyvista; r1chardj0n3s/parse; rsalmei/alive-progress; rubik/radon; rustedpy/result; scanny/python-pptx; seatgeek/thebuzz; sloria/environs; sqlfluff/sqlfluff; termcolor/termcolor; tobymao/sqlglot; tox-dev/pipdeptree; tweepy/tweepy; un33k/python-slugify; vi3k6i5/flashtext
Other	Mimino666/langdetect; PyCQA/flake8; agronholm/exceptiongroup; agronholm/typeguard; aio-libraries/async-timeout; benoitc/gunicorn; cloudpipe/cloudpickle; davidhalter/parso; django-money/django-money; gruns/furl; kurtmckee/feedparser; lincolnloop/python-qr-code; mahmoud/boltons; mahmoud/glob; mozilla/bleach; pexpect/ptyprocess; prettytable/prettytable; pwaller/pyfiglet; pydata/patsy; pydicom/pydicom; python-trio/trio; python/mypy; pyutils/line_profiler; seperman/deepdiff

Table 6: License associated with each repository as of April 8, 2025. All licenses are permissive and allow for public, nonprofit use.

Automated repository installation. The goal of this step is to first, get the installation and testing instructions for a repository, and second, create a Docker image containing the repository with the development environment set up.

We provide the system prompt given to SWE-agent that asks it to install a repository in Figure 11. Each repository installation task is initialized with a clone of the original

repository. No additional steps (e.g. pypi package downloads, conda environment setup) are performed.

We run SWE-agent with `claude-3-5-sonnet-20241022` with a maximum cost limit of \$2 and a maximum call limit of 150. The installation run terminates whenever one of these conditions is met. For every run, we record the interactions. We then manually review the trajectory, identifying the appropriate installation and testing specifications.

Each run incurs an average cost of \$0.72 and an average of 17 steps before SWE-agent issues the `submit` command. The runs typically finish within two minutes. The majority of Python repositories require fewer steps — typically, SWE-agent will view the `CONTRIBUTING.md`, runs the installation command provided verbatim in the text, and then runs `pytest`, showing all tests passing. A minority of repositories will require several steps because additional dependencies must be installed with `apt-get`. The manual review process following this requires 3 to 20 minutes. One author carried out this effort for 125 repositories; we estimate this took a total of 18 human hours to accomplish. In the process of reaching 125 repositories, the author gave up on 17 repositories at the manual review stage.

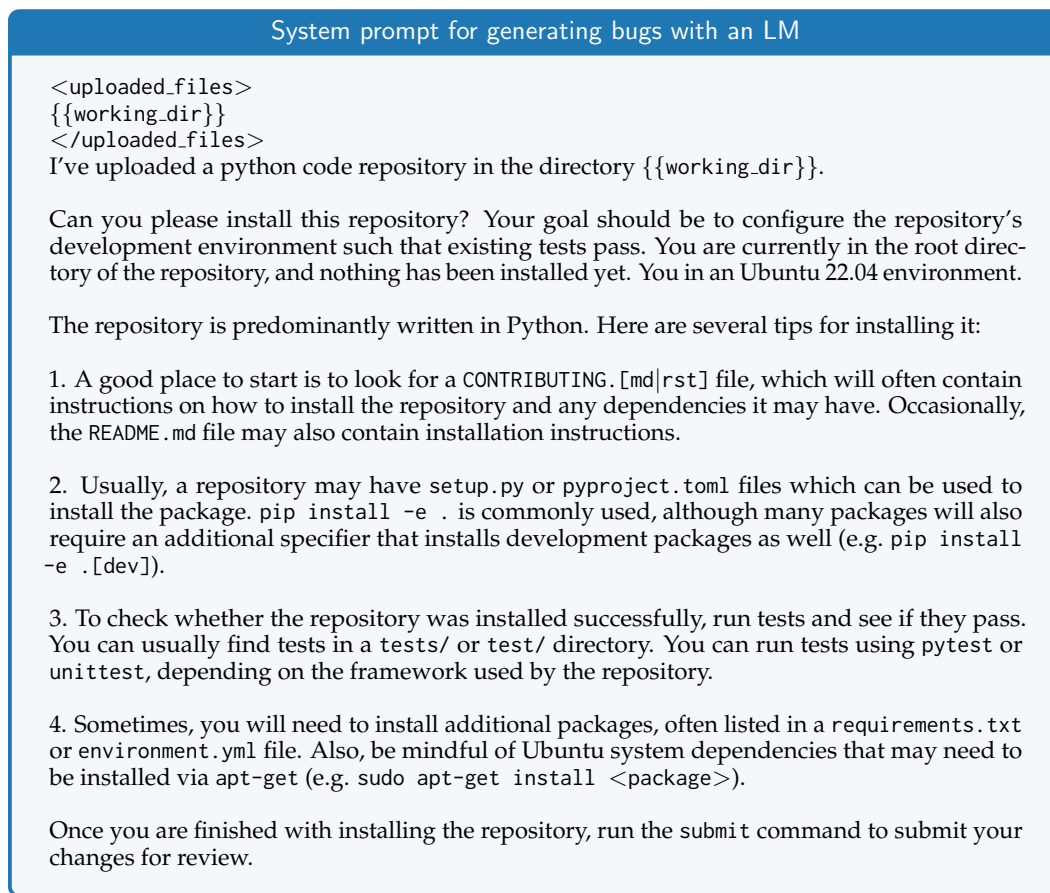


Figure 11: Prompt provided to SWE-agent + an LM asking it to install a repository.

A.3 Validation, Evaluation Harnesses

We adapt SWE-bench’s validation script to convert each bug patch into a SWE-bench style task instance. This step ensures SWE-smith can be run by existing SWE-bench solutions. The conversion involves two steps. First, the bug patch is applied and pushed as a branch to a mirror clone of the repository. Second, we create a SWE-bench style task instance from the bug patch, populating important fields such as Fail-to-Pass and Pass-to-Pass tests with information from the validation logs.

B Bug Generation Strategies

In this section, we review each of the bug generation strategies we employ in depth. While we experimented with several bug generation strategies, the ones we elect to include are those we found to satisfy several desirable properties.

1. The approach works in a codebase-agnostic manner.
2. The approach reliably yields usable task instances (meaning 1+ passing tests break).
3. The approach is controllable; via each strategy's parameters, we can affect the quantity and quality of the generated bugs.

System prompt for generating bugs with an LM

You are a software developer doing chaos monkey testing. Your job is to rewrite a function such that it introduces a logical bug that will break existing unit test(s) in a codebase. To this end, some kinds of bugs you might introduce include:

(Per inference call, only 3 of the following tips are randomly selected and shown)

- Alter calculation order for incorrect results: Rearrange the sequence of operations in a calculation to subtly change the output (e.g., change $(a + b) * c$ to $a + (b * c)$).
- Introduce subtle data transformation errors: Modify data processing logic, such as flipping a sign, truncating a value, or applying the wrong transformation function.
- Change variable assignments to alter computation state: Assign a wrong or outdated value to a variable that affects subsequent logic.
- Mishandle edge cases for specific inputs: Change handling logic to ignore or improperly handle boundary cases, like an empty array or a null input.
- Modify logic in conditionals or loops: Adjust conditions or loop boundaries (e.g., replace \leq with $<$) to change the control flow.
- Introduce off-by-one errors in indices or loop boundaries: Shift an index or iteration boundary by one, such as starting a loop at 1 instead of 0.
- Adjust default values or constants to affect behavior: Change a hardcoded value or default parameter that alters how the function behaves under normal use.
- Reorder operations while maintaining syntax: Rearrange steps in a process so the function produces incorrect intermediate results without breaking the code.
- Swallow exceptions or return defaults silently: Introduce logic that catches an error but doesn't log or handle it properly, leading to silent failures.

Tips about the bug-introducing task:

(At inference time, tips are randomly shuffled)

- It should not cause compilation errors.
- It should not be a syntax error.
- It should be subtle and challenging to detect.
- It should not modify the function signature.
- It should not modify the documentation significantly.
- For longer functions, if there is an opportunity to introduce multiple bugs, please do!" - Please DO NOT INCLUDE COMMENTS IN THE CODE indicating the bug location or the bug itself.

Your answer should be formatted as follows:

Explanation: <explanation>

Bugged Code:

```
““
<bugged.code>
““
```

Figure 12: System prompt provided to an LM to generate bugs by modifying an existing, working function. Text in textcolorredred are not included at the actual prompt.

B.1 Generating with an LM

We describe our workflows for generating bugs with an LM. For each function or class in a codebase, we prompt an LM to generate either a rewrite that introduces bugs or a complete re-implementation from scratch. This strategy is illustrated in Figure 13.

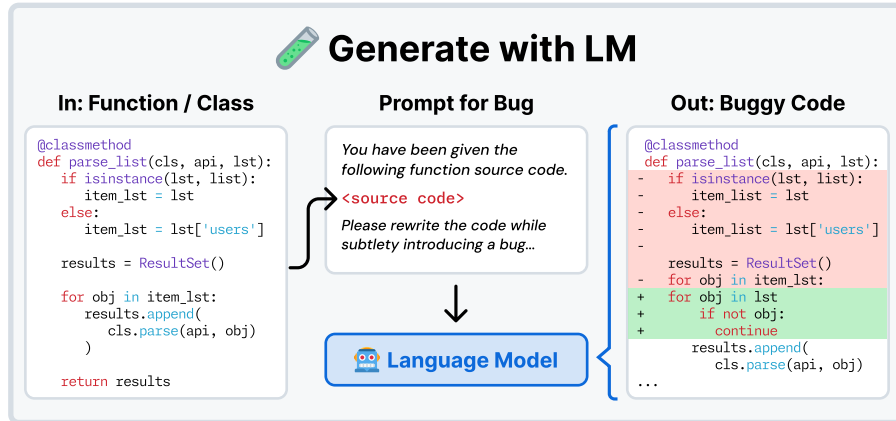


Figure 13: Workflow to generate bugs for a function or class with an LM. We first extract all functions or classes from a codebase, then enumerate across all candidates and prompt the LM to generate either a bug-laced rewrite or a re-implementation.

Modify existing functions. Given a Python codebase, we use the `ast` library to identify all unique functions, excluding any functions found under a testing related directory (e.g. `tests`, `testing`). Next, given a function, the LM is asked to write a new version that introduces logical, runtime bugs. Within the prompt, shown in Figure 12, several suggestions of types of bugs along with a demonstration of a rewrite are provided.

In our experiments, we use OpenAI’s o3 mini model (OpenAI, 2024b) (o3-mini-2025-01-31) as the main base model for bug generation. Based on our empirical observations of an LM’s tendencies, we include several explicit guidelines in the prompt about what the rewrite should not do. Notably, it is important to ask the LM to not generate any inline comments denoting the location of a bug; we observe that without explicitly specifying this, model generation outputs tend to have inline comments pointing out the bug. We also want to avoid the complexities of identifying and removing such comments from a file diff representation. Second, we state that rewrites causing compilation or syntax errors (e.g. undeclared variables, function definition modifications) should be avoided because such bugs are relatively trivial to solve.

We do not experiment extensively with different prompts or generating multiple buggy rewrites per function.

Modify existing classes. This method involves a simple amendment to the function rewriting approach. Instead of identifying unique functions (`ast.FunctionDef`), the codebase traversal logic instead looks for classes (`ast.ClassDef`). Otherwise, all other aspects of the implementation are near identical to function rewriting, with minor changes to the prompt to make bug suggestions and the demonstration more class oriented.

Rewrite existing functions. Instead of providing an LM with the original function, we explore an alternative strategy of asking an LM to re-implement a function from scratch. Similar to above, we again use the `ast` library to identify all unique functions. However, instead of directly asking for a bug, we remove the function’s implementation, then prompt the LM with the entire file containing the function (minus the original implementation). In the task description, we then explicitly ask for the LM to implement the function without changing the function signature.

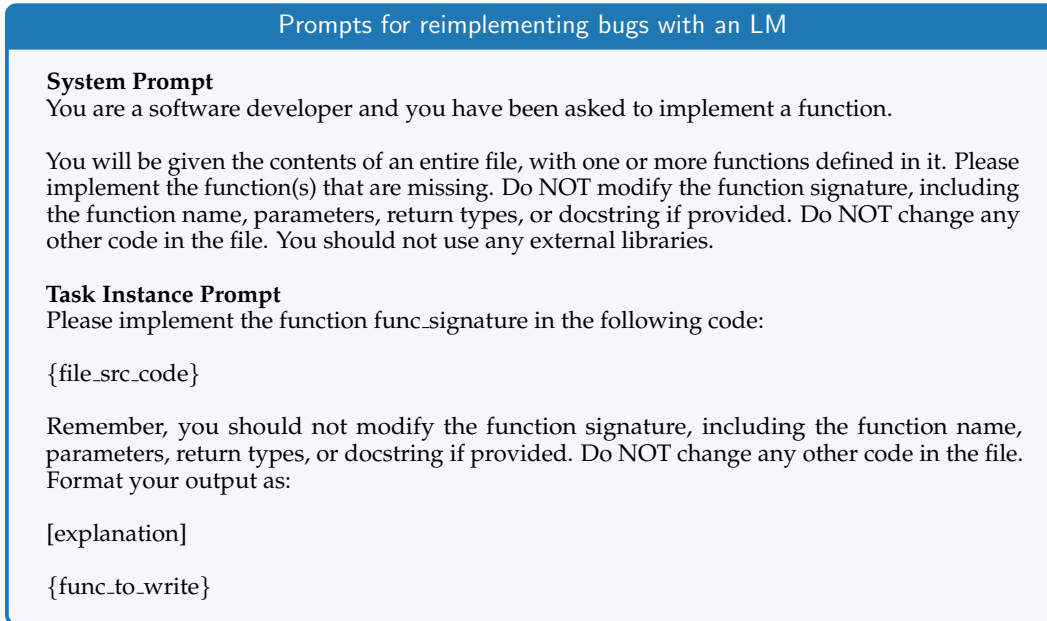


Figure 14: System prompt provided to an LM to generate bugs by re-implementing an existing target function. `file_src.code` refers to the original source file minus the target function’s original implementation. `func_to_write` refers to the signature and docstring of the target function.

B.2 Procedural Modification

We explore a zero-cost approach to create bugs by performing random modifications to the ast representation of a function or class. A “procedural modification” refers to a function that takes in an ast and applies a fixed transformation to it, such as removing a loop or swapping the blocks of an if/else clause. This strategy is illustrated in Figure 15.

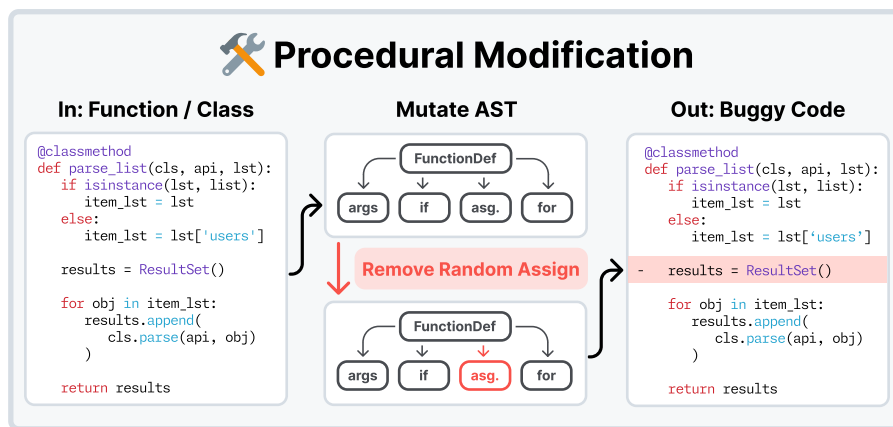


Figure 15: Workflow to generate bugs via procedural modifications. Per function/class, the source code is first converted into an ast. The modification then mutates the ast (e.g. removes an assignment statement). The modified ast is then converted back into source code with a bug of the modification type introduced.

Similar to the workflow for generating bugs with an LM, we first identify all functions or classes in a repository. Per procedural modification, we first impose a set of criteria that filters out any candidates for which the modification would be impossible. For instance, if the procedural modification removes a random conditional from a function, the

modification’s criteria will filter out any candidates that are not functions or do not have a conditional. For the remaining candidates, the procedural modification is applied with controlled likelihood, where likelihood is a fraction indicating how often the procedural modification is applied within a candidate. For example, if the procedural modification removes a random function with a likelihood of 0.5, then for every conditional declared within the function, there is a 50% chance it gets removed. We introduce likelihood so procedural modifications do not lead to changes that are too difficult. Finally, the modified ast is converted back into source code.

Table 7 is a complete list of filtering criteria that is used for any procedural modification. For the `filter_min_complexity` and `filter_max_complexity` criteria, we define a simple definition of “complexity” as a sum of the number of conditional blocks, loops, boolean operators, exception handling blocks, and comparison operators in a function. The purpose of `filter_min_complexity` is to remove both simple, uninteresting functions (e.g. `getter`, `setter` methods) from consideration. `filter_max_complexity` is occasionally used to avoid changing long, monolithic functions.

Index	Criteria	Description
1	<code>filter_functions</code>	Is the ast a function definition
2	<code>filter_classes</code>	Is the ast a class definition
3	<code>filter_classes_has_base</code>	Is the ast a class definition with parents
4	<code>filter_loops</code>	Does the ast contain a For or While loop?
5	<code>filter_conditionals</code>	Does the ast contain a conditional block?
6	<code>filter_assignments</code>	Is the ast a function def. with assignments?
7	<code>filter_wrappers</code>	Does the ast contain try or with blocks?
8	<code>filter_if_else</code>	Does the ast contain an if-else block?
9	<code>filter_operators</code>	Does the ast contain binary, boolean operators?
10	<code>filter_min_complexity</code>	Is the ast \geq a complexity score?
11	<code>filter_max_complexity</code>	Is the ast \leq a complexity score?

Table 7: Pool of criteria used to filter for functions or classes with specific properties. Per procedural modification, a subset of these criteria is first used to filter functions and/or classes from a codebase. The modification is then run on the remainder.

Table 8 is an exhaustive list of all procedural modifications used to create bugs in a codebase.

Procedural Modification		Criteria	Description
Class	Remove Functions	2, 10	Removes method(s) + reference(s).
	Remove Parent	3, 10	Removes base class from class header.
	Shuffle Methods	2, 10	Shuffles method definitions in a class.
Control Flow	Invert If/Else	8	Inverts the if-else bodies of a condition.
	Shuffle Lines	11, 12	Shuffles the lines of a function.
Expressions	Change Constants	1, 9, 10	± 1 to a constant numeric value.
	Break Chains	1, 9, 10	Removes operator(s), operator(s).
	Swap Operands	1, 9, 10	Mixes order of operands.
	Change Operator	1, 9, 10	Changes operator(s) (e.g. + to -).
Removal	Loops	1, 4, 10	Remove loops (e.g. for, while).
	Conditionals	1, 5, 10	Remove conditionals (if).
	Assignments	1, 6, 10	Remove assignment statements.
	Wrappers	1, 7, 10	Remove exception (try), context (with).

Table 8: The 13 procedural modification techniques we use to create bugs in a codebase. The “Criteria” column contains indices referencing the corresponding filter defined in Table 7. There are four informal categories — Class, Control Flow, Expressions, Removal — which indicates the general type of modification being made.

B.3 Combine Bug Patches

We discuss the two strategies we use to combine bug patches from the same file or the same module. In practice, we combine LM and procedurally generated bugs that have been validated successfully as usable task instances.

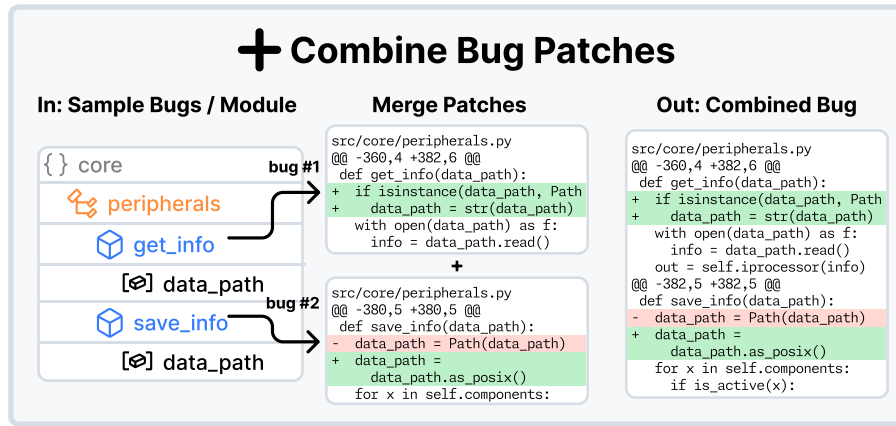


Figure 16: Workflow to generate bugs by combining bug patches. We take n patches (generated using an LM or procedural modification), then sequentially apply each bug patch to the codebase. If all individual patches apply successfully, we save the resulting single patch which now represents all n bugs combined.

From the same file. If two or more functions are defined within a single file, this strategy merges the function-level bug patches together. Given n function-level bugs and k as the number of bugs to combine, there are $\binom{n}{k}$ unique file-level candidate bug patches, which can be a large search space to cover. To make the search space tractable, ensure no single function-level bug is repeatedly used, and generate instances that reliably have 1+ Fail to Pass tests, we implement the following approach described in Algorithm 1.

Algorithm 1 Combine multiple patches from the same file.

Require: *codebase*, *bugs*; *num_bugs*, *limit_per_file*; *max_combos*

Ensure: $min_bugs \geq 2$;

$max_bugs \geq min_bugs$;

procedure COMBINEFILEBUGS

for each *file* **in** *codebase* **do**

file_bugs \leftarrow bugs that apply to *file*

combinations \leftarrow get_combos(*file_bugs*, *num_bugs*, *max_combos*)

for each *combo* **in** *combinations* **do**

 Apply *combo* to *codebase*

if success **then**

 Save *combo* to disk

if *limit_per_file* reached **then**

break

end if

combinations \leftarrow [*c* for *c* in *combinations* if $c \cap combo = \emptyset$]

end if

end for

end for

end procedure

For each file in a codebase, we first identify the function-level bugs (or bug patches) that edit that file. The pool of bugs we draw from have been *validated*, meaning we have already ensured there is 1+ Fail to Pass test(s) associated with the bug. From these pool of *file_bugs*,

the `get_combos` function then generates up to `max_combos` sets of bugs, where the size of each set is `num_bugs`. For each combo, or set of bugs, the bugs are applied to the codebase one by one. If all patches are successfully combined, this means they were successfully merged, and the merged patch, which consists of multiple function-level bugs, is saved and re-validated as a single bug. Merging patches occasionally fails if there is an overlapping conflict between two files, akin to a merge conflict with `git`; this usually happens when a function is declared within another. To ensure a function-level bug is only used once, any remaining bug sets in combinations using any patch in combo are removed.

The `limit_per_file` and `max_combos` parameters prevent any one file from being over-represented and constrains an otherwise combinatorial large search space. We run this algorithm across all codebase files, typically setting `num_bugs= [2,4]`, `limit_per_file= 3`, `max_combos= 40`. Decreasing `num_bugs` or increasing the other three parameters improves the yield.

From the same module. There are several ways one could imagine composing function-level bugs from multiple bugs, such as combining those that break the same test or have a programmatic relationship (e.g. function `a` calls function `b`). We found a relatively straightforward and effective approach to be combining files that edit the same “module”. By “module” we are referring to a subdirectory within the source code (e.g. `sklearn/feature_extraction`, `astropy/convolution`). Out of all SWE-bench instances that edit 2+ files, 75% modify files within the same submodule, suggesting a high degree of intra-module code changes. The implementation for our approach is described in Algorithm 2

Algorithm 2 Combine multiple patches from the same module.

Require: `bugs`; `num_bugs`; `limit_per_module`; `max_combos`; `depth`

Ensure: `num_bugs` ≥ 2 ;

```

procedure COMBINEMODULEBUGS
  map_path_to_bugs  $\leftarrow \{\}$ 
  for each bug in bugs do
    path  $\leftarrow$  get_path_from(bug)
    map_path_to_patches[path]  $\leftarrow$  [bug]
  end for
  Collapse nested paths based on depth
  for all (path, patches) in map_path_to_patches do
    combinations  $\leftarrow$  get_combos(patches, num_bugs, max_combos)
    for each combo in combinations do
      Apply combo to codebase
      if success and num_files.changed(combo)  $\geq$  2 then
        Save combo to disk
        if limit_per_module reached then
          break
        end if
        combinations  $\leftarrow$  [c for c in combinations if  $c \cap$  combo =  $\emptyset$ ]
      end if
    end for
  end for
end procedure

```

The implementation this approach is similar to Algorithm 1 with two key changes. First, we do not do file-by-file or folder-by-folder traversal. Instead, using the diff patches, we create a dictionary `map_path_to_bugs` that mimics the file structure of a codebase. For example, if bug modifies path `a/b/c/d.py`, it is represented as `map_path_to_bugs[a][b][c][d.py] = [bug]`. Additional bugs that modify the same path are appended to the list. Since every bug is a function-level bug, there will never be a bug registered in multiple lists. We then “collapse” up to depth indices. So for instance, at `depth = 3`, the above data structure is collapsed into `map_path_to_bugs[a/b/c][d.py] = [bug]`. Finally, any nested dictionaries are collapsed into a single list of patches (e.g. `map_path_to_bugs[a/b/c] = [bug]`). Mirroring the procedure in Algorithm 1, we then iterate across this dictionary’s values (lists of bugs).

Second, we only save patches that modify 2+ files; aggregate bugs (represented by combo modifying a single file are not considered).

Again, we run this strategy across all 100 repositories, with parameters `num_bugs= [2,5]`, `limit_per_module= 10`, `max_combos= 100`, and `depth= 2`. Reducing `num_bugs`, `depth` and increasing the other parameters yields more bugs. We choose a depth of 2 because empirically, we find that meaningful modules are usually declared as immediate sub-folders of the main source code folder (e.g. in `sklearn/feature_extraction`, `sklearn` is the source code folder while `feature_extraction` is the module). A shallower depth leads to less meaningful groupings, while yield decreases significantly for every increased level of depth, particularly for smaller repositories.

B.4 Pull Request Mirroring

We finally discuss the fourth and last strategy for generating bugs - mirroring real world pull requests (PR). We visualize this process in Figure 17.

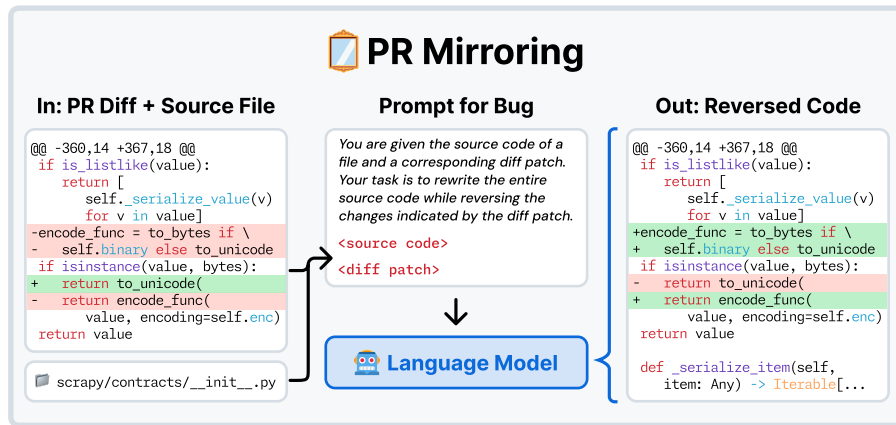


Figure 17: Workflow to generate bugs by reverting changes made in the diff patch correspond to a real GitHub pull request (PR). Given the patch and the files modified by the patch, we prompt the LM to generate a complete rewrite of each file that *reverses* the changes made in the PR. The changes are applied to the codebase, and we extract the patch, which now captures the reversal of the PR changes.

Why use an LM? When we initially implemented this approach, we attempted to directly perform a `git apply --reverse [patch]` on the codebase. However, for the large majority of patches, this fails. We performed troubleshooting by inspecting 100 PR patches on the `sqlfluff/sqlfluff` repository, leading us to two observations. First, the majority of these PRs reflect changes that remain present in the codebase today, so such a strategy would not be fruitless. Second, the reason these PR patches don't apply is usually *not* because the change introduced by the PR is no longer present; rather, it is more often the case that other parts of the file were modified, causing the exact location (e.g. lines, file) of the relevant code to shift over time. Therefore, we employ LMs to perform patch reversal.

Description of method. We follow SWE-bench's methodology for crawling PRs created January 1st, 2023 and onwards, with minor and arbitrary exceptions for some repositories where we crawl older PRs as well. Per PR, we iterate across the file(s) changed by the patch. Per file, we prompt an LM with the file-specific changes from the patch along with the file's source code in the current state of the repository (*not* the repository's state corresponding to when the PR was applied, referred to as the `base_commit` in SWE-bench). The LM is asked to generate a rewrite of the file that reverts the changes reflected in the PR. We aggregate the changes across all file(s) into a single patch.

Because we are interested in problems that our expert trajectory generation method (SWE-agent + Claude 3.7 Sonnet) has a chance of solving, we do not attempt to reproduce PRs

that change more than 8 files. This constraint is imposed because no SWE-bench instance that edits more than 6 files has ever been solved (Jimenez et al., 2024a).

How well does PR mirroring work? We scrape the PRs corresponding to 100 randomly selected SWE-bench task instances from the `django/django` GitHub repository and attempt to recreate these task instances with SWE-smith’s collection process. We successfully recovered 92 of 100 task instances. Of these, 84 break identical F2P test(s), with the remaining 8 breaking a subset because some tests were removed over time. This sanity check gives us confidence that the PR mirroring strategy lives up to its name.

Comparison to SWE-bench. This approach has several benefits and drawbacks compared to SWE-bench’s collection pipeline. First, it removes the need to create instance-specific Docker images — all PRs are mirrored against the same version of a repository. This also implies that there is no need to write installation specifications for past versions of a repository, which is typically the most laborious step in task construction with SWE-bench. Finally, this strategy also allows us to loosen the requirements on what PRs we attempt to convert into a task instance. In SWE-bench, the core requirements for what PRs to attempt to convert into a task instance include:

1. It must edit 1+ code files (e.g. not just `.md`, `.rst` files).
2. It must reference 1+ GitHub issues, which serves as the problem statement.
3. It must edit 1+ testing related files (1+ files with a test-adjacent keyword in it).

With this collection strategy and SWE-smith’s focus on training data, the second and third requirements are no longer necessary. If there is no associated issue, issue text can simply be generated. If the patch does not contain any testing related changes, this is tolerable, as the validation stage will determine whether the PR breaks any tests. With these considerations, we purport that SWE-smith’s PR mirroring strategy can re-purpose a higher percentage of real world code changes for training purposes.

The main downside is that the rest of the repository is out of sync with the state of the codebase when the PR was applied. As a result, it’s possible that changes in the behavior of the rest of the codebase may affect the issue’s reproducibility or the accuracy of the issue description (e.g. line numbers referenced in the issue text are likely somewhat off with respect to the codebase). However, a simple mitigation for this is to create a Docker image for a repository at an earlier commit that’s closer to the original creation date of the issue. While we do not carry out a targeted experiment, we hypothesize that using SWE-smith, we would be able to reproduce SWE-bench entirely with 10x less human hours with an estimated $2294 \times \$0.055 = \126.17 in costs.

C Dataset Statistics

We present additional breakdowns and analyses of the SWE-smith dataset, focusing on the kinds of repositories and bugs that are represented.

Repository categorization. We present an exhaustive list of repositories used in SWE-smith in Table 10. We categorize the repositories into seven general buckets: Data Parsing and Transformation (39), Web & API Development (11), Code Quality & Testing (12), Visualization & Presentation (8), System Tools & Protocols (17), Natural Language Processing (7), and Miscellaneous (6). The categorizations were performed by first, determining an appropriate set of categories based on manual inspection supported by the descriptions and GitHub topics associated with each repository. After settling upon the buckets, we asked GPT-4o to provide a label based on the repository’s metadata and README dump. SWE-smith represents a wider and more variegated coverage of software tools and applications compared to any prior works.

Repository	Description
<i>Code Quality and Testing</i>	
PyCQA/flake8	flake8 is a python tool that glues together pycodestyle, pyflakes, mccabe, and third-party plugins to check the style and quality of some python code.
Suor/funcy	A fancy and practical functional tools
adrienverge/yamllint	A linter for YAML files.
agronholm/typeguard	Run-time type checker for Python
cknd/stackprinter	Debugging-friendly exceptions for Python
cool-RR/PySnooper	Never use print for debugging again
getmoto/moto	A library that allows you to easily mock out tests based on AWS infrastructure.
pylint-dev/astroid	A common base representation of python source code for pylint and other projects
pytest-dev/iniconfig	None
python/mypy	Optional static typing for Python
pyupio/safety	Safety checks Python dependencies for known security vulnerabilities and suggests the proper remediations for vulnerabilities detected.
pyutils/line_profiler	Line-by-line profiling for Python
rubik/radon	Various code metrics for Python code
spulec/freezegun	Let your Python tests travel through time
sqlfluff/sqlfluff	A modular SQL linter and auto-formatter with support for multiple dialects and templated code.
<i>Data Parsing and Transformation</i>	
alecthomias/voluptuous	CONTRIBUTIONS ONLY: Voluptuous, despite the name, is a Python data validation library.
andialbrecht/sqlparse	A non-validating SQL parser module for Python
bury/python-readability	fast python port of arc90's readability tool, updated to match latest readability.js!
burnash/gspread	Google Sheets Python API
chardet/chardet	Python character encoding detector
cloudpipe/cloudpickle	Extended pickling support for Python objects
dask/dask	Parallel computing with task scheduling
datamade/usaddress	:us: a python library for parsing unstructured United States address strings into address components
davidhalter/parso	A Python Parser
erikrose/parsimonious	The fastest pure-Python PEG parser I can muster
facelessuser/soupsieve	A modern CSS selector implementation for BeautifulSoup
gawel/pyquery	A jquery-like library for python
google/textfsm	Python module for parsing semi-structured text into python tables.
gruns/furl	URL parsing and manipulation made easy.
gweis/isodate	ISO 8601 date/time parser
hukkin/tomli	A lil' TOML parser
jawah/charset_normalizer	Truly universal encoding detector in pure Python
john-kurkowski/tldextract	Accurately separates a URL's subdomain, domain, and public suffix, using the Public Suffix List (PSL).
joke2k/faker	Faker is a Python package that generates fake data for you.
jsvine/pdfplumber	Plumb a PDF for detailed information about each char, rect-angle, line, et cetera — and easily extract text and tables.

Continued on next page

Repository	Description
kayak/ypika	PyPika is a python SQL query builder that exposes the full richness of the SQL language using a syntax that reflects the resulting query. PyPika excels at all sorts of SQL queries but is especially useful for data analysis.
keleshev/schema	Schema validation just got Pythonic
kennethreitz/records	SQL for Humans™
kurtmckee/feedparser	Parse feeds in Python
lepture/mistune	A fast yet powerful Python Markdown parser with renderers and plugins.
madzak/python-json-logger	Json Formatter for the standard python logger
mahmoud/glob	Python's nested data operator (and CLI), for all your declarative restructuring needs. Got data? Glom it!
marshmallow-code/marshmallow	A lightweight library for converting complex objects to and from simple Python datatypes.
martinblech/xmltodict	Python module that makes working with XML feel like you are working with JSON
matthewwithanm/python-markdownify	Convert HTML to Markdown
mewwts/addict	The Python Dict that's better than heroin.
mido/mido	MIDI Objects for Python
modin-project/modin	Modin: Scale your Pandas workflows by changing a single line of code
mozilla/bleach	Bleach is an allowed-list-based HTML sanitizing library that escapes or strips markup and attributes
msiemens/tinydb	TinyDB is a lightweight document oriented database optimized for your happiness :)
pandas-dev/pandas	Flexible and powerful data analysis / manipulation library for Python, providing labeled data structures similar to R data.frame objects, statistical functions, and much more
pdfminer/pdfminer.six	Community maintained fork of pdfminer - we fathom PDF
pudo/dataset	Easy-to-use data handling for SQL data stores with support for implicit table creation, bulk loading, and transactions.
pydantic/pydantic	Data validation using Python type hints
pydata/patsy	Describing statistical models in Python using symbolic formulas
pydicom/pydicom	Read, modify and write DICOM files with python code
pygments/pygments	Pygments is a generic syntax highlighter written in Python
pyparsing/pyparsing	Python library for creating PEG parsers
python-jsonschema/jsonschema	An implementation of the JSON Schema specification for Python
python-openxml/python-docx	Create and modify Word documents with Python
r1chardj0n3s/parse	Parse strings using a specification based on the Python format() syntax.
scanny/python-pptx	Create Open XML PowerPoint documents in Python
scrapy/scrapy	Scrapy, a fast high-level web crawling & scraping framework for Python.
seperman/deepdiff	DeepDiff: Deep Difference and search of any Python object/data. DeepHash: Hash of any object based on its contents. Delta: Use deltas to reconstruct objects by adding deltas together.
sloria/environs	simplified environment variable parsing
sunpy/sunpy	SunPy - Python for Solar Physics
tkrajina/gpxpy	gpx-py is a python GPX parser. GPX (GPS eXchange Format) is an XML based file format for GPS tracks.

Continued on next page

Repository	Description
tobymao/sqlglot un33k/python-slugify	Python SQL Parser and Transpiler Returns unicode slugs
<i>Machine Learning and AI</i>	
facebookresearch/fvcore	Collection of common code that's shared among different research projects in FAIR computer vision team.
facebookresearch/hydra	Hydra is a framework for elegantly configuring complex applications
HIPS/autograd	Efficiently computes derivatives of NumPy code.
iterative/dvc	Data Versioning and ML Experiments
jaraco/inflect	Correctly generate plurals, ordinals, indefinite articles; convert numbers to words
life4/textdistance	Compute distance between sequences. 30+ algorithms, pure python implementation, common interface, optional external libs usage.
luozhouyang/python-string-similarity	A library implementing different string similarity and distance measures using Python.
Mimino666/langdetect	Port of Google's language-detection library to Python.
mozillazg/python-pinyin	汉字转拼音(pypinyin)
pndurette/gTTS	Python library and CLI tool to interface with Google Translate's text-to-speech API
Project-MONAI/MONAI	AI Toolkit for Healthcare Imaging
seatgeek/thefuzz	Fuzzy String Matching in Python
vi3k6i5/flashtext	Extract Keywords from sentence or Replace keywords in sentences.
<i>System Tools and Protocols</i>	
agronholm/exceptiongroup	Backport of PEP 654 (exception groups)
aio-lib/async-timeout	asyncio-compatible timeout class
arrow-py/arrow	Better dates & times for Python
borntyping/python-colorlog	A colored formatter for the python logging module
cantools/cantools	CAN bus tools.
conan-io/conan	Conan - The open-source C and C++ package manager
cookiecutter/cookiecutter	A cross-platform command-line utility that creates projects from cookiecutters (project templates), e.g. Python package projects, C projects.
dbader/schedule	Python job scheduling for humans.
gruns/icecream	Never use print() to debug again.
jd/tenacity	Retrying library for Python
mahmoud/boltions	Like builtins, but boltions. 250+ constructs, recipes, and snippets which extend (and rely on nothing but) the Python standard library. Nothing like Michael Bolton.
oauthlib/oauthlib	A generic, spec-compliant, thorough implementation of the OAuth request-signing logic
pallets/click	Python composable command line interface toolkit
paramiko/paramiko	The leading native Python SSHv2 protocol library.
pexpect/ptyprocess	Run a subprocess in a pseudo terminal
pyasn1/pyasn1	Generic ASN.1 library for Python
pyca/pyopenssl	A Python wrapper around the OpenSSL library
python-hyper/h11	A pure-Python, bring-your-own-I/O implementation of HTTP/1.1
python-trio/trio	Trio - a friendly Python library for async concurrency and I/O
rustedpy/result	NOT MAINTAINED - A simple Rust like Result type for Python 3. Fully type annotated.

Continued on next page

Repository	Description
termcolor/termcolor theskumar/python-dotenv tox-dev/pipdeptree	ANSI color formatting for output in terminal Reads key-value pairs from a .env file and can set them as environment variables. It helps in developing applications following the 12-factor principles. A command line utility to display dependency tree of the installed Python packages
<i>Visualization and Presentation</i>	
amueller/word_cloud lincolnloop/python-qrcode prettytable/prettytable pwaller/pyfiglet rsalmei/alive-progress weaveworks/grafanalib	A little word cloud generator in Python Python QR Code image generator Display tabular data in a visually appealing ASCII table format An implementation of figlet written in Python A new kind of Progress Bar, with real-time throughput, ETA, and very cool animations! Python library for building Grafana dashboards
<i>Web and API Development</i>	
Cog-Creators/Red-DiscordBot Knio/dominate alanjds/drf-nested-routers benoitc/gunicorn bottlepy/bottle django-money/django-money django/channels django/daphne encode/starlette getnikola/nikola graphql-python/graphene marshmallow-code/apispec marshmallow-code/webargs pallets/jinja pallets/markupsafe tornadoweb/tornado tweepy/tweepy	A multi-function Discord bot Dominate is a Python library for creating and manipulating HTML documents using an elegant DOM API. It allows you to write HTML pages in pure Python very concisely, which eliminate the need to learn another template language, and to take advantage of the more powerful features of Python. Nested Routers for Django Rest Framework gunicorn 'Green Unicorn' is a WSGI HTTP Server for UNIX, fast clients and sleepy applications. bottle.py is a fast and simple micro-framework for python web-applications. Money fields for Django forms and models. Developer-friendly asynchrony for Django Django Channels HTTP/WebSocket server The little ASGI framework that shines. A static website and blog generator GraphQL framework for Python A pluggable API specification generator. Currently supports the OpenAPI Specification (f.k.a. the Swagger specification).. A friendly library for parsing HTTP request arguments, with built-in support for popular web frameworks, including Flask, Django, Bottle, Tornado, Pyramid, webapp2, Falcon, and aiohttp. A very fast and expressive template engine. Safely add untrusted strings to HTML/XML markup. Tornado is a Python web framework and asynchronous networking library, originally developed at FriendFeed. Twitter for Python!

Table 10: List of all GitHub repositories represented in SWE-smith and their corresponding descriptions. We identify 7 categories generally characterizing the range of utilities for the selected codebases.

C.1 Bug Generation Statistics

We provide extensive details about different aspects of each of the bug generation strategies, including the yield rates, labor/monetary costs, and dataset characterizations.

Yield rates. In Table 11, we provide the yield rates for each bug generation method across all repositories in SWE-smith. In general, we find that the PR Mirroring has the lowest yield rate at 13.18% (although this rate is somewhat higher than SWE-bench’s yield rate of $2294/93139 = 2.46\%$). For using LMs to generated bugs, modifying functions to introduce bugs intentionally has a higher yield than asking LMs to perform a best-effort rewrite. The efficacy of Procedural Modifications varies by strategy. For instance, shuffling the functions declared in a class only breaks existing test(s) 1.93% of the time, but inverting a conditional will lead to a task instance for 47.04% of modifications. Finally, combining bug patches has an extremely high yield rate - this is to be expected because we only attempt to combine bug patches that have been validated as usable task instances breaking 1+ tests.

Strategy	# Repos	# Candidates	# Instances	Yield Rate
Combine (file)	124	6020	5865	97.43%
Combine (module)	65	4396	4227	96.16%
LM (Modify)	108	31950	17887	55.98%
LM (Rewrite)	128	11908	4173	35.04%
PR Mirroring	108	6934	2344	33.8%
Procedural (Class Rm Base)	103	1401	463	33.05%
Procedural (Class Rm Funcs)	103	2506	1180	47.09%
Procedural (Class Shuffle Funcs)	103	2504	47	1.88%
Procedural (Ctrl Invert If)	105	4695	2321	49.44%
Procedural (Ctrl Shuffle)	104	9055	4015	44.34%
Procedural (Op Break Chains)	71	747	225	30.12%
Procedural (Op Change Const)	77	723	257	35.55%
Procedural (Op Change)	81	1507	450	29.86%
Procedural (Op Swap)	87	2141	483	22.56%
Procedural (Remove Assign)	121	5470	2661	48.65%
Procedural (Remove Cond)	120	5288	2311	43.7%
Procedural (Remove Loop)	110	1945	860	44.22%
Procedural (Remove Wrapper)	80	884	368	41.63%
All	129	100074	50137	50.1%

Table 11: Yield rates for different bug generation strategies covered in Section B. We show the number of repositories that each strategy was run on, the number of bug candidates generated by each strategy, and the number of instances, or the number of candidates that were validated to have 1+ Fail to Pass test. The yield rate for

The number of repositories captured by each bug generation technique varies due to each strategy’s specific preconditions, which at times may not be effective for some repositories. For instance, the *Procedural (Class *)* set of methods only mutates Python classes. This strategy is fruitless for the minority of SWE-smith repositories that do not define any classes. The *Procedural (Op Break Chains)* method randomly removes operations and operands from expressions with two or more operations (e.g. $a + b + c \rightarrow a + b$) — such expressions are not always present in SWE-smith repositories.

The collective yield rate across SWE-smith’s bug generation strategies is significantly higher than SWE-bench’s collection strategy.

The yield rate also varies with respect to the repository it is being applied to. We provide a summary of yield rates by repository in Table 12. We generally observe that lower test coverage correlates with a lower yield rate.

Dataset characterizations. In Table 13, we provide statistics about the validated task instances produced by different bug generation strategies. Our work’s LM-based strategies

Yield Rate	# of Repositories
0-25%	10
25-50%	31
50-75%	60
75-100%	27

Table 12: Yield rates for different repositories represented in SWE-smith.

rewrite one function in one file. Procedural modifications will also only change one file, but depending on the strategy, 1+ functions or classes may be changed. Combining multiple patches from the same file always produces a patch with 2+ functions edited. Combining across modules produces a patch with 2+ files edited. The targeted nature of each of the bug creation strategies is reflected in the typical number of functions and files that the bugs produced by each strategy edits.

Strategy	# Instances	# F2P	Δ Lines	Δ Functions	Δ Files
Combine	10092	15 (5-48)	19 (12-36)	2 (2-3)	1 (1-2)
LM	22060	4 (1-17)	6 (3-15)	1 (1-1)	1 (1-1)
PR Mirroring	2344	3 (1-14)	20 (8-55)	2 (2-4)	1 (1-2)
Procedural	15641	7 (2-32)	7 (5-15)	1 (1-1)	1 (1-1)

Table 13: Statistics for different attributes of a SWE-smith task instance across different bug generation strategies, reported as *median (IQR)*, where IQR represents the inter-quartile range (25th–75th percentile).

In Figure 18, we show the distributions for different attributes of SWE-smith compared to other SWE-bench style datasets. Compared to prior works, there is a much higher proportion of task instances with more than one Fail-to-Pass test. For any one repository, we find that SWE-smith task instances collectively cause failures for a much higher percentage of the testing suit than other datasets; a potential benefit of this is that training on SWE-smith based trajectories may expose models to a much broader set of functionalities in a codebase. The number of lines and files edited by SWE-smith task instances is highly similar to the trend lines for SWE-bench Verified.




Figure 18: Comparison of cumulative distributions for Fail-to-Pass tests along with the lines and files edited by the gold patch across SWE-smith and four SWE-bench style datasets.

We note that unlike other datasets, the trend line of SWE-smith task instances is “adjustable”. In other words, the Figure 18 distributions are a capture of the task instances provided in this release of SWE-smith. However, because of SWE-smith’s flexible bug creation techniques, the distribution can be “shaped” if needed. For instance, generating more task instances using the bug patch combination method would shift all three curves in Figure 18. We make this point to highlight the fact that the attributes of SWE-bench task instances are, in a sense, constrained by real world software development behavior. On the other hand, SWE-smith

can be used to break tests and code that may not be reflected at all in any existing pull request. In this sense, we argue that LMs trained on SWE-smith have better “exposure” to a codebase compared to exclusively training on pull requests.

Continuation of scaling execution environments. The validation and evaluation procedures for SWE-smith deviate slightly from SWE-bench’s harnesses. The main reasons for these differences can largely be attributed to the granularity of installation specifications. In SWE-bench, each task instance corresponds to a unique base commit, with additional `version` and `environment_setup_commit` keys needed as indirection for mapping an instance to the correct set of installation and testing instructions. Across time, the continuous evolution of a repository and its dependencies make for an incredibly high degree of variability in how a repository should be installed correctly. To solve this variability, the community has resorted to creating an image per task instance, as done in [Chowdhury et al. \(2024\)](#). Therefore, for 2294 SWE-bench task instances, there are 2294 unique Docker images, each at a size of at least several gigabytes (\sim 5-6 GBs).

On the other hand, the simplicity and scalability of SWE-smith’s design allows one to support many task instances with comparatively much fewer Docker images. As mentioned above, installation and testing procedures are (repository, commit) specific. Therefore, when bugs are generated from each (repository, commit), all bugs can be reproduced and tested successfully from the same Docker image. In other words, if I generate 100 bugs for a repository at some commit, instead of 100 Docker images, only a single Docker image is required to run inference on any of the 100 task instances.

This design is what enables SWE-smith to be significantly more space-efficient than SWE-bench. Based on the publicly released images, for SWE-bench’s 2294 task instances, 1.2 TBs of storage are required to download all Docker images locally. For SWE-bench Multimodal’s 517 task instances, 1.2 TBs are required. The higher per-instance Docker image size for SWE-bench Multimodal is due to how JavaScript dependency management tools (e.g. `npm`) require more storage compared to equivalent Python infrastructure (e.g. `pypi`). [Pan et al. \(2024\)](#) states that each image for the 2438 instances an average of 2.6GB, totaling 6 TB of storage total. Such a storage requirement can be a significant barrier for academic practitioners.

On the other hand, with more than 20x the number of bugs, SWE-smith requires only 125 Docker images total, corresponding to the number of unique (repository, commit) pairs (in this work, for each repository, we only determine installation and test specifications for one commit). The 125 images require a total of 290.54 GBs. In summary, compared to SWE-bench’s task collection strategy, SWE-smith’s design makes it easier to not only create task instances, but also train on them as well.

C.2 Case Study: SWE-bench & SWE-smith

To better understand the differences between the SWE-bench and SWE-smith collection strategies, we perform SWE-smith collection on the `pallets/flask` GitHub repository, one of the 12 test split repositories from the original SWE-bench benchmark. We review the steps covered in Section 2.1 applied to `pallets/flask` in detail. First, we defined the installation and testing specifications for the `pallets/flask` repository at commit `bc09840`. Next, we apply the LM modification bug generation strategy to this version of the repository, generating 267 unique bugs.

We observe several differences. First, *the SWE-smith collection strategy yields a much higher number of bugs outright*. From SWE-bench, 11 task instances are from the `pallets/flask` repository. The task instances were originally filtered from 2434 pull requests (PRs), with 107 satisfying SWE-bench’s filtering criteria of (1) being linked to one or more issues and (2) featuring 1+ new tests. Out of these 107, the 11 (0.45% of 2434) task instances represent the proportion of PRs that execution environments could be successfully constructed for. On the other hand, running the function-level rewriting strategy for bug generation originally yielded 402 candidates, of which 267 were determined to be valid task instances.

Second, *SWE-smith requires significantly less human effort while only incurring minor costs*. Collecting the 11 `pallets/flask` task instances (steps include scraping PRs, determining

repository versions across time, defining version-specific installation/test specifications, running execution-based validation multiple times) took an estimated 38 hours worth of human labor. On the contrary, defining installation and testing specifications for the latest commit of pallets/flasks took 10 minutes. The subsequent function-level rewriting strategy for bugs took 23 minutes to run, incurring a total cost of just \$2.47 (\sim \$0.00613 per instance). The final execution-based validation step that filters out $402 - 267 = 135$ unqualified bug candidates ran in 14 minutes. Since both the bug and problem statement generation strategies are repository agnostic, no additional human intervention is necessary for these steps. Head to head, per instance for the pallets/flask repository, SWE-bench style collection requires $38 \times 60 / 11 = 207.27$ minutes compared to 0.176 minutes (\sim 10.6 seconds) and \$0.00613 in API costs using SWE-smith.

Third, *collectively, SWE-smith task instances break a significantly larger proportion of existing tests in a codebase.* We define “bug coverage” as the proportion of tests broken by 1+ instance across all task instances. For the SWE-bench split of pallets/flask, there are 207 unique tests across all 11 instances. Of these 207 tests, 15 are broken by 1+ instance, corresponding to a bug coverage rate of 7.25%. For the SWE-smith split of pallets/flask, there are 474 unique tests across 267 instances. The larger amount of tests due to increased test coverage in the pallets/flask repository as of Nov. 28, 2024 (when SWE-smith was collected) compared to June 2023 (when SWE-bench was collected). Of these 474 tests, 422 are broken by 1+ instance, a bug coverage rate of 89.03%. We attribute the significant difference to a consistent tendency in real world open source software development workflows, that is, the *minority* of tests are introduced to capture existing, errant behavior in the repository. The significant majority of tests are committed alongside working code, ensuring that already correct behavior is upheld. Well-maintained repositories will typically not merge commits that cause such tests to fail. This results in a large number of tests where few to no commits correspond to those tests’ failures.

Finally, *SWE-smith does not yield instances appropriate for evaluation.* The SWE-smith pipeline as presented does not produce hidden tests, a crucial difference that makes SWE-bench more suitable for evaluation. Consequently, when expert trajectories are generated, the Fail-to-Pass tests are present in the repository at inference time. Furthermore, our issue generation strategy does not include checks for known problem such as underspecified text descriptions or solution leakage (Chowdhury et al., 2024). Simple amendments could make SWE-smith task instances suitable for evaluation, such as deleting Fail-to-Pass test functions or files along with a validating procedure around the ambiguity and leakage of the issue text. Finally, thorough analyses of how faithful SWE-smith task instances are to real world issues and PRs would be necessary to justify synthetic bugs for evaluation.

D Issue Generation

We describe the four issue generation strategies we experiment with to determine the effect of a issue text on how solvable a SWE-smith task instance is along with the value of the trajectory as a training data point.

Fixed issue templates. We create a set of 7 pre-defined issue templates, listed in Table 14. Each template uses information from the bug patch or Fail-to-Pass tests associated with every task instance. Given a dataset of task instances, we randomly select one of the templates to use as the problem statement according to the probabilities listed in Table 14. The reason we assign the highest likelihood for the prompt that provides all four categories of information (bug type, files changed, functions changed, Fail-to-Pass tests) is to ensure that a higher proportion of task instances have specified issue texts.

Fail-to-Pass test code and execution logs. Another approach is showing the source code and test execution logs for a randomly selected Fail-to-Pass test. This approach is motivated by the lack of reproduction code or expected/actual behavior of code communicated with fixed issue templates. We show code and execution logs only for a single Fail-to-Pass test; if a task instances has more than one Fail-to-Pass test, we do not disclose remaining tests.

Template	Prob.	Information Provided
Basic	0.05	None
Files	0.1	States which file(s) have bug(s).
Funcs	0.15	States which file(s) and func(s) have bug(s).
Tests	0.1	States that some tests are failing.
F2P Tests	0.1	States which tests are failing.
Bug Type	0.05	States failure type.
Bug Type + Files	0.15	States failure type and which file(s) have bug(s)
Bug Type + Files + Test	0.15	States failure type, which file(s) have bug(s), and a random F2P test.
Bug Type + Files + Funcs + Test	0.15	States failure type, which file(s) and func(s) have bug(s), and a random F2P test.

Table 14: List of issue text templates we use to generate problem statements. Across all templates, four types of information are included — the files with bugs, functions with bugs, Fail-to-Pass test(s), and the type of bug. Templates that offer less information are generally assigned a lower probability.

Generated with LM. We prompt an LM with a randomly selected SWE-bench Verified problem statement, the bug patch, list of Fail-to-Pass tests, source code for one Fail-to-Pass test, and the execution logs of running all the Fail-to-Pass tests. We ask the LM to generate an issue that describes the bug conveyed in the patch in the style of the SWE-bench Verified demonstration. Figure 19 shows the system prompt for this strategy.

System prompt for generating issues with an LM

You are a software engineer helping to create a realistic dataset of synthetic GitHub issues.

You will be given the following input:

1. Demonstration: A realistic GitHub issue to mimic (included in the <demonstration> tag).
2. Patch: A git diff output/PR changes that introduces a bug (included in the <patch> tag).
3. Test output: The output of running the tests after the patch is applied (included in the <test_output> tag).
4. Test source code: Source code for one or more tests that failed (included in the <test_source_code> tag).

Output: A realistic GitHub issue for the patch.

Guidelines:

- Mimic the style and structure of the demonstration issues. If the demonstration issues are not well structured, your output should also be not well structured. If the demonstrations use improper or no markdown, your output should also use improper or no markdown. If the demonstrations are short/long, your output should also be short/long (if possible). If the demonstrations include human "flavor text" or "fluff", your output should also include human "flavor text" or "fluff". Do this even if it conflicts with your default behavior of trying to be extremely concise and helpful.
- DO NOT explain the fix/what caused the bug itself, focus on how to reproduce the issue it introduces
- Do not mention pytest or what exact test failed. Instead, generate a realistic issue.
- If possible, include information about how to reproduce the issue. An ideal reproduction script should raise an error or print an unexpected output together with the expected output.

However, still include this information in a style very similar to the demonstration issues.

Figure 19: System prompt provided to an LM to generate an issue based off the bug patch and testing information of a task instance along with a demonstration problem statement randomly selected from SWE-bench Verified.

Original issue text. This strategy works exclusively for some task instances generated using PR Mirroring. If a PR is successfully mirrored, we use the text from the associated issues as the problem statement, exactly as done in SWE-bench. Of the 2345 task instances represented in SWE-smith mirrored from real world PRs, 708 or 30.19% of these have one or more associated GitHub issue(s) to create a SWE-bench style problem statement.

E Difficulty Rating

We train a model that labels a task with one of three difficulty labels: < 15 minutes (easy), 15 minutes - 1 hour (medium), and 1+ hour (hard). This model allows us to quantify the difficulty of individual task instances and, in aggregate, the difficulty of entire datasets.

To train this model, we use 1699 annotations from [Chowdhury et al. \(2024\)](#). In their work towards curating SWE-bench Verified, a subset of 1699 SWE-bench task instances were labeled with four difficulty levels: < 15 min, 15 min - 1 hr, 1-4 hrs, and 4+ hrs. Generally, three annotators were assigned to each instance, and the difficulty annotations were ensembled by taking the majority choice for a sample, or the median if there is no majority. The distribution of annotated difficulties, from easiest to hardest, is 24.5%, 53.5%, 19.4%, and 2.8%.

Because there are very few samples in the 4+ hr category, we reclassify the 1-4 hr and 4+ hr instances into a single 1+ hr category. Next, we create corresponding train and test datasets at a 80/20% split, randomly shuffling the instances while ensuring the train and test distributions do not deviate significantly from the original. An instance’s problem statement and solution patch are provided as input, and one of the three difficulty labels serves as the target output. We perform LoRA fine-tuning ([Hu et al., 2021](#)) on a Qwen 2.5 32B Instruct model using the Unsloth ([Daniel Han & team, 2023](#)) library. The model achieves an accuracy of 68.24% on the test set. All errant predictions are off by one; in other words, the model never predicted < 15 min when the label waops 1+ hr, and visa versa.

Using this model, we can grade the difficulty of a SWE-smith instance once the bug patch and corresponding issue text have been created. To provide a succinct summary of difficulty for a dataset of SWE-bench style task instances, we propose a “difficulty score” metric. Each label corresponds to a numeric difficulty score of 1, 5, and 9, from easiest to hardest. The difficulty score is therefore the average difficulty score across all task instances.

Figure 4 summarizes our findings for difficulties across different SWE-bench style datasets. We find that different SWE-smith bug generation methods yield different levels of difficulty. Specifically, aggregating smaller functions together will typically yield harder problems. This effect aligns with our original expectations; generally, bugs that require editing more functions and files tend to be rated as more difficult.

F Experiments

In this section, we provide additional details about the configurations and parameters used to generate trajectories with an expert model and run inference on a fine-tuned model. We then provide additional ablations and analyses about the SWE-smith dataset and the behaviors of agents trained on top of our dataset.

F.1 Training Details

Rejection sampling fine-tuning. Our fine-tuning setup heavily inherits from [Pan et al. \(2024\)](#)’s work. We perform full parameter fine tuning using the torchtune ([PyTorch, 2024](#)) library, with learning rate $5e-5$, maximum 3 epochs, and max context length of 32768. Training was carried on Modal ([Modal, 2025](#)) on 2-8 NVIDIA H100 80G GPUs. As discussed in Section 3, the procedure for rejection sampling fine-tuning (RFT) is as follows. We first generate expert demonstrations/trajectories using SWE-agent and a “strong” model (e.g. Claude 3.7, GPT 4o) on SWE-smith task instances. Of these, we then only train a student model on the trajectories corresponding to resolved instances.

SWE-agent configuration. We use two different configurations, one for generating trajectories with an expert model, and a separate one for running inference on the fine-tuned Qwen, student models. The configurations are generally quite similar, with minor differences around how LMs’ responses are elicited, the parsing mechanism for an LM response, constraints around message sizes, and the system prompt.

Task Instance Prompt provided to SWE-agent

```

<uploaded_files>
{{working_dir}}
</uploaded_files>
I've uploaded a python code repository in the directory {{working_dir}}. Consider
the following PR description:

<pr_description>
{{problem_statement}}
</pr_description>

Can you help me implement the necessary changes to the repository so that the
requirements specified in the <pr_description> are met? I've already taken care of
all changes to any of the test files described in the <pr_description>. This means
you DON'T have to modify the testing logic or any of the tests in any way! Your task
is to make the minimal changes to non-tests files in the {{working_dir}} directory to
ensure the <pr_description> is satisfied. Follow these steps to resolve the issue:
1. As a first step, it might be a good idea to find and read code relevant to the
<pr_description>
2. Create a script to reproduce the error and execute it with 'python <filename.py>'
using the bash tool, to confirm the error
3. Edit the source code of the repo to resolve the issue
4. Rerun your reproduce script and confirm that the error is fixed!
5. Think about edgcases and make sure your fix handles them as well Your thinking
should be thorough and so it's fine if it's very long.

```

Figure 20: A copy of the prompt provided to an LM via SWE-agent informing the LM of the nature of the task, the task description itself, and several tips on how to proceed.

We will first review the information common to both configurations. The prompt template informing an agent of the task’s nature and problem statement is included in Figure 20. This prompt is very similar to the original SWE-agent prompt used in Yang et al. (2024a). The prompt templates for showing environment feedback are identical as well. If there is execution output, the text is simply preceded by OBSERVATION: [output]. If there is no output (e.g. `rm -r` succeeds silently), then the agent is informed “Your command ran successfully and did not produce any output”. The agent computer interface (ACI) provided is also identical; SWE-agent provides LM with access to three general tools:

- `bash`: Execute a bash command in terminal.
- `str_replace_editor`: A tool for viewing, creating, and editing files.
- `submit`: A special keyword for the LM to indicate the task is completed or if it is unable to proceed further with the task.

We briefly review the distinctions. First, the way tools are invoked is different for expert versus student models. For the Claude and GPT series models that are used as experts, we use function calling for models to invoke the aforementioned tools. On the other hand, the student model is asked to generate a response with XML tags to delineate the thought and action. Therefore, when fine-tuning on expert trajectories, a key processing step is to convert the expert trajectories’ function calling format into the XML style response — fine-tuning *directly* on the expert trajectories does not work.

For generating trajectories with expert models, we run with a maximum of 75 steps and a cost limit of \$2.00. A run terminates automatically when either of these limits are reached

Purpose	Bug Gen.	Issue Gen.	# Instances	Temp.	# Traj.
claude-3-7-sonnet-20250219					
Ablation (Bug Type)	LM (Modify)	LM	1000	0	605
	LM (Rewrite)	LM	1000	0	507
	Procedural	LM	1000	0	745
	PR Mirrors	LM	1000	0	557
Ablation (Issue Type)	PR Mirrors	Fixed	600	0	259
	PR Mirrors	F2P Test	600	0	390
	PR Mirrors	Original	600	0	328
	PR Mirrors	LM	600	0	319
Ablation (Repositories)	Procedural	LM	1000	0	721
	Procedural	LM	1000	0	709
	Procedural	LM	1000	0	723
	Procedural	LM	1000	0	707
Final Dataset Curation	LM (Rewrite)	LM	3574	0	1003
	PR Mirrors	LM	1049	0	349
claude-3-5-sonnet-20250219					
Compare with prior work	All	LM	800	0	535
gpt-4o-2024-08-06					
Compare with prior work	All	LM	200	0	89

Table 15: Breakdown of trajectories sampled from SWE-smith. Trajectories were generated from subsets of SWE-smith that were either for the purpose of ablations or performance. All trajectories were generated with a maximum of 75 steps and a \$2 cost limit.

or the context window of the expert model is exceeded. The overwhelming majority of automatic terminations are due to the 75 maximum steps limit.

For running inference with student models, we run with a maximum of 75 steps or a cost limit³ of \$2.00, where the run similarly terminates when either the steps, cost or context window limit is reached. For the student model, per LM inference call, we truncate the message history to only keep the 5 most recent tool outputs. While we occasionally sample trajectories with the expert model set at various temperatures, for the student model, the temperature is fixed at 0.0.

F.2 Trajectory Dataset Breakdown

We provide a thorough review of the dataset of SWE-agent trajectories released with this work in Table 15. The majority are generated with claude-3-7-sonnet-20250219. To compare with prior work, a minority were generated with claude-3-5-sonnet-20240620 and gpt-4o-2024-08-06. As mentioned in Section 4, to guard against the easy data bias phenomenon, we impose a per-instance cap of 3, meaning for any task instance, we include at most 3 trajectories successfully resolving that task instance in our fine-tuning dataset. From the pool of trajectories reflected in Table 15, we curate a set of 5000 trajectories that we then use to train SWE-agent-LM-32B.

Tables 16 and 17 show what repositories and bug types are represented in the final training dataset. In total, 123 repositories are represented, with at least 10 trajectories from 91 repositories. Trajectories are on average 58 turns long, meaning an LM typically takes 29 actions for a given demonstration trajectory. We visualize this distribution in Figure 21.

³We include the cost limit in addition the step limit to provide realistic behavior with respect to handling long context. To calculate a cost value for our model, we use the gpt-4o cost function as of April, 2025.

Bug Type	Count
Combine (File)	123
Combine (Module)	7
LM (Modify)	11
LM (Rewrite)	1532
Procedural	1495
PR Mirror	1848

Table 16: Bug types represented in final training dataset.

Repository	Count	Repository	Count
getmoto/moto	378	sqlfluff/sqlfluff	122
pandas-dev/pandas	320	pylint-dev/astroid	110
conan-io/conan	243	pydicom/pydicom	103
pydantic/pydantic	209	tobymao/sqlglot	101
iterative/dvc	181	pygments/pygments	99
dask/dask	139	scanny/python-pptx	98

Table 17: Top ten repositories by number of trajectories represented in final dataset for main result.




Figure 21: Distribution of number of turns for trajectories represented in the final dataset.

F.3 Agent Behavioral Studies

F.3.1 Turn counts and cost

While agents are frequently quoted with a singular cost-per-instance number, this can be very misleading in the case of SWE-agent-LM-32B. Because most of the failed instances fail due to termination by the cost or turn count limit, the average cost and turn counts depend strongly on these limits (see Fig. 22).

We can also chart the number of resolved instances vs step limits. To avoid reevaluating the agent with multiple step limits, we use one run with step limit 75 and then assume that a successful agent run that terminates after step n would have failed when restricted by a limit smaller than n . This chart corroborates the point made in section 3: SWE-agent-LM-32B has a higher resolution rate for very low step limits.




Figure 22: The average step count depends strongly on the prescribed step limit.




Figure 23: Number of successful instances submitted before a given step limit.

F.3.2 Analysis of agent action space

Reduction to base commands. In addition to the dedicated tools provided to the agent as part of the agent computer interface (Section F.1), the agent can execute arbitrary bash commands. This makes quantitative analyses of the agent action space challenging. For example, the agent might issue commands like `PYTHONPATH=/testbed/repo cd /testbed/repo && python3 reproduce.py`. We have found the following procedure to determine a *base command* effective to meaningfully describe the action:

1. Strip any environment variable manipulation from the beginning of the command
2. When multiple commands are chained with `&&` or semicolons, only consider the last one
3. Remove all arguments. Because some commands have subcommands (e.g., `git checkout`), we apply several basic heuristics to determine whether to keep the first or the first two words.

Repetitive actions. We determine the longest repetitive sequence of actions by determining the longest sequence of identical base commands within the agent actions. Note that this means that e.g., `str_replace_editor view` actions that target different files are considered to be identical actions as far as this analysis is concerned.

F.3.3 Failure mode analysis

Categorizing the failure mode proceeds as shown in Figure 24:

1. **Error conditions:** If the agent terminates due to an error (environment errors, inability of the LM to correctly format its messages, etc.) or because it exceeded its maximum context window, we return the **error** or **context** category.
2. **Early termination:** If the agent was terminated because of a step or cost limit, we return one of the **stuck ...** subcategories. Note that the SWE-agent still attempts to extract a submission (list of changes/patch). We determine the subcategory based on which part of the workflow agentic loop was terminated:
 - (a) If no source (i.e., non-test) file was modified⁴ and no attempt at testing was made, we return **stuck at localization**. If test commands were run (i.e., `python`, `pytest`, ..., or similar commands), we return **stuck at reproduction**.
 - (b) If source files *were* modified, we check whether the changes include changes to all source files that are modified in the gold patch. If not, we return **incorrect localization (stuck)**, else **incorrect edit (stuck)**.
3. **Successful submission:** If the agent terminated and submitted a solution naturally, we return **incorrect localization** or **incorrect edit**, depending on whether the changes from the submitted patch included changes to all files from the SWE-bench gold patch.

G Miscellaneous

Teaser figure description. We briefly describe how the left hand graph of Figure 1, which depicts scaling of task instance collection for the SWE-smith vs. SWE-bench, was created. For SWE-smith, we simply collected the number of task instances for each repository. For SWE-bench, we ran SWE-bench task instance candidate collection script on all 128 repositories, which first crawls all PRs from a given repository. Then, each PR that edits at least one or more Python files and changes at least one or more testing related files is converted into a candidate task instance. Finally, based on the average task instance yield rate reported in Jimenez et al. (2024b), we estimate the number of viable task instances to be 20% of the candidates. We then determine the number of task instances for n repositories

⁴We exclude added files because solving SWE-bench instances always requires *changes* to existing files.

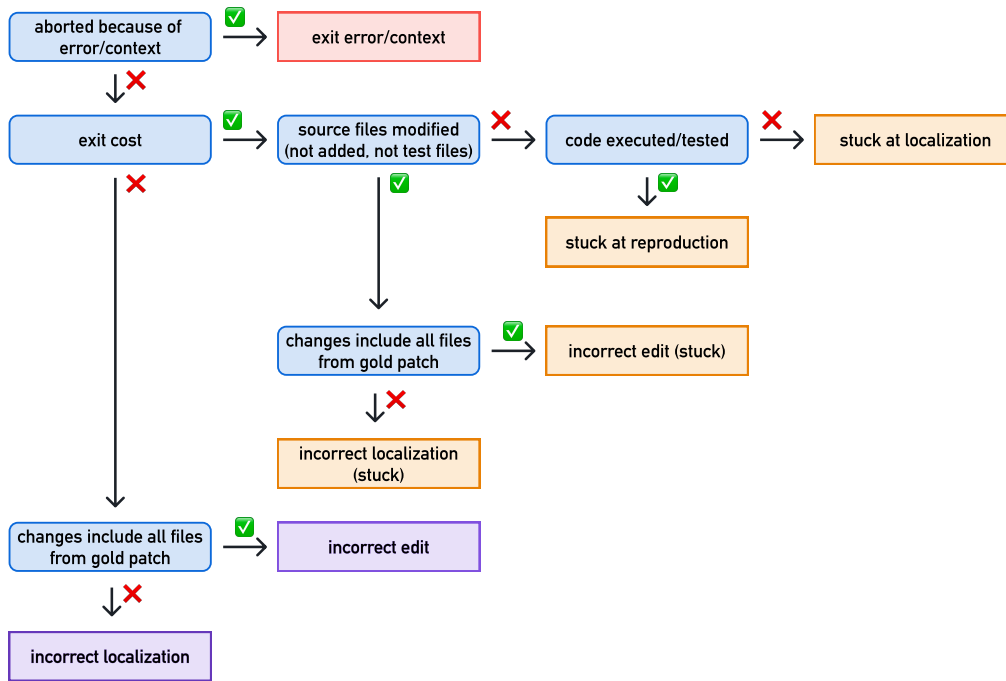


Figure 24: Categorizing failure modes

at intervals of 5 repositories ranging from 5 to 250, where the repositories are sorted by number of stars. In other words, the first five repositories we account for in the figure are the five with the fewest number of stars out of the 128 repositories used.

Extended related works. We discuss additional related works briefly, primarily about similar work towards synthesizing trajectories for training LM agents, but for the domain of web tasks. To improve the interactive capabilities of open source LMs (Chen et al., 2023), prior works have also explored trajectory generation techniques for web benchmarks and settings (Xie et al., 2024; Yao et al., 2023a; Zhou et al., 2024). For web navigation, existing strategies rely on (1) performing random walks which are then labeled retroactively with instructions (Xiang et al., 2023; Murty et al., 2024), (2) using online web tutorials as a source of indirect supervision for generating synthetic trajectories (Ou et al., 2024), or (3) collecting human demonstrations (Shen et al., 2024; Xu et al., 2024). These procedures do not translate well to the software engineering setting; random sequences of command line interactions usually do not achieve meaningful effects on a codebase. Our cursory efforts around replaying trajectories synthesized from online code edit sequences (e.g. GitHub commit histories) were unsuccessful due to the limited information available, which primarily capture file-level changes without reflecting the underlying skills, decision-making, or the broader context of a software development process.

Our exploration of using SWE-agent to automatically determine installation and testing specifications for a repository is heavily influenced by two research directions - automatic execution environment construction using LMs (Bogin et al., 2024; Eliseeva et al., 2025; Vergopoulos et al., 2025), and generating unit tests using LMs (Mündler et al., 2025). Although relatively much less than SWE-bench style collection, SWE-smith still requires minimal amounts of human labor (around 8 minutes total per repository). As we expand SWE-smith to more repositories and languages, we are continuing to consider how to completely automate the environment construction process end to end.