
SWE-Replay: Efficient Test-Time Scaling for Software Engineering Agents

Yifeng Ding¹ Lingming Zhang¹

Abstract

Test-time scaling has been widely adopted to enhance the capabilities of Large Language Model (LLM) agents in software engineering (SWE) tasks. However, the standard approach of repeatedly sampling trajectories from scratch is computationally expensive. While recent methods have attempted to mitigate costs using specialized value agents, they can suffer from model miscalibration and fail to generalize to modern agents that synthesize custom bash scripts as tools. In this paper, we introduce SWE-Replay, the first efficient and generalizable test-time scaling technique for modern agents without reliance on potentially noisy value estimates. SWE-Replay optimizes the scaling process by recycling trajectories from prior trials, dynamically choosing to either explore from scratch or exploit archived experience by branching at critical intermediate steps. This selection of intermediate steps is driven by the potential and reasoning significance of repository exploration, rather than external LLM-based quality estimates. Our evaluation shows that, on SWE-Bench Verified, SWE-Replay consistently outperforms naive scaling, reducing costs by up to 17.4% while maintaining or even improving performance by up to 3.8%. Further evaluation on SWE-Bench Pro and Multilingual validates the generalizability of SWE-Replay, establishing it as a robust foundation for efficient test-time scaling of software engineering agents.

1. Introduction

Large Language Models (LLMs) have evolved significantly, progressing from basic code completion tools (Austin et al., 2021; Chen et al., 2021; Li et al., 2022; Ding et al., 2024; 2025a) to sophisticated interactive agents capable of navigating repositories, executing tests, and submitting patches

¹Siebel School of Computing and Data Science, University of Illinois Urbana-Champaign, USA. Correspondence to: Yifeng Ding <yifeng6@illinois.edu>.

end-to-end (Yang et al., 2024; Wang et al., 2024; Zhang et al., 2024; Liu et al., 2024; Xia et al., 2024; Carbonneau et al., 2025). Modern agentic frameworks, such as SWE-agent (Yang et al., 2024) and OpenHands (Wang et al., 2024), equip LLMs with tools like terminals, editors, and search engines to tackle complex repositories. To further enhance capabilities in software engineering (SWE) tasks, a primary driver of progress has been *test-time scaling* (Zhang et al., 2025b), which increases inference-time computation to yield higher-quality solutions. Existing works have demonstrated that generating multiple candidate solutions with a positive temperature and selecting a final answer via test-based feedback improves performance log-linearly with the number of samples, establishing a promising scaling law for SWE tasks (Brown et al., 2024; Ehrlich et al., 2025).

Despite this promise, test-time scaling is computationally expensive due to the cost of repeated sampling from scratch (Kim et al., 2026). Consequently, several existing works have been proposed to mitigate these costs. SWE-Search (Antoniades et al., 2024) proposed using a modified Monte Carlo Tree Search (MCTS) to prune unpromising trajectories early, relying on a value agent to estimate the quality of each step. Similarly, Satori-SWE (Zeng et al., 2025) proposes to achieve sample-efficient test-time scaling by prompting agents to self-improve the scores of their prior generations estimated by a reward model. However, these approaches have the following limitations: quality scores from value agents and reward models can be compromised by model miscalibration (Son et al., 2024), introducing noise during scaling; more importantly, existing approaches are tailored to pipeline-based scaffolds (*e.g.*, Moatless¹) and cannot generalize to modern agentic frameworks like SWE-agent (Yang et al., 2024). Specifically, existing works design specific prompts to evaluate pre-defined tools in the pipeline, such as a structured search/retrieval tool. However, modern agents are designed to synthesize custom bash scripts without templates. This flexibility makes it infeasible to design tool-specific evaluation prompts, rendering existing works incompatible with modern agents.

To bridge this gap, we introduce SWE-Replay, the first *efficient* and *generalizable* test-time scaling technique designed to reduce trajectory sampling costs while improving

¹<https://github.com/aorwall/moatless-tools>

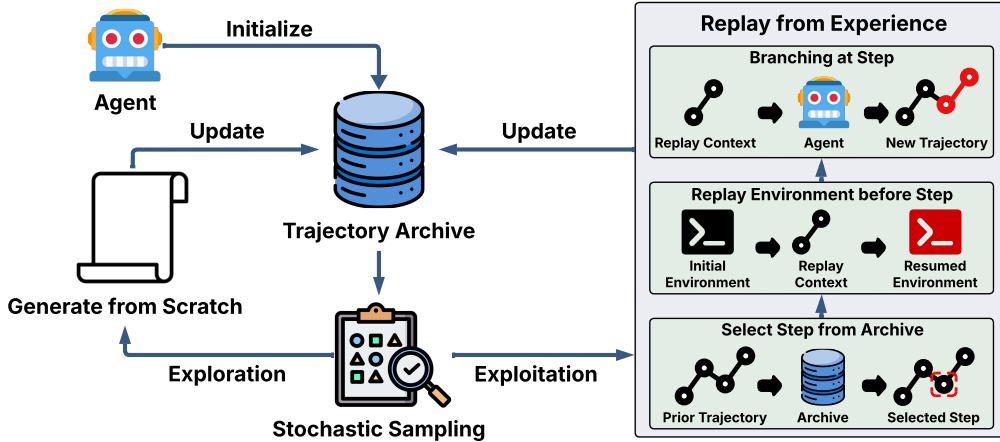


Figure 1. Overview of SWE-Replay.

trajectory quality for modern agents, *without reliance on LLM-as-a-Judge*. Our core insight is to recycle previously sampled trajectories by resuming exploration at carefully selected intermediate steps, rather than generating every trajectory from scratch. As shown in Figure 1, SWE-Replay maintains an archive of sampled trajectories and performs stochastic sampling iteratively: either explore by sampling a new trajectory from scratch or exploit by replaying from the middle of an existing trajectory. In the exploitation phase, SWE-Replay identifies a critical step s_t based on its potential to explore new search space and its reasoning intensity, efficiently restores its environment state, and branches at step s_t by sampling a new step s'_t to replace s_t while continuing exploration. This mechanism bypasses reliance on potentially inaccurate LLM-as-a-Judge, generalizes naturally to modern agentic scaffolds, and employs a streamlined select-and-replay mechanism to ensure scalability.

We evaluate SWE-Replay on the widely used SWE-Bench Verified², and the more complex SWE-Bench Pro (Deng et al., 2025) and Multilingual³. Our results show that, on SWE-Bench Verified, SWE-Replay consistently reduces the cost of naive test-time scaling by up to 17.4% while maintaining or even improving performance by up to 3.8%, across three different LLM backends and two different agentic scaffolds. On SWE-Bench Pro and Multilingual, SWE-Replay further demonstrates its consistent generalizability to diverse SWE problems. In addition, our analysis reveals that SWE-Replay successfully shifts exploration to the long-tail of repository files, as visualized in Figure 6. Finally, we provide an interesting theoretical intuition on how replaying optimizes the quality of exploration in SWE-Replay.

In summary, we make the following contributions:

- We present SWE-Replay, the first efficient test-time scaling technique for software engineering agents that reduces trajectory sampling costs while improving trajec-

tory quality, **without any reliance on LLM-as-a-Judge and generalizable to any modern agent scaffold**.

- On SWE-Bench Verified, SWE-Replay consistently reduces the sampling cost by up to 17.4% while maintaining or even improving the resolve rate by up to 3.8%, across different agent scaffolds and LLM backends. We further validate the generalization of SWE-Replay to diverse software issues on SWE-Bench Pro and Multilingual.
- Our in-depth analysis provides both *empirical* results showing that SWE-Replay enables agents to explore more diverse repository spaces, and *theoretical* intuition that bridges the gap between the performance gain and the effectiveness of step selection in SWE-Replay.

2. SWE-Replay

To enable efficient and effective exploration for software engineering agents, we propose SWE-Replay, a novel algorithm that reduces the cost of trajectory sampling while improving trajectory quality. The key idea of SWE-Replay is to reuse previously sampled trajectories by resuming exploration at carefully selected intermediate steps, rather than repeatedly generating trajectories from scratch.

As shown in Figure 1, SWE-Replay maintains an archive of previously sampled trajectories and updates the archive with newly generated ones. Specifically, SWE-Replay initializes the archive with a single trajectory generated from scratch. It then iteratively conducts a stochastic sampling to decide whether to explore by generating a new trajectory from scratch or to exploit existing trajectories by resuming from intermediate steps in the archive. To exploit existing trajectories, SWE-Replay will select a critical step s_t from the archive based on its potential to explore new repository space and its reasoning importance (§2.1), restore the environment state before this step (*i.e.*, $s_{1,2,\dots,t-1}$) efficiently (§2.2), and resume exploration by generating a new step s'_t to replace s_t for new branches (§2.3).

²<https://openai.com/index/introducing-swe-bench-verified/>

³<https://www.swebench.com/multilingual.html>

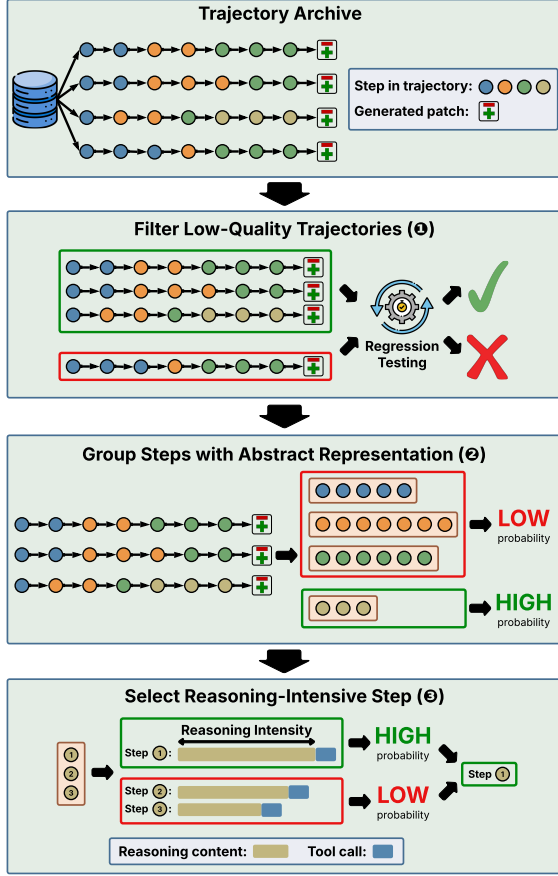


Figure 2. Overview of step selection in SWE-Replay.

2.1. Selecting Steps

Reliable identification of critical intermediate steps is central to efficient exploration. Intuitively, resuming at such steps allows the agent to revisit valuable regions of the search space that **are not explored comprehensively** and **require more reasoning efforts** from the agent.

To this end, SWE-Replay employs a hierarchical selection pipeline consisting of four components, as shown in Figure 2: (1) filtering low-quality trajectories from the archive (§2.1.1), (2) grouping steps using abstract state representations and selecting the critical group (§2.1.2), (3) selecting concrete steps within a group based on reasoning intensity (§2.1.3), and (4) balancing exploration and exploitation through stochastic sampling (§2.1.4).

2.1.1. FILTERING LOW-QUALITY TRAJECTORIES

To ensure that resumed exploration originates from reliable trajectories, SWE-Replay first removes low-quality trajectories from the archive. Trajectory quality is evaluated based on the final generated patch. In particular, we check whether the patch causes any existing regression tests to fail, and **discard all trajectories whose final patches introduce regression failures**. The motivation is that trajectories leading to regression failures often contain misleading intermediate

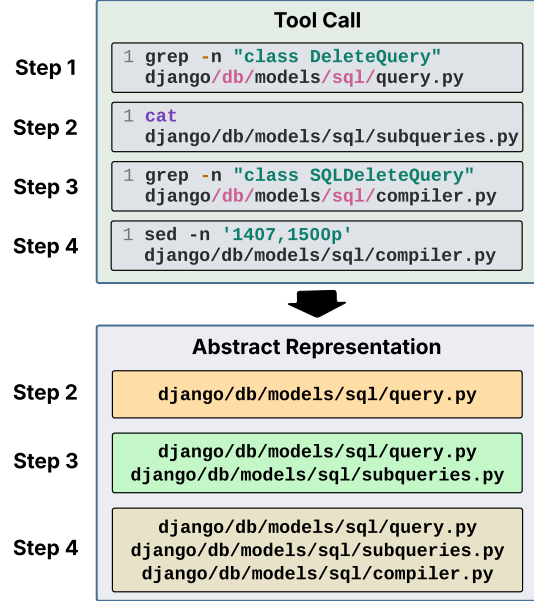


Figure 3. Example of tool calls and their file-level representations. Step 1 is skipped as no file has been explored before it.

steps, such as exploring irrelevant files, which can propagate the same errors when reused. In practice, we follow Agentless (Xia et al., 2024) to obtain regression tests for each issue. Since the original repositories do not provide reproduction tests and LLM-generated tests may be unreliable, SWE-Replay exclusively relies on existing regression tests from the repository, as shown in Figure 2 (1).

2.1.2. GROUPING STEPS VIA ABSTRACTION

A key challenge in selecting critical steps is to distinguish meaningfully different exploration steps while avoiding fragmentation caused by superficial differences. Although one could represent each step using the raw environment output, such representations are typically long and nearly unique, making similarity assessment ineffective. Instead, SWE-Replay introduces a lightweight yet effective abstraction by **representing each step using the set of repository files explored before that step**, as shown in Figure 3.

We refer to this representation as a *state*. Each state corresponds to a subset of steps that have explored the same files. Let $\{s_1, s_2, \dots, s_n\}$ denote the set of distinct abstract states extracted from all steps, where state s_i corresponds to v_i concrete steps. To prioritize states that access rarely explored regions of the repository, SWE-Replay assigns a sampling probability p_i to each state based on its rarity:

$$p_i = \text{softmax}\left(\frac{1}{v_i}\right) = \frac{e^{\frac{1}{v_i}}}{\sum_{j=1}^n e^{\frac{1}{v_j}}} \quad (1)$$

This formulation favors states reached by fewer steps, but still maintains a balanced trade-off between exploration and exploitation through sampling rather than greedily choosing

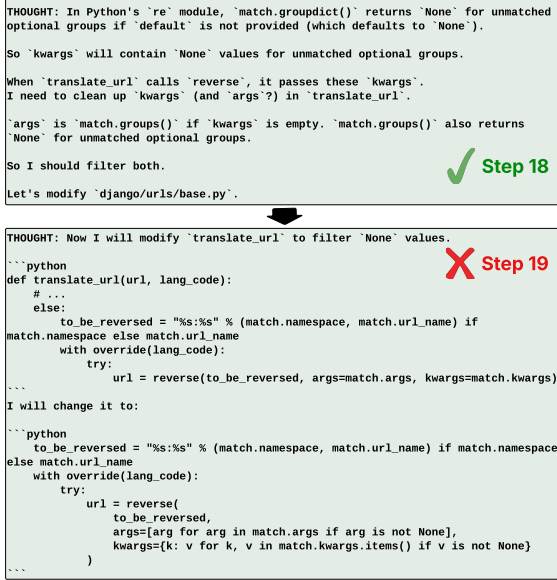


Figure 4. Example of reasoning contents with different intensity.

the most rare state. Softmax normalization is adopted to maintain the sampling diversity for **worst** performance with a flatter distribution, as validated in §A.7.

2.1.3. DETECTING REASONING-INTENSIVE STEPS

After selecting a critical abstract state, SWE-Replay must choose a concrete step at which to resume exploration. Typically, modern agents are designed to perform explicit reasoning at each step. Inspired by this, we select the final step by **prioritizing reasoning-intensive ones**, as shown in Figure 2 (⊕). It is based on the intuition that steps requiring more deliberation are often *critical decision points* in the task. Therefore, SWE-Replay branches at such steps to explore alternative solutions by sampling new steps to replace them.

To measure the reasoning intensity, one straightforward idea is to use the raw reasoning length. However, the raw length cannot **ineffectively** convey the reasoning intensity, as the rea-soning content may include copies of raw code snippets or concrete inputs, which significantly extend the length with-out providing any meaningful reasoning gains. As shown in Figure 4, while step 19 has longer reasoning content, its reasoning consists of two detailed but less informative edit drafts; in contrast, step 18 contains multiple short, in-sightful paragraphs analyzing the program logic. Based on this observation, SWE-Replay instead uses the number of paragraphs in the reasoning content as a *robust structural proxy* to measure the reasoning intensity. Empirical ablation (§A.6) confirms that paragraph count outperforms raw token length (60.0% vs 58.0%), validating that structural segmen-tation captures reasoning quality better than verbosity.

Given the v_i steps $\{s_{i,1}, \dots, s_{i,v_i}\}$ within state s_i , with corresponding numbers of reasoning paragraphs

$\{l_{i,1}, \dots, l_{i,v_i}\}$, similar to §2.1.2, SWE-Replay assigns sampling probabilities as follows:

$$p_{i,j} = \text{softmax}(l_{i,j}) = \frac{e^{l_{i,j}}}{\sum_{j=1}^{v_i} e^{l_{i,j}}} \tag{2}$$

2.1.4. BALANCING EXPLORATION AND EXPLOITATION

To support effective test-time scaling, SWE-Replay balances exploration (sampling trajectories from scratch) and exploitation (resuming from prior trajectories) as follows:

- **Initialization (Exploration):** The first run always generates a trajectory from scratch using only the original issue description, forming the initial archive.
- **Scaling (Exploration vs. Exploitation):** For subsequent runs, exploration and exploitation are treated as independent Bernoulli trials with probability $p = 0.5$. The strategy for each run is then sampled stochastically.

This design enables SWE-Replay to naturally trade off coverage and reuse as scaling proceeds.

2.2. Replaying Environment before Steps

Once a target step is selected, SWE-Replay must restore the replay context and environment state *before* this step, as shown in Figure 1. Since prior trajectories are recorded in the archive, the replay context can be retrieved directly. For the environment state, storing full snapshots is storage-intensive, and replaying entire action sequences can be inefficient, as some actions (*e.g.*, running test suites) are costly. To address this, we observe that software agents modify the environment mostly through text edits and file creation/deletion in the repository. Consequently, we start by checking with pattern matching whether agents’ actions have mutated non-repo state (*e.g.*, package installs): if unchanged, we will only record the repository diff for each step: when resuming from a step, the environment is restored by applying the stored diffs to the repository, avoiding action replay and improving efficiency; otherwise, SWE-Replay will replay the original action sequence to restore the environment.

2.3. Branching at Steps

After restoration, the agent resumes exploration by sampling a new step to *replace the selected step* and then continuing to complete a new trajectory, as shown in Figure 1. The resulting new trajectory, which combines the replay context and generated suffix, is added to the archive for future reuse.

Overhead analysis. Unlike existing works (Antoniades et al., 2024; Zeng et al., 2025) which require costly tree search or model training, SWE-Replay introduces negligible overhead. Firstly, step selection only requires lightweight system operations (*e.g.*, loading trajectory files) and simple numeric computations, which do not introduce any meaningful overhead. Likewise, the overhead of returning to the

Table 1. Results on SWE-Bench Verified.

Agent Scaffold	LLM Backend	Scaling Method	Performance		Efficiency	
			% Resolved \uparrow	# Input (k) \downarrow	# Output (k) \downarrow	Avg. Cost (\$) \downarrow
mini-SWE-agent	◆ Gemini-2.5-Pro	Naive Scaling	58.0%	3464.5	98.3	1.52
		SWE-Replay	60.2%	2980.4	86.0	1.32
		Improvement	\uparrow 3.8%	\downarrow 14.0%	\downarrow 12.5%	\downarrow 13.2%
	🏠 Devstral-Small-2	Naive Scaling	62.2%	14914.9	130.0	1.53
		SWE-Replay	63.2%	13221.4	111.2	1.36
		Improvement	\uparrow 1.6%	\downarrow 11.4%	\downarrow 14.5%	\downarrow 11.1%
Live-SWE-agent	◆ Gemini-3-Pro	Naive Scaling	75.4%	6849.2	106.4	2.88
		SWE-Replay	75.6%	5874.3	85.4	2.38
		Improvement	\uparrow 0.3%	\downarrow 14.2%	\downarrow 19.7%	\downarrow 17.4%
	🏠 Devstral-Small-2	Naive Scaling	63.2%	14993.9	134.9	1.54
		SWE-Replay	65.0%	13233.0	115.7	1.36
		Improvement	\uparrow 2.8%	\downarrow 11.7%	\downarrow 14.2%	\downarrow 11.7%
◆ Gemini-3-Pro	Naive Scaling	75.6%	6948.7	107.0	2.93	
	SWE-Replay	75.8%	6395.0	91.4	2.59	
	Improvement	\uparrow 0.3%	\downarrow 8.0%	\downarrow 14.6%	\downarrow 11.6%	

Table 2. Results on SWE-Bench Pro and SWE-Bench Multilingual.

Benchmark	Scaling Method	% Resolved \uparrow	Avg. Cost (\$) \downarrow
Pro	Naive Scaling	28.73%	1.29
	SWE-Replay	28.97%	1.25
	Improvement	\uparrow 0.8%	\downarrow 3.1%
Multilingual	Naive Scaling	31.0%	1.22
	SWE-Replay	38.0%	1.11
	Improvement	\uparrow 22.6%	\downarrow 9.0%

steps is negligible: it only requires applying a diff file, which takes on the order of seconds and is insignificant compared to the cost of generating full trajectories.

3. Evaluation

3.1. Experimental Setup

Implementation. SWE-Replay is generalizable across various agent scaffolds. In our experiments, we implement SWE-Replay on top of two popular frameworks: mini-SWE-agent (Yang et al., 2024) and Live-SWE-agent (Xia et al., 2025). We evaluate performance using two closed-source and one open-source LLM backends: Gemini-2.5-Pro (Comanici et al., 2025), Gemini-3-Pro⁴, and Devstral-Small-2⁵. We sample ten patches per issue and employ an Agentless-style selection mechanism to submit the final solution (Xia et al., 2024). More details are provided in Appendix A.1.

Datasets. We evaluate SWE-Replay on three primary datasets: SWE-Bench Verified, SWE-Bench Pro (Deng et al., 2025), and SWE-Bench Multilingual. Following prior works (Zhang et al., 2025a; Xia et al., 2025), we utilize

SWE-Bench Verified mini⁶, a 50-problem subset of SWE-Bench Verified, for our ablation studies in §3.3. Further dataset details are presented in Appendix A.2.

Baselines and Metrics. Our primary baseline is *naive scaling*, where all ten trajectories are sampled from scratch. In contrast, SWE-Replay generates the initial trajectory from scratch and performs exploration-exploitation sampling (§2.1.4) for the remaining nine. We do not compare directly with SWE-Search or Satori-SWE as they rely on pre-defined tool templates and are incompatible with the open-ended action spaces of modern agents. Instead, we consider an LLM-as-a-Judge baseline (Table 4) where we replace step selection in SWE-Replay with LLM-as-a-Judge to estimate step quality and select steps based on the score.

Following standard practices (Xia et al., 2024; 2025), we report two metrics: (1) **Performance**, measured by the resolve rate; and (2) **Efficiency**, measured by the average number of input/output tokens and the cost per problem. Costs account for *explicit* prompt caching, as SWE-Replay maximizes prefix sharing. See Appendix A.3 for cost modeling details.

3.2. Main Results

SWE-Bench Verified. Table 1 presents the comparative results of SWE-Replay and naive test-time scaling on SWE-Bench Verified. We observe that SWE-Replay consistently achieves a higher resolve rate while incurring lower average costs across all agent scaffolds and LLM backends. Specifically, SWE-Replay reduces the computational cost of naive scaling by up to 17.4% while improving performance by up to 3.8%. These results highlight the ability of SWE-Replay

⁴<https://deepmind.google/models/gemini/>

⁵<https://mistral.ai/news/devstral-2-vibe-cli>

⁶<https://github.com/mariushobbahn/SWEBench-verified-mini>

Table 3. Ablation on SWE-Replay designs. When selection strategies are disabled (X), the method defaults to uniform random selection.

Naive Scaling	Trajectory Filtering	Selection #1: Representation	Selection #2: Reasoning	Performance		Efficiency	
				% Resolved \uparrow	# Input (k) \downarrow	# Output (k) \downarrow	Avg. Cost (\$) \downarrow
✓	-	-	-	52.0%	16814.7	149.4	1.73
-	X	X	X	56.0%	17373.3	132.1	1.78
-	✓	X	X	56.0%	14952.7	117.4	1.53
-	✓	✓	X	58.0%	15083.4	126.7	1.55
-	✓	✓	✓	60.0%	14795.3	131.8	1.52

Table 4. Ablation on different step selection methods.

Step Selection	% Resolved \uparrow	Avg. Cost (\$) \downarrow		
		Agent	Reward	Total
Random	56.0%	1.53	0	1.53
LLM-as-a-Judge	54.0%	1.44	1.87	3.31
SWE-Replay	60.0%	1.52	0	1.52

to deliver consistent efficiency gains and performance improvements without requiring specialized reward models.

SWE-Bench Pro and Multilingual. We conduct additional evaluations on the more complex SWE-Bench Pro and Multilingual benchmarks, using Devstral-Small-2 with mini-SWE-agent. Table 2 shows that SWE-Replay maintains its superiority over naive scaling, achieving an up to 22.6% performance gain with an up to 9.0% reduction in cost. These results confirm that SWE-Replay generalizes effectively to enterprise-level multilingual SWE problems.

3.3. Ablation Experiments

Ablation of SWE-Replay Components. We perform ablation studies to validate the design choices of SWE-Replay, focusing on three core components: trajectory filtering, step grouping via abstraction, and step selection based on reasoning intensity. Following the setup in §3.1, all ablations use the Devstral-Small-2 backend with Live-SWE-agent as it achieves a better performance on SWE-Bench Verified. As shown in Table 3, the full SWE-Replay implementation achieves the best results, with a 60.0% resolve rate and an average cost of \$1.52 per problem. Progressively disabling these components leads to degradation in both performance and efficiency. These findings underscore the necessity of each component in the SWE-Replay pipeline.

Comparison of LLM-as-a-Judge and SWE-Replay. We experiment to study whether we can replace the step selection mechanism of SWE-Replay with LLM-as-a-Judge. Specifically, we explore directly using LLM-as-a-Judge to estimate the quality of each step in the archive and then sampling one based on the LLM-generated quality score. More details are included in Appendix A.4. As shown in Table 4, LLM-as-a-Judge performs worse than SWE-Replay while requiring additional costs to estimate the quality score using




Figure 5. Scaling curve of SWE-Replay for Devstral-Small-2.

an LLM, demonstrating that step selection in SWE-Replay is crucial to boost the scaling performance without being impacted by potentially inaccurate LLM-as-a-Judge.

Scaling analysis of SWE-Replay. We evaluate how well the resolve rate on SWE-Bench Verified scales with the number of samples in SWE-Replay. Figure 5 shows that increasing the number of samples enhances its performance steadily, while outperforming naive scaling consistently.

More ablations on the impact of different step representations (§A.5), reasoning intensity representations (§A.6), and normalizations (§A.7) are detailed in the Appendix.

4. Discussion

4.1. SWE-Replay Diversifies Repository Exploration

A core intuition motivating the step selection mechanism of SWE-Replay (§2.1.2) is to explore repository spaces that are largely ignored in prior trials. To verify whether SWE-Replay can diversify the search space, we analyze the distribution of repository files visited by the agent on SWE-Bench Verified. Specifically, for each problem, we extract the set of files explored during the scaling process. To aggregate distributions across problems with varying file structures, we normalize file identities by ranking them based on visitation frequency. For a given problem, the most frequently visited file is assigned rank 1, the second most visited rank 2, and so on. This normalization allows us to average the exploration frequency of files across differ-

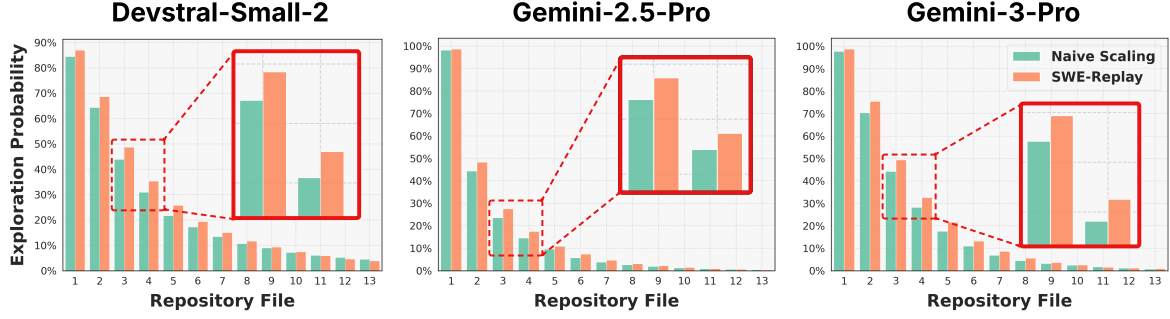


Figure 6. Distribution of repository files explored by SWE-Replay and naive scaling.



Figure 7. Example of fixing failure due to inadequate exploration. Naive scaling fixated on `message.py` (incorrect), whereas SWE-Replay successfully branched out to explore `utils.py` (correct).

ent problems. Thus, effectively accessing the "tail" of the distribution indicates greater diversity.

Figure 6 illustrates the distribution of repository files across different LLMs. A key observation is that while naive scaling and SWE-Replay exhibit similar probabilities of exploring the most frequent (top-ranked) files, SWE-Replay consistently increases the probability of visiting less frequent files (the long tail). These results confirm that SWE-Replay successfully encourages the agent to investigate less frequently visited files. The case study in Figure 7 demonstrates the practical advantage of this diversity. In this instance, the issue description explicitly points to `message.py` as the potential location of the bug. Consequently, naive scaling fixates on this file, generating a patch that fails to address the root cause. In contrast, SWE-Replay breaks away from this local optimum to explore `utils.py`, successfully fixing the issue at a lower level. This highlights the critical importance of exploring a diverse repository space to locate the correct patch.

4.2. Theoretical Intuition of SWE-Replay

We present the theoretical intuition of SWE-Replay by considering a simplified scenario to model the efficiency of SWE-Replay compared to naive scaling. We assume a challenging SWE task characterized by the following properties:

1. All sampled trajectories consist of exactly N steps.
2. There exists a single optimal path (the "correct trajectory") that solves the issue.
3. **Recoverability**: If the archive contains the correct trajectory and SWE-Replay branches from any step within it, the agent is guaranteed to fix the issue.
4. **Hardness**: The probability p of generating this correct trajectory from scratch in a single trial is small ($p \ll 1$).

Let us define the probabilities for the t -th trial ($t \geq 2$):

1. $P_{\text{replay}}^{(t)}$: The probability of SWE-Replay generating the correct solution at trial t .
2. $P_{\text{select}}^{(t)}$: The probability that the step selection mechanism chooses a step belonging to the correct trajectory, if the correct trajectory already exists in the archive.
3. $P'_{\text{select}}(t)$: The probability of selecting a step from an incorrect trajectory but still randomly stumbling upon the solution (negligible for hard tasks).

We aim to show: **provided $P_{\text{select}}^{(t)}$ exceeds the probability of random selection, SWE-Replay satisfies $P_{\text{replay}}^{(t)} \geq p$.**

At each trial, SWE-Replay generates from scratch with probability 0.5 and also replays from archive with probability 0.5. The marginal probability of success at trial t is:

$$\begin{aligned}
 P_{\text{replay}}^{(t)} &= \frac{1}{2}p + \frac{1}{2} \left[\underbrace{\mathbb{P}(\text{Archive has correct})}_{=1-(1-p)^{t-1}} \cdot P_{\text{select}}^{(t)} \right. \\
 &\quad \left. + \underbrace{\mathbb{P}(\text{Archive has only incorrect})}_{=(1-p)^{t-1}} \cdot P'_{\text{select}}(t) \right]. \quad (3)
 \end{aligned}$$

Dropping the negligible term $P'_{\text{select}}(t)$ and requiring $P_{\text{replay}}^{(t)} \geq p$, we obtain the inequality:

$$\frac{1}{2}p + \frac{1}{2} (1 - (1-p)^{t-1}) P_{\text{select}}^{(t)} \geq p. \quad (4)$$

Rearranging for $P_{\text{select}}^{(t)}$:

$$P_{\text{select}}^{(t)} \geq \frac{p}{1 - (1 - p)^{t-1}}. \quad (5)$$

Since p is small (Assumption 4), we can approximate the denominator using the first-order Taylor expansion, $(1 - p)^{t-1} \approx 1 - (t - 1)p$. This simplifies the condition to:

$$P_{\text{select}}^{(t)} \gtrsim \frac{p}{(t - 1)p} = \frac{1}{t - 1}. \quad (6)$$

This lower bound, $\frac{1}{t-1}$, corresponds precisely to the probability of *random selection*. At trial t , the archive contains $t - 1$ trajectories, and the probability of uniformly selecting a step from the correct trajectory is $\frac{N}{(t-1)N} = \frac{1}{t-1}$. Consequently, inequality (6) implies that as long as the archive has contained the correct trajectory and the SWE-Replay selection mechanism assigns higher probability to steps on the correct trajectory than uniform random sampling, SWE-Replay will outperform naive scaling ($P_{\text{replay}}^{(t)} \geq p$). In practice, by filtering low-quality trajectories and prioritizing reasoning-intensive steps, SWE-Replay biases selection toward the correct trajectory, as validated in §3.3.

5. Related Work

5.1. Scaffolding for Software Engineering Agents

Drawing inspiration from human debugging loops, where developers learn from environmental feedback such as test failures, early interactive solutions (Xia & Zhang, 2024; Chen et al., 2023) pioneered bug-fixing via Large Language Models (LLMs). Since those initial efforts, a significant body of research regarding bug fixing and general coding tasks has focused on automatically providing LLMs with greater context through multi-turn conversations (Yang et al., 2024; Wang et al., 2024). More recently, foundation LLMs have demonstrated substantial improvements in reasoning and tool usage, enabling the development of feedback-driven solutions that leverage these emergent capabilities (Carbonneaux et al., 2025; Ding et al., 2025b). In March 2024, the field saw the release of Devin AI⁷, the first software engineer agent designed to autonomously complete end-to-end software tasks, such as resolving GitHub issues. The initial release of Devin demonstrated impressive performance on the SWE-Bench dataset, which consists of thousands of real-world GitHub issues (Jimenez et al., 2023). Following this milestone, numerous dedicated software agent scaffolds have been proposed, including SWE-agent (Yang et al., 2024), OpenHands (Wang et al., 2024), AutoCodeRover (Zhang et al., 2024), Trae Agent (Gao et al., 2025), and Live-SWE-agent (Xia et al., 2025). These systems typically equip LLMs with a suite of coding tools, enabling the models to determine the necessary actions to complete real-world software tasks autonomously.

⁷<https://cognition.ai/blog/introducing-devin>

5.2. Test-Time Scaling for Software Engineering Agents

Allocating test-time computation to sample multiple candidate patches has emerged as a highly effective strategy for software engineering agents. Existing works demonstrate that repeated sampling with smaller LLMs often yields superior solution coverage compared to single-shot outputs from larger models (Brown et al., 2024). This approach is particularly advantageous in software engineering tasks, where execution-based feedback facilitates the rigorous selection of valid code (Ehrlich et al., 2025; Ma et al., 2025). To mitigate the computational costs associated with such scaling, SWE-Search (Antoniades et al., 2024) has integrated tree search algorithms to prune unpromising trajectories early, relying on a separate agent to estimate quality. Similarly, Satori-SWE (Zeng et al., 2025) proposes to achieve sample-efficient test-time scaling by using LLM to self-improve the scores of its prior generations assigned by a reward model. However, this work presents two significant drawbacks: (1) quality scores are often skewed by model overconfidence and miscalibration (Son et al., 2024), introducing noise into the scaling process; (2) SWE-Search and Satori-SWE are designed specifically for the tool designs of pipeline-based scaffolds (e.g., Moatless) and not directly generalizable to modern agents. In contrast, SWE-Replay bypasses reliance on potentially inaccurate predictions, is naturally generalizable to all agentic scaffolds, and employs a streamlined select-and-replay mechanism to ensure scalability.

6. Conclusion

In this work, we presented SWE-Replay, a novel efficient test-time scaling framework that reconciles the tension between computational efficiency and solution quality for software engineering agents by recycling previously sampled trajectories, dynamically choosing to either explore from scratch or exploit archived experience by branching at critical intermediate steps. Critical intermediate steps are selected based on the potential and reasoning significance of repository exploration, rather than LLM-as-a-Judge.

Through evaluation on SWE-Bench Verified, SWE-Bench Pro, and SWE-Bench Multilingual across multiple LLM backends and agentic scaffolds, we demonstrated that: (1) SWE-Replay reduces computational costs by up to 17.4% while improving resolve rates by up to 3.8% on SWE-Bench Verified, and (2) SWE-Replay showcases consistent generalizability towards diverse types of problems in SWE-Bench Pro and SWE-Bench Multilingual. We further validate empirically that, compared with naive scaling, SWE-Replay enables agents to explore more diverse, less visited repository spaces. These empirical gains are supported by our theoretical intuition, which provides a foundational understanding of how replaying optimizes the quality of exploration.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Antoniades, A., Örwall, A., Zhang, K., Xie, Y., Goyal, A., and Wang, W. Swe-search: Enhancing software agents with monte carlo tree search and iterative refinement. *arXiv preprint arXiv:2410.20285*, 2024.
- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Brown, B., Juravsky, J., Ehrlich, R., Clark, R., Le, Q. V., Ré, C., and Mirhoseini, A. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*, 2024.
- Carbonneaux, Q., Cohen, G., Gehring, J., Kahn, J., Kossen, J., Kreuk, F., McMilin, E., Meyer, M., Wei, Y., Zhang, D., et al. Cwm: An open-weights llm for research on code generation with world models. *arXiv preprint arXiv:2510.02387*, 2025.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto, H. P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F. P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W. H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A. N., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., and Zaremba, W. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Chen, X., Lin, M., Schärli, N., and Zhou, D. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.
- Comanici, G., Bieber, E., Schaekermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Deng, X., Da, J., Pan, E., He, Y. Y., Ide, C., Garg, K., Lauffer, N., Park, A., Pasari, N., Rane, C., et al. Swe-bench pro: Can ai agents solve long-horizon software engineering tasks? *arXiv preprint arXiv:2509.16941*, 2025.
- Ding, Y., Liu, J., Wei, Y., and Zhang, L. XFT: Unlocking the power of code instruction tuning by simply merging upcycled mixture-of-experts. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12941–12955, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.699. URL <https://aclanthology.org/2024.acl-long.699/>.
- Ding, Y., Ding, H., Wang, S., Sun, Q., Kumar, V., and Wang, Z. Planning-aware code infilling via horizon-length prediction. In Christodoulopoulos, C., Chakraborty, T., Rose, C., and Peng, V. (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 32930–32942, Suzhou, China, November 2025a. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.1672. URL <https://aclanthology.org/2025.emnlp-main.1672/>.
- Ding, Y., Le, H., Han, S., Ruan, K., Jin, Z., Kumar, V., Wang, Z., and Deoras, A. Empowering multi-turn tool-integrated reasoning with group turn policy optimization, 2025b. URL <https://arxiv.org/abs/2511.14846>.
- Ehrlich, R., Brown, B., Juravsky, J., Clark, R., Ré, C., and Mirhoseini, A. Codemonkeys: Scaling test-time compute for software engineering. *arXiv preprint arXiv:2501.14723*, 2025.
- Gao, P., Tian, Z., Meng, X., Wang, X., Hu, R., Xiao, Y., Liu, Y., Zhang, Z., Chen, J., Gao, C., et al. Trae agent: An llm-based agent for software engineering with test-time scaling. *arXiv preprint arXiv:2507.23370*, 2025.
- Jimenez, C. E., Yang, J., Wettig, A., Yao, S., Pei, K., Press, O., and Narasimhan, K. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023.
- Kim, J., Shin, B., Chung, J., and Rhu, M. The cost of dynamic reasoning: Demystifying ai agents and test-time scaling from an ai infrastructure perspective, 2026. URL <https://arxiv.org/abs/2506.04301>.
- Li, Y., Choi, D., Chung, J., Kushman, N., Schrittwieser, J., Leblond, R., Eccles, T., Keeling, J., Gimeno, F., Lago, A. D., Hubert, T., Choy, P., de Masson d’Autume, C., Babuschkin, I., Chen, X., Huang, P.-S., Welbl, J., Goyal,

- S., Cherepanov, A., Molloy, J., Mankowitz, D. J., Robson, E. S., Kohli, P., de Freitas, N., Kavukcuoglu, K., and Vinyals, O. Competition-level code generation with alphacode. *arXiv preprint arXiv:2203.07814*, 2022.
- Liu, J., Wang, K., Chen, Y., Peng, X., Chen, Z., Zhang, L., and Lou, Y. Large language model-based agents for software engineering: A survey. *arXiv preprint arXiv:2409.02977*, 2024.
- Ma, Y., Li, Y., Dong, Y., Jiang, X., Cao, R., Chen, J., Huang, F., and Li, B. Thinking longer, not larger: Enhancing software engineering agents via scaling test-time compute, 2025. URL <https://arxiv.org/abs/2503.23803>.
- Son, G., Ko, H., Lee, H., Kim, Y., and Hong, S. Llm-as-a-judge & reward model: What they can and cannot do. *arXiv preprint arXiv:2409.11239*, 2024.
- Wang, X., Li, B., Song, Y., Xu, F. F., Tang, X., Zhuge, M., Pan, J., Song, Y., Li, B., Singh, J., Tran, H. H., Li, F., Ma, R., Zheng, M., Qian, B., Shao, Y., Muennighoff, N., Zhang, Y., Hui, B., Lin, J., Brennan, R., Peng, H., Ji, H., and Neubig, G. Openhands: An open platform for ai software developers as generalist agents. *arXiv preprint arXiv:2407.16741*, 2024.
- Xia, C. S. and Zhang, L. Automated program repair via conversation: Fixing 162 out of 337 bugs for \$0.42 each using chatgpt. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 819–831, 2024.
- Xia, C. S., Deng, Y., Dunn, S., and Zhang, L. Agentless: Demystifying llm-based software engineering agents. *arXiv preprint arXiv:2407.01489*, 2024.
- Xia, C. S., Wang, Z., Yang, Y., Wei, Y., and Zhang, L. Live-swe-agent: Can software engineering agents self-evolve on the fly? *arXiv preprint arXiv:2511.13646*, 2025.
- Yang, J., Jimenez, C. E., Wettig, A., Lieret, K., Yao, S., Narasimhan, K., and Press, O. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.
- Zeng, G., Shen, M., Chen, D., Qi, Z., Das, S., Gutfreund, D., Cox, D., Wornell, G., Lu, W., Hong, Z.-W., and Gan, C. Satori-swe: Evolutionary test-time scaling for sample-efficient software engineering, 2025. URL <https://arxiv.org/abs/2505.23604>.
- Zhang, J., Hu, S., Lu, C., Lange, R., and Clune, J. Darwin godel machine: Open-ended evolution of self-improving agents, 2025a. URL <https://arxiv.org/abs/2505.22954>.
- Zhang, Q., Lyu, F., Sun, Z., Wang, L., Zhang, W., Hua, W., Wu, H., Guo, Z., Wang, Y., Muennighoff, N., et al. A survey on test-time scaling in large language models: What, how, where, and how well? *arXiv preprint arXiv:2503.24235*, 2025b.
- Zhang, Y., Ruan, H., Fan, Z., and Roychoudhury, A. Autocoderover: Autonomous program improvement. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 1592–1604, 2024.

A. Appendix

Algorithm 1 SWE-Replay

Require: Issue Description D , Agent \mathcal{A} , Budget N

Ensure: Final Patch P^*

```

1: Initialize Trajectory Archive  $\mathcal{T} \leftarrow \emptyset$ 
2: for  $i = 1$  to  $N$  do
3:    $mode \leftarrow \text{EXPLORE}$ 
4:   if  $i > 1$  then
5:      $mode \leftarrow \text{Bernoulli}(0.5) ? \text{EXPLORE} : \text{EXPLOIT}$ 
6:   end if
7:   if  $mode = \text{EXPLORE}$  then
8:      $S_{start} \leftarrow \text{InitializeEnv}(D)$ 
9:      $\tau_{new} \leftarrow \mathcal{A}.\text{Run}(S_{start}, \text{context} = \emptyset)$ 
10:  else
11:     $s_{selected} \leftarrow \text{SelectStep}(\mathcal{T})$ 
12:     $S_{resumed} \leftarrow \text{RestoreEnv}(s_{selected})$ 
13:     $C_{resumed} \leftarrow \text{GetContext}(s_{selected})$ 
14:     $\tau_{suffix} \leftarrow \mathcal{A}.\text{Run}(S_{resumed}, C_{resumed})$ 
15:     $\tau_{new} \leftarrow \text{Concatenate}(C_{resumed}, \tau_{suffix})$ 
16:  end if
17:   $\mathcal{T} \leftarrow \mathcal{T} \cup \{\tau_{new}\}$ 
18: end for
19:  $\mathcal{P}_{candidates} \leftarrow \{\text{GetPatch}(\tau) \mid \tau \in \mathcal{T}\}$ 
20:  $\mathcal{P}_{valid} \leftarrow \text{FilterTestFailures}(\mathcal{P}_{candidates})$ 
21:  $P^* \leftarrow \text{MajorityVote}(\mathcal{P}_{valid})$ 

```

A.1. Implementation Details

Following the default setting of mini-SWE-agent (Yang et al., 2024), we set a maximum step limit of 250 for Gemini-2.5-Pro and Gemini-3-Pro. For Devstral-Small-2, we set a maximum step limit of 128 to accelerate the experiments. We ignore the cost limit in our experiments to **unsure** that the agents have enough opportunities to explore the repositories without being affected by the external model pricing. Based on our experience, we set a temperature of 0.8 for Gemini-2.5-Pro and a temperature of 0.2 for Devstral-Small-2 and Gemini-3-Pro to achieve the best performance. This also gives us a chance to observe the performance of SWE-Replay under different levels of temperatures. We sample ten trajectories per problem during scaling for all the experiments on SWE-Bench Verified and SWE-Bench Multilingual. Due to the large size of SWE-Bench Pro, we sample five trajectories per problem on this benchmark.

While we follow the Agentless (Xia et al., 2024) pipeline to select the final patch, we remove LLM-generated reproduction tests in the pipeline to remove any potential noise. So our final selection pipeline is: (1) run the regression tests for all the candidate patches, (2) filter out all the candidate patches with regression failures, and (3) select the final patch out of the remaining ones using majority voting.

A.2. Benchmark Details

The SWE-Bench Verified benchmark comprises 500 SWE tasks, where the objective is to modify a repository based on a provided description. These instances are human-validated to ensure that the problem descriptions contain sufficient information for a solution. To further test the agent on realistic, enterprise-level complexity, we also employ SWE-Bench Pro and SWE-Bench Multilingual. SWE-Bench Pro comprises 731 publicly available problems and differs from SWE-Bench Verified by presenting higher difficulty levels across a diverse range of repositories and programming languages. Similarly, SWE-Bench Multilingual is a benchmark of 300 software engineering tasks across 9 programming languages (JavaScript, TypeScript, Rust, Ruby, Go, C/C++, PHP, and Java).

SWE-Bench Verified mini is a subset of SWE-Bench Verified that uses 50 instead of 500 data points, requires 5GB instead of 130GB of storage, and has approximately the same distribution of performance, test pass rate, and difficulty as the original

dataset. It includes the following problems from SWE-Bench Verified:

- django__django-11790
- django__django-11815
- django__django-11848
- django__django-11880
- django__django-11885
- django__django-11951
- django__django-11964
- django__django-11999
- django__django-12039
- django__django-12050
- django__django-12143
- django__django-12155
- django__django-12193
- django__django-12209
- django__django-12262
- django__django-12273
- django__django-12276
- django__django-12304
- django__django-12308
- django__django-12325
- django__django-12406
- django__django-12708
- django__django-12713
- django__django-12774
- django__django-9296
- sphinx-doc__sphinx-10323
- sphinx-doc__sphinx-10435
- sphinx-doc__sphinx-10466
- sphinx-doc__sphinx-10673
- sphinx-doc__sphinx-11510
- sphinx-doc__sphinx-7590
- sphinx-doc__sphinx-7748
- sphinx-doc__sphinx-7757
- sphinx-doc__sphinx-7985
- sphinx-doc__sphinx-8035
- sphinx-doc__sphinx-8056
- sphinx-doc__sphinx-8265
- sphinx-doc__sphinx-8269
- sphinx-doc__sphinx-8475
- sphinx-doc__sphinx-8548
- sphinx-doc__sphinx-8551
- sphinx-doc__sphinx-8638
- sphinx-doc__sphinx-8721
- sphinx-doc__sphinx-9229
- sphinx-doc__sphinx-9230
- sphinx-doc__sphinx-9281
- sphinx-doc__sphinx-9320
- sphinx-doc__sphinx-9367
- sphinx-doc__sphinx-9461
- sphinx-doc__sphinx-9698

A.3. Metric Details

To evaluate computational efficiency, we record the cumulative number of input and output tokens across all ten trajectories for each problem. However, raw token counts fail to capture the practical benefits of *prompt caching*, an inference optimization that stores the processed states of repeated text sequences to reduce latency and cost.

The branching mechanism of SWE-Replay, which resumes execution from intermediate steps, naturally yields substantial shared prefixes, thereby maximizing the utility of prompt caching. Although current commercial LLM backends (e.g., Gemini-2.5-Pro, Gemini-3-Pro) support caching, their default *implicit* prompt caching policies do not always guarantee deterministic cost savings. Consequently, to provide a precise assessment of the cost reductions enabled by our method, we report the cost after enabling *explicit* prompt caching. Specifically, we store all the prior generations in the prompt cache during scaling and report the cost of input/output tokens from LLMs.

A.4. Ablation Details for LLM-as-a-Judge Variant

To replace the original step selection mechanism with LLM-as-a-Judge in SWE-Replay, we use the same LLM that the agent is using (i.e., Devstral-Small-2 in this experiment) to generate a quality score $q_i \in [0, 1]$ for each step s_i in the archive. Then, similar to §2.1.3, we assign sampling probabilities for each step as follows:

$$p(s_i) = \text{softmax}(q_i) = \frac{e^{q_i}}{\sum_j e^{q_j}}. \tag{7}$$

In our experiment, we use Devstral-Small-2 to estimate a quality score for each step based on the following prompt:

Quality estimate prompt for LLM-as-a-Judge

```

Your task is to evaluate the quality of a middle step generated by an AI
↪ agent in a long trajectory. Your evaluation will focus on whether the
↪ step was well-constructed, whether the resulting execution outcome is
↪ relevant and useful for solving the problem at hand, and whether the
↪ analysis in the step is appropriate and helpful.

Here is the full trajectory:
"""
{trajectory}
"""

Here is the step you need to evaluate:
"""
{step}
"""

On the scale of 0 to 1 with a scale of 0.01, where 0 is extremely bad and 1
↪ is extremely good, rate the quality of this step in the format of "score:
↪ ".

```

Figure 8. Prompt for Devstral-Small-2 to estimate the quality score for each step in the mini-SWE-agent trajectories.

A.5. Effectiveness of File-Level Abstraction

We evaluate different levels of granularity for the abstract representation of steps (§2.1.2): (1) **Line**: the set of all raw observations before the step; (2) **Method**: the set of all explored methods before the step; and (3) **File**: the set of all explored repository files before the step. As shown in Table 5, the file-level abstraction yields the highest resolve rate (60.0%) and the lowest cost (\$1.52). In contrast, both method-level and line-level abstraction perform worse than file-level. This suggests that line-level and method-level representations are too granular and sparse, leading to fragmented groups that hinder effective reuse. File-level abstraction offers the optimal balance, grouping semantically related steps to maximize the benefits of exploration-exploitation.

Table 5. Ablation on different step representations.

Abstraction Level	% Resolved ↑	Avg. Cost (\$) ↓
Line	58.0%	1.58
Method	60.0%	1.57
File	60.0%	1.52

A.6. Impact of Reasoning Intensity Representation

We verify our design choice for representing reasoning intensity (§2.1.3) by comparing two variants: using the raw length of reasoning content and our proposed method, which uses the number of reasoning paragraphs. As shown in Table 6, using the number of paragraphs yields the largest performance gains (60.0%) while minimizing cost (\$1.52). This suggests that raw length is a noisy proxy for reasoning quality. In contrast, paragraph count serves as a lightweight yet accurate heuristic for identifying reasoning-intensive steps that are worthy of replay.

Table 6. Ablation on different reasoning intensity representations.

Reasoning Intensity	% Resolved ↑	Avg. Cost (\$) ↓
Length	58.0%	1.55
# Paragraphs	60.0%	1.52

A.7. Impact of Softmax Normalization

SWE-Replay utilizes a normalized probability distribution to balance exploration and exploitation during step selection. We compare two normalization techniques: *Unit Sum* ($\frac{v_i}{\sum v_i}$) and *Softmax* ($\frac{e^{v_i}}{\sum e^{v_i}}$). In the context of our step selection mechanism, Unit Sum tends to produce a sharper distribution that may prematurely converge on specific steps. In contrast, Softmax produces a flatter distribution that better preserves diversity. As shown in Table 7, replacing Softmax with Unit Sum results in a marked reduction in both performance (56.0% vs. 60.0%) and efficiency (\$1.63 vs. \$1.52), confirming that Softmax normalization is critical for maintaining an effective diversity of sampling candidates.

Table 7. Ablation on normalization methods for step selection.

Normalization	% Resolved \uparrow	Avg. Cost (\$) \downarrow
Unit Sum	56.0%	1.63
Softmax	60.0%	1.52