

EvoScale: Evolutionary Test-Time Scaling for Software Engineering

Anonymous authors

Paper under double-blind review

ABSTRACT

Language models (LMs) perform well on standardized coding benchmarks but struggle with real-world software engineering tasks such as resolving GitHub issues in SWE-Bench—especially when model parameters are less than 100B. While smaller models are preferable in practice due to their lower computational cost, improving their performance remains challenging. Existing approaches primarily rely on supervised fine-tuning (SFT) with high-quality data, which is expensive to curate at scale. An alternative is test-time scaling: generating multiple outputs, scoring them using a verifier, and selecting the best one. Although effective, this strategy often requires excessive sampling and costly scoring, limiting its practical application. We propose Evolutionary Test-Time Scaling (EvoScale), a sample-efficient method that treats generation as an evolutionary process. By iteratively refining outputs via selection and mutation, EvoScale shifts the output distribution toward higher-scoring regions, reducing the number of samples needed to find correct solutions. To reduce the overhead from repeatedly sampling and selection, we train the model to *self-evolve* using reinforcement learning (RL). Rather than relying on external verifiers at inference time, the model learns to self-improve the scores of its own generations across iterations. Evaluated on SWE-Bench-Verified, EvoScale enables a 32B model to match or exceed the performance of models with over 100B parameters while using a few samples. Code, data, and models will be fully open-sourced.

1 INTRODUCTION

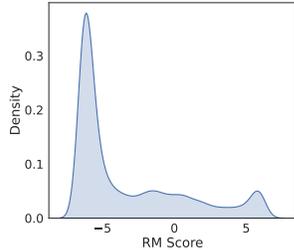
Language models (LMs) perform well on coding benchmarks like HumanEval (Chen et al., 2021) or LiveCodeBench (Jain et al., 2025a) but struggle with real-world software engineering (SWE) tasks (Jimenez et al., 2024). Unlike standardized coding problems, real issues—such as GitHub issues (Jimenez et al., 2024)—are often under-specified and require reasoning across multiple files. Even large models like Claude reach only around 60% accuracy on SWE-bench (Jimenez et al., 2024), despite using carefully engineered prompting pipelines (Xia et al., 2024). Smaller models (under 100B parameters) perform significantly worse, typically scoring below 10% in zero-shot settings and plateauing around 30% after supervised fine-tuning (SFT) (Xie et al., 2025; Pan et al., 2025) on GitHub issue datasets. Improving the performance of these models remains a key challenge for practical deployment, where repeatedly querying large models is often too costly or inefficient.

Recent and concurrent works to improve the performance of small LMs on SWE tasks have mainly focused on expanding SFT datasets—either through expert annotation or distillation from larger models (Yang et al., 2025; Xie et al., 2025; Pan et al., 2025). These approaches show that performance improves as the quality and quantity of training data increase. However, collecting such data is both costly and time-consuming.

An alternative is *test-time scaling*, which improves performance by generating multiple outputs at inference and selecting the best one using a verifier, such as a reward model (Cobbe et al., 2021; Lightman et al., 2023). While widely applied in math and logical reasoning (Hoffmann et al., 2022; Snell et al., 2025), test-time scaling remains underexplored in SWE. Yet it shows strong potential: prior works (Pan et al., 2025; Brown et al., 2024) demonstrate that small models can generate correct solutions when sampled many times. Specifically, their $\text{pass}@N$, the probability that at least one of N samples is correct, is close to the $\text{pass}@1$ performance of larger models. This indicates that small models *can* produce correct solutions; the challenge lies in efficiently identifying them.

054 Test-time scaling assumes that among many sampled outputs, at least one will be correct. However,
 055 when correct solutions are rare, these methods often require a large number of samples to succeed.
 056 This is particularly costly in SWE tasks, where generating each sample is slow due to long code
 057 contexts, and scoring is expensive when unit tests execution is needed (Xia et al., 2024). Recent
 058 work (Wei et al., 2025) uses reinforcement learning (RL) (Wei et al., 2025) to enhance the reasoning
 059 capabilities of LMs for improved output quality but still requires hundreds of code edits (i.e., patch
 060 samples) per issue. Also, Pan et al. (2025) depends on slow interactions with the runtime environment
 061 in agentic workflows. This motivates the need for *sample-efficient* test-time scaling methods that can
 062 identify correct solutions with fewer samples.

063 In this paper, we propose **Evolutionary Test-Time Scaling**
 064 (**EvoScale**), a sample-efficient method for improving test-time per-
 065 formance on SWE tasks. Existing test-time scaling methods often
 066 require an excessive number of samples because model outputs are
 067 highly dispersed—correct solutions exist but are rare, as shown in
 068 Figure 1. EvoScale mitigates this by progressively steering genera-
 069 tion toward higher-scoring regions, reducing the number of samples
 070 needed to find correct outputs. Inspired by evolutionary algorithms
 071 (Shen et al., 2023; Wierstra et al., 2014; Hansen, 2016; Salimans
 072 et al., 2017), EvoScale iteratively refines candidate patches through
 073 *selection* and *mutation*. Instead of consuming the sample budget in a
 074 single pass, EvoScale amortizes it over multiple iterations: the model
 075 generates a batch of outputs, a scoring function selects the top ones,
 076 and the next batch is generated by conditioning on these—effectively
 077 mutating prior outputs. Early iterations focuses on exploration; later
 078 ones focus on exploitation.



078 Figure 1: Reward score distribution of SFT outputs, with
 079 high-scoring outputs concentrated in the long tail.

078 Although EvoScale improves sample efficiency, the selection step still incurs overhead: like standard
 079 evolutionary algorithms (Wierstra et al., 2014), it generates more outputs than needed and filters only
 080 the high-scoring ones, increasing sampling and computation costs. To eliminate this, we use RL to
 081 internalize the reward model’s guidance into the model itself, enabling it to *self-evolve*—refining
 082 its own outputs without external reward models at inference. We formulate this as a potential-
 083 based reward maximization problem (Ng et al., 1999), where the model learns to improve output
 084 scores across iterations based on score differences. This avoids discarding low-scoring outputs and
 085 reduces sample usage per iteration. Our theoretical analysis shows that this RL objective ensures
 086 monotonic score improvement across iterations. We evaluate the proposed EvoScale method on
 087 SWE-Bench-Verified (Jimenez et al., 2024), and summarize our key contributions as follows:

- 088 • A new perspective of formulating test-time scaling as an evolutionary process, improving sample
 089 efficiency for software engineering tasks.
- 090 • A novel RL training approach that enables self-evolution, eliminating the need for external reward
 091 models or verifiers at inference time.
- 092 • Empirical results showing that a 32B model with EvoScale achieves performance comparable to
 093 models exceeding 100B parameters, while requiring only a small number of samples

094 2 RELATED WORK

096 **Dataset Curation for SWE.** Prior works (Ma et al., 2024; Pan et al., 2025) and concurrent ef-
 097 forts (Yang et al., 2025; Jain et al., 2025b) use proprietary LLMs (e.g., Claude, GPT-4) as autonomous
 098 agents to collect SFT data by recording step-by-step interactions in sandboxed runtime environ-
 099 ments. While this automates the data collection process for agent-style training (Yang et al., 2024),
 100 it involves substantial engineering overhead (e.g., Docker setup, sandboxing) and high inference
 101 costs. In contrast, Xie et al. (2025) uses a pipeline-based framework (Xia et al., 2024), collecting
 102 real pull-request-issue pairs and prompting GPT-4 to generate CoT traces and ground-truth patches
 103 without runtime interaction. Though easier to collect, this data requires careful noise filtering. Our
 104 approach instead improves small models’ performance by scaling the computation at test time.

105 **Test-time scaling for SWE.** Xia et al. (2024) showed that sampling multiple patches and selecting
 106 the best one based on unit test results in sandboxed environments improves performance. Unit tests
 107 have since been widely adopted in SWE tasks (Wei et al., 2025; Ehrlich et al., 2025; Jain et al.,
 2025b; Brown et al., 2024). Other works (Pan et al., 2025; Jain et al., 2025b; Ma et al., 2025) train

verifiers or reward models to score and select patches. To reduce the cost of interacting with runtime environments in agentic frameworks (Yang et al., 2024), some methods (Ma et al., 2025; Antoniadou et al., 2025) integrate tree search, pruning unpromising interaction paths early. While prior works improve patch ranking or interaction efficiency, our focus is on reducing the number of samples needed for effective test-time scaling.

RL for SWE. Pan et al. (2025) used a basic RL approach for SWE tasks, applying rejection sampling to fine-tune models on their own successful trajectories. Wei et al. (2025) later used policy gradient RL (Shao et al., 2024), with rewards based on string similarity to ground truth patches, showing gains over SFT. In contrast, our method trains the model to iteratively refine its past outputs, improving scores over time. We also use a learned reward model that classifies patches as correct or incorrect, which outperforms string similarity as shown in Appendix A.

Improve Test-time Scaling Efficiency. Recent work has explored making test-time scaling more efficient by improving accuracy under limited sampling budgets. Snell et al. (2025) study optimal strategies for allocating sampling budgets across different scaling methods. Puri et al. (2025) introduce probabilistic resampling techniques that concentrate compute on promising candidates. Tan et al. (2025) leverage process-level reward models to encourage models to “think longer” only when needed. Other approaches reduce redundant sampling through calibrated stopping or confidence-based pruning (Huang et al., 2025; Wu et al., 2025). While these techniques improve sample efficiency using a fixed LLM, our work instead trains the model to learn iteratively refine and improve its outputs, shifting the sampling distribution toward high-scoring regions.

3 PRELIMINARIES



Figure 2: **Pipeline for SWE Tasks.** Given a GitHub issue, the retriever identifies the code files most relevant to the issue. The code editor then generates a code patch to resolve it.

Software engineering (SWE) tasks. We study the problem of using LMs to resolve real-world GitHub issues, where each issue consists of a textual description and a corresponding code repository. Since issues are not self-contained, solving them requires identifying and modifying relevant parts of the codebase. There are two main paradigms for solving SWE tasks with LMs: agentic (Yang et al., 2024) and pipeline-based (Xia et al., 2024; Wei et al., 2025). Agentic methods allow the model to interact with the runtime environment, such as browsing files, running shell commands, and editing code through tool use. While flexible, these approaches are computationally intensive and rely on long-context reasoning, making them less practical for small models. In contrast, pipeline-based methods decompose the task into subtasks, typically retrieval and editing, and solve each without runtime interaction, which is more computationally efficient and suited for small models. Retrieval refers to identifying the files or functions relevant to the issue, while editing involves generating the code changes needed to resolve it.

Formally, given an issue description x , the goal is to produce a code edit (i.e., patch) y that fixes the bug or implements requested changes. A retrieval model selects a subset code context $C(x) \subseteq \mathcal{C}$ from the full codebase \mathcal{C} , and an editing model π generates the patch $y = \pi(x, C(x))$ by modifying $C(x)$. While retrieval has reached around 70% accuracy in prior work Xie et al. (2025); Xia et al. (2024), editing remains the main bottleneck. This work focuses on improving editing performance in pipeline-based settings, using off-the-shelf localization methods in experiments. In this setup, the dominant cost comes from sampling and scoring outputs from the editing model at test time.

Goal: Sample-efficient test-time scaling. Test-time scaling improves performance by selecting the best output from multiple samples (Hoffmann et al., 2022; Snell et al., 2025), but often requires a large number of generations, especially in SWE tasks (Wei et al., 2025). Our goal is to enable test-time scaling more sample-efficient, achieving stronger performance with fewer samples.

Why is test-time scaling sample-inefficient in SWE task? Given a sample budget N , typical test-time scaling methods in SWE (Xia et al., 2024; Wei et al., 2025; Jain et al., 2025b; Pan et al.,

2025) draw N outputs (patches) $\{y_i\}_{i=1}^N$ from a frozen editor model π , score them with a score function R (e.g., reward model or unit tests), and selects the best one $\arg \max_{y_i} R(x, y_i)$. While high-scoring outputs near the mode could be sampled easily, the challenge of test-time scaling is to identify high-scoring outputs from the tail of $\pi(\cdot | x, C(x))$, especially for hard issues. However, doing so typically requires a large sample size N , making the process sample-inefficient.

4 METHOD: EVOLUTIONARY TEST-TIME SCALING

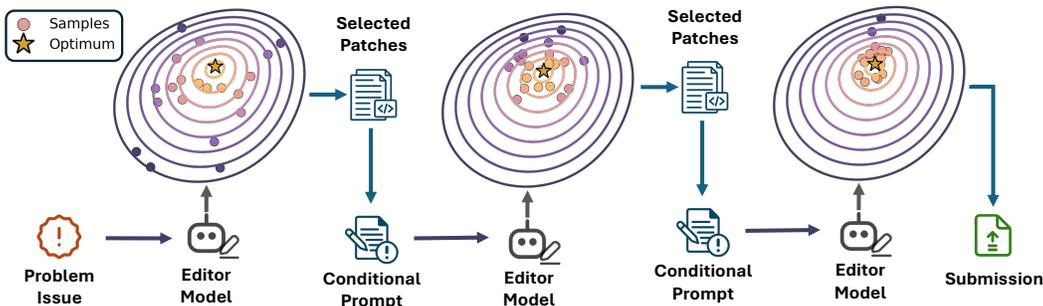


Figure 3: **An Overview of Evolutionary Test-Time Scaling.** Given a GitHub issue x and its code context $C(x)$, the editor model π first generates a batch of candidate patches \mathcal{Y}^t . The reward landscape is illustrated with contour lines, where brighter contours indicate a higher score of a scoring function R (e.g., reward model or unit tests). A set of patches \mathcal{E}^t is selected (e.g., via a scoring function R) and combined with x and $C(x)$ to form a conditional prompt (see Section 4.1), which guides the model to generate the next batch $\mathcal{Y}^{t+1} = \{y_1^{t+1}, \dots, y_M^{t+1}\}$, increasingly concentrated around the optimum. The process continues under a fixed sampling budget until convergence, after which the final patch is submitted.

This motivates our method, Evolutionary Test-Time Scaling (EvoScale), which iteratively refines generation by using earlier outputs to guide subsequent sampling. We recast patch generation for a GitHub issue as an evolutionary process. The objective is to explore the patch space with a small number of samples, identify high-scoring patches, and iteratively refine the generated patches. As shown in Figure 3, initial samples are scattered and far from the correct solutions (denoted by stars), but over iterations, the distribution shifts closer to the correct solution. Through evolution, EvoScale more efficiently uncovers high-scoring outputs in long tails. We formulate the problem in Section 4.1 and detail the training procedure in Sections 4.2 and 4.3.

4.1 FORMULATION: PATCH GENERATION AS EVOLUTION

We amortize the sampling budget over T iterations by generating $M < N$ samples per iteration, rather than sampling all N at once. The goal is to progressively improve sample quality across iterations. A key challenge lies in effectively using early samples to guide later ones. Typical evolutionary strategies select top-scoring candidates and mutate them—often by adding random noise—to steer future samples toward high-scoring regions. However, in SWE tasks, where patches are structured code edits, random perturbations often break syntax or semantics (e.g., undefined variables, etc).

Algorithm. Instead of using random noise for mutation, we use a language model (LM) as a *mutation operator*, leveraging its ability to produce syntactically and semantically valid patches. At each iteration t , the LM generates a batch of patches $\mathcal{Y}^{t+1} = \{y_1^{t+1}, \dots, y_M^{t+1}\}$ conditioned on a set of prior patches \mathcal{E}^t : $\mathcal{Y}^{t+1} \sim \pi(\cdot | x, C(x), \mathcal{E}^t)$. We refer to \mathcal{E}^t as *conditioning examples* consisting of patches generated at iteration t . Following the selection step in evolutionary algorithms, \mathcal{E}^t could be selected as the top- K patches ranked by a scoring function R (i.e., fitness function in evolutionary algorithms). Note that we find that our model after training can self-evolve without this selector (see Section 4.3 and Section 5.2), so this step is optional. The full procedure is detailed in Algorithm 1.

Question: Can a language model naturally perform mutation? Ideally, the mutation operator should generate patches that improve scores. However, as shown in Section 5.2, models trained with classical SFT—conditioned only on the issue and code context—struggle to refine existing patches. In the next section, we present our approach to overcome this limitation.

4.2 SMALL-SCALE MUTATION SUPERVISED FINE-TUNING

Classical supervised fine-tuning (SFT) fails at mutation because it never learns to condition on previous patches. To train the model for mutation, it must observe *conditioning examples*—patches from previous iterations—so it can learn to refine them. In EvoScale, conditioning examples are drawn from the model’s earlier outputs. We introduce a two-stage supervised fine-tuning (SFT) process: classical SFT followed by mutation SFT. The classical SFT model is first trained and then used to generate conditioning examples for training the mutation SFT model.

Stage 1 — Classical SFT. We fine-tune a base model on inputs consisting of the issue description x and code context $C(x)$, with targets that include a chain-of-thought (CoT) trace and the ground-truth patch, jointly denoted as y_{SFT}^* . Following prior work on dataset curation (Yang et al., 2025; Xie et al., 2025), we use a teacher model μ (e.g., a larger LLM; see Section 5.1) to generate CoT traces. The training objective is:

$$\max_{\pi_{\text{SFT}}} \mathbb{E}_{x \sim \mathcal{D}, y_{\text{SFT}}^* \sim \mu(\cdot | x, C(x))} [\log \pi_{\text{SFT}}(y_{\text{SFT}}^* | x, C(x))]. \quad (1)$$

We refer to the resulting model π_{SFT} as the classical SFT model.

Stage 2 — Mutation SFT. We fine-tune a second model, initialized from the same base model, using inputs x , $C(x)$, and a set of conditioning examples \mathcal{E} consisting of patches sampled from the classical SFT model π_{SFT} . The target $y_{\text{M-SFT}}^*$ includes a CoT trace generated by the teacher model μ conditioned on \mathcal{E} , along with the ground-truth patch. The training objective is:

$$\max_{\pi_{\text{M-SFT}}} \mathbb{E}_{x \sim \mathcal{D}, \mathcal{E} \sim \pi_{\text{SFT}}(\cdot | x, C(x)), y_{\text{M-SFT}}^* \sim \mu(\cdot | x, C(x), \mathcal{E})} [\log \pi_{\text{M-SFT}}(y_{\text{M-SFT}}^* | x, C(x), \mathcal{E})]. \quad (2)$$

We refer to the resulting model $\pi_{\text{M-SFT}}$ as the mutation SFT model.

Training on small-scale datasets.

EvoScale targets issues where one-shot generation often fails, but high-scoring patches can still be found through sufficient sampling. This means the model generates a mix of high- and low-scoring patches, so conditioning examples should reflect this diversity. If all examples were already high-scoring, test-time scaling would offer limited benefit. Training a classical SFT model on the full dataset, however, leads to memorization, reducing output diversity and making it difficult to construct diverse conditioning examples for mutation SFT. To preserve diversity, we collect y_{SFT}^* and $y_{\text{M-SFT}}^*$ on disjoint subsets of the data. See Appendix D for details.

Limitation of SFT in *self-evolving*. The mutation SFT model $\pi_{\text{M-SFT}}$ is trained on conditioning examples from the classical SFT model π_{SFT} , which include both low- and high-scoring patches. This raises a natural question: can $\pi_{\text{M-SFT}}$ learn to improve low-scoring patches on its own—i.e., *self-evolve*—without relying on reward models to select high-scoring examples? If so, we could eliminate the selection step (Line 3 in Algorithm 1), reducing scoring costs and sample usage. However, we find that SFT alone cannot enable self-evolution. Section 4.3 introduces a reinforcement learning approach that trains the model to self-evolve without scoring or filtering.

Algorithm 1 Evolutionary Test-Time Scaling (EvoScale)

Require: Issue description x , code context $C(x)$, editor model π , number of iterations T , samples per iteration M , optional selection size K

- 1: Generate initial outputs $\mathcal{Y}^0 := \{y_1^0, \dots, y_M^0\} \sim \pi(\cdot | x, C(x))$
- 2: **for** $t = 1$ to T **do**
- 3: (Optional) Select conditioning examples $\mathcal{E}^{t-1} := \{\bar{y}_1^{t-1}, \dots, \bar{y}_K^{t-1}\} = \text{Select}(\mathcal{Y}^{t-1})$
- 4: Generate new outputs $\mathcal{Y}^t := \{y_1^t, \dots, y_M^t\} \sim \pi(\cdot | x, C(x), \mathcal{E}^{t-1})$
- 5: **end for**

4.3 LEARNING TO SELF-EVOLVE VIA LARGE-SCALE REINFORCEMENT LEARNING (RL)

To *self-evolve*, the model must generate patches that maximize a scoring function R , given conditioning examples \mathcal{E} from previous patches. This setup naturally aligns with the reinforcement learning (RL) (Sutton and Barto, 2018), where a policy π is optimized to maximize expected rewards over time. Since our goal is to maximize the reward at the final iteration T , a naïve RL objective is:

$$\max_{\pi} \mathbb{E}_{y^t \sim \pi(\cdot | x, C(x), \mathcal{E}^{t-1})} \left[\sum_{t=0}^T r_t \right], \quad \text{where } r_t = \begin{cases} R(x, y^t), & t = T \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

This objective focuses solely on maximizing the final reward. However, it presents two key challenges: (1) rewards are sparse, with feedback only at iteration T , making learning inefficient (Lee et al., 2024; Shen et al., 2025); and (2) generating full T -step trajectories is computationally expensive (Snell et al., 2025).

Potential shaping alleviates sparse rewards. We address the sparse reward challenge using *potential-based reward shaping* (Ng et al., 1999), where the potential function is defined as $\Phi(y) = R(x, y)$. The potential reward at step t is:

$$r_t = \Phi(y^t) - \Phi(y^{t-1}) = R(x, y^t) - R(x, y^{t-1}). \quad (4)$$

Unlike the naïve formulation (Equation 3), this provides non-zero potential rewards at every step, mitigating the sparse reward challenge. The cumulative potential reward forms a telescoping sum: $\sum_{t=1}^T r_t = R(x, y^T) - R(x, y^0)$. Since y^0 is fixed, maximizing this sum is equivalent to maximizing the final score as shown by Ng et al. (1999).

Monotonic improvement via local optimization. While optimizing Equation 3 achieves the optimal final reward, it is computationally expensive due to the need for full T -step trajectories. As a more efficient alternative, we train the model to maximize the potential reward at each individual iteration t (Equation 4), avoiding the cost of generating full T -step trajectories. This local optimization reduces computation and runtime while ensuring monotonic reward improvement (see Section 5.2), which is sufficient for improving patch scores over iterations. We formally show this property in Section 4.4.

Implementation. Using the full dataset, we fine-tune the mutation SFT model $\pi_{\text{M-SFT}}$ to maximize the expected potential rewards (Equation 4) in score between a newly generated patch y and a previous patch y' drawn from the conditioning examples \mathcal{E} :

$$\max_{\pi_{\text{RL}}} \mathbb{E}_{y \sim \pi_{\text{RL}}(\cdot | x, C(x), \mathcal{E}), y' \sim \mathcal{E}} [R(x, y) - R(x, y') - \lambda F(y)]. \quad (5)$$

This objective encourages the model to generate patches that consistently improve upon previous ones. To ensure the outputs follow the required syntax, we incorporate a formatting penalty term F into the reward function (see Appendix D for details). The conditioning patch y' is sampled from conditioning examples constructed using patches generated by earlier models, such as π_{SFT} or intermediate checkpoints of π_{RL} .

4.4 THEORETICAL ANALYSIS

We analyze the RL objective in Equation 5, which leverages potential-based reward shaping (Ng et al., 1999), and show that the induced policy yields non-decreasing scores at each iteration.

Assumption 1 (Φ -monotonicity). *Let \mathbb{Y} be the set of all patches and $\Phi: \mathbb{Y} \rightarrow \mathbb{R}$ a potential function. For every $y \in \mathbb{Y}$, there exists a finite sequence $y = y_0, y_1, \dots, y_k$ such that $\Phi(y_{t+1}) \geq \Phi(y_t)$ for all $0 \leq t < k$.*

This ensures that from any initial patch one can reach higher-scoring patches without decreasing Φ .

Definition 1 (Myopic Policy). *Define the one-step action-value $Q_0(y, y') = \Phi(y') - \Phi(y)$, $y, y' \in \mathbb{Y}$. The myopic policy π_0 selects, at each state y , any successor that maximizes Q_0 : $\pi_0(y) \in \arg \max_{y' \in \mathbb{Y}} [\Phi(y') - \Phi(y)]$.*

Proposition 1 (Monotonic Improvement). *Under Assumption 1, any trajectory $\{y^t\}_{t \geq 0}$ generated by the myopic policy π_0 satisfies $\Phi(y^t) \geq \Phi(y^{t-1})$ and $r_t = \Phi(y^t) - \Phi(y^{t-1}) \geq 0 \quad \forall t \geq 1$.*

Proof. By definition of π_0 , at each step $y^t \in \arg \max_{y'} [\Phi(y') - \Phi(y^{t-1})]$. Hence $\Phi(y^t) - \Phi(y^{t-1}) \geq 0$, which immediately gives $\Phi(y^t) \geq \Phi(y^{t-1})$ and $r_t \geq 0$. In particular, training with the potential reward in Equation equation 5 guarantees that

$$R(x, y^t) = \Phi(y^t) \geq \Phi(y^{t-1}) = R(x, y^{t-1}) \quad \forall t.$$

Thus the learned policy produces non-decreasing scores over iterations.

5 EXPERIMENTS

5.1 SETTINGS

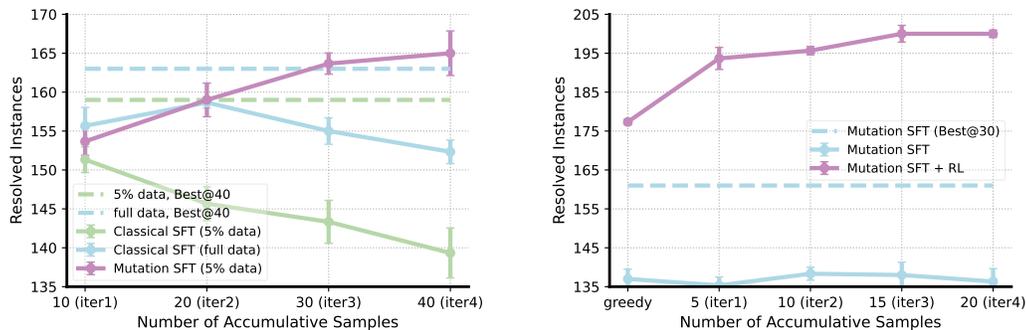
Implementation Details. We adopt a pipeline-based scaffold consisting of a retriever and a code editing model (see Appendix C). Both components are trained using small-scale SFT and large-scale RL. We use the Qwen2.5-Coder-32B-Instruct model Hui et al. (2024) as our base model due to its strong code reasoning capabilities. Our training data is sourced from SWE-Fixer (Xie et al., 2025) and SWE-Gym (Pan et al., 2025). After filtering and deduplication, we obtain a total of 29,404 high-quality instances. For RL training of the code editing model, we rely on a reward model trained on data collected from open source data¹ with 1,889 unique instances. Additional experimental details are provided in Appendix D.

Evaluation and Metrics. We consider two metrics in our evaluation: (1) Greedy: zero-shot pass@1 accuracy, which measures the number of correctly solved instances using greedy generation with syntax retrieval (i.e., random sampling up to five times until syntactically correct); (2) Best@ N : accuracy of the optimal sample selected by the verifier among N randomly generated samples. Greedy evaluates the model’s budget-efficient performance, while Best@ N represents the model’s potential for test-time scaling.

Test-time Scaling Methods. We evaluate the following test-time scaling methods: (1) Reward Model Selection: selects the optimal patch sample with the highest reward model score; (2) Unit Tests Selection: selects the optimal patch sample based on whether it passes unit tests, including both regression and reproduction tests. If multiple samples pass, one is selected at random; (3) EvoScale: at each evolution iteration, the model generates M patch samples and selects $K \leq M$ samples as the conditional prompt for the next generation. The selection of the K samples is guided by the reward model. In our experiments, we set $M = 10$, $K = 5$, and perform up to four iterations of evolution.

5.2 ANALYSIS

In this section, we present a comprehensive analysis of the proposed EvoScale approach. To simplify our analysis, we use ground-truth localization (retrieval) and focus on the code editing part. All reported results are averaged over three random trials. More results are provided in Appendix A.



(a) **RM as selector:** Classical SFT v.s. Mutation SFT

(b) **Self-evolve:** Mutation SFT v.s. RL

Figure 4: Evolutionary Capability of Different Stages of SFT and RL Models. (a) Reward Model selects the top-5 patch candidates from 10 samples from the previous iteration, and the model iteratively evolves by generating new 10 samples conditioned on the candidates. Performance of the top-1 sample selected by RM is reported. Without the additional mutation SFT training, the model fails to exhibit evolutionary behavior, even when scaling up the training set. (b) Without RM selection, the model only iteratively evolves by conditioning on 5 random samples from the last iteration. RL training improves the model’s initial performance and incentivizes the self-evolution capability, while the SFT model fails to self-evolve without guidance from RM.

Can LLMs Iteratively Evolve without Mutation SFT Training? First, we investigate whether the mutation SFT is necessary for LLMs to learn how to iteratively improve their generations. Specifically,

¹<https://huggingface.co/nebius>

we fine-tune base LLMs using either classical SFT (without conditional generation) or mutation SFT. As shown in Figure 4(a), models trained with classical SFT fail to improve their outputs when conditioned on previous samples. In contrast, mutation SFT enables the model to iteratively improve under the guidance of a reward model. The performance of the mutation SFT model at later iterations can surpass the classical SFT model by scaling up the samples (e.g., Best@40). Moreover, this iterative refinement capability can be learned effectively even with a small number of training data.

RL Enables Self-evolve Capability. While mutation SFT model demonstrates evolutionary behavior when guided by a reward model, we further examine whether it can self-evolve without such guidance. Specifically, instead of selecting the top- K candidates to ensure generation quality, we allow the model to generate $M = K = 5$ random samples for the next iteration of conditional generation. However, as shown in Figure 4(b), the SFT model fails to learn self-evolution without reward model selection. Interestingly, RL training significantly improves the SFT model in two key aspects. First, RL substantially boosts the model’s greedy performance, surpassing even the Best@ N performance of 30 randomly generated samples from the SFT model. Second, we observe that the RL-trained model exhibits strong self-evolution capability: even when conditioned on its random outputs, the model can self-refine and improve performance across iterations without reward model guidance. We provide further analysis of the model’s behavior through demo examples in Appendix B.1.

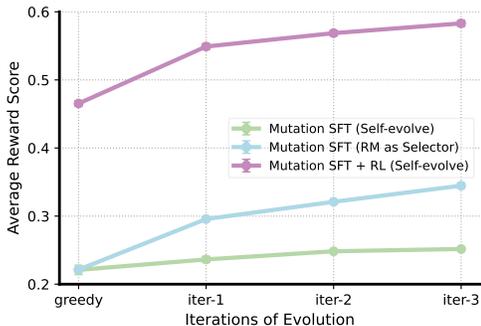


Figure 5: **Average Reward Score of Patch Samples at Each Evolution Iteration.** Reward scores are normalized via a sigmoid function before average. The SFT model struggles to improve reward scores without the guidance of a reward model to select top- K conditional patch samples, while the RL model consistently self-improves its reward score across iterations without external guidance, validating our theoretical results of monotonic improvement in Section 4.4

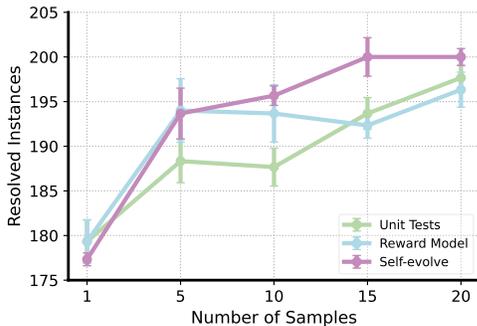


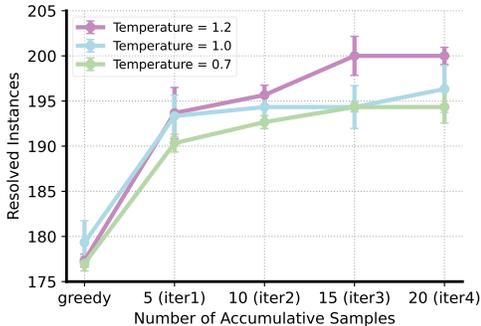
Figure 6: **Comparison with Other Test-Time Scaling Methods.** Reward model selection requires deploying an additional model at test time and can become unstable as the number of samples increases. Unit test selection is computationally expensive and performs poorly with a small sample size. In contrast, self-evolution demonstrates high sample efficiency and strong test-time scaling performance.

Do our SFT and RL Models Monotonically Improve Reward Scores over Iterations? We further analyze the evolutionary behavior of the SFT and RL models by measuring the average reward score of the patch samples generated at each iteration. As shown in Figure 5, although the SFT model learns to iteratively improve reward scores, it relies on the reward model to select high-quality conditioning examples to achieve significant improvements. In contrast, the RL model trained with potential-based reward, naturally learns to self-evolve without any external guidance. Its reward scores improve monotonically across iterations, aligns with our theoretical analysis in Section 4.4.

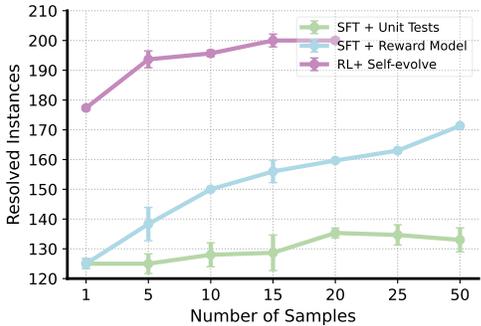
Evolutionary Test-time Scaling v.s. Other Test-time Scaling Methods. Next, we further compare evolutionary test-time scaling with other test-time scaling methods. Starting from the RL model, we first randomly sample $N = 5, 10, 15, 20$ patch samples and let the reward model and unit tests select the best sample among the subsets. In addition, we let the model perform self-evolution with $K = 5$ samples per iteration, up to four iterations (20 samples in total). The test-time scaling results presented in Figure 6 demonstrate both efficiency and effectiveness of evolutionary test-time scaling.

EvoScale Prefers Higher Mutation Sampling Temperature. Mutation sampling plays a critical role in Evolutionary Test-Time Scaling. To investigate its impact, we vary the model’s sampling temperature across 0.7, 1.0, 1.2 and perform self-evolution over four iterations. As shown in Figure 7, higher temperatures demonstrate better performance. Intuitively, a larger temperature increases the

432 diversity of generated patch samples, providing richer information for the mutation operator to
 433 produce improved patches in subsequent iterations. In contrast, lower temperatures tend to result in
 434 repetitive patch samples and may lead the model to converge quickly to suboptimal solutions.
 435



436
437
438
439
440
441
442
443
444
445
446
447
448 **Figure 7: Impact of Mutation Sampling Temperature.** Higher sampling temperatures in EvoScale
 449 encourage greater diversity among mutation samples, leading to more effective iterative improvements.
 450
451
452



453
454
455
456
457
458
459 **Figure 8: Classical SFT + Test-time Scaling v.s. Mutation RL + Self-evolve.** RL Model with self-evolve
 460 capability is more effective than classical SFT model using other test-time scaling methods.
 461
462
463
464
465
466

467 **SFT + Test-time Scaling v.s. RL + Self-evolve.** In Figure 6, we demonstrated the superior
 468 performance of Evolutionary Test-Time Scaling using our RL-trained model. To further investigate
 469 this result, we compare against other test-time scaling methods applied to a classical SFT model
 470 trained on the full dataset (30K instances), since the typical procedure in most existing SWE work
 471 Pan et al. (2025); Xia et al. (2024) trains a SFT model and applying verifiers (e.g., reward models or unit
 472 tests) for test-time scaling.
 473

474 However, as shown in Figure 8, this approach proves to be less effective: (1) With 50 samples, the
 475 SFT model’s Best@50 performance is still outperformed by the greedy decoding of the RL model,
 476 despite both being trained on the same dataset. (2) The SFT model is relatively sensitive to the choice
 477 of verifier. When using unit tests (including both reproduction and regression tests) as the verifier,
 478 increasing the number of samples results in only marginal performance gains. These observations
 479 support our hypothesis: while correct solutions do exist in the SFT model’s output distribution, they
 480 are rarely sampled due to its dispersed sampling distribution. In contrast, the RL model learns to
 481 self-refine the sampling distribution towards high-scoring region.
 482

483 **More Results.** We include additional analysis in Appendix A, including different reward modeling
 484 methods in RL, the impact of mutation sampling temperature, the impact of RL training, and runtime
 485 comparison of different test-time scaling methods.

470
471 **5.3 RESULTS IN THE WILD: SWE-BENCH PERFORMANCE**

472 We present the main results of our RL-trained model, EvoScale-32B, on the SWE-bench Verified
 473 benchmark (Jimenez et al., 2024) and compare its performance against both open-source and prop-
 474 rietary systems. We report results for both greedy decoding and Best@N metrics, using our own
 475 retrieval framework (see details of retrieval in Appendix C). For test-time scaling, we apply iterative
 476 self-evolution, allowing the RL model to generate $M = 25$ samples per iteration. We observe that
 477 the initial iterations produce more diverse candidate patches, while later iterations generate higher-quality,
 478 more refined patches. To balance diversity and refinement, we aggregate all generated samples across
 479 iterations into a combined pool of $N = 50$ candidates. As discussed in Section 5.2, different verifiers
 480 provide complementary strengths. We therefore combine both the reward model and unit tests to
 481 select the best patch from the candidate pool.

482 As shown in Table 1, EvoScale-32B achieves a greedy accuracy of 35.8, outperforming all existing
 483 small-scale models under greedy decoding. Additionally, it achieves a Best@50 score of 41.6,
 484 matching the performance of the current state-of-the-art Llama3-SWE-RL-70B Wei et al. (2025),
 485 which requires Best@500 decoding—incurring over 10× higher sampling cost. It is also worth
 noting that agent-based methods incur even higher test-time computational cost, as each generation

Table 1: **Results on SWE-bench Verified.** EvoScale-32B outperforms all small-scale models under greedy decoding, while achieving comparable performance with current SOTA SWE-RL with much fewer training data and test-time scaling samples.

Model Scale	Model/Methods	Scaffold	SWE-Verified Resolved Rate
Large	GPT-4o Yang et al. (2024)	SWE-agent	23.0
	GPT-4o Xia et al. (2024)	Agentless	38.8
	GPT-4o Zhang et al. (2024)	AutoCodeRover	28.8
	GPT-4o Ma et al. (2024)	SWE-SynInfer	31.8
	OpenAI o1 Xia et al. (2024)	Agentless	48.0
	Claude 3.5 Sonnet Yang et al. (2024)	SWE-agent	33.6
	Claude 3.5 Sonnet Wang et al. (2025)	OpenHands	53.0
	Claude 3.5 Sonnet Xia et al. (2024)	Agentless	50.8
	Claude 3.5 Sonnet Zhang et al. (2024)	AutoCodeRover	46.2
	Claude 3.7 Sonnet Anthropic (2025)	SWE-agent	58.2
	DeepSeek-R1 Guo et al. (2025)	Agentless	49.2
	DeepSeek-V3 Liu et al. (2024)	Agentless	42.0
	Small	Lingma-SWE-GPT-7B (Greedy) Ma et al. (2024)	SWE-SynInfer
Lingma-SWE-GPT-72B (Greedy) Ma et al. (2024)		SWE-SynInfer	28.8
SWE-Fixer-72B (Greedy) Xie et al. (2025)		SWE-Fixer	30.2
SWE-Gym-32B (Greedy) Pan et al. (2025)		OpenHands	20.6
SWE-Gym-32B (Best@16) Pan et al. (2025)		OpenHands	32.0
Llama3-SWE-RL-70B (Best@80) Wei et al. (2025)		Agentless Mini	37.0
Llama3-SWE-RL-70B (Best@160) Wei et al. (2025)		Agentless Mini	40.0
Llama3-SWE-RL-70B (Best@500) Wei et al. (2025)		Agentless Mini	41.0
EvoScale-32B (Greedy)	EvoScale-32B (Greedy)	EvoScale	35.8
	EvoScale-32B (Self-evolve@10)	EvoScale	36.4
	EvoScale-32B (Self-evolve@20)	EvoScale	38.6
	EvoScale-32B (Best@10)	EvoScale	38.9
	EvoScale-32B (Best@25)	EvoScale	40.2
	EvoScale-32B (Best@50)	EvoScale	41.6

corresponds to a full rollout trajectory with multiple interactions. In contrast, EvoScale-32B achieves state-of-the-art performance with significantly lower inference cost and is trained on fewer than 30K open-source samples, compared to millions of proprietary data used to train Llama3-SWE-RL-70B.

6 CONCLUDING REMARKS

We propose Evolutionary Test-time Scaling (EvoScale), a sample-efficient inference-time method that enables small language models to approach the performance of 100B+ parameter models using just 50 code patch samples—without requiring interaction trajectories with the runtime environment. EvoScale opens up a new direction for sample-efficient test-time scaling in real-world software engineering tasks: **(1) Evolution improves sample efficiency.** Our results show that evolutionary strategies, which iteratively refine generations, can drastically reduce the number of required samples. This contrasts with prior work that primarily focuses on improving verifiers (e.g., reward models, test cases); **(2) RL enables self-evolution.** We show that reinforcement learning (RL) can train models to refine their outputs without relying on external verifiers at inference. While our current method optimizes local reward differences, future work may explore optimizing cumulative potential rewards over entire trajectories. Compared to Snell et al. (2025), who maintains all prior outputs in the prompt during revision, our method retains only the most recent output—making it more suitable for SWE tasks with long context windows; **(3) Limitations and future work.** This work focuses on a pipeline-based (agentless) setup. Extending EvoScale to agentic settings where models interact with code and runtime environments, remains an interesting future work. While our study focuses on SWE due to its realism and difficulty, we believe the core ideas behind EvoScale could generalize to other agentic tasks, an exciting direction for future work.

REFERENCES

- Anthropic. Claude 3.7 sonnet and claude code, 2025. URL [Claude3.7SonnetandClaudeCode](#). 10
- Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, and William Yang Wang. SWE-search: Enhancing software agents with monte carlo tree search and iterative refinement. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=G7sIFXugTX>. 3

- 540 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and
541 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling,
542 2024. URL <https://arxiv.org/abs/2407.21787>. 1, 2
543
- 544 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared
545 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,
546 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,
547 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,
548 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios
549 Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino,
550 Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,
551 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa,
552 Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob
553 McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating
554 large language models trained on code. 2021. 1
- 555 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
556 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reichiro Nakano, et al. Training verifiers to solve
557 math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 1
- 558 Ryan Ehrlich, Bradley Brown, Jordan Juravsky, Ronald Clark, Christopher Ré, and Azalia Mirho-
559 seini. Codemonkeys: Scaling test-time compute for software engineering. *arXiv preprint*
560 *arXiv:2501.14723*, 2025. 2
561
- 562 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
563 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
564 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 10
565
- 566 Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016. 2
567
- 568 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
569 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al.
570 Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022. 1, 3
- 571 Jian Hu, Xibin Wu, Zilin Zhu, Xianyu, Weixun Wang, Dehao Zhang, and Yu Cao. Openrlhf: An
572 easy-to-use, scalable and high-performance rlhf framework. *arXiv preprint arXiv:2405.11143*,
573 2024. 34
574
- 575 Jian Hu, Jason Klein Liu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with robustness to
576 both prompt and reward models, 2025. URL <https://arxiv.org/abs/2501.03262>. 35,
577 36
- 578 Baihe Huang, Shanda Li, Tianhao Wu, Yiming Yang, Ameet Talwalkar, Kannan Ramchandran,
579 Michael I Jordan, and Jiantao Jiao. Sample complexity and representation ability of test-time
580 scaling paradigms. *arXiv preprint arXiv:2506.05295*, 2025. 3
581
- 582 Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun
583 Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei
584 Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng
585 Ren, Jingren Zhou, and Junyang Lin. Qwen2.5-coder technical report, 2024. URL <https://arxiv.org/abs/2409.12186>. 7, 34
586
- 587 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
588 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free eval-
589 uation of large language models for code. In *The Thirteenth International Conference on Learning*
590 *Representations*, 2025a. URL <https://openreview.net/forum?id=chfJJYC3iL>. 1
591
- 592 Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2e-gym:
593 Procedural environments and hybrid verifiers for scaling open-weights swe agents. *arXiv preprint*
arXiv:2504.07164, 2025b. 2, 3

- 594 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R
595 Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth*
596 *International Conference on Learning Representations*, 2024. URL [https://openreview.](https://openreview.net/forum?id=VTF8yNQM66)
597 [net/forum?id=VTF8yNQM66](https://openreview.net/forum?id=VTF8yNQM66). 1, 2, 9
- 598 Chi-Chang Lee, Zhang-Wei Hong, and Pulkit Agrawal. Going beyond heuristics by imposing
599 policy improvement as a constraint. *Advances in Neural Information Processing Systems*, 37:
600 138032–138087, 2024. 6
- 601 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
602 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth*
603 *International Conference on Learning Representations*, 2023. 1
- 604 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
605 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*
606 *arXiv:2412.19437*, 2024. 10
- 607 Yingwei Ma, Rongyu Cao, Yongchang Cao, Yue Zhang, Jue Chen, Yibo Liu, Yuchen Liu, Binhua
608 Li, Fei Huang, and Yongbin Li. Lingma swe-gpt: An open development-process-centric language
609 model for automated software improvement. *arXiv preprint arXiv:2411.00622*, 2024. 2, 10
- 610 Yingwei Ma, Yongbin Li, Yihong Dong, Xue Jiang, Rongyu Cao, Jue Chen, Fei Huang, and Binhua
611 Li. Thinking longer, not larger: Enhancing software engineering agents via scaling test-time
612 compute, 2025. URL <https://arxiv.org/abs/2503.23803>. 2, 3
- 613 Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations:
614 Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287, 1999. 2, 6
- 615 Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang.
616 Training software engineering agents and verifiers with SWE-gym. In *ICLR 2025 Third Work-*
617 *shop on Deep Learning for Code*, 2025. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=lpFFpTbi9s)
618 [lpFFpTbi9s](https://openreview.net/forum?id=lpFFpTbi9s). 1, 2, 3, 7, 9, 10, 14
- 619 Isha Puri, Shivchander Sudalairaj, Guangxuan Xu, Kai Xu, and Akash Srivastava. A probabilistic
620 inference approach to inference-time scaling of llms using particle-based monte carlo methods.
621 *arXiv preprint arXiv:2502.01618*, 2025. 3
- 622 Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a
623 scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017. 2
- 624 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
625 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
626 mathematical reasoning in open language models, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2402.03300)
627 [2402.03300](https://arxiv.org/abs/2402.03300). 3
- 628 Maohao Shen, Soumya Ghosh, Prasanna Sattigeri, Subhro Das, Yuheng Bu, and Gregory Wornell.
629 Reliable gradient-free and likelihood-free prompt tuning. In *Findings of the Association for*
630 *Computational Linguistics: EACL 2023*. Association for Computational Linguistics, 2023. URL
631 <https://aclanthology.org/2023.findings-eacl.183/>. 2
- 632 Maohao Shen, Guangtao Zeng, Zhenting Qi, Zhang-Wei Hong, Zhenfang Chen, Wei Lu, Gre-
633 gory Wornell, Subhro Das, David Cox, and Chuang Gan. Satori: Reinforcement learning with
634 Chain-of-Action-Thought enhances llm reasoning via autoregressive search. *arXiv preprint*
635 *arXiv:2502.02508*, 2025. 6
- 636 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
637 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient RLHF framework. In *Proceedings*
638 *of the Twentieth European Conference on Computer Systems*. ACM, 2025. 34
- 639 Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute
640 optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth Inter-*
641 *national Conference on Learning Representations*, 2025. URL [https://openreview.net/](https://openreview.net/forum?id=4FWAwZtd2n)
642 [forum?id=4FWAwZtd2n](https://openreview.net/forum?id=4FWAwZtd2n). 1, 3, 6, 10

- 648 Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. 2018. 5
- 649
- 650 Zhendong Tan, Xingjun Zhang, Chaoyi Hu, Yancheng Pan, and Shaoxun Wang. Adaptive rectification
651 sampling for test-time compute scaling. *arXiv preprint arXiv:2504.01317*, 2025. 3
- 652
- 653 Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan,
654 Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill
655 Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan,
656 Hao Peng, Heng Ji, and Graham Neubig. Openhands: An open platform for AI software developers
657 as generalist agents. In *The Thirteenth International Conference on Learning Representations*,
2025. URL <https://openreview.net/forum?id=OJd3ayDDoF>. 10
- 658
- 659 Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried,
660 Gabriel Synnaeve, Rishabh Singh, and Sida I Wang. Swe-rl: Advancing llm reasoning via
661 reinforcement learning on open software evolution. *arXiv preprint arXiv:2502.18449*, 2025. 2, 3,
662 9, 10, 15, 36
- 663
- 664 Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber.
665 Natural evolution strategies. *The Journal of Machine Learning Research*, 15(1):949–980, 2014. 2
- 666
- 667 Menghua Wu, Cai Zhou, Stephen Bates, and Tommi Jaakkola. Thought calibration: Efficient and
668 confident test-time scaling. *arXiv preprint arXiv:2505.18404*, 2025. 3
- 669
- 670 Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Agentless: Demystifying llm-
671 based software engineering agents, 2024. URL <https://arxiv.org/abs/2407.01489>.
672 1, 2, 3, 9, 10, 33, 36
- 673
- 674 Chengxing Xie, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou, and Kai Chen. Swe-fixer:
675 Training open-source llms for effective and efficient github issue resolution, 2025. URL <https://arxiv.org/abs/2501.05040>. 1, 2, 3, 5, 7, 10
- 676
- 677 John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik R Narasimhan,
678 and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering.
679 In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL
<https://openreview.net/forum?id=mXpq6ut8J3>. 2, 3, 10, 14
- 680
- 681 John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang,
682 Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. Swe-smith: Scaling data for software
683 engineering agents, 2025. URL <https://arxiv.org/abs/2504.21798>. 1, 2, 5
- 684
- 685 Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. Autocoderover: Autonomous
686 program improvement. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on*
687 *Software Testing and Analysis*, pages 1592–1604, 2024. 10
- 688
- 689
- 690
- 691
- 692
- 693
- 694
- 695
- 696
- 697
- 698
- 699
- 700
- 701

702	APPENDIX	
703		
704	A Additional Experiments	14
705	A.1 Extending EvoScale to an Agentic SWE System	14
706	A.2 Additional Anlysis	15
707		
708	B Demo Examples	17
709	B.1 Type 1: Prior patches are all wrong	17
710	B.2 Type 2: Prior patches are partially wrong	20
711	B.3 Type 3: Prior patches are all correct	25
712		
713	C Scaffold of EvoScale	32
714	C.1 Retriever	32
715	C.2 Code Editing Model	33
716	C.3 Verifier	33
717	D Implementation Details	34
718	D.1 Dataset Collection	34
719	D.2 Training Pipeline and Hardware	34
720	D.3 Retrieval Model	34
721	D.4 Retrieval Reward Model	35
722	D.5 Code Editing Model	35
723	D.6 Code Editing Reward Model	36
724	D.7 Reproduction Test Generator	36
725	E Prompt Template	38
726		

DISCLAIMER ON THE USE OF LARGE LANGUAGE MODELS (LLMs)

In preparing this manuscript, Large Language Models (LLMs) were used exclusively for polishing the writing. They were not involved in research ideation, methodology, experiments, analysis, or paper content generation. The authors take full responsibility for all claims and results presented in this paper.

A ADDITIONAL EXPERIMENTS

A.1 EXTENDING EVOSCALE TO AN AGENTIC SWE SYSTEM

While our main experiments adopt an Agentless scaffold for simplicity of the training framework, we show that the proposed method naturally extends to an agentic scaffold without modification.

Experimental Setting. We integrate EvoScale into the SWE-Agent scaffold (Yang et al., 2024) and utilize the standard tool set: `edit anthropic` (file viewing and editing), `bash` (restricted shell commands), and `submit`. For execution, we adopt `Seed-OSS-36B-Instruct` as the LLM due to its strong tool-use and instruction-following capability.

Within the SWE-Agent environment, the agent can perform multiple interactions inside the sandbox before submitting the final diff, including searching files, editing code, and generating unit tests. Following the same setup as the RL-trained EvoScale-32B, we allow `Seed-OSS-36B-Instruct` to self-evolve for four iterations. At each iteration, the model generates five random samples, which serve as the candidates for the next round of conditional generation.

Results. Since SWE-Agent can self-generate unit tests during the problem-solving process, we primarily compare EvoScale with reward-model selection (Pan et al., 2025) as the standard test-time scaling baseline. The results shown in Figure 9 reveal a surprising observation: even without any training, `Seed-OSS-36B-Instruct` exhibits strong self-evolution capabilities and surpasses reward-model selection in both performance and sample efficiency.

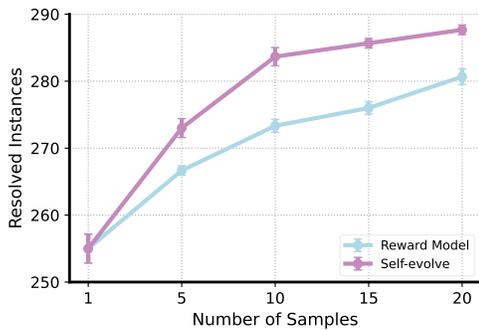


Figure 9: **EvoScale v.s. Standard Test-Time Scaling using Agentic Workflows.** On Agentic SWE task, EvoScale still demonstrates high sample efficiency and strong test-time scaling performance.

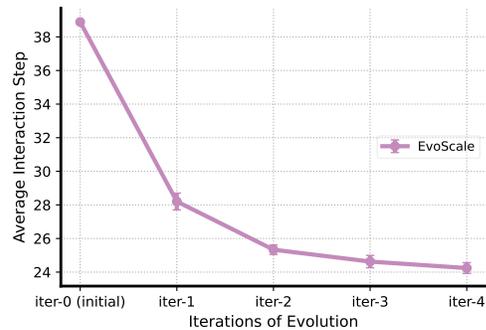


Figure 10: **Agentic Interaction Steps of EvoScale over Iterations.** Agentic workflow often suffers from long-horizon interaction, but leveraging past experience, EvoScale effectively reduces the interaction steps across iterations.

We further analyze how EvoScale leverages previously generated patches to self-evolve. In Figure 10, we plot the average number of interaction steps used by SWE-Agent across evolution iterations. Interestingly, the agent requires fewer steps over time to solve the problem. This finding suggests that leveraging past experience is essential for agentic tasks, and that our proposed evolutionary test-time scaling serves as a key mechanism for enabling such in-situ adaptation.

A.2 ADDITIONAL ANALYSIS

In this section, we provide additional analytical experiments of EvoScale and model training.

RL Reward Modeling: Reward Model is More Effective than String-matching. A reliable reward signal is the key to drive RL training. To better understand the impact of different components in reward modeling, we conduct an ablation study comparing three variants: using only the reward model score, using only string-matching rewards proposed in (Wei et al., 2025), and using both. As shown in Table 2, models trained with a single reward component show degraded greedy decoding performance compared to the model trained with the hybrid reward. In particular, the reward model plays a crucial role in boosting the performance, while the string-matching reward helps the model learn better syntactic structure. However, the results also suggest that naïve string-matching Wei et al. (2025) alone may not serve as a reliable reward signal for SWE tasks.

Table 2: **Ablation Study on Reward Modeling.** The total number of instances is 500. Compared to the SFT model, RL using the RM reward significantly improves performance but introduces more syntax errors. In contrast, RL with a string-matching reward reduces syntax errors but fails to improve reasoning capability. A hybrid reward signal effectively balances both aspects, achieving superior performance.

Metrics	Mutation SFT (5% data)	RM RL	String-matching RL	Hybrid Reward RL
Num of Resolved Instances	137 ± 2.5	171 ± 1.7	140.7 ± 0.5	179.3 ± 2.4
Num of Syntax-correct Instances	427	404	478	471

Runtime Comparison of Different Test-time Scaling Methods. To evaluate the efficiency of different test-time scaling methods, we measure the average runtime per instance. For our proposed EvoScale approach, the runtime consists solely of iteratively prompting the RL model to generate samples. Reward model selection incurs additional computational cost due to running the reward model to score each sample, and unit test selection requires executing each patch in a sandbox environment. EvoScale achieves better performance with highest efficiency compared with the other two approaches (see Figure 11). while unit test selection is effective when scaling to larger sample sizes, it comes at a cost around 6× slower than EvoScale.

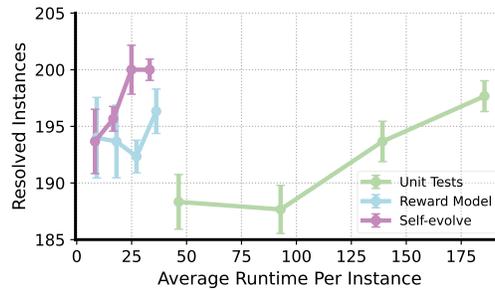


Figure 11: **Performance vs. Average Runtime per Instance for Different Test-Time Scaling Methods.** EvoScale achieves better performance while incurs much less runtime overhead.

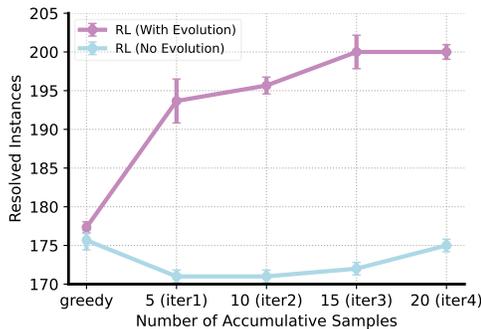


Figure 12: **RL with vs. without Self-Evolution Training.** Removing evolution training during the RL stage results in a model that lacks iterative self-improvement capabilities.

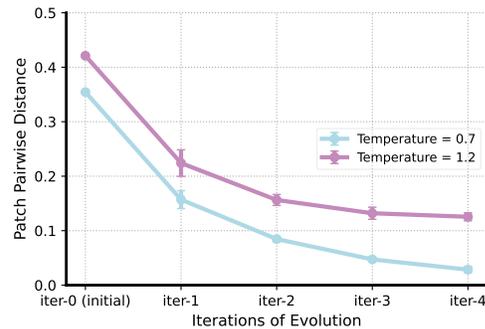


Figure 13: **Diversity of Generated Patches over Iterations.** The generated patches gradually converge toward a high-scoring distribution, while higher temperatures encourage greater diversity.

Would RL without Evolution Training still Work? We consider a simplified training setup for the code editing model, where the base model is trained using classical SFT followed by RL without incorporating mutation data or potential-based rewards. As shown in Figure 2, although this simplified RL approach can still improve the SFT model’s greedy performance, it fails to equip the model with iterative self-improvement ability. This finding demonstrates the importance of evolution training, particularly the use of potential-based rewards, in incentivizing the model to learn how to self-refine over multiple iterations.

Analyzing Diversity of Generated Patches. To better understand the behavior of the evolutionary generation, we measure how the generated patch diversity changes over iterations. To quantify this, we compute the average pairwise distance between the generated patches for each instance and report the mean over the entire dataset. Lower values indicate reduced diversity, i.e., more similar patches. We evaluate our RL model at two sampling temperatures (0.7 and 1.2). As shown in Figure 13, diversity decreases over iterations, which aligns with our goal of gradually converging towards high-quality solutions (illustrated in Fig. 3). However, rapid convergence risks local optima. As shown in Figure 7, using higher temperatures helps preserve diversity and can increase performance.

Impact of Population Size. While our RL model can self-evolve without relying on a reward model, it is important to note that population size plays an important role only when a selection mechanism, such as a reward model, is available to choose among candidates. To investigate this, we conduct the experiment shown in Figure 4(a) using the Mutation SFT model. Figure 14 reports the Best-of-N (BoN) performance of the SFT model under different population sizes ($M = 10, 20, 30$). We observe that larger population sizes yield higher performance upper bounds after multiple evolution iterations, but at the price of increased inference costs. Therefore, we adopt $M = 10$ in our main experiments, as it achieves strong performance after four iterations while maintaining relatively low cost.

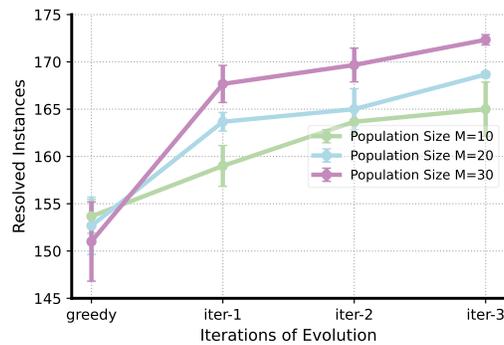


Figure 14: **Performance of Mutation SFT (RM as Selector) vs. Population Size (M).** Increasing the population size raises the performance upper bound through evolution, but also results in higher inference costs.

B DEMO EXAMPLES

The core idea of EvoScale is to use a mix of correct and incorrect patches as context to generate a correct patch. To do this effectively, the model must demonstrate the following capabilities:

1. Generate correct patches even when all prior patches are incorrect—essential for making progress when initial generations fail.
2. Generate correct patches when most prior patches are incorrect—common in practice, requiring the model to identify and build on the few correct ones.
3. Preserve correctness when all prior patches are correct—ensuring the model doesn’t over-edit and degrade already correct solutions.

We present examples below to illustrate each of these capabilities.

B.1 TYPE 1: PRIOR PATCHES ARE ALL WRONG

Problem Statement 1

```
sum: min_count is not available for reduction with more than one dimensions
**Is your feature request related to a problem? Please describe.**

`sum` with `min_count` errors when passing more than one dim:

```python
import xarray as xr
da = xr.DataArray([[1., 2, 3], [4, 5, 6]])
da.sum(["dim_0", "dim_1"], min_count=1)
```

**Describe the solution you'd like**
The logic to calculate the number of valid elements is here:
https://github.com/pydata/xarray/blob/1be777fe725a85b8cc0f65a2bc41f4bc2ba18043/xarray/core/nanops.py#L35

I think this can be fixed by replacing

`mask.shape[axis]` with `np.take(a.shape, axis).prod()`

**Additional context**
nongPotentially relevant for #4351
```

Conditional patch 1 (Wrong)

```
diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
index 74e25c7..3e6c5db 100644
--- a/xarray/core/nanops.py
```

```

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932

```

```

+++ b/xarray/core/nanops.py
@@ -32,7 +32,16 @@ def _maybe_null_out(result, axis, mask, min_count=1):
    )

    if axis is not None and getattr(result, "ndim", False):
-       null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
+       if np.isscalar(axis):
+           # in the scalar case we can just use the shape of that axis
+           size = mask.shape[axis]
+       else:
+           # in the case of multiple axis we need to take the product
+           # of all the axis sizes since we want the total number of
+           # elements we are reducing, not just the size of one axis
+           size = np.take(mask.shape, axis).prod()

+       null_mask = (size - mask.sum(axis) - min_count) < 0
        if null_mask.any():
            dtype, fill_value = dtypes.maybe_promote(result.dtype)
            result = result.astype(dtype)

```

Conditional patch 2 (Wrong)

```

933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949

```

```

diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
index 74e25c7..f533eed 100644
--- a/xarray/core/nanops.py
+++ b/xarray/core/nanops.py
@@ -32,7 +32,15 @@ def _maybe_null_out(result, axis, mask, min_count=1):
    )

    if axis is not None and getattr(result, "ndim", False):
-       null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
+       # handle the case where we are reducing on multiple axis
+       # we need to consider the product of their sizes, not just any
+       # one of them.
+       if np.isscalar(axis):
+           size = mask.shape[axis]
+       else:
+           size = np.take(mask.shape, axis).prod()

+       null_mask = (size - mask.sum(axis) - min_count) < 0
        if null_mask.any():
            dtype, fill_value = dtypes.maybe_promote(result.dtype)
            result = result.astype(dtype)

```

Conditional patch 3 (Wrong)

```

950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966

```

```

diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
index 74e25c7..f533eed 100644
--- a/xarray/core/nanops.py
+++ b/xarray/core/nanops.py
@@ -32,7 +32,15 @@ def _maybe_null_out(result, axis, mask, min_count=1):
    )

    if axis is not None and getattr(result, "ndim", False):
-       null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
+       # handle the case where we are reducing on multiple axis
+       # we need to consider the product of their sizes, not just any
+       # one of them.
+       if np.isscalar(axis):
+           size = mask.shape[axis]
+       else:
+           size = np.take(mask.shape, axis).prod()

+       null_mask = (size - mask.sum(axis) - min_count) < 0
        if null_mask.any():
            dtype, fill_value = dtypes.maybe_promote(result.dtype)
            result = result.astype(dtype)

```

Conditional patch 4 (Wrong)

```

967
968
969
970
971

```

```

diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
index 74e25c7..f533eed 100644
--- a/xarray/core/nanops.py
+++ b/xarray/core/nanops.py

```

```

972
973 @@ -32,7 +32,15 @@ def _maybe_null_out(result, axis, mask, min_count=1):
974     )
975     if axis is not None and getattr(result, "ndim", False):
976         - null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
977         + # handle the case where we are reducing on multiple axis
978         + # we need to consider the product of their sizes, not just any
979         + # one of them.
980         + if np.isscalar(axis):
981             + size = mask.shape[axis]
982         + else:
983             + size = np.take(mask.shape, axis).prod()
984         +
985         + null_mask = (size - mask.sum(axis) - min_count) < 0
986         + if null_mask.any():
987             + dtype, fill_value = dtypes.maybe_promote(result.dtype)
988             + result = result.astype(dtype)
989

```

Conditional patch 5 (Wrong)

```

987 diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
988 index 74e25c7..f533eed 100644
989 --- a/xarray/core/nanops.py
990 +++ b/xarray/core/nanops.py
991 @@ -32,7 +32,15 @@ def _maybe_null_out(result, axis, mask, min_count=1):
992     )
993     if axis is not None and getattr(result, "ndim", False):
994         - null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
995         + # handle the case where we are reducing on multiple axis
996         + # we need to consider the product of their sizes, not just any
997         + # one of them.
998         + if np.isscalar(axis):
999             + size = mask.shape[axis]
1000         + else:
1001             + size = np.take(mask.shape, axis).prod()
1002         +
1003         + null_mask = (size - mask.sum(axis) - min_count) < 0
1004         + if null_mask.any():
1005             + dtype, fill_value = dtypes.maybe_promote(result.dtype)
1006             + result = result.astype(dtype)
1007

```

Thinking Process

```

1004 1. Peer Patch Review:
1005
1006 - All five patches are essentially identical in their approach to solving the issue. They
1007 ↪ all modify the same section of code in _maybe_null_out() to handle the case where
1008 ↪ axis is a list/tuple by calculating the product of the sizes of the reduced
1009 ↪ dimensions instead of just using the size of a single axis.
1010
1011 - The key insight is correct - when reducing over multiple dimensions, the total number
1012 ↪ of elements being reduced is the product of the sizes of those dimensions, not just
1013 ↪ the size of one dimension.
1014
1015 - All patches correctly identify that when axis is a scalar, we should use the single
1016 ↪ dimension size as before, but when axis is a list/tuple, we should use the product of
1017 ↪ the sizes of the specified dimensions.
1018
1019 - The solution is consistent with the issue description which suggests using
1020 ↪ np.take(a.shape, axis).prod() to get the total number of elements being reduced.
1021
1022 2. Issue Analysis:
1023
1024 - The problem occurs when using sum() with min_count on multiple dimensions because the
1025 ↪ current code only checks the size of one dimension (the first in the axis list)
1026 ↪ rather than the total number of elements being reduced.
1027
1028 - The intended behavior is that min_count should consider the total number of valid
1029 ↪ elements across all reduced dimensions, not just the number of elements along a
1030 ↪ single dimension.
1031
1032 - The fix requires modifying the logic that calculates the number of valid elements when
1033 ↪ reducing over multiple dimensions to use the product of the sizes of all reduced
1034 ↪ dimensions.
1035

```

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

3. Solution Approach:

- The key change needed is in the `_maybe_null_out()` function where it calculates `↪ null_mask` for multi-dimensional reductions.
- When axis is a list/tuple, we should calculate the total number of elements being `↪ reduced` by taking the product of the sizes of the specified dimensions (using `↪ np.take(mask.shape, axis).prod()`).
- For single dimension reductions, we can keep the existing behavior of using just `↪ mask.shape[axis]`.
- The rest of the function's logic (checking `min_count` against the number of valid `↪ elements`) can remain the same.

4. Code Localization:

- The issue is in `xarray/core/nanops.py` in the `_maybe_null_out()` function.
- The specific location is where it calculates `null_mask` for multi-dimensional reductions `↪ (when axis is not None and result has ndim)`.
- The current code uses `mask.shape[axis]` which is incorrect for multi-dimensional `↪ reductions`. We need to modify this to use the product of the sizes of all reduced `↪ dimensions` when axis is a list/tuple.

Evolved Patch

```
diff --git a/xarray/core/nanops.py b/xarray/core/nanops.py
index 74e25c7..c642ed6 100644
--- a/xarray/core/nanops.py
+++ b/xarray/core/nanops.py
@@ -27,12 +27,22 @@ def _maybe_null_out(result, axis, mask, min_count=1):
     xarray version of pandas.core.nanops._maybe_null_out
     """
     if hasattr(axis, "__len__"): # if tuple or list
-        raise ValueError(
-            "min_count is not available for reduction with more than one dimensions."
-        )
+        # we still need to handle the case where axis is a list/tuple
+        # by calculating the product of the sizes of the reduced
+        # dimensions rather than just using one dimension's size
+        pass

     if axis is not None and getattr(result, "ndim", False):
-        null_mask = (mask.shape[axis] - mask.sum(axis) - min_count) < 0
+        if np.isscalar(axis):
+            # in the scalar case we can just use the shape of that axis
+            size = mask.shape[axis]
+        else:
+            # in the case of multiple axis we need to take the product
+            # of all the axis sizes since we want the total number of
+            # elements we are reducing, not just the size of one axis
+            size = np.take(mask.shape, axis).prod()
+        null_mask = (size - mask.sum(axis) - min_count) < 0
     if null_mask.any():
         dtype, fill_value = dtypes.maybe_promote(result.dtype)
         result = result.astype(dtype)
```

B.2 TYPE 2: PRIOR PATCHES ARE PARTIALLY WRONG

Problem Statement 1

Incorrect parenthesizing of Subs
Here is an example.

```
python
>>> from sympy import Subs
>>> from sympy.abc import x,y
>>> 3*Subs(-x+y, (x,),(1,))
...

LaTeX printing of this gives:
python
'3 \left. - x + y \right|_{\substack{x=1}}
```

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

```

! [image] (https://quicklatex.com/cache3/76/ql_9672fd7e62c909ff3d9ac8543c2e2576_13.png)

It would be better to be parenthesized to:
python
'3 \left. \left(- x + y\right) \right|_{\substack{x=1}}'
python

! [image] (https://quicklatex.com/cache3/bf/ql_936ffdb876e784206d4c54bb93d28dbf_13.png)

```

Conditional patch 1 (Wrong)

```

diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..8b39416 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,12 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
-       latex_expr = self._print(expr)
+       # get the expression in proper parenthesized form
+       # note: we need to use a lower precedence than MUL
+       # to get parens for things like 3*Subs(3+2,...) --> 3*(3+2)|...
+       # the MUL prec level is the highest that always adds the
+       # needed parens
+       latex_expr = self.parenthesize(expr, PRECEDENCE["Mul"] - 1)
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ ' .join(

```

Conditional patch 2 (Correct)

```

diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..2d262c4 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,10 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
-       latex_expr = self._print(expr)
+       # issue 12165: need to parenthesize the expression if it is
+       # an Add (or lower precedence) to get proper grouping
+       # when a coefficient multiplies a substituted expression
+       latex_expr = self.parenthesize(expr, PRECEDENCE["Mul"])
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ ' .join(

```

Conditional patch 3 (Correct)

```

diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..b5b4285 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,11 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
-       latex_expr = self._print(expr)
+       # always parenthesize the expr so Subs is properly grouped
+       # when multiplied, e.g. 3*Subs(-x+y, (x,),(1,)) vs 3*(-x+y)|_x=1
+       # we use Mul precedence since we want to group it at the same
+       # level as args to Mul (but not as high as Pow which wouldn't group)
+       latex_expr = self.parenthesize(expr, PRECEDENCE["Mul"])
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ ' .join(

```

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Conditional patch 4 (Wrong)

```
diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..868ab44 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,13 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
-       latex_expr = self._print(expr)
+       # we want to always parenthesize the expression, otherwise
+       # something like 3*Subs(-x+y, (x,),(1,)) will have
+       # incorrect LaTeX output
+       # to see the issue, compare 3*Subs(-x+y, (x,),(1,)) with
+       # 3*(-x+y).subs(x, 1) in IPython with init_printing()
+       # so we sub in a low precedence level to be safe
        latex_expr = self.parenthesize(expr, 0)
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ \ '.join(
```

Conditional patch 5 (Wrong)

```
diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..dab5b75 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,13 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
-       latex_expr = self._print(expr)
+
+       # issue 15210: always parenthesize the expression because
+       # otherwise the wrong precedence may give incorrect results
+       # For example 3*Subs(x+y, x, 2) would be 3*x+y|_{x=2} rather
+       # than the correct 3*(x+y)|_{x=2}
        latex_expr = self.parenthesize(expr, 0)
+
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ \ '.join(
```

Thinking Process

1. Peer Patch Review:

- Patch 1:
 - Correctly identifies the need to use a lower precedence than MUL to force parentheses
 - The solution is reasonable since MUL precedence is high enough to always add needed
 - ↪ parens
 - However, using PRECEDENCE["Mul"] - 1 might be too low since we want to group at the
 - ↪ same level as MUL args
- Patch 2:
 - Uses PRECEDENCE["Mul"] which is a reasonable choice since we want to group at the
 - ↪ same level as MUL args
 - Correctly identifies the need to parenthesize Add expressions when multiplied
 - Simple and correct solution that addresses the core issue
- Patch 3:
 - Similar to patch 2 but with clearer explanation of why Mul precedence is appropriate
 - Also correctly identifies the need to always parenthesize the expression
 - Good solution that matches the intended behavior
- Patch 4:
 - Uses precedence 0 which is too low - this would always add parentheses even when not
 - ↪ needed
 - While it would fix the issue, it would also introduce unnecessary parentheses in many
 - ↪ cases
 - Not the optimal solution since it's too aggressive with parenthesizing
- Patch 5:
 - Similar to patch 4 but with a clearer explanation of the issue
 - Also uses precedence 0 which has the same problem as patch 4

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

```

- While the intent is correct, the solution is too broad in its application

2. Issue Analysis:
- The issue shows that when a Subs expression is multiplied by a coefficient (like
↳ 3*Subs(-x+y, (x),(1,))), the expression inside Subs isn't properly parenthesized in
↳ LaTeX output
- The current code just prints the expression directly without considering its precedence
↳ relative to multiplication
- The correct behavior should always parenthesize the Subs expression when it appears in
↳ a context where precedence matters (like multiplication)

3. Solution Approach:
- The key is to modify the _print Subs method to always parenthesize the expression using
↳ the MUL precedence level
- This ensures that when Subs appears in a multiplication context, it will be properly
↳ grouped with parentheses
- Using MUL precedence is appropriate because we want it to group at the same level as
↳ other MUL arguments but not higher than necessary

4. Code Localization:
- The issue is in sympy/printing/latex.py in the _print Subs method
- The current implementation directly prints the expression without any parenthesizing
- We need to modify it to use the parenthesize method with MUL precedence

```

Evolved Patch

```

diff --git a/sympy/printing/latex.py b/sympy/printing/latex.py
index 675a3c4..6dd89a9 100644
--- a/sympy/printing/latex.py
+++ b/sympy/printing/latex.py
@@ -698,7 +698,14 @@ class LatexPrinter(Printer):

    def _print Subs(self, subs):
        expr, old, new = subs.args
        latex_expr = self._print(expr)
        # issue 15210: always parenthesize the expression because
        # otherwise the wrong precedence may give incorrect results
        # For example 3*Subs(x+y, x, 2) would be 3*x+y|_{x=2} rather
        # than the correct 3*(x+y)|_{x=2}
        # Using MUL precedence ensures we group at the same level as
        # other MUL args but not higher than necessary (we don't want
        # to group for Pow since Pow has higher precedence)
        latex_expr = self.parenthesize(expr, PRECEDENCE["Mul"])
        latex_old = (self._print(e) for e in old)
        latex_new = (self._print(e) for e in new)
        latex_subs = r'\ \ '.join(

```

Problem Statement 2

```

Contains.as_set returns Contains
```py
>>> Contains(x, Reals).as_set()
Contains(x, Reals)
```

This is wrong because Contains is not a set (it's a boolean). It results in failures in
↳ other places because it doesn't have as_relational (since it isn't a set). For
↳ instance, from https://github.com/sympy/sympy/pull/14965#discussion\_r205281989

```pytb
>>> Piecewise((6, Contains(x, Reals)), (7, True))
Traceback (most recent call last):
 File "<stdin>", line 1, in <module>
 File "./sympy/functions/elementary/piecewise.py", line 136, in __new__
 r = cls.eval(*newargs)
 File "./sympy/functions/elementary/piecewise.py", line 185, in eval
 c = c.as_set().as_relational(x)
AttributeError: 'Contains' object has no attribute 'as_relational'
```

```

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

Conditional patch 1 (Correct)

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..3fd780b 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,6 @@ class Contains(BooleanFunction):
     def as_set(self):
-         raise NotImplementedError()
+         # The set being contained in is the only set we can return
+         # from a Contains object in any meaningful way
+         return self.args[1]
```

Conditional patch 2 (Correct)

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..4c089aa 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,5 @@ class Contains(BooleanFunction):
     def as_set(self):
-         raise NotImplementedError()
+         # The set that is being tested for contains
+         return self.args[1]
```

Conditional patch 3 (Correct)

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..985c88e 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,6 @@ class Contains(BooleanFunction):
     def as_set(self):
-         raise NotImplementedError()
+         # The set being contained in is the set representation
+         # of Contains objects
+         return self.args[1]
```

Conditional patch 4 (Correct)

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..e732791 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,6 @@ class Contains(BooleanFunction):
     def as_set(self):
-         raise NotImplementedError()
+         # Contains(x, S) means "x is in S", so as_set returns S
+         return self.args[1]
```

Conditional patch 5 (Wrong)

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..4980364 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,22 @@ class Contains(BooleanFunction):
     def as_set(self):
+         # try to evaluate Contains(x, S)
+         try:
+             res = self._eval_evalf()
+             if res == S.true:
+                 return S.UniversalSet
+             elif res == S.false:
+                 return S.EmptySet
+         except (TypeError, AttributeError):
+             pass
+         # fallback
+         raise NotImplementedError()
```

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Thinking Process

1. Peer Patch Review:
 - Patch 1: This patch correctly identifies that the set being contained in `(self.args[1])` is the most meaningful set representation of a `Contains` object. The implementation is simple and correct.
 - Patch 2: Similar to Patch 1 but with slightly different wording in the comment. The functionality is correct but the comment could be more precise about why this is the right approach.
 - Patch 3: This patch adds a note about potential invalid sets, which isn't strictly necessary since the issue is about the basic behavior. The core implementation is correct but the comment is overly cautious.
 - Patch 4: This patch provides a clear explanation of why returning the set `(self.args[1])` is the correct behavior - because `Contains(x, S)` means "x is in S", so the set representation is just `S`. This is a correct and well-explained solution.
 - Patch 5: This patch takes a different approach by trying to evaluate the `Contains` expression and return either the universal or empty set based on the result. While creative, this doesn't match the simple and correct solution of just returning the set being contained in. The approach is overly complex for this simple issue.
2. Issue Analysis:
 - The issue occurs when `Contains.as_set()` returns the `Contains` object itself instead of a proper set.
 - This causes problems because other code expects a set with `as_relational()` method.
 - The correct behavior should be to return the set that `x` is being tested for containment in `(self.args[1])`, since that's the only meaningful set representation of a `Contains` object.
3. Solution Approach:
 - The simplest and most correct solution is to modify `as_set()` to return `self.args[1]`, which is the set being tested for containment.
 - This matches the mathematical interpretation that `Contains(x, S)` means "x is in S", so the set representation is just `S`.
 - No other evaluation or special cases are needed since the set being contained in is the only meaningful set representation of a `Contains` object.
4. Code Localization:
 - The issue is in `sympy/sets/contains.py` in the `Contains` class.
 - The problematic method is `as_set()` which currently raises `NotImplementedError`.
 - The fix is to modify this method to return `self.args[1]`.

Evolved Patch

```
diff --git a/sympy/sets/contains.py b/sympy/sets/contains.py
index 0b1e933..e732791 100644
--- a/sympy/sets/contains.py
+++ b/sympy/sets/contains.py
@@ -45,4 +45,6 @@ class Contains(BooleanFunction):
     def as_set(self):
-         raise NotImplementedError()
+         # Contains(x, S) means "x is in S", so as_set returns S
+         return self.args[1]
```

B.3 TYPE 3: PRIOR PATCHES ARE ALL CORRECT

Problem Statement 1

```
_pytest.capture.EncodedFile mode should not include `b` (binary)
<!--
Thanks for submitting an issue!

Here's a quick checklist for what to provide:
-->
- [x] a detailed description of the bug or suggestion
```

```

1350
1351 Exception when youtube-dl logs to pytest captured output. Youtube-dl looks for `b` in
↳ `out.mode` to decide whether to writes `bytes` or `str`. `_pytest.capture.EncodedFile`
1352 ↳ incorrectly advertises `rb+`, the mode of the underlying stream. Its `write()` method
1353 ↳ raises an exception when passed `bytes`.
1354
1355 ...
1356 (pytest-issue-ve3) 01:11:48:nlevitt@Internets-Air-2:/tmp$ py.test test.py
===== test
1357 ↳ session starts
1358 ↳ =====
1359 platform darwin -- Python 3.7.3, pytest-4.5.0, py-1.8.0, pluggy-0.11.0
1360 rootdir: /private/tmp
1361 collected 1 item
1362
1363 test.py F
1364 ↳ [100%]
1365
1366 =====
1367 ↳ FAILURES
1368 ↳ =====
1369
1370 ↳ test_foo
1371 ↳ -----
1372
1373     def test_foo():
1374 >         youtube_dl.YouTubeDL().extract_info('http://example.com/')
1375
1376 test.py:4:
1377 -----
1378 ↳ -----
1379 ↳ _
1380 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/YouTubeDL.py:796: in extract_info
1381     ie_result = ie.extract(url)
1382 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/extractor/common.py:529: in
1383     extract
1384     ie_result = self._real_extract(url)
1385 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/extractor/generic.py:2245: in
1386     _real_extract
1387     self.to_screen('%s: Requesting header' % video_id)
1388 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/extractor/common.py:913: in
1389     to_screen
1390     self._downloader.to_screen('[%s] %s' % (self.IE_NAME, msg))
1391 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/YouTubeDL.py:502: in to_screen
1392     return self.to_stdout(message, skip_eol, check_quiet=True)
1393 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/YouTubeDL.py:516: in to_stdout
1394     self._write_string(output, self._screen_file)
1395 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/YouTubeDL.py:505: in
1396     _write_string
1397     write_string(s, out=out, encoding=self.params.get('encoding'))
1398 pytest-issue-ve3/lib/python3.7/site-packages/youtube_dl/utils.py:1496: in write_string
1399     out.write(byt)
1400 -----
1401 ↳ -----
1402 ↳ _
1403
1404 self = <_pytest.capture.EncodedFile object at 0x10df124a8>, obj = b'[generic] example:
1405 ↳ Requesting header\n'
1406
1407     def write(self, obj):
1408         if isinstance(obj, six.text_type):
1409             obj = obj.encode(self.encoding, "replace")
1410         elif _PY3:
1411             raise TypeError(
1412 >                 "write() argument must be str, not {}".format(type(obj).__name__)
1413             )
1414     E         TypeError: write() argument must be str, not bytes
1415
1416 pytest-issue-ve3/lib/python3.7/site-packages/_pytest/capture.py:437: TypeError
1417 ===== 1 failed in
1418 ↳ 2.74 seconds
1419 ↳ =====
1420 ...
1421
1422 - [x] output of `pip list` from the virtual environment you are using
1423 ...
1424
1425 Package      Version
1426 -----
1427 atomicwrites 1.3.0
1428 attrs        19.1.0

```

```

1404
1405 more-itertools 7.0.0
1406 pip 19.1.1
1407 pluggy 0.11.0
1408 py 1.8.0
1409 pytest 4.5.0
1410 setuptools 41.0.1
1411 six 1.12.0
1412 wcwidth 0.1.7
1413 wheel 0.33.4
1414 youtube-dl 2019.5.11
1415 ...
1416
1417 - [x] pytest and operating system versions
1418 ...
1419 This is pytest version 4.5.0, imported from
1420 ↪ /private/tmp/pytest-issue-ve3/lib/python3.7/site-packages/pytest.py
1421 ...
1422 ...
1423 macOS 10.14.4 (18E226)
1424 ...
1425
1426 - [x] minimal example if possible
1427 ...
1428 pip install pytest youtube-dl
1429 py.test test.py
1430 ...
1431
1432 test.py:
1433 ...
1434 import youtube_dl
1435 def test_foo():
1436     youtube_dl.YouTubeDL().extract_info('http://example.com/')
1437 ...

```

Conditional patch 1 (Correct)

```

1431 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1432 index 6b2d076..aa57c0d 100644
1433 --- a/src/_pytest/capture.py
1434 +++ b/src/_pytest/capture.py
1435 @@ -448,6 +448,17 @@ class EncodedFile(object):
1436     return repr(self.buffer)
1437
1438     def __getattr__(self, name):
1439 + # when code inspects the mode attribute (like youtube-dl does)
1440 + # we need to hide the 'b' flag from it - we act as a text stream
1441 + # even though we write to a bytes buffer
1442 + if name == "mode":
1443 +     mode = getattr(self.buffer, name, None)
1444 +     if mode and "b" in mode:
1445 +         # return a copy of the mode without 'b'
1446 +         # this makes e.g. youtube-dl think this is a text stream
1447 +         # and write str instead of bytes
1448 +         return mode.replace("b", "")
1449 + # fall back to regular attribute access
1450     return getattr(object.__getattr__(self, "buffer"), name)

```

Conditional patch 2 (Correct)

```

1447
1448 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1449 index 6b2d076..01e03c8 100644
1450 --- a/src/_pytest/capture.py
1451 +++ b/src/_pytest/capture.py
1452 @@ -448,6 +448,20 @@ class EncodedFile(object):
1453     return repr(self.buffer)
1454
1455     def __getattr__(self, name):
1456 + # youtube-dl specifically looks at out.mode to decide whether to
1457 + # write bytes or str. Since we're a text interface over a bytes
1458 + # interface, we should make it look like we're text.
1459 + #
1460 ↪ https://github.com/ytdl-org/youtube-dl/blob/2019.05.11/youtube_dl/utils.py#L1496
1461 + # See: https://github.com/pytest-dev/pytest/issues/5691
1462 + if name == "mode":

```

```

1458
1459 +         mode = getattr(self.buffer, name, None)
1460 +         if mode and "b" in mode:
1461 +             # remove the 'b' from the mode, but otherwise leave it
1462 +             # unchanged. This makes youtube-dl think it can write text
1463 +             # to us when in reality we'll encode it to bytes for the
1464 +             # underlying buffer.
1465 +             mode = "".join(ch for ch in mode if ch != "b")
1466 +         return mode
1467
1468 return getattr(object.__getattribute__(self, "buffer"), name)

```

Conditional patch 3 (Correct)

```

1468 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1469 index 6b2d076..14dfb3c 100644
1470 --- a/src/_pytest/capture.py
1471 +++ b/src/_pytest/capture.py
1472 @@ -447,6 +447,21 @@ class EncodedFile(object):
1473     """Ensure that file.name is a string."""
1474     return repr(self.buffer)
1475
1476 + @property
1477 + def mode(self):
1478 +     """
1479 +     Get the mode of the underlying buffer but remove 'b' from it.
1480 +     This is needed because some code (like youtube-dl) checks
1481 +     for 'b' in mode to decide whether to write bytes or str.
1482 +     Since this is an encoded (text) file, we need to hide that
1483 +     it's wrapping a binary stream.
1484 +     """
1485 +     mode = getattr(self.buffer, "mode", None)
1486 +     if mode and "b" in mode:
1487 +         # remove 'b' but keep other flags like '+' or 'r'
1488 +         mode = "".join(c for c in mode if c != "b")
1489 +     return mode
1490
1491 def __getattr__(self, name):
1492     return getattr(object.__getattribute__(self, "buffer"), name)

```

Conditional patch 4 (Correct)

```

1488 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1489 index 6b2d076..df7fb36 100644
1490 --- a/src/_pytest/capture.py
1491 +++ b/src/_pytest/capture.py
1492 @@ -447,6 +447,22 @@ class EncodedFile(object):
1493     """Ensure that file.name is a string."""
1494     return repr(self.buffer)
1495
1496 + @property
1497 + def mode(self):
1498 +     """Lie about our mode - remove 'b' if present.
1499 +
1500 +     Some libraries (like youtube-dl) inspect the mode to decide whether to
1501 +     write bytes or str to the stream. Since we handle that in .write() and
1502 +     .writelines() already, we should pretend to be a text stream even if our
1503 +     underlying buffer is a byte stream.
1504 +     """
1505 +     # just in case the buffer doesn't even have a mode...
1506 +     mode = getattr(self.buffer, "mode", "")
1507 +     # ... and remove the 'b'
1508 +     mode = mode.replace("b", "")
1509 +     return mode
1510
1511 def __getattr__(self, name):
1512     return getattr(object.__getattribute__(self, "buffer"), name)

```

Conditional patch 5 (Correct)

```

1508 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1509 index 6b2d076..f94e9dc 100644
1510 --- a/src/_pytest/capture.py
1511 +++ b/src/_pytest/capture.py
1512 @@ -447,6 +447,14 @@ class EncodedFile(object):

```

```

1512         """Ensure that file.name is a string."""
1513         return repr(self.buffer)
1514
1515     + @property
1516     + def mode(self):
1517     +     # youtube-dl checks for 'b' in mode to decide whether to write bytes or str.
1518     +     # since we are a text interface (we encode str to bytes), we should not
1519     +     # appear as a binary stream, regardless of the underlying buffer's mode.
1520     +     # https://github.com/pytest-dev/pytest/issues/5295
1521     +     return self.buffer.mode.replace("b", "")
1522
1523     def __getattr__(self, name):
1524         return getattr(object.__getattr__(self, "buffer"), name)

```

Thinking Process

1. Peer Patch Review:

- Patch 1:
 - Correctly identifies the need to modify the mode attribute when accessed
 - Uses `__getattr__` to intercept the mode attribute and remove 'b'
 - Simple and effective solution that matches the issue requirements
 - Correctly maintains all other attribute behavior
- Patch 2:
 - Similar to patch 1 but uses a different string manipulation method
 - Also correctly intercepts the mode attribute and removes 'b'
 - Works just as well as patch 1 but with slightly different implementation
- Patch 3:
 - Takes a different approach by adding a mode property directly
 - This is a cleaner solution since it explicitly defines the behavior
 - Correctly handles the case where the buffer doesn't have a mode
 - Maintains all other attribute behavior through `__getattr__`
- Patch 4:
 - Similar to patch 3 but with more detailed docstring
 - Also adds the mode property to remove 'b' from the underlying mode
 - Correctly handles the case where the buffer doesn't have a mode
 - Maintains all other attribute behavior through `__getattr__`
- Patch 5:
 - Simplest solution that just adds a mode property
 - Directly returns the underlying mode with 'b' removed
 - Doesn't handle the case where the buffer might not have a mode attribute
 - Still solves the core issue but is less robust than patches 3-4

2. Issue Analysis:

- The problem occurs when youtube-dl checks the mode of the captured output stream
- It sees 'b' in the mode and tries to write bytes when the EncodedFile expects str
- The solution needs to make the EncodedFile appear as a text stream (no 'b' in mode)
- While still maintaining all other file-like behavior from the underlying buffer

3. Solution Approach:

- The key is to intercept access to the 'mode' attribute of the underlying buffer
- When 'mode' is accessed, we should return the mode string with 'b' removed
- All other attribute access should continue to work normally from the buffer
- This matches exactly what youtube-dl expects while maintaining proper functionality

4. Code Localization:

- The issue is in the EncodedFile class in `src/_pytest/capture.py`
- Need to modify the class to add a mode property that removes 'b' from the underlying
 - ↪ buffer's mode
- All other behavior should remain unchanged

Evolved Patch

```

1558
1559
1560 diff --git a/src/_pytest/capture.py b/src/_pytest/capture.py
1561 index 6b2d076..d3b3473 100644
1562 --- a/src/_pytest/capture.py
1563 +++ b/src/_pytest/capture.py
1564 @@ -447,6 +447,23 @@ class EncodedFile(object):
1565         """Ensure that file.name is a string."""
1566         return repr(self.buffer)
1567
1568     + @property

```

```

1566
1567 + def mode(self):
1568 +     """When code inspects our mode (like youtube-dl does), lie and
1569 +     remove the 'b' if present. This makes youtube-dl think we're a text
1570 +     stream and write str instead of bytes, which we can then encode to
1571 +     bytes for the underlying buffer.
1572 +
1573 +     Note: we don't need to handle the case where the buffer doesn't have
1574 +     a mode attribute because in that case, normal attribute access will
1575 +     raise AttributeError like it should.
1576 +     """
1577 +     mode = self.buffer.mode
1578 +     if "b" in mode:
1579 +         # remove the 'b' but otherwise preserve the mode
1580 +         mode = mode.replace("b", "")
1581 +     return mode
1582 +
1583 + def __getattr__(self, name):
1584 +     return getattr(object.__getattr__(self, "buffer"), name)
1585

```

Problem Statement 2

DateFormat.Y() is not zero-padded.
Description

The Y specifier for django.utils.dateformat.DateFormat is supposed to always return a
↔ four-digit year padded with zeros. This doesn't seem to be the case for year < 1000.

Conditional patch 1 (Correct)

```

1588 diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
1589 index 38cb1c1..823e893 100644
1590 --- a/django/utils/dateformat.py
1591 +++ b/django/utils/dateformat.py
1592 @@ -314,7 +314,9 @@ class DateFormat(TimeFormat):
1593
1594     def Y(self):
1595         "Year, 4 digits; e.g. '1999'"
1596         - return self.data.year
1597         + # We want to zero-pad the year (this is only useful for year < 1000)
1598         + # so that we get 4 digits.
1599         + return '%04d' % self.data.year
1600
1601     def z(self):
1602         """Day of the year, i.e. 1 to 366."""

```

Conditional patch 2 (Correct)

```

1601 diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
1602 index 38cb1c1..15090a4 100644
1603 --- a/django/utils/dateformat.py
1604 +++ b/django/utils/dateformat.py
1605 @@ -314,7 +314,9 @@ class DateFormat(TimeFormat):
1606
1607     def Y(self):
1608         "Year, 4 digits; e.g. '1999'"
1609         - return self.data.year
1610         + # We want to zero-pad the year, so can't just return the number --
1611         + # hence the "%04d" dance.
1612         + return '%04d' % self.data.year
1613
1614     def z(self):
1615         """Day of the year, i.e. 1 to 366."""

```

Conditional patch 3 (Correct)

```

1614 diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
1615 index 38cb1c1..abab54f 100644
1616 --- a/django/utils/dateformat.py
1617 +++ b/django/utils/dateformat.py
1618 @@ -314,7 +314,8 @@ class DateFormat(TimeFormat):
1619
1620     def Y(self):
1621         "Year, 4 digits; e.g. '1999'"

```

```

1620
1621 -         return self.data.year
1622 +         # We want a four-digit year; 0-padding if necessary
1623 +         return '%04d' % self.data.year
1624
1625     def z(self):
1626         """Day of the year, i.e. 1 to 366."""

```

Conditional patch 4 (Correct)

```

1627
1628 diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
1629 index 38cblc1..ecf7449 100644
1630 --- a/django/utils/dateformat.py
1631 +++ b/django/utils/dateformat.py
1632 @@ -314,7 +314,11 @@ class DateFormat(TimeFormat):
1633
1634     def Y(self):
1635         "Year, 4 digits; e.g. '1999'"
1636 -         return self.data.year
1637 +         # We hardcode 4 here instead of using a settings value or
1638 +         # self.data.year.__len__() because we want to always return
1639 +         # a 4-digit year for consistency with PHP's Y behavior.
1640 +         # See also: https://bugs.php.net/bug.php?id=67077
1641 +         return '%04d' % self.data.year
1642
1643     def z(self):
1644         """Day of the year, i.e. 1 to 366."""

```

Conditional patch 5 (Correct)

```

1645
1646 diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
1647 index 38cblc1..acd0ae0 100644
1648 --- a/django/utils/dateformat.py
1649 +++ b/django/utils/dateformat.py
1650 @@ -314,7 +314,9 @@ class DateFormat(TimeFormat):
1651
1652     def Y(self):
1653         "Year, 4 digits; e.g. '1999'"
1654 -         return self.data.year
1655 +         # We want to always return a 4-digit year, zero-padded if necessary
1656 +         # This ensures years < 1000 are properly formatted
1657 +         return '%04d' % self.data.year
1658
1659     def z(self):
1660         """Day of the year, i.e. 1 to 366."""

```

Thinking Process

```

1655 1. Peer Patch Review:
1656
1657 - Patch 1:
1658 - Correctly identifies the need to zero-pad the year using '%04d'
1659 - The comment explains the issue well for years < 1000
1660 - Simple and correct solution that matches the issue requirements
1661
1662 - Patch 2:
1663 - Similar to Patch 1 but with slightly different comment wording
1664 - Also correctly uses '%04d' to zero-pad the year
1665 - Correct solution that addresses the issue
1666
1667 - Patch 3:
1668 - Uses '%04d' like the other patches but with simpler comment
1669 - Correct solution that matches the issue requirements
1670
1671 - Patch 4:
1672 - Adds more detailed comment explaining the PHP compatibility
1673 - While the solution is correct, the additional comment about PHP behavior isn't
1674   ↪ necessary for this Django-specific issue
1675 - Still a correct solution to the issue
1676
1677 - Patch 5:
1678 - Simple and correct solution using '%04d' with clear comment
1679 - Correctly addresses the issue requirements
1680
1681 2. Issue Analysis:

```

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686

```

- The issue states that DateFormat.Y() should always return a 4-digit year padded with
↳ zeros
- Currently it returns the raw year number which doesn't zero-pad for years < 1000
- The fix requires modifying the Y() method to format the year as a 4-digit string with
↳ zero-padding
- This matches the behavior described in the method's docstring which says "4 digits"

3. Solution Approach:
- The solution is to modify the Y() method to use string formatting with '%04d'
- This will ensure the year is always returned as a 4-character string with leading zeros
↳ if needed
- The change is simple and focused on just the Y() method where the issue occurs

4. Code Localization:
- The issue is in django/utils/dateformat.py in the DateFormat class
- The problematic method is Y() which currently just returns self.data.year
- Need to modify it to return '%04d' % self.data.year instead

```

1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700

Evolved Patch

```

diff --git a/django/utils/dateformat.py b/django/utils/dateformat.py
index 38cb1c1..d0e66c3 100644
--- a/django/utils/dateformat.py
+++ b/django/utils/dateformat.py
@@ -314,7 +314,9 @@ class DateFormat(TimeFormat):

    def Y(self):
        "Year, 4 digits; e.g. '1999'"
-       return self.data.year
+       # We want to zero-pad the year to ensure we always get
+       # 4 digits, even for years < 1000.
+       return '%04d' % self.data.year

    def z(self):
        ""Day of the year, i.e. 1 to 366.""

```

C SCAFFOLD OF EVOSCALE

1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711

We adopt a pipeline-based scaffold consisting of two main components: a retriever and a code editing model. The retriever takes either the GitHub repository structure or the contents of individual code files as input and identifies the files most relevant to the given issue. Once retrieval is complete, the full content of the selected code files is passed to the code editing model.

The code editing model receives both the issue description and the retrieved code content as input and generates a patch to resolve the issue. Additionally, there is an optional verifier component, which can be used to select the best patch from a large pool of candidate samples. We describe each component in detail below.

1712
1713

C.1 RETRIEVER

1714
1715

Our retriever is entirely LLM-based and consists of two components: a retrieval model and a retrieval reward model.

1716
1717
1718
1719
1720
1721
1722
1723
1724



1725
1726
1727

Figure 15: **Retrieval Pipeline.** Given the repository’s file structure, the retrieval model first selects the top-5 candidate files. These candidates are then re-scored by the retrieval reward model based on file content, and the top-ranked (Top-1) file is returned as the final result.

Retrieval Model The first stage of our retriever uses a retrieval model to identify the top 5 most relevant files based on the repository structure and the GitHub issue description. We adopt the same format as Agentless Xia et al. (2024) to represent the repository’s file structure. Given this representation and the issue description, the retrieval model performs a reasoning process and outputs five file paths from the repository. The model is trained using a combination of small-scale supervised fine-tuning (SFT) and large-scale reinforcement learning (RL), see Appendix D.3 for details.

Retrieval Reward Model The retrieval reward model is designed to refine retrieval results in a more fine-grained manner. After the initial top-5 files are retrieved by the retrieval model, the reward model evaluates each one by considering both the file’s code content and the issue description. It then outputs a relevance score for each file, and the file with the highest score is selected as the final target for code editing. The retrieval reward model is a classifier-style LLM trained with a binary classification objective, see Appendix D.4 for training details.

C.2 CODE EDITING MODEL

The code editing model receives a prompt formed by concatenating the issue statement with the code content of the retrieved target file. It performs iterative sampling to enable self-evolution during generation.

In the first iteration, given the issue statement and code context, the model generates five diverse responses, each corresponding to a different patch candidate. These five patch candidates are then appended to the input as a conditional prompt for the next iteration of generation. This iterative process allows the model to progressively refine its outputs.

As discussed in Section 4, the code editing model is trained using a combination of small-scale SFT and large-scale RL. Additional training details are provided in Appendix D.5.

C.3 VERIFIER

As discussed in Section 5.2, our code editing model demonstrates the ability to self-evolve by iteratively refining its own generations. While this process improves the quality of patch samples, incorporating external verifiers to select the optimal patch can further boost performance. For software engineering tasks, we consider two primary types of verifiers: an LLM-based reward model and unit tests.

Code Editing Reward Model The code editing reward model is designed to select the best patch from a pool of candidates. It takes as input the retrieved file’s code content, the issue description, and a patch candidate in git diff format. The model then outputs a score indicating the quality of the patch. This reward model is implemented as a classifier-based LLM trained with a binary classification objective (see Appendix D.6 for details).

Unit Tests Unit tests consist of two components: (1) Reproduction tests, which validate whether the original GitHub issue can be reproduced and resolved by the patch; (2) Regression tests, which check whether the patch preserves the existing functionality of the codebase. To construct the regression tests, we extract existing test files from the repository using the Agentless Xia et al. (2024) pipeline. For the reproduction tests, we use a trained test generation model that takes the issue description as input and generates tests aimed at reproducing the issue. This reproduction test generator is trained using supervised fine-tuning (see Appendix D.7). For each instance, we sample 100 reproduction tests and retain 5 valid patches to serve as reproduction tests.

Hybrid Verifiers We combine multiple verifiers to select the most promising patch candidates. The selection process is as follows: (1) Regression tests are applied first. Any patch that fails is discarded; (2) Reproduction tests are then executed on the remaining patches. Candidates are ranked based on how many of the five tests they pass. (3) The top- k ($k = 2$) unique patches are retained per instance. (4) If no patch passes both regression and reproduction tests, we fall back to using all generated candidates without filtering. (5) Finally, among the remaining patches, the code editing reward model is used to select the candidate with the highest reward score for submission.

D IMPLEMENTATION DETAILS

D.1 DATASET COLLECTION

Our primary training data is sourced from SWE-Fixer and SWE-Gym. To ensure data quality, we apply a comprehensive filtering and deduplication process. Specifically, we discard instances that meet any of the following criteria: (1) extremely short or excessively long issue statements; (2) multimodal or non-text content (e.g., images, videos, LaTeX); (3) presence of unrelated external links; (4) inclusion of commit hashes; (5) patches requiring modifications to more than three files. After applying these filters, we obtain 29,404 high-quality training instances, which we use to train both the retrieval and code editing models.

D.2 TRAINING PIPELINE AND HARDWARE

Both the retrieval and code editing models are trained in two stages: (1) small-scale supervised fine-tuning (SFT) for warm-up, and (2) large-scale reinforcement learning (RL) for self-improvement. The SFT trajectories are generated via chain-of-thought prompting using Deepseek-V3-0324. We adopt the Qwen2.5-Coder-32B-Instruct model Hui et al. (2024) as the base model for all components due to its strong code reasoning capabilities. We utilize OpenRLHF Hu et al. (2024) as the training framework for SFT, and use VERL Sheng et al. (2025) as the training framework for RL. All training runs are conducted on NVIDIA H100 GPUs with 80 GB of memory. For evaluation and model inference, we serve models using the sglang framework², employing tensor parallelism with a parallel size of 8 on NVIDIA H100 GPUs.

D.3 RETRIEVAL MODEL

Dataset Construction To ensure cost-efficient synthetic data generation, we randomly select 1,470 instances (5% of the full dataset) as the small-scale SFT dataset for training the retrieval model.

To train the model to perform step-by-step reasoning and generate relevant file paths conditioned on both the issue statement and the repository structure, we require training data that includes both intermediate reasoning and final retrieved files. However, the raw data only provides the issue descriptions and the ground-truth retrieved files (extracted from the final patch), without any intermediate reasoning.

To address this, we use Deepseek-V3-0324 with a custom retrieval model prompt (see Appendix E) and apply rejection sampling to collect high-quality chain-of-thought (CoT) reasoning traces. Specifically, we check whether the model’s top-5 retrieved files include the ground-truth retrieval files. If so, we retain the response as part of our synthetic SFT data.

For each selected instance, we generate one response via greedy decoding and three additional responses using random sampling (temperature = 0.7), as long as they include the correct retrieval files. This results in four responses per instance.

For RL training, we use the remaining 27,598 instances (95% of the dataset), filtering out prompts whose total sequence length exceeds 16,384 tokens. The RL dataset consists of prompts (issue statement + repo structure) as input and the corresponding ground-truth retrieval files as the final answer, without requiring intermediate reasoning.

Supervised Fine-Tuning We perform supervised fine-tuning (SFT) on the Qwen2.5-Coder-32B-Instruct model using the synthetic SFT dataset described above. The prompt template used for training is identical to the one used to construct the synthetic data (see the retrieval model prompt in Appendix E). We train the model using a cosine learning rate scheduler with an initial learning rate of 1e-5. The model is fine-tuned for one epoch with a batch size of 128 and a maximum sequence length of 32,768 tokens.

Reinforcement Learning To further push the limit of the model’s retrieval performance, we apply a large-scale reinforcement learning (RL) stage after the SFT stage. After SFT, the model has learned to

²<https://github.com/sgl-project/sglang>

reason step-by-step and generates a set of candidate retrieval files, denoted as $\mathcal{Y} = \{y_1, y_2, \dots, y_k\}$, where $k = 5$. Given the ground-truth set of target files \mathcal{F} , we define the reward as the proportion of correctly retrieved files:

$$\text{Reward} = \frac{|\mathcal{Y} \cap \mathcal{F}|}{|\mathcal{F}|}$$

Given the prompt in the RL dataset, we let the model generate response and self-improve through this reward signal. We use the REINFORCE++ [Hu et al. \(2025\)](#) algorithm with a fixed learning rate of $1e-6$ for the actor model. During training, we sample 8 rollouts per prompt. The training batch size is 64, and the rollout batch size is 256. The model is trained for 3 epochs, with a maximum prompt length of 16k tokens and generation length of 4k tokens. Additional hyperparameters include a KL divergence coefficient of 0.0, entropy coefficient of 0.001 and a sampling temperature of 1.0.

D.4 RETRIEVAL REWARD MODEL

To train a reward model capable of reliably identifying the most relevant code files for modification, we construct a reward model dataset derived from our main dataset. The final reward model dataset consists of 112,378 samples corresponding to 25,363 unique instances. For each instance, the prompt is constructed using the retrieval reward model prompt template (see Appendix E), incorporating the issue statement along with the code content of each of the top-5 candidate retrieval files. Each data point is labeled with a binary value $\in \{0, 1\}$, indicating whether the provided code content belongs to the ground-truth retrieval files. The model is initialized from the Qwen2.5-Coder-32B-Instruct and trained as a binary classifier using cross-entropy loss. Training is conducted with a batch size of 128, a learning rate of $5e-6$, and a maximum sequence length of 32,768 tokens, over two epochs.

D.5 CODE EDITING MODEL

Dataset construction As described in Section 4, our code editing model is trained in two stages using supervised fine-tuning (SFT)—classical SFT and mutation SFT—followed by large-scale reinforcement learning (RL). We randomly select 1,470 instances (5%) from the full dataset for the classical SFT set and a separate 1,470 instances (5%) for the mutation SFT set. These two subsets are kept disjoint to ensure that the model learns to self-refine without direct exposure to the ground-truth solutions. For RL training, we use the remaining 22,102 instances (90% of the dataset), filtering out any prompts with sequence lengths exceeding 16,384 tokens. The RL dataset contains only the prompt (issue + code context) as input and the corresponding ground-truth patch as the output.

To synthesize the reasoning chain-of-thought (CoT) for classical SFT, we prompt Deepseek-V3-0324. Unlike the retrieval setting, we do not use rejection sampling, as Deepseek-V3-0324 often fails to generate the correct patch even after multiple samples. Instead, we adopt a more efficient approach by designing a “role-playing” prompt that provides the model access to the ground-truth patch and instructs it to explain the reasoning process behind it (see the “Generating Reasoning CoT for Code Editing Model (Classical SFT)” prompt in Appendix E). This ensures that the generated reasoning is both accurate and reflects an independent thought process. We then synthesize the classical SFT dataset using the “Code Editing Model (Classical SFT)” prompt template in Appendix E.

We first fine-tune the base model on the classical SFT dataset. This fine-tuned model is then used to generate five random patch candidates per instance with a sampling temperature of 1.0. These candidate patches are used to construct the mutation SFT dataset. For each instance, we prompt Deepseek-V3-0324 with: the issue statement, the content of the target file, the five candidate patches, and the ground-truth patch. Using the “Generating Reasoning CoT for Code Editing Model (Mutation SFT)” prompt (see Appendix E), the model is instructed to review each patch, critique their strengths and weaknesses, and propose an improved solution. We then extract the reasoning process and synthesize the mutation SFT dataset using the “Code Editing Model (Mutation SFT)” prompt template.

Supervised Fine-Tuning We perform supervised fine-tuning (SFT) on the Qwen2.5-Coder-32B-Instruct model using the synthetic SFT datasets described above. The prompt templates used for training are the same as those used to construct the two-stage SFT datasets (classical and mutation SFT). We employ a cosine learning rate scheduler with an initial learning rate of $1e-5$. Training is conducted for one epoch, with a batch size of 128 and a maximum sequence length of 32,768 tokens.

Reinforcement Learning We fine-tune the mutation SFT model on the full dataset using REINFORCE++ [Hu et al. \(2025\)](#) with the following reward function:

$$r = \underbrace{R(x, y)}_{\text{Bonus}} + \underbrace{R(x, y) - \sum_{i=1}^K R(x, \bar{y}_i)}_{\text{Potential}} - \lambda \underbrace{F(y)}_{\text{Format}}, \quad (6)$$

where each term is defined as follows:

- $R(x, y)$ (Bonus): Encourages the model to produce high-reward outputs. Although similar in effect to the potential term, including this bonus stabilizes training and consistently improves the model’s average reward.
- $R(x, y) - \sum_{i=1}^K R(x, \bar{y}_i)$ (Potential): Measures the improvement of the current patch y over the average reward of the K conditioning patches \bar{y}_i . See Section 4.3 for details.
- $F(y)$ (Format): Penalizes outputs that violate format or syntax constraints. It consists of:
 - **String matching:** Rewards outputs that closely match the ground-truth patch y^* using sequence similarity, following [Wei et al. \(2025\)](#).
 - **Syntax check:** Ensures the output can be parsed into the expected search-replace format, passes Python’s `ast` syntax check, and satisfies `flake8` static analysis. If any check fails, the format reward is set to zero.

The RL model is trained on a mix of data with and without conditioning examples. Conditioning examples are generated not only using the classical SFT model but also using a checkpoint of an RL-trained model at the first epoch.

As for implementation, we use REINFORCE++ [Hu et al. \(2025\)](#) algorithm with a fixed learning rate of $1e-6$ for the actor model. During training, we sample 8 rollouts per prompt. The training batch size is 64, and the rollout batch size is 256. The model is trained for only 1 epochs, with a maximum prompt length of 16k tokens and generation length of 8k tokens. Additional hyperparameters include a KL divergence coefficient of 0.0, entropy coefficient of 0.001 and a sampling temperature of 1.0.

D.6 CODE EDITING REWARD MODEL

The code editing reward model is designed to provide a more accurate reward signal, addressing the limitations of using simple string-matching scores. The training setup is similar to that of the retrieval reward model (see Appendix D.4), with the main difference being in the data collection process. We construct the reward model training dataset using data collected from `nebius/SWE-agent-trajectories` and `nebius/SWE-bench-extra`, resulting in 56,797 samples across 1,889 unique instances. For each instance, the prompt is constructed using the code editing reward model prompt template (see Appendix E), and includes the issue statement, the code content of the target file to be modified, and a candidate patch. Each sample is labeled with a binary value $\in \{0, 1\}$, indicating whether the candidate patch successfully resolves the issue. The model is trained as a binary classifier using the same training settings as the retrieval reward model.

D.7 REPRODUCTION TEST GENERATOR

Following a similar approach to that used for code editing, we generate intermediate reasoning steps for reproduction test generation using the `Deepseek-V3-0324` model. Given the issue description and the corresponding ground-truth test patch, the model is prompted to produce a response that includes the reasoning behind constructing a valid test in a chain-of-thought format.

To support automated verification, we follow the strategy used in `Agentless` [Xia et al. \(2024\)](#), employing a test script template that prints clear diagnostic messages indicating whether the issue has been successfully reproduced or resolved. Specifically, If the test successfully triggers the target error (e.g., raises an `AssertionError`), it prints `"Issue reproduced"`; If the test completes without errors, it prints `"Issue resolved"`. An example template for this diagnostic test script is shown below:

1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

Reproduction Test Template

```
def test_<meaningful_name>() -> None:
    try:
        # Minimal code that triggers the bug
        ...
    except AssertionError:
        print("Issue reproduced")
        return
    except Exception:
        print("Other issues")
        return

    print("Issue resolved")
    return

if __name__ == "__main__":
    test_<meaningful_name>()
```

Starting from our filtered dataset, we generate one response per instance using greedy decoding and three additional responses via sampling with a temperature of 0.7. These synthetic examples are then used to fine-tune the Qwen2.5-Coder-32B-Instruct model over three epochs, resulting in our reproduction test generation model. The prompt templates used for generating intermediate reasoning and for supervised fine-tuning are provided in Appendix E.

E PROMPT TEMPLATE

Prompt Template — Retrieval Model

Please look through the following GitHub problem description and Repository structure. Determine the files most likely to be edited to fix the problem. Identify 5 most important files.

```

### GitHub Problem Description ###
{problem_statement}

### Repository Structure ###
{structure}

### Format Instruction ###
1. Enclose reasoning process within `...</think>`.
2. Please only provide the full path and return 5 most important files. Always return exactly 5 files, Do Not output less than 5 or more than 5 files.
3. The returned files should be separated by new lines ordered by most to least important.
4. Do not include any explanations after `...</file>`.

### Examples ###
<think>
1. Analyze the issue...
2. Check the files in provided repository structure for relevance...
3. Confirm that the issue might be most relevant to 5 relevant files...
</think>

<file>
file1.py
file2.py
file3.py
file4.py
file5.py
</file>

---
```

Please provide your response below.

Prompt Template — Retrieval Reward Model

You are an expert software engineer and seasoned code reviewer, specializing in bug localization and
↳ code optimization.

You will be presented with a GitHub issue and a source code file.
Your task is to decide if the code file is relevant to the issue.

```

# Issue Statement
{problem_statement}

# File to be Modified
{file_content}
```

Prompt Template — Generating Reasoning CoT for Code Editing Model (Classical SFT)

You are a student striving to become an expert software engineer and seasoned code reviewer,
↳ specializing in bug localization and code optimization within real-world code repositories. Your
↳ strengths lie in understanding complex codebase structures and precisely identifying and modifying
↳ the relevant parts of the code to resolve issues. You also excel at articulating your reasoning
↳ process in a coherent, step-by-step manner that leads to efficient and correct
bug fixes.

You are now taking an exam to evaluate your capabilities. You will be provided with a codebase and an
↳ issue description. Your task is to simulate a complete reasoning process--step-by-step--as if
↳ solving the issue from scratch, followed by the code modifications to resolve the issue.

To evaluate your correctness, an oracle code modification patch will also be provided. You must ensure
↳ that your final code modifications MATCH the oracle patch EXACTLY. However, your reasoning process
↳ must appear fully self-derived and **must NOT reference, suggest awareness of, or appear to be
↳ influenced by** the oracle patch. You must solve the problem as if you are unaware of the oracle
↳ solution.

```

---
```

```

# Issue Statement
{problem_statement}

---
```

```

2052 # Files to be Modified
2053 Below are some code files that might be relevant to the issue above. One or more of these files may
2054 ↪ contain bugs.
2055
2056 {file_content}
2057
2058 ---
2059 # Oracle Code Modification Patch (For Evaluation Only):
2060 {oracle_patch}
2061
2062 ---
2063 # Reasoning Guidelines
2064 Your reasoning process should generally follow these steps, with flexibility to adjust as needed for
2065 ↪ clarity and accuracy:
2066
2067 1. Issue Analysis: Start by thoroughly analyzing the issue. Explain what the problem is, why it
2068 ↪ matters, and what the intended behavior should be. Identify the key goals and constraints that must
2069 ↪ be addressed in your solution.
2070
2071 2. Task Decomposition: Break down the issue into smaller, manageable sub-tasks. Describe the
2072 ↪ purpose of each sub-task and how it contributes to solving the overall problem.
2073
2074 3. Code Localization and Editing: For each sub-task:
2075 ↪ - Identify relevant code snippets by file path and code location.
2076 ↪ - Explain how each snippet relates to the sub-task.
2077 ↪ - Describe how the code should be changed and justify your reasoning.
2078 ↪ - After thorough explanation, provide the corresponding edited code.
2079
2080 Your final output must precisely match the oracle patch, but your thinking must remain fully grounded
2081 ↪ in the issue description and provided code files.
2082
2083 ---
2084 # General Requirements
2085 1. Independent and Evidence-Based Reasoning: Your reasoning must be constructed as if independently
2086 ↪ derived, based solely on the issue and code. Do not reference or imply knowledge of the oracle
2087 ↪ patch.
2088 2. Clarity and Justification: Ensure that each reasoning step is clear, well-justified, and easy to
2089 ↪ follow.
2090 3. Comprehensiveness with Focus: Address all relevant components of the issue while remaining
2091 ↪ concise and focused.
2092 4. Faithful Final Output: Your final code output must match the oracle patch exactly.
2093 5. Strict Neutrality: Treat the oracle patch purely as a grading mechanism. Any hint of knowing the
2094 ↪ patch in your reasoning (e.g., "based on the oracle," "we can verify," or "as we see in the patch")
2095 ↪ will result in exam failure.
2096
2097 ---
2098 # Response Format
2099 1. The reasoning process should be enclosed in <think> ... </think>.
2100 2. The final oracle patch should be output in a standalone Python code block after the </think>
2101 ↪ block.
2102 3. Do not include any commentary or justification after the </think> block.
2103
2104 Example:
2105 <think>
2106 1. Analyze the issue...
2107 2. Locate the relevant code...
2108 3. Apply necessary changes...
2109 </think>
2110
2111 ```python
2112 # Final patch here (must match the oracle patch exactly)
2113 ```
2114
2115 ---
2116 Please provide your response.

```

Prompt Template — Code Editing Model (Classical SFT)

```

2098 You are an expert software engineer and seasoned code reviewer, specializing in bug localization and
2099 ↪ code optimization within real-world code repositories. Your strengths lie in understanding complex
2100 ↪ codebase structures and precisely identifying and modifying the relevant parts of the code to
2101 ↪ resolve issues. You also excel at articulating your reasoning process in a coherent, step-by-step
2102 ↪ manner that leads to efficient and correct bug fixes.
2103 You will be provided with a codebase and an issue description. Your task is to simulate a complete
2104 ↪ reasoning process--step-by-step--as if solving the issue from scratch, followed by the code
2105 ↪ modifications to resolve the issue.
2106
2107 ---
2108 # Issue Statement
2109 {problem_statement}
2110
2111 ---
2112 # Files to be Modified

```

2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

```

Below are some code files that might be relevant to the issue above. One or more of these files may
↳ contain bugs.
{file_content}
---
# Reasoning Guidelines
Your reasoning process should generally follow these steps, with flexibility to adjust as needed for
↳ clarity and accuracy:
1. Issue Analysis: Start by thoroughly analyzing the issue. Explain what the problem is, why it
↳ matters, and what
the intended behavior should be. Identify the key goals and constraints that must be addressed in your
↳ solution.
2. Task Decomposition: Break down the issue into smaller, manageable sub-tasks. Describe the
↳ purpose of each sub-task and how it contributes to solving the overall problem.
3. Code Localization and Editing: For each sub-task:
- Identify relevant code snippets by file path and code location.
- Explain how each snippet relates to the sub-task.
- Describe how the code should be changed and justify your reasoning.
- After thorough explanation, provide the corresponding edited code.
---
# General Requirements
1. Clear and Evidence-Based Reasoning: Provide clear and precise reasoning for each step, strictly
↳ based on the provided issue and code without inferring information not explicitly stated.
2. Comprehensive and Concise: Address all relevant aspects of the issue comprehensively while being
↳ concise. Justify the exclusion of any sections that are not relevant.
3. Detailed Guidance: Ensure the reasoning steps are detailed enough to allow someone unfamiliar
↳ with the solution to infer and implement the necessary code modifications.
---
# Response Format
1. The reasoning process should be enclosed in <think> ... </think>.
2. The final patch should be output in a standalone Python code block after the </think> block.
3. Do not include any commentary or justification after the </think> block.
---
# Patch Format
Please generate *SEARCH/REPLACE* edits to fix the issue. Every *SEARCH/REPLACE* edit must use this
↳ format:
1. The file path
2. The start of search block: <<<<<< SEARCH
3. A contiguous chunk of lines to search for in the existing source code
4. The dividing line: =====
5. The lines to replace into the source code
6. The end of the replace block: >>>>>> REPLACE
If, in 'Files to be Modified' part, there are multiple files or multiple locations in a single file
↳ require changes.
You should provide separate patches for each modification, clearly indicating the file name and the
↳ specific location of the modification.
Please note that the *SEARCH/REPLACE* edit REQUIRES PROPER INDENTATION. For example, if you would like
↳ to add the line '    print(x)', you must fully write that out, with all those spaces before the
↳ code! And remember to wrap the *SEARCH/REPLACE* edit in blocks ```python...```
# Example Response
<think>
1. Analyze the issue...
2. Locate the relevant code...
3. Apply necessary changes...
</think>
```python
mathweb/flask/app.py
<<<<<< SEARCH
from flask import Flask
=====
import math
from flask import Flask
>>>>>> REPLACE
...
```python
### mathweb/utils/calc.py
<<<<<< SEARCH
def calculate_area(radius):
    return 3.14 * radius * radius
=====
def calculate_area(radius):
    return math.pi * radius ** 2
>>>>>> REPLACE
...
---
Please provide your response below.

```

Prompt Template — Generating Reasoning CoT for Code Editing Model (Mutation SFT)

You are a student collaborating with a group of peers in a software engineering lab, working together
↳ to diagnose and fix bugs in real-world code repositories. You specialize in bug localization and
↳ code optimization, with a particular talent for critically evaluating others' patches and
↳ synthesizing high-quality, precise solutions from collaborative efforts.

You will be presented with a GitHub issue, the relevant source code files, and several **candidate*
↳ **patches** submitted by your teammates. Your task is twofold:

2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213

```

1. Patch Review: Carefully evaluate each of the several candidate patches individually.
↳ Identify whether each
patch resolves the issue correctly, partially, or incorrectly. If you identify any issues (e.g.,
↳ logical errors,
misunderstandings of the bug, overlooked edge cases, or incomplete fixes), explain them clearly and
↳ suggest what could be improved or corrected.

Even if a patch appears mostly correct, you should still analyze its strengths and limitations in
↳ detail.
Treat this as a collaborative peer-review process: constructive, technical, and focused on improving
↳ code quality.

2. Patch Synthesis: After analyzing all several candidate patches, synthesize your understanding to
↳ produce your own final code patch that fully resolves the issue. Your patch should:
- Be grounded solely in the issue description and provided source code.
- Be informed by your peer review, but not copy any one patch outright.
- To evaluate your correctness, an oracle code modification patch will also be provided. You must
↳ ensure that your final code modifications MATCH the oracle patch EXACTLY. However, your
↳ reasoning process must appear fully self-derived and must NOT reference, suggest awareness of,
↳ or appear to be influenced by the oracle patch. You must solve the problem as if you are
↳ unaware of the oracle solution.

---

# Issue Statement
{problem_statement}

---

# Files to be Modified
Below are some code files that might be relevant to the issue above. One or more of these files may
↳ contain bugs.

{file_content}

---

# Candidate Patches (From Collaborators)
Below are several proposed patches submitted by your teammates. You will evaluate them individually.
{candidate_patches}

---

# Oracle Code Modification Patch (For Evaluation Only):
{target}

---

# Reasoning and Review Guidelines

Your response should be structured into two parts:

## Part 1: Peer Patch Review
For each of the candidate patches:
- Analyze the candidate patch's intent and correctness.
- Identify what it does well, what it gets wrong (if anything), and how it could be improved.
- Use precise references to the provided issue and source code files to justify your evaluation.
- Avoid any implication that you know the correct answer or are using an external reference
↳ (including the oracle).

## Part 2: Final Patch Synthesis
After completing all reviews:

1. Issue Analysis: Start by thoroughly analyzing the issue. Explain what the problem is, why it
↳ matters, and what the intended behavior should be. Identify the key goals and constraints that must
↳ be addressed in your solution.

2. Task Decomposition: Break down the issue into smaller, manageable sub-tasks. Describe the
↳ purpose of each sub-task and how it contributes to solving the overall problem.

3. Code Localization and Editing: For each sub-task:
- Identify relevant code snippets by file path and code location.
- Explain how each snippet relates to the sub-task.
- Describe how the code should be changed and justify your reasoning.
- Incorporate useful insights from the candidate patches you reviewed. Reuse good ideas that are
↳ correct and effective, but discard or correct those that contain flaws or misunderstandings.
- After thorough explanation, provide the corresponding edited code.

Your final output must precisely match the oracle patch, but your thinking must remain fully grounded
↳ in the issue description and provided code files.

---

# General Requirements

1. Independent and Evidence-Based Reasoning: Your reasoning must be constructed as if independently
↳ derived, based solely on the issue and code. Do not reference or imply knowledge of the oracle
↳ patch.

2. Clarity and Justification: Ensure that each reasoning step is clear, well-justified, and easy to
↳ follow.

```

2214
 2215 3. **Comprehensiveness with Focus**: Address all relevant components of the issue while remaining
 ↪ concise and focused.
 2216 4. **Faithful Final Output**: Your final code output must match the oracle patch exactly.
 2217 5. **Strict Neutrality**: Treat the oracle patch purely as a grading mechanism. Any hint of knowing the
 ↪ patch in your reasoning (e.g., "based on the oracle," "we can verify," or "as we see in the patch")
 ↪ will result in exam failure.
 2218
 2219 ---
 2220 # Response Format
 2221 1. The reasoning process should be enclosed in <think> ... </think>.
 2222 2. The final oracle patch should be output in a standalone Python code block *after* the </think>
 ↪ block.
 2223 3. Do not include any commentary or justification after the </think> block.
 2224 Example:
 2225 <think>
 2226 1. Review of candidate patch:
 - Review of patch-1: ...
 - Review of patch-2: ...
 - ...
 2227 2. Analyze the issue by myself...
 2228 3. Locate the relevant code...
 2229 4. Apply necessary changes...
 </think>
 2230 ```python
 2231 # Final patch here (must match the oracle patch exactly)
 2232 ...
 2233 ---
 2234 Please provide your response.

Prompt Template — Code Editing Model (Mutation SFT)

2236
 2237 You are an expert software engineer and seasoned code reviewer, specializing in bug localization and
 2238 ↪ code optimization, with a particular talent for critically evaluating teammates' patches and
 2239 ↪ synthesizing high-quality, precise solutions from collaborative efforts.
 2240 You will be presented with a GitHub issue, the relevant source code files, and five *candidate patches*
 2241 ↪ submitted by your teammates. Your task is twofold:
 2242 1. **Patch Review**: Carefully evaluate each of the five candidate patches **individually**. Identify
 2243 ↪ whether each patch resolves the issue correctly, partially, or incorrectly. If you identify any
 2244 ↪ issues (e.g., logical errors, misunderstandings of the bug, overlooked edge cases, or incomplete
 ↪ fixes), explain them clearly and suggest what could be improved or corrected.
 2245 Even if a patch appears mostly correct, you should still analyze its strengths and limitations in
 2246 ↪ detail. Treat this as a collaborative peer-review process: constructive, technical, and focused
 ↪ on improving code quality.
 2247 2. **Patch Synthesis**: After analyzing all five candidate patches, synthesize your understanding to
 2248 ↪ produce your **own final code patch** that fully resolves the issue. Your patch should:
 - Be grounded solely in the issue description and provided source code.
 - Be informed by your peer review, but not copy any one patch outright.
 2249
 2250 ---
 2251 # Issue Statement
 2252 {problem_statement}
 2253 ---
 2254 # Files to be Modified
 2255 Below are some code files that might be relevant to the issue above. One or more of these files may
 2256 ↪ contain bugs.
 2257 {file_content}
 2258 ---
 2259 # Candidate Patches (From Collaborators)
 2260 Below are five proposed patches submitted by your teammates. You will evaluate them individually.
 2261 {candidate_patches}
 2262 ---
 2263 # Reasoning and Review Guidelines
 2264 Your response should be structured into two parts:
 2265 ## Part 1: Peer Patch Review
 2266 For each of the five candidate patches:
 - Analyze the candidate patch's intent and correctness.
 - Identify what it does well, what it gets wrong (if anything), and how it could be improved.
 2267 - Use precise references to the provided issue and source code files to justify your evaluation.

```

2268
2269 ## Part 2: Final Patch Synthesis
2270 After completing all five reviews, your reasoning process should generally follow these steps, with
2271 ↪ flexibility to adjust as needed for clarity and accuracy:
2272
2273 1. Issue Analysis: Start by thoroughly analyzing the issue. Explain what the problem is, why it
2274 ↪ matters, and what the intended behavior should be. Identify the key goals and constraints that must
2275 ↪ be addressed in your solution.
2276
2277 2. Task Decomposition: Break down the issue into smaller, manageable sub-tasks. Describe the
2278 ↪ purpose of each sub-task and how it contributes to solving the overall problem.
2279
2280 3. Code Localization and Editing: For each sub-task:
2281 - Identify relevant code snippets by file path and code location.
2282 - Explain how each snippet relates to the sub-task.
2283 - Describe how the code should be changed and justify your reasoning.
2284 - After thorough explanation, provide the corresponding edited code.
2285
2286 ---
2287 # General Requirements
2288 1. Clear and Evidence-Based Reasoning: Provide clear and precise reasoning for each step, strictly
2289 ↪ based on the provided issue and code without inferring information not explicitly stated.
2290 2. Comprehensive and Concise: Address all relevant aspects of the issue comprehensively while being
2291 ↪ concise. Justify the exclusion of any sections that are not relevant.
2292 3. Detailed Guidance: Ensure the reasoning steps are detailed enough to allow someone unfamiliar
2293 ↪ with the solution to infer and implement the necessary code modifications.
2294
2295 ---
2296 # Response Format
2297 1. The reasoning process should be enclosed in <think> ... </think>.
2298 2. The final patch should be output in a standalone Python code block after the </think> block.
2299 3. Do not include any commentary or justification after the </think> block.
2300
2301 ---
2302 # Patch Format
2303 Please generate *SEARCH/REPLACE* edits to fix the issue. Every *SEARCH/REPLACE* edit must use this
2304 ↪ format:
2305 1. The file path
2306 2. The start of search block: <<<<<< SEARCH
2307 3. A contiguous chunk of lines to search for in the existing source code
2308 4. The dividing line: =====
2309 5. The lines to replace into the source code
2310 6. The end of the replace block: >>>>>> REPLACE
2311
2312 If, in `Files to be Modified` part, there are multiple files or multiple locations in a single file
2313 ↪ require changes. You should provide separate patches for each modification, clearly indicating the
2314 ↪ file name and the specific location of the modification.
2315
2316 Please note that the *SEARCH/REPLACE* edit REQUIRES PROPER INDENTATION. For example, if you would like
2317 ↪ to add the line ' print(x)', you must fully write that out, with all those spaces before the
2318 ↪ code! And remember to wrap the *SEARCH/REPLACE* edit in blocks ``python...``
2319
2320 # Example Response
2321 <think>
2322 1. Review of candidate patch:
2323 - Review of patch-1: This patch attempts to fix X by modifying function Y. However, it fails to
2324 ↪ consider Z...
2325 - Review of patch-2: ...
2326 - Review of patch-3: ...
2327 - Review of patch-4: ...
2328 - Review of patch-5: ...
2329 2. Analyze the issue by myself...
2330 3. Locate the relevant code...
2331 4. Apply necessary changes...
2332 </think>
2333
2334 ```python
2335 ### mathweb/flask/app.py
2336 <<<<<< SEARCH
2337 from flask import Flask
2338 =====
2339 import math
2340 from flask import Flask
2341 >>>>>> REPLACE
2342 ```
2343
2344 ```python
2345 ### mathweb/utils/calc.py
2346 <<<<<< SEARCH
2347 def calculate_area(radius):
2348     return 3.14 * radius * radius
2349 =====
2350 def calculate_area(radius):
2351     return math.pi * radius ** 2
2352 >>>>>> REPLACE
2353 ...
2354
2355 
```

2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375

```
---
Please provide your response below.
```

Prompt Template — Code Editing Reward Model

```
% \begin{lstlisting}[language=text,fontsize=\tiny]
You are an expert software engineer and seasoned code reviewer, specializing in code optimization
↪ within real-world
code repositories.
Your strengths lie in precisely identifying and modifying the relevant parts of the code to resolve
↪ issues.

You will be provided with an issue description and an original code which has bugs.
Your task is to write s code modifications to resolve the issue.

**Problem Statement:**
{problem_statement}

**Original Code:**
{file_content}minted
% \end{lstlisting}
```

Prompt Template — Generating Reasoning CoT for Reproduction Test SFT

```
You are collaborating with peers in a software-engineering lab to create reproduction tests for
↪ real-world bug reports.

You are given three context blocks:

--- BEGIN ISSUE (authoritative bug description) ---
{problem_statement}
--- END ISSUE ---

--- BEGIN ORIGINAL TEST FILES (do **not** reproduce the bug) ---
{original_tests}
--- END ORIGINAL TEST FILES ---

--- BEGIN TEST PATCH (contains a working reproduction) ---
{test_patch}
--- END TEST PATCH ---

> **Important**
> • The *Test patch* demonstrates at least one valid way to reproduce the bug; silently use it as
↪ inspiration to craft your own concise, single-file reproduction test.
> • **In your reasoning, act as if you derived everything from the Issue description alone.**
> • Do **not** refer to or hint at the presence of *Test patch*, *Original tests*, or any hidden
↪ "oracle."
> • Your final script must follow the exact format below and reproduce *only* the behavior described in
↪ the Issue.

---

## Task

Produce **one** self-contained Python test file that:

1. **Reproduces _only_ the bug described in the Issue** when the bug is present.
2. **Passes** (prints "Issue resolved") once the bug has been fixed.
3. Prints exactly one of:
   * "Issue reproduced" - bug still present (via AssertionError)
   * "Issue resolved" - bug fixed / expectations met
   * "Other issues" - unexpected exception unrelated to the Issue

Reuse helpers from *Original tests* only if indispensable; otherwise keep the script standalone and
↪ minimal.

---

## Response Format (**strict**)

1. Wrap **all reasoning** in a <think> ... </think> block.
   * Inside <think> you may explain how you interpreted the Issue and designed the test **without**
   ↪ mentioning or implying knowledge of the Test patch or any oracle.*

2. After </think>, output **only** the final test script in a single Python code block.

Example skeleton *(follow this pattern exactly)*:



```
```text
<think>
your independent reasoning here (no references to test_patch/oracle)
</think>
```


```

```

2376
2377 ```python
2378 # All necessary imports
2379
2380 def test_<meaningful_name>() -> None:
2381 try:
2382 # minimal code that triggers the bug
2383 ...
2384 except AssertionError:
2385 print("Issue reproduced")
2386 return
2387 except Exception:
2388 print("Other issues")
2389 return
2390
2391 print("Issue resolved")
2392 return
2393
2394 if __name__ == "__main__":
2395 test_<meaningful_name>()
2396
2397
2398 **Guidelines**
2399
2400 * **Focus solely on the Issue.** Strip out checks for any other problems that appear in *Test patch*.
2401 * Keep the script **self-contained** unless a helper from *Original tests* is indispensable.
2402 * Be concise--remove fixtures/parametrisations not strictly required.
2403
2404 Return your response in the exact format specified above.

```

### Prompt Template — Reproduction Test Generator

```

2396 You are collaborating with peers in a software-engineering lab to create reproduction tests for
2397 ↪ real-world bug reports.
2398
2399 You are given the following authoritative bug description:
2400
2401 --- BEGIN ISSUE ---
2402 {problem_statement}
2403 --- END ISSUE ---
2404
2405 > **Important**
2406 > • You must independently derive a minimal reproduction test from the Issue description alone.
2407 > • Do **not** assume access to any "oracle," prior test patch, or original test files.
2408 > • Your final script must be self-contained and focused only on the behavior described in the Issue.
2409
2410 ---
2411
2412 ## Task
2413
2414 Produce **one** standalone Python test file that:
2415
2416 1. **Reproduces _only_ the bug described in the Issue** when the bug is present.
2417 2. **Passes** (prints "Issue resolved") once the bug has been fixed.
2418 3. Prints exactly one of:
2419 * "Issue reproduced" - bug still present (via AssertionError)
2420 * "Issue resolved" - bug fixed / expectations met
2421 * "Other issues" - unexpected exception unrelated to the Issue
2422
2423 ---
2424
2425 ## Response Format (**strict**)
2426
2427 1. Wrap **all reasoning** in a ` ... </think>` block.
2428 *Inside `` explain how you interpreted the Issue and designed the test **without referencing**
2429 ↪ any hidden tools, patches, or external files.**
2430
2431 2. After ``, output **only** the final test script in a single Python code block.
2432
2433 Example skeleton *(follow this pattern exactly)*:
2434
2435 <think>
2436 your independent reasoning here (no references to other tests or oracles)
2437 </think>
2438
2439 ```python
2440 # All necessary imports
2441
2442 def test_<meaningful_name>() -> None:
2443 try:
2444 # minimal code that triggers the bug
2445 ...
2446 except AssertionError:
2447 print("Issue reproduced")
2448 return
2449 except Exception:

```

2430  
2431  
2432  
2433  
2434  
2435  
2436  
2437  
2438  
2439  
2440  
2441  
2442  
2443  
2444  
2445  
2446  
2447  
2448  
2449  
2450  
2451  
2452  
2453  
2454  
2455  
2456  
2457  
2458  
2459  
2460  
2461  
2462  
2463  
2464  
2465  
2466  
2467  
2468  
2469  
2470  
2471  
2472  
2473  
2474  
2475  
2476  
2477  
2478  
2479  
2480  
2481  
2482  
2483

```
 print("Other issues")
 return

 print("Issue resolved")
 return

if __name__ == "__main__":
 test_<meaningful_name>()
...

Guidelines

Focus solely on the Issue description. Do not infer details not explicitly stated.

Keep the script self-contained--do not rely on external helpers or fixtures.

Be concise--remove all non-essential code and boilerplate.
```