
SWE-RL: Advancing LLM Reasoning via Reinforcement Learning on Open Software Evolution

Yuxiang Wei¹² Olivier Duchenne¹ Jade Copet¹ Quentin Carbonneaux¹
Lingming Zhang² Daniel Fried¹³ Gabriel Synnaeve¹ Rishabh Singh¹ Sida I. Wang¹

¹Meta AI ²University of Illinois Urbana-Champaign ³Carnegie Mellon University

✉ ywei40@illinois.edu ✉ sida@meta.com

<https://github.com/facebookresearch/swe-rl>

Abstract

The recent DeepSeek-R1 release has demonstrated the immense potential of reinforcement learning (RL) in enhancing the general reasoning capabilities of large language models (LLMs). While DeepSeek-R1 and other follow-up work primarily focus on applying RL to competitive coding and math problems, this paper introduces SWE-RL, the first approach to scale RL-based LLM reasoning for real-world software engineering. Leveraging a lightweight rule-based reward (e.g., the similarity score between ground-truth and LLM-generated solutions), SWE-RL enables LLMs to autonomously recover a developer’s reasoning processes and solutions by learning from extensive open-source software evolution data—the record of entire software development cycles, including code snapshots, code changes, and events such as issues and pull requests. Trained on top of Llama 3, our resulting reasoning model, Llama3-SWE-RL-70B, achieves a **41.0%** solve rate on SWE-bench Verified—a human-verified collection of real-world GitHub issues. To our knowledge, this is the best performance reported for medium-sized (<100B) LLMs to date, even comparable to leading proprietary LLMs like GPT-4o. Surprisingly, despite performing RL solely on software evolution data, Llama3-SWE-RL has even emerged with generalized reasoning skills. For example, it shows improved results on five out-of-domain tasks, namely, function coding, library use, code reasoning, mathematics, and general language understanding, whereas a supervised-finetuning baseline even leads to performance degradation on average. Overall, SWE-RL opens up a new direction to improve the reasoning capabilities of LLMs through reinforcement learning on massive software engineering data.

1 Introduction

The application of large language models (LLMs) to software engineering (SE) tasks has received significant attention, with researchers exploring their potential to automate various complex SE tasks, such as library-level and complex code generation (Zhuo et al., 2024; Zhao et al., 2024), real-world bug/issue resolution (Xia and Zhang, 2022; Jimenez et al., 2023; Yang et al., 2024c), and software testing (Deng et al., 2023; Jain et al., 2024a). Among these tasks, SWE-bench (Jimenez et al., 2023)—a benchmark for solving real-world software issues—has emerged as a focal point of research efforts, and researchers have proposed various agentic (Yang et al., 2024b; Wang et al., 2024b; Gauthier, 2024) and pipeline-based (Xia et al., 2024; Örwall, 2024; Zhang et al., 2024) methods to push LLMs’ real-world issue solving capability. However, most current techniques rely on powerful proprietary LLMs like GPT-4o (OpenAI, 2024a) or Claude-3.5-Sonnet (Anthropic, 2024a), where advancements are driven more by enhanced prompting strategies than by improvements in the underlying LLM.

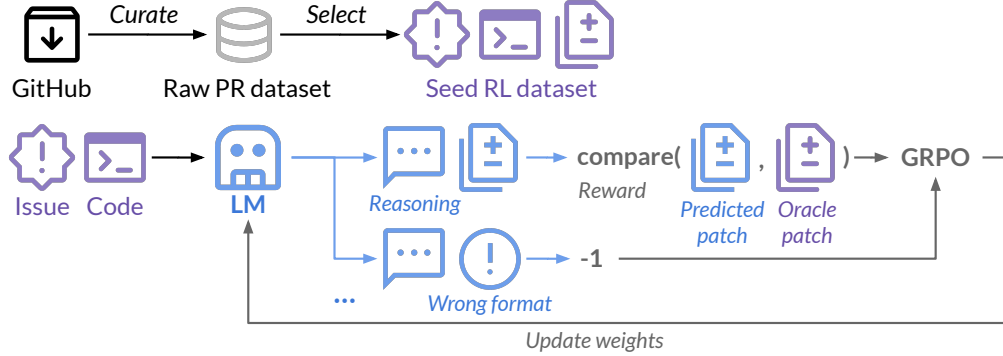


Figure 1: **Overview of SWE-RL.** We create a seed RL dataset from GitHub PRs, including issue descriptions, code context, and oracle patches. A policy LLM generates code edits via reasoning. Rewards are based on the similarity to oracle patches, with penalties for formatting errors. We optimize the policy using GRPO.

With the release of DeepSeek-R1 (DeepSeek-AI, 2025), reinforcement learning (RL) using rule-based rewards has become a crucial technique for enhancing the reasoning capabilities of LLMs across various downstream tasks, including coding (Zeng et al., 2025) and mathematics (Yeo et al., 2025). However, their effectiveness in SE tasks remains limited (DeepSeek-AI, 2025), and their substantial total parameter size (671B) poses challenges for researchers attempting to train them. For mathematics, the reward is generally defined as whether the answer predicted by the LLM can exactly match the ground truth (DeepSeek-AI, 2025). While in coding, existing RL research typically utilizes execution feedback (DeepSeek-AI, 2025; Gehring et al., 2025) as the reward signal and is limited to competitive programming tasks (Li et al., 2022; Jain et al., 2024b), where code is self-contained and easily executable. This is challenging to apply to real-world SE tasks due to the execution cost and lack of executable environments (Pan et al., 2024). Meanwhile, previous research (Ma et al., 2024; Pan et al., 2024; Xie et al., 2025) relied on proprietary teacher models and focused primarily on supervised fine-tuning (SFT), which, as we show in the paper, is less effective and less generalizable.

To address these limitations, we propose SWE-RL, the first RL method to improve LLMs on SE tasks by directly using rule-based rewards and software evolution data—the record of entire software lifecycle, including all code snapshots, changes, and events like PRs and issues. As shown in Figure 1, we begin by curating a comprehensive dataset of GitHub pull requests (PRs), which is then transformed into the seed dataset for RL. Each data item includes an issue, the corresponding code context, and the oracle patch merged by the PR. During RL, the policy LLM is tasked with solving a given issue through reasoning and producing the code changes. The code changes are then converted into a consistent patch format for reward calculation. If the response is incorrectly formatted, the reward will be -1 ; otherwise, the reward is a similarity score (between 0 and 1) of the predicted and the oracle patch calculated by Python’s `difflib.SequenceMatcher` (Ratcliff and Metzener, 1988). Notably, we provide the complete content of each file in the input prompt, which implicitly teaches the model to reason about the precise fault locations before suggesting repair edits.

Applying SWE-RL to Llama-3.3-70B-Instruct (Dubey et al., 2024), our model Llama3-SWE-RL-70B solves **41.0%** of the issues in SWE-bench Verified (OpenAI, 2024), a human-verified subset of SWE-bench, with Agentless Mini, our pipeline-based scaffold built upon Agentless (Xia et al., 2024), featuring simplifications to match our RL process and enhancements for scaling. This performance is comparable to leading proprietary LLMs like GPT-4o (OpenAI, 2024a) and state-of-the-art among medium-sized LLMs with less than 100B total parameters. Our ablation studies demonstrate that Llama3-SWE-RL-70B significantly outperforms its Llama baseline. Additionally, we developed a competitive supervised fine-tuning (SFT) model from Llama-3.3-70B-Instruct using synthetic data generated in the Magicoder (Wei et al., 2024) style to enhance the chain-of-thought (Wei et al., 2022) process, employing the same seed as SWE-RL. We show that Llama3-SWE-RL-70B, trained with SWE-RL solely for solving issues, not only surpasses the SFT model in SWE-bench but also excels in other out-of-domain (OOD) tasks, including function-level coding (Chen et al., 2021; Liu et al., 2023), practical code generation with library use (Zhuo et al., 2024), code reasoning (Gu et al., 2024), mathematics (Hendrycks et al., 2021c), and general language understanding (Hendrycks et al., 2021b).

In these OOD tasks, Llama3-SWE-RL-70B even outperforms Llama-3.3-70B-Instruct, whereas the SFT model results in decreased performance.

In summary, our contributions are as follows:

- We introduce SWE-RL, the first RL approach specifically designed to enhance LLMs for SE tasks using software evolution data (e.g., PRs) and rule-based rewards.
- We develop Llama3-SWE-RL-70B, trained with SWE-RL on Llama-3.3-70B-Instruct. It achieves **41.0%** on SWE-bench Verified, the best performance among medium-sized language models (<100B) and even comparable to leading proprietary models like GPT-4o.
- We show for the first time that applying RL solely to real-world SE tasks, such as issue solving, can already enhance an LLM’s general reasoning abilities, enabling it to improve on out-of-domain tasks like math, code generation, and general language understanding.

2 SWE-RL

To prepare the initial dataset for RL, we extract 273k high-quality PR seeds from the raw PR dataset we collected from GitHub. These seeds are selected based on specific heuristics. For example, a PR instance should include at least one linked issue, the issue should describe a bug-fixing request, and the code changes should involve programming files. For each seed, we extract the issue descriptions and code context, including all changed files and some relevant but unchanged files. These are converted to input prompts for the policy LLM. We also take the oracle patch from each seed, which will be used for reward calculation. More details are explained in Appendix A.

2.1 Reward modeling

[System] A user will ask you to solve a task. You should first draft your thinking process (inner monologue). Then, generate the solution. Your response format must follow the template below:

<think>
Your thoughts or/and draft, like working through an exercise on scratch paper. Be as casual and as long as you want until you are confident to generate a correct solution.
</think>

<solution>
Final solution presented to the user.
</solution>

[User] ...here is the issue text: {issue}
 ...here is the code context: {context}
 ...generate *search/replace* edits to fix the issue...

[Assistant]

Figure 2: **Prompt template used to train Llama3-SWE-RL with SWE-RL.** Given an issue description and the corresponding code context, the policy LLM needs to generate search/replace edits (Xia et al., 2024) to fix this issue through reasoning. This is the only subtask we incorporate in the RL training. During inference, the LLM can generalize to tasks outside the training domain (e.g, file-level localization and test generation). For conciseness, we exclude certain prompt details, with the complete prompt template available in Appendix D.

We bootstrap the policy LLM with the prompt template shown in Figure 2. Assuming that the LLM has generated a rollout o given an issue and its code context, the reward function is defined as follows:

$$\mathcal{R}(o) = \begin{cases} -1, & \text{if } o \text{ has wrong format,} \\ \text{compare}(\text{patch}_{\text{pred}}, \text{patch}_{\text{gt}}), & \text{otherwise.} \end{cases} \quad (1)$$

Here, $\text{patch}_{\text{pred}}$ is the patch extracted from the LLM generation o if it is correctly formatted, and patch_{gt} means the oracle patch for this issue.

In the implementation, we use Python’s `difflib.SequenceMatcher` as the *compare* function, which returns a floating point between 0 and 1 indicating the sequence similarity, and adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024) for policy optimization.

Given a seed RL dataset $\mathcal{D}_{\text{seed}}$ where each item contains (1) issue, the issue description, (2) ctx, the code context required to solve this issue, including both files to repair and relevant files not intended for editing, and (3) `patchgt`, which is the oracle patch. The input prompt for each data item is formed by instantiating the prompt template in Figure 2 with the issue description and code context, which we denote as $q = \text{form-prompt}(\text{issue}, \text{ctx})$. The policy LLM π_θ tries to solve the issue by generating code changes through reasoning, and produces multiple outputs o_i for each input prompt q given the group size G . Then, the policy LLM aims to maximize the following GRPO objective:

$$\mathcal{J}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \left(\min \left(\frac{\pi_\theta(o_i | q)}{\pi_{\theta_{\text{old}}}(o_i | q)} A_i, \text{clip} \left(\frac{\pi_\theta(o_i | q)}{\pi_{\theta_{\text{old}}}(o_i | q)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right) \right],$$

where $(\text{issue}, \text{ctx}, \text{patch}_{\text{gt}}) \sim \mathcal{D}_{\text{seed}}$, $q = \text{form-prompt}(\text{issue}, \text{ctx})$, and $\{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)$. (2)

ϵ and β are hyperparameters, and $\pi_{\theta_{\text{old}}}$ and π_{ref} are the old and reference policy. The rewards $r_i = \mathcal{R}(o_i)$ and advantages $A_i = \frac{r_i - \text{mean}(r_1, \dots, r_G)}{\text{std}(r_1, \dots, r_G)}$ are calculated using the normalized rewards within each group following GRPO. D_{KL} denotes the estimated KL-divergence (Schulman, 2020).

Our training approach conditions the model on the complete context of each file, implicitly forcing it to identify detailed fault locations before generating repair edits. This process inherently teaches the model both bug diagnosis and repair generation. However, during the evaluation on SWE-bench, Agentless Mini requires capabilities beyond generating repair edits, such as file-level fault localization, reproduction test generation, and regression test selection. Remarkably, even without explicit training on these subtasks, Llama3-SWE-RL can generalize to them through the RL process.

2.2 Aha moments and generalized reasoning capabilities

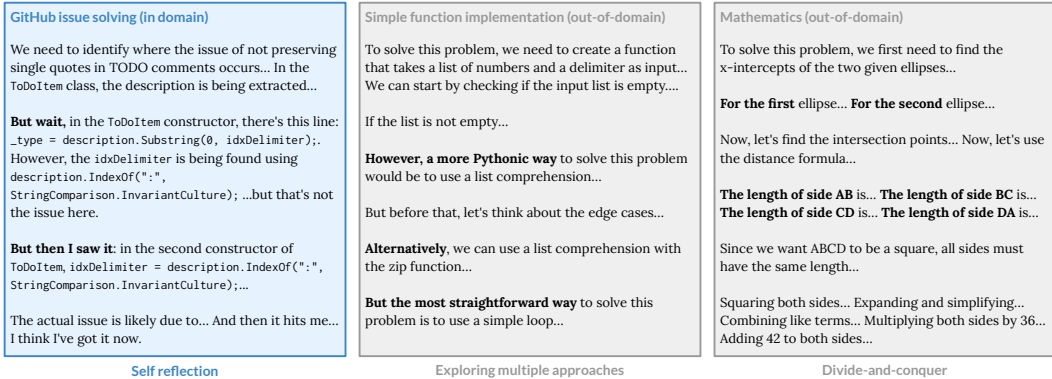


Figure 3: **Reasoning skills emerged from Llama3-SWE-RL-70B following the application of SWE-RL.** RL helps the model develop reasoning skills like self-reflection, exploring alternatives, and divide-and-conquer strategies, for both in-domain (e.g., issue-solving) and out-of-domain tasks (e.g., function implementation and mathematics).

“Aha moments” on software engineering. With the application of SWE-RL, we observe “aha moments” (DeepSeek-AI, 2025) where Llama3-SWE-RL exhibits emergent reasoning skills. To our knowledge, this is the first study demonstrating the existence of such aha moments in the realm of real-world software engineering tasks, confirming the findings of DeepSeek-R1 (DeepSeek-AI, 2025), which mainly focuses on competition coding and math. As shown in Figure 3, through RL, Llama3-SWE-RL-70B can allocate more thinking time to reflect on its initial assumptions during the issue-solving process. This behavior emerges naturally from the model’s interaction with RL, rather than being explicitly programmed.

General reasoning capabilities. Surprisingly, we have identified additional aha moments where Llama3-SWE-RL acquires general reasoning abilities that are transferrable to various out-of-domain

tasks, such as function-level code generation and mathematics, although RL is applied exclusively to software issue solving. Figure 3 demonstrates that Llama3-SWE-RL is capable of reasoning through self-reflection, exploring alternative approaches, and solving complex problems by breaking them down into smaller subtasks. In §3.5, we demonstrate that Llama3-SWE-RL improves over even more out-of-domain tasks, including library use, code reasoning, and general language understanding.

3 Evaluation

3.1 Experimental setup

Training configs. Llama3-SWE-RL-70B is trained on top of Llama-3.3-70B-Instruct (Dubey et al., 2024) using SWE-RL for 1,600 steps with a 16k context window. We use a global batch size of 512, sampling 16 rollouts from each of the 32 problems in every batch. For every global step, a single optimization step is performed using Adam (Kingma and Ba, 2017). We train our models on 512 NVIDIA H100 GPUs; a training run takes approximately 32 wall-time hours.

Scaffolding. We have developed Agentless Mini on top of Agentless (Xia et al., 2024) as the underlying scaffold. Different from Agentless’s multi-step localization, Agentless Mini focuses solely on file-level localization, delegating detailed reasoning to the repair step by providing the entire file contents in the input. This enables Llama3-SWE-RL to do more reasoning during SWE-RL and simplifies the RL process to focus on only one issue-solving task. Despite this simplification, Llama3-SWE-RL can still seamlessly complete other pipeline steps, benefiting from out-of-domain generalization through RL. Furthermore, while Agentless only employs one reproduction test per issue for reranking, Agentless Mini can use multiple reproduction tests, which has proven effective in our scaling analysis (§3.4). More details about Agentless Mini can be found in Appendix B.

Evaluation setup. We conduct evaluation on SWE-bench Verified (OpenAI, 2024), a subset of SWE-bench with 500 human-verified problems that can more reliably evaluate AI models’ capability in solving real-world software issues. In the main evaluation (§3.2), we generate 500 patches for each problem using a 1.0 temperature, and use the top 30 reproduction tests for execution and reranking. Ultimately, only the highest-ranked patch for each issue will be submitted for SWE-bench evaluation to calculate the pass@1 score.

SFT baseline. To understand the advantages of SWE-RL, we also trained an SFT baseline, named Llama3-SWE-SFT-70B, for experiments in §3.3, §3.4, and §3.5. It is trained on top of Llama-3.3-70B-Instruct (Dubey et al., 2024) using a mixture of synthetic code editing data, Llama 3 (Dubey et al., 2024) coding SFT data, and Llama 3 general SFT data. The synthetic data is generated using an approach inspired by Magicoder (Wei et al., 2024), where high-quality PR data serves as the seeds for creating chain-of-thoughts and subsequent editing, as well as serving as the oracle for filtering. Llama-3.3-70B-Instruct is employed for this generation process. In contrast to our RL model, which only requires a seed PR dataset to trigger the RL loop, the SFT baseline needs synthetic data generation for chain-of-thoughts and additional data mix to ensure the dataset diversity and model generalizability. More details are explained in Appendix C.

3.2 Main results

Table 1 presents the pass@1 results on SWE-bench Verified for models that utilize open-source scaffolds, with models categorized by their size. From the table, we observe that Llama3-SWE-RL-70B achieves state-of-the-art results among small and medium-sized language models (<100B) by resolving 41.0% of the issues. Additionally, all other open-source baselines we compare, such as Lingma-SWE-GPT (Ma et al., 2024), SWE-Gym (Pan et al., 2024), and SWE-Fixer (Xie et al., 2025), include distilled outputs from GPT-4o or Claude-3.5-Sonnet in the training data. In contrast, Llama3-SWE-RL is trained solely with publicly available data through our reinforcement learning technique SWE-RL, without relying on any proprietary LLMs in the pipeline. Llama3-SWE-RL also sets a new record for Llama-based methods on SWE-bench. Additionally, we show that our RL-ed model is significantly better than the SFT baseline in this end-to-end setup, indicating that SWE-RL enhances the general reasoning ability of LLMs to solve issues.

Table 1: **Main results on SWE-bench Verified.** We include representative methods with open-source scaffolds. The scores are either collected from the SWE-bench Leaderboard (Jimenez et al., 2024) or from the corresponding reference.

Model	Scaffold	SWE-bench Verified	Reference
Model closed-source or size \gg 100B			
GPT-4o	SWE-agent	23.2	Yang et al. (2024b)
Claude-3.5-Sonnet	SWE-agent	33.6	Yang et al. (2024b)
GPT-4o	Agentless	38.8	Xia et al. (2024)
o1-preview	Agentless	41.3	OpenAI (2024b)
DeepSeek-V3 ¹	Agentless	42.0	DeepSeek-AI (2024)
Claude-3.5-Sonnet	AutoCodeRover-v2.0	46.2	Zhang et al. (2024)
Claude-3.5-Sonnet	Tools	49.0	Anthropic (2024b)
DeepSeek-R1 ¹	Agentless	49.2	DeepSeek-AI (2025)
Claude-3.5-Sonnet	Agentless	50.8	Xia et al. (2024)
Claude-3.5-Sonnet	OpenHands	53.0	Wang et al. (2024b)
Model size \leq 100B			
SWE-Llama-13B	RAG	1.2	Jimenez et al. (2023)
SWE-Llama-7B	RAG	1.4	Jimenez et al. (2023)
Lingma-SWE-GPT-7B	SWE-SynInfer	18.2	Ma et al. (2024)
Lingma-SWE-GPT-72B	SWE-SynInfer	28.8	Ma et al. (2024)
SWE-Gym-32B	OpenHands	32.0	Pan et al. (2024)
SWE-Fixer-72B	SWE-Fixer	32.8	Xie et al. (2025)
Llama3-SWE-SFT-70B	Agentless Mini	36.2	This paper
Llama3-SWE-RL-70B	Agentless Mini	41.0	This paper

¹Open-source Mixture-of-Experts model with 671B total and 37B active parameters

Table 2: Baseline comparison on SWE-bench Verified. In this experiment, we compare the repair-only performance of baseline LLMs by providing oracle localized files in the input context, without doing test generation and execution. We use greedy decoding by default, but for Llama-3.3-70B-Instruct, we include a 20-sample majority voting result at a temperature of 0.6 to improve formatting accuracy.

Model	Setting	Correct format	Repair performance (oracle)
Llama-3.3-70B-Instruct	Greedy decoding	12.2%	5.4
Llama-3.3-70B-Instruct	Majority voting	44.6%	16.6
Llama3-SWE-SFT-70B	Greedy decoding	96.2%	29.6
Llama3-SWE-RL-70B	Greedy decoding	95.6%	34.8

3.3 Baseline comparison

To understand how much SWE-RL improves LLMs in solving software issues, we compare Llama3-SWE-RL with the corresponding Llama-3 and SFT baseline in Table 2, using Agentless Mini as the underlying scaffold. In this experiment, we also evaluate on SWE-bench Verified but focus on the models’ repair ability. To achieve this, we provide oracle files in the context and let the model generate a single repair edit using greedy decoding, without incorporating additional pipeline steps such as localization and test generation. The table reveals that the base Llama-3.3 model struggles to produce correctly formatted code edits, even when using a 20-sample majority voting approach, where outputs with incorrect formats are pre-filtered. With SFT, most code edits generated by the language model are correctly formatted, and the repair performance shows significant improvement. However, Llama3-SWE-RL-70B demonstrates even greater enhancement in repair capabilities, although its format accuracy is slightly lower than that of the SFT version. This indicates that SWE-RL aids the LLM in better reasoning about issue solving and code editing.

3.4 Scaling analysis with more samples

Agentless Mini supports scaling both the number of repair samples and the number of generated reproduction tests. The difference in the number of samples may affect the reranking accuracy. In this section, we evaluate how the final pass@1 performance on SWE-bench Verified scales with the two factors.




Figure 4: **Scaling analysis with more repair samples and more reproduction tests.** The figure on the left illustrates the resolve rate on SWE-bench Verified in relation to the number of repair samples, while maintaining a constant 30 test samples. Conversely, the figure on the right depicts the resolution rate as it varies with the number of reproduction test samples, with a fixed 500 repair samples.

Figure 4 shows that increasing both the number of repair samples and test samples enhances performance on SWE-bench. Notably, for repair samples, there is a significant score increase from 33.6 to 40.0 when the sample size is expanded from 20 to 160. However, beyond 160 samples, the improvement trend begins to plateau, with scores only rising slightly from 40.0 to 41.0 as the sample size increases to 320 and 500. Although the impact of adding more reproduction test samples is less obvious, there is still a gradual score improvement from 38.8 to 41.0 as the number of test samples increases up to 20. There is no difference between 20 and 30 test samples, suggesting a performance saturation point has been reached.

3.5 Generalizability of RL

Table 3: **Generalizability of Llama3-SWE-RL-70B beyond SWE-bench.** This table compares Llama-3.3-70B-Instruct, the SFT variant, and the RL model on five out-of-domain tasks, highlighting RL improvements and SFT declines. All experiments are done in a consistent setting using zero-shot greedy decoding. We report the macro average over category accuracy for MMLU and pass@1 for the others. In MATH, we use simple-evals’s “Answer: . . .” prompt format (OpenAI, 2024). However, only the RL model consistently follows the format requirements, so we also report MATH (lenient) to relax the constraint to include “\boxed. . .”.

Category	Llama-3.3-70B-Instruct	Llama3-SWE-SFT-70B	Llama3-SWE-RL-70B
Function coding			
HumanEval+	76.2	73.2	79.9
Library use			
BigCodeBench-Hard (I)	28.4	25.7	28.4
BigCodeBench-Hard (C)	29.1	24.3	29.1
Code reasoning			
CRUXEval-I	60.5	68.4	71.6
CRUXEval-O	61.9	75.1	75.5
Math			
MATH (strict)	63.2	54.0	73.7
MATH (lenient)	70.9	71.7	73.7
General			
MMLU	86.49	85.26	86.82

Llama3-SWE-RL is only trained with SWE-RL on issue-solving data. This raises a natural question whether such domain-specific training harms the performance on other tasks. To address this, we conduct an experiment in Table 3, evaluating the LLMs on five out-of-domain benchmarks, i.e., HumanEval+ (Chen et al., 2021; Liu et al., 2023) for function-level code generation, BigCodeBench (Zhuo et al., 2024) for practical code generation with library use, CRUXEval (Gu et al., 2024) for code execution reasoning, MATH (Hendrycks et al., 2021c) for mathematical reasoning, and MMLU (Hendrycks et al., 2021b) for general language understanding. We also include the SFT baseline, which is finetuned on the same Llama-3.3-70B-Instruct model using issue-solving data, combined with general coding and dialog data.

From the table, it is evident that Llama3-SWE-RL-70B, trained with RL, outperforms both its base model and the SFT baseline. There are notable improvements in CRUXEval and MATH, where significant reasoning efforts are required to arrive at the final answer. Through SWE-RL, the model enhances its reasoning skills and dedicates more thinking effort to solving problems compared to other baselines. Although trained on a single task, SWE-RL enables the model to generalize its reasoning capabilities across various domains. In contrast, the SFT version, on average, underperforms relative to the original model.

Overall, our results suggest for the first time that reinforcement learning on real-world software data like PRs enables the model to acquire generalized reasoning skills, whereas supervised finetuning steers the language model towards a specific task distribution, leading to performance declines on tasks with lower emphasis, even when a meticulously curated data mix is used.

Statistical significance analysis. While small absolute gains (1–2 percentage points) on HumanEval or CRUXEval are, by themselves, unlikely to be statistically significant, we evaluate across multiple benchmarks and aggregate evidence, observing improvements that consistently favor RL. With Eval Arena (Wang et al., 2024a), we indicate that improvements of > 0.8 percentage points on MMLU, 3 points on CRUXEval, and > 3 points on the full MATH dataset are significant on their own; taken together, these results achieve significance at the 0.05 level.

3.6 Reward ablation

According to Equation (1), the reward design of SWE-RL allows different instantiations of the *compare* function. Throughout the paper, we adopt the sequence similarity between the predicted and the oracle patch, which is a continuous value from 0 to 1. We denote this type of reward as continuous. It is natural to compare this continuous reward with a discrete reward, where the *compare* function outputs 1 if the predicted and oracle patches exactly match each other, and 0 otherwise. We trained a variant of Llama3-SWE-RL with the discrete reward function, using the same training setup as in the continuous reward case.

Reward type	Correct format	Repair (oracle)
Discrete	94.2%	29.0
Continuous	95.6%	34.8




Figure 5: **Ablation on SWE-RL’s reward functions and their training dynamics.** We compare SWE-RL using the default continuous reward function against a discrete reward. Repair (oracle) evaluates the repair-only performance using greedy decoding, with oracle files in the input context.

As shown in Figure 5, while the discrete and continuous reward functions lead to similar format accuracy, the continuous reward is more effective in enhancing the repair performance. From the training dynamics, we can see that discrete rewards grow slower than continuous rewards. Additionally, the average discrete reward remains approximately zero upon the completion of training, meaning it struggles to obtain patches exactly matching the oracles. This is because real-world patches are highly diverse and often cannot be easily matched. The continuous reward function

better captures partial correctness and incremental improvements, allowing the model to learn more nuanced and effective repair strategies.

4 Related work

4.1 Language models for software engineering

Large language models (LLMs), trained with billions to trillions of code tokens, have demonstrated outstanding performance in a wide range of coding tasks, including code generation (Chen et al., 2021; Li et al., 2023; Rozière et al., 2023; Guo et al., 2024; Wei et al., 2024; Lozhkov et al., 2024; DeepSeek-AI et al., 2024), code optimization (Cummins et al., 2024; Liu et al., 2024a), program repair (Xia and Zhang, 2022; Xia et al., 2023b; Jiang et al., 2023; Wei et al., 2023), and software testing (Xia et al., 2023a; Deng et al., 2023; Yuan et al., 2023; Schäfer et al., 2023; Lemieux et al., 2023). Initially, researchers primarily focused on single-shot code generation tasks, such as function-level (Chen et al., 2021; Austin et al., 2021; Hendrycks et al., 2021a; Li et al., 2022; Jain et al., 2024b), class-level (Du et al., 2023), and repository-level code completion (Ding et al., 2023; Liu et al., 2024b; Zhang et al., 2023). However, with the rapid development of LLMs, the performance on many popular single-shot code generation benchmarks like HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), and EvalPlus (Liu et al., 2023) has become saturated. Since the development of SWE-bench (Jimenez et al., 2023), which requires solving real-world GitHub issues, researchers start to work on improving LLMs’ real-world issue-solving capability and have designed various scaffolds for SWE-bench. Two general types are (1) agentic scaffolds (Yang et al., 2024b; Wang et al., 2024b; Gauthier, 2024), where an LLM drives the decision-making process based on its past actions and observations through tool-based interaction with the environment; and (2) pipeline-based scaffolds (Xia et al., 2024; Örwall, 2024; Zhang et al., 2024), where an LLM goes through human-defined stages to solve a given issue. Generally, agentic methods are more general but require strong instruction-following and capable LLMs to drive the autonomous process, and can be computationally intensive due to multi-round interactions. In contrast, pipeline-based approaches are more specialized but efficient, with a focus on LLMs’ pure code editing capability. Therefore, we designed our minimalist pipeline-based scaffold, Agentless Mini, to focus on the enhancements of Llama3-SWE-RL’s core code editing capability.

4.2 Training software agents

While existing scaffolds have successfully leveraged proprietary language models to tackle real-world software engineering tasks, open models typically yield subpar results in these settings. Moreover, the most effective approach to enhancing real-world software engineering capabilities through training remains unclear. Recently, researchers have begun exploring the possibility of training open LLMs specifically for software engineering tasks, aiming to improve performance on benchmarks such as SWE-bench. For instance, Lingma-SWE-GPT (Ma et al., 2024) introduces 7B and 72B model variants that build on top of Qwen2.5-Coder-7B (Hui et al., 2024) and Qwen2.5-72B-Instruct (Yang et al., 2024a), using an iterative development-process-centric approach. SWE-Gym (Pan et al., 2024) presents the first open training environment for software engineering agents, significantly improving the performance of Qwen2.5-Coder’s 7B and 32B variants on SWE-bench. More recently, SWE-Fixer (Xie et al., 2025) finetunes the Qwen2.5 base series, resulting in a 7B code retriever and a 72B code editor focused on efficient issue resolution, achieving notable best@1 improvements. Notably, all these works incorporate distilled samples from either GPT-4o or Claude-3.5-Sonnet in their training data and are built upon Qwen2.5 models. Their training objectives are all based on supervised finetuning. On the contrary, Llama3-SWE-RL is based on Llama 3 (Dubey et al., 2024) and trained through reinforcement learning (RL) using SWE-RL. The seed dataset for RL is sourced exclusively from publicly available repositories, allowing Llama3-SWE-RL to self-improve its issue-solving capabilities through the RL incentive. Remarkably, Llama3-SWE-RL achieves the best performance among these models with a 41.0% solve rate on SWE-bench Verified (OpenAI, 2024), demonstrating for the first time that LLMs can already effectively address real-world issues through RL on real-world software artifacts.

5 Conclusion

We introduce SWE-RL, the first reinforcement learning (RL) approach to improve language models (LLMs) on software engineering tasks using software evolution data (e.g., PRs) and rule-based rewards. The resulting model, Llama3-SWE-RL-70B, achieves a **41.0%** solve rate on SWE-bench Verified, a human-verified collection of high-quality GitHub issues. This performance is state-of-the-art among medium-sized models and comparable to many proprietary LLMs such as GPT-4o. While SWE-RL is specifically applied to the issue-solving task, Llama3-SWE-RL has developed generalized reasoning skills through RL, demonstrating improved performance on out-of-domain tasks such as code reasoning, mathematics, and general language understanding. Overall, SWE-RL opens a new direction for enhancing the software engineering capabilities of LLMs through RL.

Limitations. Despite the promising results, our approach has several limitations. First, our reward implementation compares the sequence similarity between the predicted and oracle patch rather than their semantic equivalence. This may prevent the policy LLM from exploring alternative, functional equivalent solutions. Additionally, in Agentless Mini, the localization process is simplified to mapping repository structures to file paths, which lacks comprehensive context. Moreover, as a pipeline-based approach, Agentless Mini divides all steps into distinct inference stages. This “external structure” prevents the model from learning through interaction feedback and hinders its ability to consider the entire problem holistically.

Acknowledgement

We thank Chris Waterson, Chris Cummins, and Volker Seeker for their assistance in debugging and testing Agentless Mini; Kunhao Zheng, David Zhang, and Taco Cohen, Jonas Gehring, Vegard Mella for their assistance in the online RL infra; Pierre Chambon, Mihir Sanjay Kale, Parth Thakkar, Michael Jiang, Sinong Wang, and Jingyue Shen for their involvement in the PR dataset discussion; Jannik Kossen for his helpful discussions in model training; Kamila Benzina, Rachel Kim, Megan Miller, Hannah Schubloom, and Madeline Hepler for their support in the review process; Steven Xia and Yinlin Deng for their aid in setting up Agentless; Jiawei Liu, Yifeng Ding, Terry Yue Zhuo, and Naman Jain for their valuable discussions; Albert Örwall for their help in resolving issues with Moatless EvalTools.

References

- Anthropic. Claude 3.5 sonnet model card addendum. *Claude-3.5 Model Card*, 2024a.
- Anthropic. Raising the bar on SWE-bench Verified with Claude 3.5 Sonnet. <https://www.anthropic.com/research/swe-bench-sonnet>, 2024b.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models, 2021. <https://arxiv.org/abs/2108.07732>.
- Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. Codet: Code generation with generated tests. In *The Eleventh International Conference on Learning Representations*, 2023. <https://openreview.net/forum?id=ktrw68Cmu9c>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021.

- Chris Cummins, Volker Seeker, Dejan Grubisic, Baptiste Roziere, Jonas Gehring, Gabriel Synnaeve, and Hugh Leather. Meta large language model compiler: Foundation models of compiler optimization, 2024. <https://arxiv.org/abs/2407.02524>.
- DeepSeek-AI. Deepseek-v3 technical report, 2024. <https://arxiv.org/abs/2412.19437>.
- DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. <https://arxiv.org/abs/2501.12948>.
- DeepSeek-AI, Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y. Wu, Yukun Li, Huazuo Gao, Shirong Ma, Wangding Zeng, Xiao Bi, Zihui Gu, Hanwei Xu, Damai Dai, Kai Dong, Liyue Zhang, Yishi Piao, Zhibin Gou, Zhenda Xie, Zhewen Hao, Bingxuan Wang, Junxiao Song, Deli Chen, Xin Xie, Kang Guan, Yuxiang You, Aixin Liu, Qiushi Du, Wenjun Gao, Xuan Lu, Qinyu Chen, Yaohui Wang, Chengqi Deng, Jiashi Li, Chenggang Zhao, Chong Ruan, Fuli Luo, and Wenfeng Liang. Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence, 2024. <https://arxiv.org/abs/2406.11931>.
- Yinlin Deng, Chunqiu Steven Xia, Haoran Peng, Chenyuan Yang, and Lingming Zhang. Large language models are zero-shot fuzzers: Fuzzing deep-learning libraries via large language models, 2023.
- Yanguibo Ding, Zijian Wang, Wasi Uddin Ahmad, Hantian Ding, Ming Tan, Nihal Jain, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, Dan Roth, and Bing Xiang. Crosscodeeval: A diverse and multilingual benchmark for cross-file code completion. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. <https://openreview.net/forum?id=wgDcbBMSfh>.
- Xueying Du, Mingwei Liu, Kaixin Wang, Hanlin Wang, Junwei Liu, Yixuan Chen, Jiayi Feng, Chaofeng Sha, Xin Peng, and Yiling Lou. Classeval: A manually-crafted benchmark for evaluating llms on class-level code generation. *arXiv preprint arXiv:2308.01861*, 2023.
- AI@Meta: Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, and Angela Fan et al. The llama 3 herd of models, 2024. <https://arxiv.org/abs/2407.21783>.
- Paul Gauthier. Aider is ai pair programming in your terminal. <https://aider.chat/>, 2024.
- Jonas Gehring, Kunhao Zheng, Jade Copet, Vegard Mella, Quentin Carbonneaux, Taco Cohen, and Gabriel Synnaeve. Rlef: Grounding code llms in execution feedback with reinforcement learning, 2025. <https://arxiv.org/abs/2410.02089>.
- GitHub. Github rest api documentation. <https://docs.github.com/en/rest?apiVersion=2022-11-28>, 2022. Accessed: 2025-02-24.
- GitHub. Github. <https://github.com>, 2025.
- Ilya Grigorik. Gh archive. <https://www.gharchive.org/>, 2025. Accessed: 2025-02-23.
- Alex Gu, Baptiste Roziere, Hugh James Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida Wang. CRUXEval: A benchmark for code reasoning, understanding and execution. In *Proceedings of the 41st ICML*, volume 235 of *Proceedings of Machine Learning Research*, pages 16568–16621. PMLR, 21–27 Jul 2024. <https://proceedings.mlr.press/v235/gu24c.html>.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming – the rise of code intelligence, 2024.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. Measuring coding challenge competence with apps, 2021a.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021b. <https://openreview.net/forum?id=d7KBjmI3GmQ>.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021c. <https://openreview.net/forum?id=7Bywt2mQsCe>.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. Qwen2.5-coder technical report, 2024. <https://arxiv.org/abs/2409.12186>.
- Kush Jain, Gabriel Synnaeve, and Baptiste Rozière. Testgeneval: A real world unit test generation and test completion benchmark, 2024a. <https://arxiv.org/abs/2410.00752>.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code, 2024b.
- Nan Jiang, Kevin Liu, Thibaud Lutellier, and Lin Tan. Impact of code language models on automated program repair, 2023.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues?, 2023.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench leaderboard. <https://www.swebench.com>, 2024. Accessed: 2025-02-04.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. <https://arxiv.org/abs/1412.6980>.
- Caroline Lemieux, Jeevana Priya Inala, Shuvendu K Lahiri, and Siddhartha Sen. Codamosa: Escaping coverage plateaus in test generation with pre-trained large language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 919–931. IEEE, 2023.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umabathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. Starcoder: may the source be with you!, 2023.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, December 2022. ISSN 1095-9203. doi: 10.1126/science.abq1158. <http://dx.doi.org/10.1126/science.abq1158>.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatGPT really correct? rigorous evaluation of large language models for code generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. <https://openreview.net/forum?id=1qvx610Cu7>.
- Jiawei Liu, Songrun Xie, Junhao Wang, Yuxiang Wei, Yifeng Ding, and Lingming Zhang. Evaluating language models for efficient code generation. In *First Conference on Language Modeling*, 2024a. <https://openreview.net/forum?id=IBCBMeAhmC>.

- Tianyang Liu, Canwen Xu, and Julian McAuley. Repobench: Benchmarking repository-level code auto-completion systems. In *The Twelfth International Conference on Learning Representations*, 2024b. <https://openreview.net/forum?id=pPjZIOuQuF>.
- Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osaе Osaе Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun Chai, Niklas Muennighoff, Xiangru Tang, Muhtasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sebastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. Starcoder 2 and the stack v2: The next generation, 2024.
- Yingwei Ma, Rongyu Cao, Yongchang Cao, Yue Zhang, Jue Chen, Yibo Liu, Yuchen Liu, Binhua Li, Fei Huang, and Yongbin Li. Lingma swe-gpt: An open development-process-centric language model for automated software improvement, 2024. <https://arxiv.org/abs/2411.00622>.
- OpenAI. Gpt-4o system card, 2024a. <https://arxiv.org/abs/2410.21276>.
- OpenAI. Openai o1 system card, 2024b. <https://arxiv.org/abs/2412.16720>.
- OpenAI. simple-evals. <https://github.com/openai/simple-evals>, 2024. Accessed: 2025-02-23.
- OpenAI. Introducing swe-bench verified. <https://openai.com/index/introducing-swe-bench-verified>, 2024. Accessed: 2025-02-04.
- Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with swe-gym, 2024. <https://arxiv.org/abs/2412.21139>.
- John W. Ratcliff and David E. Metzener. Pattern Matching: The Gestalt Approach. *Dr. Dobbs's Journal*, page 46, July 1988. <https://www.drdoobbs.com/database/pattern-matching-the-gestalt-approach/184407970>.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code, 2023.
- Max Schäfer, Sarah Nadi, Aryaz Eghbali, and Frank Tip. An empirical evaluation of using large language models for automated unit test generation. *IEEE Transactions on Software Engineering*, 2023.
- John Schulman. Approximating kl divergence. <http://joschu.net/blog/kl-approx.html>, 2020. Accessed: 2025-02-22.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. <https://arxiv.org/abs/2402.03300>.
- Sida I. Wang, Alex Gu, Lovish Madaan, Dieuwke Hupkes, Jiawei Liu, Yuxiang Wei, Naman Jain, Yuhang Lai, Sten Sootla, Ofir Press, Baptiste Rozière, and Gabriel Synnaeve. Eval-Arena: noise and errors on llm evaluations. <https://github.com/crux-eval/eval-arena>, 2024a.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. Openhands: An open platform for ai software developers as generalist agents, 2024b. <https://arxiv.org/abs/2407.16741>.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. https://openreview.net/forum?id=_VjQlMeSB_J.
- Yuxiang Wei, Chunqiu Steven Xia, and Lingming Zhang. Copiloting the copilots: Fusing large language models with completion engines for automated program repair, 2023.
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. Magicoder: Empowering code generation with OSS-instruct. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 52632–52657. PMLR, 21–27 Jul 2024. <https://proceedings.mlr.press/v235/wei24h.html>.
- Chunqiu Steven Xia and Lingming Zhang. Less training, more repairing please: Revisiting automated program repair via zero-shot learning, 2022.
- Chunqiu Steven Xia, Matteo Paltenghi, Jia Le Tian, Michael Pradel, and Lingming Zhang. Universal fuzzing via large language models, 2023a.
- Chunqiu Steven Xia, Yuxiang Wei, and Lingming Zhang. Automated program repair in the era of large pre-trained language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 1482–1494, 2023b. doi: 10.1109/ICSE48619.2023.00129.
- Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Agentless: Demystifying llm-based software engineering agents, 2024. <https://arxiv.org/abs/2407.01489>.
- Chengxing Xie, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou, and Kai Chen. Swe-fixer: Training open-source llms for effective and efficient github issue resolution, 2025.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2024a.
- John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik R Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. <https://openreview.net/forum?id=mXpq6ut8J3>.
- John Yang, Carlos E. Jimenez, Alex L. Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press, Niklas Muennighoff, Gabriel Synnaeve, Karthik R. Narasimhan, Diyi Yang, Sida I. Wang, and Ofir Press. Swe-bench multimodal: Do ai systems generalize to visual software domains?, 2024c. <https://arxiv.org/abs/2410.03859>.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-of-thought reasoning in llms, 2025. <https://arxiv.org/abs/2502.03373>.
- Zhiqiang Yuan, Yiling Lou, Mingwei Liu, Shiji Ding, Kaixin Wang, Yixuan Chen, and Xin Peng. No more manual tests? evaluating and improving chatgpt for unit test generation. *arXiv preprint arXiv:2305.04207*, 2023.
- Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen, and Wenhui Chen. Acecoder: Acing coder rl via automated test-case synthesis. *ArXiv*, 2502.01718, 2025.
- Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. Repocoder: Repository-level code completion through iterative retrieval and generation, 2023.
- Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. Autocoderover: Autonomous program improvement, 2024. <https://arxiv.org/abs/2404.05427>.

Wenting Zhao, Nan Jiang, Celine Lee, Justin T Chiu, Claire Cardie, Matthias Gallé, and Alexander M Rush. Commit0: Library generation from scratch, 2024. <https://arxiv.org/abs/2412.01769>.

Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*, 2024.

Albert Örwall. Moatless tools. <https://github.com/aorwall/moatless-tools>, 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We listed the core contributions and key results in the last paragraph of the introduction section (§1).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discussed the limitations of our work in §5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our work studies improving LLMs’ reasoning ability through reinforcement learning on open software data. Therefore, theoretical results are not applicable here. Instead, we performed a comprehensive set of evaluations (§3) in an empirical fashion.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We detailed our data curation (Appendix A), technique (§2), and experimental configurations (§3). We also included the reward implementation and evaluation pipeline in the supplemental material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We included the reward implementation and evaluation code in the supplemental material, along with detailed instructions for using the artifact to ensure transparency and reproducibility of our results. At this time, we cannot share the training data or training pipeline due to privacy and proprietary considerations. However, once the necessary reviews are completed and associated risks are addressed, we will make these resources available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We detailed the configurations and rationales for model training in §3.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We acknowledge that we did not include error bars for all evaluations on SWE-bench. This is because SWE-bench comprises real-world software engineering problems, and each instance is very expensive to run. Also, we tried to follow the design of SWE-bench where all accepted submissions are evaluated using a single attempt per example.

Guidelines:

- The answer NA means that the paper does not include experiments.

- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We reported our compute configurations in §3.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have read the NeurIPS Code of Ethics and believe our work does not violate the terms.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Our technique is neutral in not implying clear positive or negative impacts on society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work does not involve the release of models or datasets that pose a high risk of misuse, so safeguards are not applicable in this context.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cited the datasets and models for training, synthetic data generation, and evaluations, and specified their versions to our best efforts. We also licensed our code in the supplementary material.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We clearly described the license of our asset and respected the license of all derived assets in the supplementary material.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: LLMs are used to fix grammar mistakes and awkward phrasing in writing.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Raw pull request data curation



Figure 6: **Overview of SWE-RL’s raw pull request data curation process.** The collected git clones and GitHub events are transformed into self-contained PR instances via decontamination, aggregation, relevant files prediction, and filtering.

Figure 6 provides a high-level overview of our process for curating the raw PR dataset for Llama3-SWE-RL. In the following paragraphs, we detail each step of the curation process. During data processing, we exclude all the repositories used by SWE-bench (Jimenez et al., 2023) to prevent data contamination.

GitHub events and clones. The goal of this stage is to recover all pull request details that human developers can inspect on GitHub. To achieve this, we need two sources of information: (1) all events that occur within a PR and (2) the source code of a repo before the changes introduced by the PR are merged. We derive all GitHub (GitHub, 2025) events from GHArchive (Grigorik, 2025), which contains all activity events data from GitHub. Our collection includes all GitHub events from Jan 1, 2015 to Aug 31, 2024.

To obtain source code, since pull requests often occur at different commit stages of a repository, we opt to use `git clone` to retrieve the entire repository with its commit history, rather than relying on the GitHub API (GitHub, 2022) to download specific code snapshots. Eventually, we successfully cloned and processed 4.6M repositories.

PR data aggregation. The collected events and git clones are disparate entities that require further processing before they can be used for training. At this stage, we focus on each PR individually and aggregate all pertinent information associated with it. This includes mentioned issues, user discussions, review comments, initial code contents, and subsequent commits and code changes.

First, we keep only merged PRs and gather all related conversational events for each PR, sorting them in chronological order. Next, using the `base_commit` and `head_commit` hashes of a PR, we retrieve the contents of all modified files indicated by its patch at the merge base of the two commits. The reason is that many PRs aim to merge back into the main branch, which may have undergone changes since the PR was created. By considering the merge base as the actual starting point for a developer working on the PR, we can more accurately understand the context of the changes. We save all intermediate commits and code changes made between the merge base and the head commit. Additionally, we extract the complete patch representing cumulative changes from start to finish. Finally, we scan each aggregated PR to identify patterns that resemble issues, and associate the matched issues with the corresponding PR. In the end, we have 24M aggregated PR instances.

Relevant files prediction. Currently, each pull request includes only the code files that have been modified. In our earlier experiments, we noticed that this approach let LLMs learn a bias: the model consistently generated edits for every code file presented and was unable to handle noisy files presented in the context. This issue was also mentioned in the finetuning experiment discussed in the SWE-bench paper (Jimenez et al., 2023). To mitigate this problem, we prompt Llama-3.1-70B-Instruct (Dubey et al., 2024) to generate a list of relevant but unmodified files given each PR description and the changed files, and include the contents of these files in our final dataset.

Data filtering. GitHub PRs can be quite noisy, so we implemented various filtering strategies to eliminate potentially harmful PRs. In designing these filtering rules, our goal is to maximize the recall of high-quality PRs while permitting a certain level of noise. First, we remove the bot-generated PRs whose title, description, or username contains keywords “[bot]”, “dependabot”, “renovate”, “bump”, or “automerger”. Also, we remove PRs with empty changes or with extremely large number of changes (e.g., in some PRs, the developer uploaded a directory of data files by mistake). Additionally, we implemented a more fine-grained set of filters used in CodeLlama (Rozière et al., 2023) to examine each code change hunk. We then removed any PRs where these filters flagged all code changes. For example, this can exclude PRs with only lock file changes or version updates. Finally, this gives us around 11M unique PR instances. Before applying reinforcement learning, we select 273k high-quality PRs from these raw PRs, as discussed in §2.

B Agentless Mini

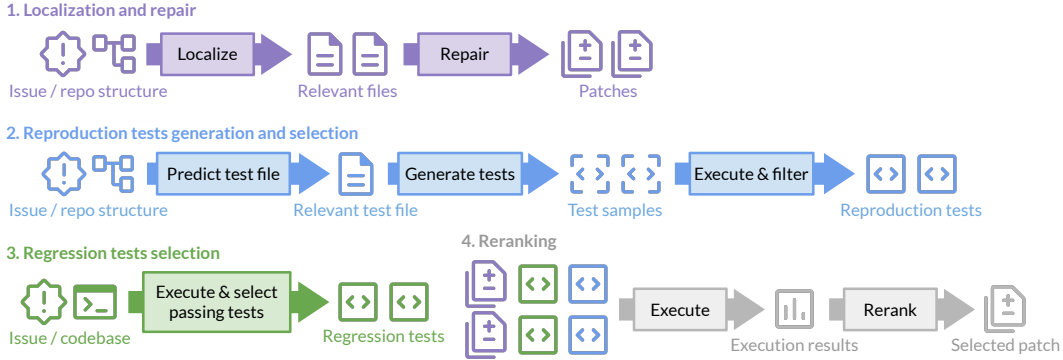


Figure 7: **The Agentless Mini scaffold.** The design emphasizes easy decomposition, parallelization, and scalability.

In addition to a model proficient in code editing, effectively tackling software engineering tasks, such as those found in SWE-bench (Jimenez et al., 2023), also requires a robust scaffold. Agentless (Xia et al., 2024) is one of the state-of-the-art scaffolds for SWE-bench at the time of writing. Building upon Agentless with various simplifications and enhancements, we developed Agentless Mini, a framework that prioritizes straightforward component decomposition, parallelization, and scalability. With Agentless Mini, each step’s inference or execution compute can be independently scaled to enhance SWE-bench performance. In Figure 7, we present a detailed illustration of Agentless Mini’s working principles. The following paragraphs will elaborate on each step and highlight the differences from Agentless.

Localization and repair. For localization, we employ a prompting-based approach that enables the model to predict relevant file paths based on a given issue and the repository’s structure. Unlike Agentless, which involves two additional detailed steps to identify related elements and files, as well as a separate embedding model, Agentless Mini simplifies the process. It generates multiple samples of potentially problematic files from the model and consolidates them into unique sets for repair.

During the repair phase, the LLM is conditioned on the full content of the files to predict search/replace edits. We generate multiple repair samples from different location sets, ensuring a comprehensive exploration of the patch search space.

Reproduction tests generation and selection. Agentless samples reproduction tests for patch selection. Initially, multiple reproduction tests are generated based on an issue description, and one majority sample is selected after filtering. These tests must have distinct logic to output "Issue reproduced" and "Issue resolved" when the issue is reproduced or resolved, respectively. They are filtered based on whether they correctly output "Issue reproduced" when executed in the original codebase. Agentless Mini enhances this pipeline with two key improvements. First, instead of relying solely on the issue description, the model retrieves a relevant test file to guide test generation. Additionally, rather than selecting just one majority sample, Agentless Mini allows for the selection of multiple top test samples based on voting results. In our evaluation, using more test samples has proven beneficial for reranking (§3.4).

Regression tests selection. We select regression tests in the same manner as Agentless. Initially, we gather all passing tests for each issue by executing the code before any modifications. This step does not require model inference and needs to be performed only once. Subsequently, an additional, optional inference step is conducted to filter out passing tests that are supposed to fail after the issue is fixed. The rest of the tests will be marked as regression tests.

Reranking. Agentless Mini utilizes both regression and reproduction tests for reranking. For each issue, every applied patch is executed against the regression tests. The patches that result in the minimum number of existing test failures are selected. For each issue, we choose the top- N reproduction tests and run each patch against these tests. If the execution outputs "Issue resolved", we mark the patch passing this test. We then adopt the dual execution agreement objective from CodeT (Chen et al., 2023). Specifically, patches \mathcal{P} that pass the same set of reproduction tests \mathcal{T} are denoted as a

consensus group. Each consensus group is scored using the formula $|\mathcal{P}| \times |\mathcal{T}|^2$. With this objective, consensus groups with patches passing more tests receive higher scores. Additionally, groups where more patches pass the tests are scored higher, although passing more tests is prioritized over having more patches. Finally, we identify the consensus group with the highest score and select the best patch from this group using majority voting.

C Synthesizing supervised-finetuning data




Figure 8: **Synthetic data pipeline for constructing SFT data.** We start by collecting high-quality seed PRs using heuristics, then generate synthetic localization and code-editing samples, and finally use the ground-truth edited files and patches to filter out incorrect samples.

Figure 8 shows our method of generating synthetic supervised-finetuning (SFT) data. The data generation pipeline is inspired by Magicoder (Wei et al., 2024), where the OSS-Instruct technique generates high-quality code instruction data from open-source seed snippets. We apply a similar methodology to generate both fault localization and code editing data using high-quality PR seeds.

Collecting high-quality seed PRs. To begin with, we extract high-quality PR seeds from the raw dataset we collected, as detailed in Appendix A. These seeds are chosen based on specific heuristics. For example, a PR instance should include at least one linked issue, the issue should describe a bug fix request, and the code changes should involve programming files.

Localization and editing data synthesis. We adopt Llama-3.3-70B-Instruct (Dubey et al., 2024) for data synthesis. For localization data, we prompt the model with the issue description, repository structure, and the paths of edited and relevant files as hints. We then ask the model to identify the relevant files for modification or review by generating a thought process, followed by a prioritized list of file paths. During filtering, we ensure that the model’s response includes all files that are genuinely edited in the PR, and these files should be prioritized in the ranking.

Similarly, in terms of code editing, we synthesize code edits for a given issue, in search/replace format (Xia et al., 2024), by providing the ground-truth PR and patch as guidance to Llama-3.3-70B-Instruct. During the filtering process, we ensure that all search/replace blocks adhere to the correct format and that all search paths and blocks can be accurately matched within the input files’ context.

SFT baseline. As discussed in §3.1, we meticulously constructed Llama3-SWE-SFT-70B as a strong SFT baseline. This SFT baseline is trained on a 16k context window for 2B tokens, where the training data consists of a mix of the aforementioned synthetic localization and editing data, as well as coding and general SFT datasets from Llama 3 (Dubey et al., 2024).

D Complete prompt

The following is a complete version of Figure 2:

```
PROMPT_TEMPLATE = ""<|begin_of_text|><|start_header_id|>system<|end_header_id|>
```

```
A user will ask you to solve a task. You should first draft your thinking process (inner monologue). Then, generate the solution.
```

```
Your response format must follow the template below:
```

```
<think>
```

```
Your thoughts or/and draft, like working through an exercise on scratch paper. Be as casual and as long as you want until you are confident to generate a correct solution.
```

```
</think>
```

```
<solution>
```

Final solution presented to the user.
</solution><|eot_id|><|start_header_id|>user<|end_header_id|>

We are currently solving the following issue within our repository. Here is the issue text:

```
--- BEGIN ISSUE ---  
{problem_statement}  
--- END ISSUE ---
```

Below are some code segments, each from a relevant file. One or more of these files may contain bugs.

```
--- BEGIN FILE ---  
...  
{content}  
...  
--- END FILE ---
```

Please first localize the bug based on the issue statement, and then 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

Here is an example:

```
```python  
mathweb/flask/app.py
<<<<<< SEARCH
from flask import Flask
=====
import math
from flask import Flask
>>>>>> REPLACE
```
```

Please note that the *SEARCH/REPLACE* edit REQUIRES PROPER INDENTATION. If you would like to add the line ' print(x)', you must fully write that out, with all those spaces before the code!

Wrap each *SEARCH/REPLACE* edit in a code block as shown in the example above. If you have multiple *SEARCH/REPLACE* edits, use a separate code block for each one.<|eot_id|><|start_header_id|>assistant<|end_header_id|>

"""
