

---

# SWE-Spot: Building Small Repo-Experts with Repository-Centric Learning

---

Jinjun Peng<sup>\*1</sup> Magnus Saebo<sup>\*1</sup> Tianjun Zhong<sup>1</sup> Yi-Jie Cheng<sup>1</sup> Junfeng Yang<sup>1</sup> Baishakhi Ray<sup>1</sup>  
 Simin Chen<sup>1</sup> Yangruibo Ding<sup>2</sup>

## Abstract

The deployment of coding agents in privacy-sensitive and resource-constrained environments drives the demand for capable open-weight Small Language Models (SLMs). However, they suffer from a fundamental capability gap: unlike frontier large models, they lack the inference-time strong generalization to work with complicated, unfamiliar codebases. We identify that the prevailing Task-Centric Learning (TCL) paradigm, which scales exposure across disparate repositories, fails to address this limitation. In response, we propose Repository-Centric Learning (RCL), a paradigm shift that prioritizes vertical repository depth over horizontal task breadth, suggesting SLMs must internalize the “physics” of a target software environment through parametric knowledge acquisition, rather than only attempting to recover it via costly inference-time search. Following this new paradigm, we design a four-unit Repository-Centric Experience, transforming static codebases into interactive learning signals, to train SWE-SPOT-4B, a family of highly compact models built as repo-specialized experts that breaks established scaling trends, outperforming open-weight models up to 8× larger (e.g., CWM by Meta, Qwen3-Coder-30B) and surpassing/matching efficiency-focused commercial models (e.g., GPT-4.1-mini, GPT-5-nano) across multiple SWE tasks. Further analysis reveals that RCL yields higher training sample efficiency and lower inference costs, emphasizing that for building efficient intelligence, repository mastery is a distinct and necessary dimension that complements general coding capability.

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Computer Science, Columbia University, New York, USA <sup>2</sup>Computer Science Department, University of California, Los Angeles, California, USA. Correspondence to: Jinjun Peng <jinjun@cs.columbia.edu>, Yangruibo Ding <yrbding@cs.ucla.edu>.

Preprint. January 30, 2026.

## 1 Introduction

Large language model (LLM) agents have significantly improved productivity across different kinds of software engineering workflows. Yet today’s most widely adopted systems (Anthropic, 2025b; Google, 2025; Cursor, 2025a; GitHub, 2025; Amazon, 2025) primarily rely on small, fully open-source models with only a few million parameters, making them inexpensive to operate and ideal for local deployment in privacy-sensitive settings. As agentic workflows transition from experimental use to production across organizations, there is a growing demand for capable *small language models* (SLMs) that have fewer than 10B parameters and can be deployed locally and cost-effectively (Apple, 2025), preserving code confidentiality while avoiding dependence on proprietary APIs and their evolving economic and access constraints (Cursor, 2025b).

**The current paradigm** To widen the capability gap between small open-weight models and much larger proprietary systems, recent work has abandoned Task-Centric Learning (TCL) and explicitly avoided framing software engineering as GitHub issue resolution (Jimenez et al., 2024). Early approaches relied on fixed-workflow scaffolds (Xia et al., 2024; Li et al., 2025b), decomposing issue resolution into predefined stages, including retrieval, localization, patch generation, and validation, and training specialized models for each subtask (Tang et al., 2025; Xie et al., 2025). More recent work, driven by Reinforcement Learning with Verifiable Rewards (RLVR) (Guo et al., 2025), retains rigid fixed pipelines and removes agent-driven exploration, ignoring the correctness of the final patch as supervision (Wei et al., 2025). To scale this paradigm, several “gym” infrastructures have been introduced to support large-scale TCL training and evaluation across diverse, unrelated repositories (Luo et al., 2025; He et al., 2025; Zan et al., 2025a; Deng et al., 2025). The underlying assumption is the exposure to massive numbers of task trajectories will yield broadly generalizable bug-fixing skills, producing increasingly capable open-weight *issue resolvers*. (See Appendix A for more related work.)

**The limitations** While Task-Centric Learning (TCL) has shown promise for medium-to-large models (e.g., 30B+ pa-

rameters), it exposes fundamental limitations when applied to small language models (SLMs). The core issue is an inference-time capability gap:

frontier models can compensate for missing prior knowledge in unfamiliar environments through extensive inference-time search, trial-and-error, and reflection. In contrast, SLMs lack this capacity: when trained as task specialists across disparate codebases, they tend to overfit surface-level patterns (e.g., patch templates, shell syntax) (Wang et al., 2023) rather than internalizing repository-specific design and semantics. Without a grounded understanding of the target codebase, SLMs struggle to navigate through and adapt to unseen environments, leading to inefficient search and brittle generalization at deployment time, as demonstrated in Section 4.2.

**Our proposal** These observations motivate a fundamental shift in perspective. If small language models (SLMs) cannot reliably acquire repo-specific knowledge at inference time, that knowledge must be internalized during training and carried as prior into deployment. Our key insight is that, for SLMs, the horizontal breadth of task-centric learning must be complemented with the vertical depth of **Repository-Centric Learning (RCL)**. This training–inference asymmetry mirrors how human developers build expertise: they do not gain proficiency by fixing isolated bugs across thousands of unrelated projects; instead, they develop expertise through sustained interaction with a target repository (or a small set of related repositories). Such interaction spans multiple learning signals, including code reading, execution and testing, debugging, and iterative modification, which together expose the repository’s structure and behavior. Over time, these experiences induce an implicit representation of the repository’s *operational semantics*, including its dependency structure, execution dynamics, and behavioral invariants.

Analogously, RCL trains SLMs to acquire repository-specific representations *before* deployment. By grounding learning in repeated, multi-task interactions within a single codebase, RCL encourages SLMs to encode the repository’s operational semantics as a reusable prior. This enables more efficient reasoning over complex software engineering tasks, reducing over-reliance on inference-time search and mitigating brittle generalization to unseen repository dynamics. Figure 1 contrasts RCL with task-centric learning, showing that when performance is evaluated across multiple tasks *within* the same repository, RCL yields stronger and more stable generalization.

We emphasize that RCL complements rather than replaces task-centric learning (TCL). While TCL optimizes a *single task*, typically issue resolution, across many unrelated repositories to train a general-purpose issue resolver, RCL targets *multiple software engineering tasks* within a *fixed repository*,




Figure 1. **Conceptual Illustration on Repository-Centric Learning (RCL) vs. Task-Centric Learning (TCL)**. TCL scales across repos to learn shared task skills, but struggles on complex codebases due to limited inference-time generalization. RCL instead scales within a repository, diversifying repo-centric experience to enable durable knowledge acquisition and repository mastery.

yielding a repository-specialized model. Using Django as a representative example, TCL may resolve individual issues in isolation, whereas RCL enables consistent performance across tasks within the same codebase, including issue resolution, feature implementation, test development, etc. This distinction reflects a key requirement of repository-level SWE performance: combining task-agnostic capability with repository-specific knowledge. While large models may recover the latter at inference time through strong generalization, RCL is necessary for SLMs to internalize this knowledge during training.

To operationalize RCL, we train SLMs not merely as issue resolvers, but as *repository experts*, learning from diverse signals that span the full software engineering lifecycle of a target codebase. We structure this training around four units of **Repository-Centric Experience (RCX)** (Figure 2): (i) *Software design*, where the agent infers architectural intent and module responsibilities through *active* code analysis; (ii) *Contextual Implementation*, where it models cross-file dependencies, internal APIs, and implementation logic via context-aware completion; (iii) *Evolutionary Replay*, where it learns historical changes to understand design trade-offs and constraints shaping the codebase over time. (iv) *Semantic-Runtime Alignment*, where it designs tests to capture nuanced discrepancies between specifications and runtime behaviors; Together, these experiences enable SLMs to acquire repo-specific knowledge that supports robust, multi-task performance within a single codebase.

**Results** We validate RCL by training SWE-SPOT, a family of *only 4B-parameter* models specialized

as Small Repo-Experts in target codebases. Under a **repository-centric evaluation** with a strict temporal protocol which spans four software engineering tasks (issue resolution, feature implementation, test generation, and codebase QA), SWE-SPOT consistently breaks the scaling trends established by TCL. Despite its compact size, SWE-

SPOT outperforms open-weight TCL-trained models up to 8x larger and matches or exceeds efficiency-comparable commercial models such as GPT-5-nano and GPT-4.1-mini.

Beyond end-to-end accuracy, RCL induces a qualitative shift in agent behavior. By internalizing repository-specific structure and dynamics, SWE-SPOT achieves *higher training sample efficiency* and substantially *lower inference-time cost*. Crucially, RCL remains necessary even in the combination with strong context retrieval and test-time scaling, establishing a clear performance gap that indicates *deep parametric knowledge acquisition* beyond context memorization or superficial pattern matching. Ablation studies further show that the proposed units of RCX are synergistic: for example, training on test generation improves issue resolution, confirming that repository mastery emerges from holistic exposure rather than isolated tasks.

**Contributions** In summary, our paper makes the following contributions:

- **The Paradigm:** We introduce *Repository-Centric Learning (RCL)* as a new dimension for training coding agents. We show that for SLMs, task breadth alone is insufficient: repository-specific knowledge must be internalized during training to compensate limited inference-time capacity.
- **The Framework:** We present a practical RCL framework that converts static repositories into interactive learning signals via four-unit *Repository-Centric Experiences*, mirroring how human experts acquire codebase mastery.
- **The Results:** We train SWE-SPOT-4B using RCL and demonstrate a clear break from task-centric scaling trends. Across multiple software engineering tasks, SWE-SPOT-4B outperforms much larger open-weight models (30B+) and exceeds/matches efficiency-comparable commercial APIs (e.g., GPT-5-nano, GPT-4.1-mini).
- **The Analysis:** We show that RCL yields higher sample efficiency, lower inference cost, and strong cross-task transfer. Ablations confirm that these gains persist beyond context retrieval and test-time scaling, indicating deep parametric knowledge acquisition rather than surface-level memorization.

We open-source our code, data, and model weights to support follow-on research into efficient, repo-centric software agents: [GitHub: SWE-Spot/swespot](#).

## 2 Repository-Centric Learning

Existing approaches to software agent training face a critical dichotomy. As discussed in Section 1, Task-Centric Learning (TCL) builds generalizable task capabilities but lacks the deep, repository-specific grounding required for true mastery of a target codebase. Conversely, our pilot studies indicate that classical domain adaptation, such as continued pre-training on static source files including its structural

variants like fill-in-the-middle (FIM), yields negligible improvements for agentic tasks, because such static objectives limit the model to merely *reciting* the code, failing to teach it how to effectively *work with* the software environment.

To bridge this gap, we propose Repository-Centric Learning (RCL) from diverse types of **Repository-Centric Experience (RCX)**. Inspired by the recent trend toward experience-driven AI (Silver & Sutton, 2025), we prioritize dynamic interaction over static ingestion. By learning from experience rather than static text, the model internalizes the repository’s “physics” during its interactions. RCX is made concrete as learnable agentic trajectories in the target environment. These trajectories include observing search outputs, reasoning about runtime behaviors, leveraging grounded knowledge for planning, and so on. This approach transcends simple recitation by fusing repository-specific knowledge with active agentic thinking, planning, and actions.

To ensure a comprehensive mastery of the target environment, we identify four distinct units of RCX, listed below and visualized in Figure 2, that cover the full spectrum of agentic software engineering.

**Software Design** Static code encodes rich but implicit information about architectural intent, design philosophy, and cross-file dependencies, knowledge critical for navigating a complex codebase. Traditional static analysis captures *syntactic* relationships but misses this *semantic* layer, which partially lies in natural language (e.g., variables/files names) beyond the capability of static analysis. Therefore, we extract this knowledge through *active code analysis*, inspired by Lin et al. (2025): given a target location in the codebase (a module, file, or function/class), an agent interactively explores the repository to produce a structured report articulating the component’s functionality, design rationale, and interactions with the broader system.

This experience teaches the agent to reason about why code is structured as it is, not only what it does.

**Contextual Implementation** Implementing functionality that complies with an existing codebase requires awareness of global conventions, abstractions, and cross-file dependencies. While existing fill-in-the-middle (FIM) pretraining teaches models to complete missing code spans, most formulations provides cross-file contexts obtained from retrieval systems, reducing the task to *passive* completion. We propose *agentic FIM* as a more realistic formulation: given only a functionality to implement,

the agent must *actively* explore the repository to discover relevant context before producing a compliant implementation. To maximize the density of repository-level learning signals rather than trivial, local ones, we use static analysis as a cost-efficient way to assist with task construction, prioritizing implementation targets involving cross-file dependencies.

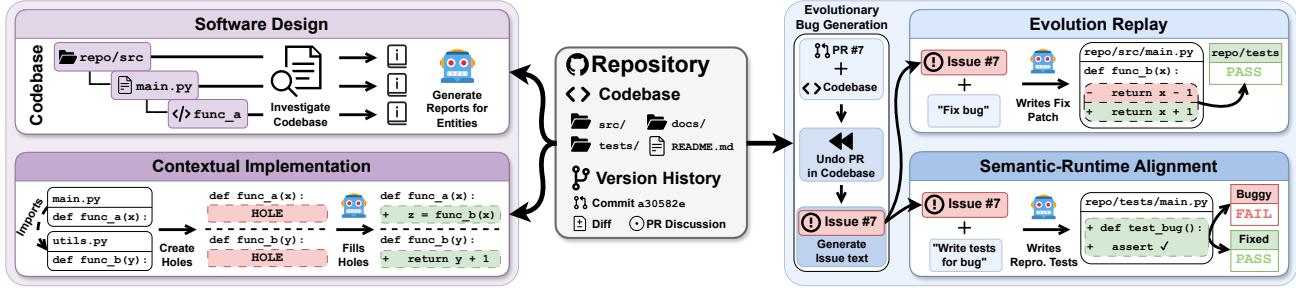


Figure 2. A Four-Unit Design of Repository-Centric Experience (RCX)

However, the resulting RCX tasks crucially do not include these analysis results: the learning agent must discover the relevant context through its own exploration, building internalized structural knowledge rather than relying on external tooling.

**Evolutionary Replay** A core maintainer’s expertise emerges not from reading one version of the codebase alone, but from experiencing the repository’s evolution over time. We replicate this learning signal by mining pull request history to reintroduce historical bugs into the current codebase, tasking the agent with resolving realistic issues during the repository’s evolution. While the PR Mirror method from Yang et al. (2025) is designed for TCL, which may fall short on intra-repository diversity, we adapt it for RCL by relaxing their strict filtering, leveraging agent-generated issue text, and more to maximize coverage, yielding roughly an 18-fold increase in per-repository task instances. We further augment this reality-grounded data with wholly synthesized tasks to compensate for repositories with fewer commits. Debugging historical faults demands intimate knowledge of the codebase’s structure and behavioral features, teaching the agent not only what the code does, but how it has failed and been repaired over time.

**Semantic-Runtime Alignment** Bugs arise from misalignment between expected program behavior and actual runtime execution. We train our agent to detect these misalignments through experiences of generating test cases that specify expected behavior for historical faults in the codebase. Concretely, we reuse the issue instances generated through Evolutionary Replay, but instead of resolving the bug, the agent must produce *reproduction tests* that formally specify the expected semantics: failing when the historical bug is applied, and passing on the fixed codebase. This dual use of the same underlying bugs maximizes the learning extracted from the repository’s evolution history, and additionally trains the agent to identify and formalize semantic-runtime misalignments based on real past cases.

In this work, we instantiate RCL by synthesizing RCX trajectories with a teacher model (Gemini-2.5-Pro) and training a student SLM (Qwen3-4B-Instruct-2507) via supervised

fine-tuning (SFT).

While this implementation suffices to demonstrate the benefits and necessity of RCL in the following sections, the paradigm is not tied to any particular training recipe, such as alternatives discussed in Section 5. Appendix B has more implementation details of RCX.

### 3 Experimental Setup

#### 3.1 Repository-Centric Evaluation for Agentic Coding

To rigorously assess the efficacy of RCL, we design a comprehensive **Repository-Centric Evaluation (RCE)** suite. Unlike traditional Task-Centric Evaluation (TCE), which focuses on scaling the repository diversity within a single task to test generalization to wild repositories, we propose RCE to validate whether RCL enables multi-spectrum mastery in the target codebase rather than only fitting to specific task formats (e.g., issue resolving).

Specifically, we compile several existing TCE benchmarks across four agentic coding tasks as our RCE benchmark:

- **Issue Resolution:** We employ SWE-bench-Verified (Chowdhury et al.) to assess the agent’s ability to resolve realistic GitHub issues.
- **Test Generation:** Derived from SWE-bench-Verified, TDD-Bench-Verified (Ahmed et al., 2025a) evaluates the capability of writing reproduction tests, i.e., test cases that fail on a buggy codebase but pass on the patched version.
- **Feature Implementation:** FEA-Bench (Li et al., 2025c) selects realistic Pull Requests with a focus on new feature development. It can be quite challenging as it often requires scattered cross-file modifications with compliance.
- **Codebase QA:** Separate from explicit coding tasks, SWE-QA (Peng et al., 2025) evaluates the ability to answer codebase-specific questions, a practical skill in human-AI collaborative development.

**Temporal Evaluation and Repository Selection** A critical challenge in evaluating repository experts is data leakage. To ensure the model is not evaluated on changes that it has already learned during RCL, we implement a strict **Tem-**

**poral Evaluation Protocol** with a knowledge-cutoff date of December 31, 2020. We construct our training data using only repository knowledge available prior to this date, and utilize instances created after this cutoff as our RCE instances. To manage computational costs without loss generality, we select the top-7 repositories with the highest density of post-cutoff instances across the selected four benchmarks above.

Our final RCE suite consists of 231 issue resolving, 211 test generation, 152 feature implementation, and 336 codebase QA instances.

**Metrics** For the three coding tasks (issue resolution, test generation, feature implementation), we report the pass rate (pass@1) based on execution results. For codebase QA, following the protocol by SWE-QA (Peng et al., 2025), we utilize LLM-as-a-judge (based on GPT-5.2) to score responses by agents, and report the average scaled to 0-100.

### 3.2 Models and Baselines

We compare SWE-SPOT-4B post-trained using RCL against three categories of models:

- **Efficency-focused Commercial Models:** We evaluate Gemini-2.5-Flash-Lite, GPT-4.1-mini (the 2nd fastest variant in the best non-reasoning GPT series, and GPT-5-nano (the fastest version of GPT-5).
- **Open-Weight Generalists:** We include Gemma-3-27B-it, Qwen3-Coder-30B-Instruct, and our base model Qwen3-4B-Instruct-2507 as general-purpose open-weight models.
- **Task-Centric Agents:** We also compare against CWM by Meta, a 32B model trained with 3M agentic coding trajectories from 3.15k repositories (Meta FAIR CodeGen, 2025), and Mini-Coder-4B, a model post-trained with 400k SWE-Smith trajectories (Olmedo, 2025). Both are extensively trained within the paradigm of TCL.

### 3.3 Implementation Details

**Inference Details** To ensure a fair comparison focused on model capability rather than agentic scaffold engineering, we utilize mini-SWE-agent, which is designed for *model-centric evaluation* (Mini-SWE-agent, 2025) across all evaluations. We make minimal changes to adapt the default prompt to the four evaluation tasks to ensure models can fully understand the task requirements. To mitigate the “looping” issue detailed in Appendix C, for all open-weight models including our trained ones, we apply the following sampling parameters changed from the recommendation of the base model (Team, 2025): a temperature of 1.0, top- $p$  of 0.8, top- $k$  of 20, and a repetition penalty of 1.05. For commercial LLM APIs, we keep the default settings and disable or minimize the reasoning effort to focus on non-reasoning models in our evaluation. Since we do not use

greedy decoding, all experiments are run 3 times to report the average. Furthermore, to finish the costly long-trajectory evaluation within a reasonable time, we set the maximum trajectory steps to 64 and the context limit to 48k tokens. Trial runs confirmed that this configuration help prevent getting stuck in super-long trajectories without degrading most normally advancing trajectories.

**Training Details** We perform Supervised Fine-Tuning (SFT) using the ms-swift library (Zhao et al., 2024) with Megatron (Shoeybi et al., 2019) on  $2\times$ H200 GPUs. By default, we train 2 epochs for all experiments with a global batch size of 16 and a maximum sequence length of 32,768 tokens. We use a learning rate of  $1e-5$  with cosine decay to  $1e-6$  and a linear warmup fraction of 0.03.

## 4 Evaluation and Results

To evaluate the effectiveness of RCL,

we ask the following four research questions:

1. What is the overall end-to-end performance of RCL?
2. What are the advantages of RCL over TCL?
3. What is the nature of RCL and is it replaceable by other optimization techniques?
4. What is the contribution of each unit of RCX?

### 4.1 RQ1: Overall Performance of RCL

Table 1 shows that SWE-SPOT-4B trained with RCL decisively breaks the scaling barriers observed in task-centric learning. Despite having only 4B parameters, SWE-SPOT-4B achieves an average execution-based score of 17.12%, outperforming open-weight models up to  $8\times$  larger with a clear margin, including Qwen3-Coder-30B-A3B-Instruct, Gemma-3-27B-it, and CWM-32B. Notably, SWE-SPOT-4B also surpasses or effectively matches efficiency-comparable commercial models such as GPT-4.1-mini, GPT-5-nano, and Gemini-2.5-Flash-Lite.

This advantage extends beyond issue resolution to several other software engineering tasks, where SWE-SPOT-4B continues to punch well above its weight. Comparisons with Mini-Coder-4B and CWM-32B are particularly revealing: although these TCL-trained models achieve comparable performance on issue resolution, they fail to generalize across tasks. On test generation, for example, Mini-Coder-4B collapses to 0.63%, and CWM trails SWE-SPOT-4B by more than 5%, despite extensive task-centric training.

Overall, the consistent performance of SWE-SPOT-4B across all evaluated tasks demonstrates that RCL enables SLMs to internalize repo-specific knowledge that transfers across software engineering tasks, whereas TCL tends to induce brittle specialization around an individual task.

Table 1. End-to-end Repository-Centric Evaluation on Four Agentic SWE Tasks. Exec Avg denotes the average pass rate across the three tasks with execution-based evaluation (Issue Resolving, Test Gen, Feature Impl). Trained by RCL, SWE-SPOT-4B achieves comparable or superior performance to open-weight models up to  $8\times$  its size and coding agents trained by TCL with significantly more data, and also surpasses or matches efficient commercial models.

Model	Size	SWE Training Data	Issue Resolving (Pass %)	Test Gen (Pass %)	Feature Impl (Pass %)	Exec Avg (Pass %)	Codebase QA (Score 0-100)
<i>Efficiency-focused Commercial LLM APIs</i>							
GPT-4.1-mini	[10, 100]B	-	21.79	22.27	5.70	<b>17.85</b>	<b>80.28</b>
Gemini-2.5-Flash-Lite	[0.1, 10]B	-	8.80	10.74	0.88	7.46	73.36
GPT-5-nano	[0.1, 10]B	-	6.49	0.00	1.75	2.97	35.64
<i>Open-Weight Generalist Models</i>							
Gemma-3-27b-it	27B	-	1.29	2.37	0.00	1.34	65.19
Qwen3-Coder-30B	32B	-	16.74	11.85	3.29	11.56	65.48
Qwen3-4B-Instruct-2507	4B	-	3.32	5.37	0.00	3.20	61.79
<i>Coding Agents Trained with TCL</i>							
CWM (Meta)	32B	3M trajs from 3.15K repos	<b>22.22</b>	17.38	4.17	15.88	73.09
Mini-Coder-4B	4B	400K trajs from 128 repos	18.76	0.63	4.61	8.70	57.30
<i>RCL (ours)</i>							
SWE-SPOT	<b>4B</b>	<b>8K</b> trajs in <b>one</b> target repo	19.34	<b>22.75</b>	<b>5.92</b>	17.12	78.05

### 4.2 RQ2: Advantages of RCL over TCL

To rigorously assess the benefits of RCL vs. repo-agnostic TCL, we conduct a controlled comparison between the two, where we use exactly the same setting to ensure fairness

#### Superior Performance under Controlled Data Budgets

We compare the performance of RCL against TCL on the three execution-based evaluated tasks: issue resolving (8k training samples each, evaluated on all 7 selected repos), test generation (2k training samples each, evaluated on Django), and feature implementation (2k training samples each, evaluated on Sympy). We synthesize TCL data by directly leveraging SWE-Smith (Yang et al., 2025) for issue resolution and test generation; for feature implementation, we use instances of FEA-Bench (Li et al., 2025c) in repositories other than the seven ones selected for our evaluation. Note that for the latter two tasks, training budget is reduced to 2k samples for each learning strategy and evaluation is restricted to the repository with the most instances in the corresponding benchmark due to limited computational budget.

As shown in Table 2, RCL consistently outperforms TCL across all evaluated tasks, suggesting that for SLMs with limited capacity deep mastery of the specific target environment is more valuable than broad but shallow exposure to diverse, unrelated codebases.

Table 2. Controlled comparison between RCL and TCL. RCL achieves consistently higher pass rates through repository mastery.

Task	Eval Scope	TCL	RCL
Issue Resolving	All Selected Repos	14.86%	<b>19.34%</b>
Test Generation	Django	4.09%	<b>11.01%</b>
Feature Implementation	Sympy	2.15%	<b>4.84%</b>




Figure 3. Controlled Performance Scaling with Respect to the Amount of Training Samples on Django Issue Resolution.

For both the pass rate and the NLL on teacher trajectories, RCL surpasses the peak performance of TCL using significantly less data. RCL uses 10k multi-dimensional mixed samples and trains for 2 epochs, while TCL uses 20k SWE-Smith issue resolving samples and trains for only one epoch.

#### Higher Sample Efficiency in Training

Beyond final performance, we evaluate intermediate checkpoints during training, in which RCL exhibits significantly higher sample efficiency, converging to optimal performance much faster than TCL. Figure 3 illustrates the scaling trends on Django issue resolution: RCL surpasses the peak performance of TCL using approximately  $2.5\times$  fewer training budget, despite its data consists of samples across different dimensions rather than issue resolving alone. This indicates that, by grounding the learning in the target repository’s specific logic with diverse experience, rather than trying to generalize from disparate wild environments with one task, RCL is both more effective and efficient than TCL.

#### Reduced Inference Cost

Another critical advantage of RCL is its efficiency. We compare the inference trajectories

**Table 3. Trajectory-length Distributions on Django Issue Resolving.** Values are reported as Mean  $\pm$  Standard Deviation. The RCL model is more inference-time efficient than TCL, despite trained by longer trajectories. Post-SFT models show increased length with the "looping" phenomenon.

Setting	Avg. Turns	Avg. Tokens (k)
<i>Inference Cost</i>		
Base Model (Qwen3-4B)	9.22 $\pm$ 9.82	7.25 $\pm$ 7.62
TCL Model (Baseline)	41.62 $\pm$ 19.07	33.19 $\pm$ 13.33
RCL Model (Ours)	<b>32.06</b> $\pm$ 19.16	<b>29.49</b> $\pm$ 14.69
<i>Training Data</i>		
TCL Data (Wild Repos)	33.70 $\pm$ 11.12	15.47 $\pm$ 8.13
RCL Data (Django)	38.43 $\pm$ 11.03	17.70 $\pm$ 8.70

of the base model, the RCL model, and the TCL model on the issue-resolving benchmark. As detailed in Table 3, the RCL model reduces the average number of turns by **23.0%** and total token consumption by **11.1%** compared to the TCL model. Counter-intuitively, this efficiency gain occurs even though the training trajectories of RCL are on average 12.6% longer by token count than the TCL ones. We attribute this to the qualitative difference in model behavior: because a RCL model possesses more knowledge about the repository, it requires less extensive exploration and trial-and-error at inference time. In contrast, the TCL model must perform costly search operations and more iterative corrections due to the lack of environment-specific knowledge. In addition, we observe a looping phenomenon across all SFT models that results in longer trajectories than the base model and the training data, which is detailed in Appendix C.

### 4.3 RQ3: The Necessity and Nature of RCL

Having demonstrated the performance and efficiency benefits of RCL, we now investigate the possible reasons behind these improvements. Specifically, we question whether RCL simply acts as a proxy for better retrieval, whether the acquired capability is shallow (style alignment) or deep (parametric knowledge), and if its benefits persist considering other optimizations like test-time scaling.

#### Parametric Knowledge Injection vs. Context Retrieval

A natural counter-hypothesis to our approach is that RCL results in better performance simply because it merely memorizes file locations, acting as an implicit retriever to alleviate the burden of codebase exploration. If this were true, a TCL model equipped with perfect contextual knowledge should match RCL. We test this by providing "Oracle Localization"—manually injecting the locations of ground-truth patched files and functions as strong contextual hints into the prompt—thereby removing the need of navigation and proxying the presence of an excellent retrieval system.

As shown in Table 4, oracle localization indeed improves the base and TCL model, while it provides no further gain for the RCL model, suggesting that RCL does implicitly

**Table 4. Comparison of Optimization Strategies on Django Issue Resolution. Oracle Loc.** adds ground-truth patching locations to the input. **LoRA** uses rank  $r = 128$ . The results show that RCL (Full FT) renders Oracle Localization redundant, whereas LoRA fails to match Full FT even with Oracle help, indicating the necessity of high learning capacity for deep knowledge acquisition.

Model	Configuration	Pass Rate (%)	
		Standard	+ Oracle Loc.
Base Model	Zero-shot	3.95	5.37
Task-Centric Learning	Full SFT	17.51	19.21
Repo-Centric Learning	LoRA ( $r = 128$ )	17.23	18.36
	Full SFT	<b>24.29</b>	23.45

improve localization through the effectively internalized repository structure. However, this is not the complete role of RCL, because the TCL model with oracle context (19.21%) still underperforms the RCL model without any localization hints (24.29%) by a large gap, indicating RCL introduces more complex optimizations, such as not only teaching the model *where* to look, but also *how* to reason about the task and modify code in compliance with the target repository’s logic, which is likely driven by *deep parametric* knowledge injection beyond simple memorization.

**Requirements of Sufficient Learning Capacity: Knowledge Acquisition vs. Style Alignment** We further probe the nature of RCL by examining whether full-parameter fine-tuning is necessary for RCL, or if parameter-efficient methods like LoRA (Hu et al., 2022) suffice.

Because if LoRA performs competitively, then it indicates RCL primarily adapts the model to the repository’s superficial coding styles or to the agentic scaffold’s response format, achievable with much lower learning capacity than deep knowledge acquisition.

However, as shown in Table 4, the LoRA-tuned RCL model (even with a high rank 128, updating  $\approx 5.38\%$  of full parameters) significantly lags behind the fully fine-tuned model (17.23% vs. 24.29%), dropping to a performance level similar to the TCL model. Notably, this limitation persists even when perfect localization is provided (18.36%), confirming that the performance gap is not due to navigational errors but a fundamental lack of capacity. This implies that the role of repo-agnostic TCL on SLMs is likely fitting shared task-specific formats, achievable by low-capacity learning, whereas RCL involves substantial experience and knowledge storage regarding the repository’s specifications, requiring robust representations from full-weight updates.

**Orthogonality to Test-Time Scaling** Lastly, we question about the necessity of RCL when other optimizations present, such as test-time scaling. We measure  $\text{pass}@k$  performance of the base, TCL, and RCL models on Django issue resolution. As illustrated in Figure 4, the RCL model consistently outperforms baseline models at every  $k$ . The




Figure 4. Test-time scaling on Django issue resolution. The RCL model consistently outperforms the baselines across all sample budgets, indicating that repository mastery enables foundational improvements orthogonal to test-time scaling.

Table 5. Leave-one-out Ablation on Multi-spectrum RCL, Evaluated on Django. Pass rates (%) and LLM-as-a-judge scores are reported accordingly. Full-unit RCL yields the best performance, demonstrating that removing any single repository-centric experience unit degrades performance across multiple downstream tasks.

Configuration	Issue Res.	Test Gen.	Feat. Impl.	Code QA
<b>Full-unit RCL</b>	<b>24.29</b>	<b>17.30</b>	<b>14.29</b>	<b>83.74</b>
– Software Design	22.32	14.15	14.29	75.77
– Contextual Impl.	20.62	16.38	4.76	79.44
– Evolutionary Replay	23.45	3.14	9.52	80.60
– Sem.-Rt. Alignment	16.10	15.72	14.29	78.58

gap does not diminish with more samples; rather, RCL effectively halves the inference cost of TCL for roughly equivalent performance.

This suggests that RCL provides additional, irreplaceable capabilities that are orthogonal to, and synergistic with, test-time scaling.

#### 4.4 RQ4: Benefits of Multi-Spectrum RCX

To validate the contribution of each unit in the repository-centric experience supporting RCL, we conduct a comprehensive leave-one-out ablation study on the Django repository, by comparing the full-spectrum RCL (2k samples for each of the four units) against variants omitting one unit.

As shown in Table 5, RCL with full-spectrum repository-centric experience consistently outperforms all ablated variants. Notably, removing any single experience unit can result in performance degradation not just in its “mirror” task, but also in others. This suggests that for SLMs, repository expertise is possibly not a collection of isolated skills but consists of synergistic representations that can affect each other, possibly leveraging “superposition” (Elhage et al., 2022) to learn with limited capacity.

## 5 Discussions, Limitations, and Future Work

**Inter-Repository Learning Dynamics of SLMs** While our main results establish that RCL outperforms TCL on a target repository, a critical question remains regarding the interaction between different codebases. To investigate

Table 6. Inter-Repository Learning Dynamics of SLMs on Issue Resolution. Repo-specific RCL aggregates the performance of three separate specialist models evaluated on their respective target repositories. Joint RCL represents a single model trained on all three repositories simultaneously. The co-occurrence of both Synergy and Interference confirms the necessity of RCL.

Model Strategy	Overall (%)	Pass Rate on Target Repo (%)		
		django	matplotlib	sympy
TCL	16.86	17.51	11.46	21.21
Django-specific RCL	19.57	24.29	9.38	9.09
Repo-specific RCL	<b>23.64</b>	24.29	<b>19.79</b>	25.76
Joint RCL	<b>23.64</b>	<b>25.99<sup>S</sup></b>	<b>11.46<sup>I</sup></b>	<b>28.79<sup>S</sup></b>

this, we conduct an ablation study comparing our single-repository experts against a multi-repository expert jointly trained on a cluster of three codebases: django, matplotlib, and sympy. The results, presented in Table 6, reveal that the relationship between repository diversity and agentic performance is complex and non-monotonic. First, the django-specific expert suffers severe performance degradation on non-targeted repositories (e.g., dropping to 9.4% on matplotlib), confirming that RCL effectively internalizes repo-specific “physics” that may not generalize across repositories.

Second, when training on the three repositories simultaneously, we observe divergent dynamics: while the joint model achieves synergy on django and sympy, outperforming their respective dedicated experts, it suffers from a significant negative transfer on matplotlib (dropping from 19.8% to 11.5%). This fluctuation suggests that blind scaling of repository diversity—the core premise of TCL—is not always optimal, as *knowledge and experience of distinct repositories can contradict or suppress one another during SLMs’ training, and repository diversity may dilute the highly-specialized knowledge that is critical for specific, high-value targets*. This further validates RCL as a necessary new dimension: in many production contexts, we prioritize maximizing performance within a specific and limited working scope over maintaining average performance across all codebases in the entire world. We leave a few directions inspired from these observations for future work, including when and how repository characteristics lead to synergy or interference, and the applicability to larger models.

**Extensions to the RCL Instantiation** Our implementation of RCL is a SLM student learning from teacher-produced RCX through SFT, due to the limited compute budget. However, RCL can be implemented with more advanced training techniques, such as on-policy distillation (Lu & Lab, 2025), or RLVR which allows the agent to directly learn from self-produced RCX. Besides, recent advances on new model architectures that better support continual learning (Behrouz et al., 2025; Tandon et al., 2025) may help to reduce the cost of RCL to unlock more efficient

expert specialization, unifying training and inference to enable agents that keep evolving in targeted environments.

## 6 Conclusion

In this work, we introduced Repository-Centric Learning to advance efficient coding agents, a paradigm shift designed to bridge the inference-time capability gap in Small Language Models by parametrically internalizing the dynamics of target codebases. Using this approach, our SWE-SPOT-4B model breaks the established scaling barriers, demonstrating that repository mastery is a distinct, necessary new dimension that complements general coding capability.

## Impact Statements

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

## References

- Ahmed, M., Dorrah, M., Ashraf, A., Adel, Y., Elatrozy, A., Mohamed, B. E., and Gomaa, W. Codeqa: Advanced programming question-answering using llm agent and rag. In *2024 6th Novel Intelligent and Leading Emerging Sciences Conference (NILES)*, pp. 494–499. IEEE, 2024.
- Ahmed, T., Ganhotra, J., Pan, R., Shinnar, A., Sinha, S., and Hirzel, M. Otter: Generating tests from issues to validate SWE patches. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 752–771. PMLR, 13–19 Jul 2025a. URL <https://proceedings.mlr.press/v267/ahmed25b.html>.
- Ahmed, T., Ganhotra, J., Shinnar, A., and Hirzel, M. Heterogeneous prompting and execution feedback for swe issue test generation and selection. *arXiv preprint arXiv:2508.06365*, 2025b.
- Aleithan, R., Xue, H., Mohajer, M. M., Nnorom, E., Uddin, G., and Wang, S. Swe-bench+: Enhanced coding benchmark for llms. *arXiv preprint arXiv:2410.06992*, 2024.
- Amazon. Kiro: Agentic ai development from prototype to production, 2025. URL <https://kiro.dev>.
- Anthropic. Introducing claude sonnet 4.5, 2025a. URL <https://www.anthropic.com/news/claude-sonnet-4-5>.
- Anthropic. Claude code. GitHub, 2025b. URL <https://github.com/anthropics/claude-code>.
- Apple. Apple intelligence foundation language models: Tech report 2025, 2025. URL <https://arxiv.org/abs/2507.13575>.
- Behrouz, A., Razaviyayn, M., Zhong, P., and Mirrokni, V. Nested learning: The illusion of deep learning architectures, 2025. URL <https://arxiv.org/abs/2512.24695>.
- Chen, J., Zhao, K., Liu, J., Peng, C., Liu, J., Zhu, H., Gao, P., Yang, P., and Deng, S. Coreqa: uncovering potentials of language models in code repository question answering. *arXiv preprint arXiv:2501.03447*, 2025a.
- Chen, S., Lin, S., Gu, X., Shi, Y., Lian, H., Yun, L., Chen, D., Sun, W., Cao, L., and Wang, Q. Swe-exp: Experience-driven software issue resolution. *arXiv preprint arXiv:2507.23361*, 2025b.
- Chowdhury, N., Aung, J., Shern, C. J., Jaffe, O., Sherburn, D., Starace, G., Mays, E., Dias, R., Aljubeih, M., Glaese, M., Jimenez, C. E., Yang, J., Ho, L., Patwardhan, T., Liu, K., and Madry, A. Introducing swe-bench verified — openai. URL <https://openai.com/index/introducing-swe-bench-verified/>. [Online; accessed 2026-01-27].
- Cursor. Cursor agent, 2025a. URL <https://cursor.com/agents>.
- Cursor. Composer: Building a fast frontier model with rl, 2025b. URL <https://cursor.com/blog/composer>.
- Deng, X., Da, J., Pan, E., He, Y. Y., Ide, C., Garg, K., Lauffer, N., Park, A., Pasari, N., Rane, C., Sampath, K., Krishnan, M., Kundurthy, S., Hendryx, S., Wang, Z., Bharadwaj, V., Holm, J., Aluri, R., Zhang, C. B. C., Jacobson, N., Liu, B., and Kenstler, B. Swe-bench pro: Can ai agents solve long-horizon software engineering tasks?, 2025. URL <https://arxiv.org/abs/2509.16941>.
- Elhage, N., Hume, T., Olsson, C., Schiefer, N., Henighan, T., Kravec, S., Hatfield-Dodds, Z., Lasenby, R., Drain, D., Chen, C., Grosse, R., McCandlish, S., Kaplan, J., Amodei, D., Wattenberg, M., and Olah, C. Toy models of superposition. *Transformer Circuits Thread*, 2022. [https://transformer-circuits.pub/2022/toy\\_model/index.html](https://transformer-circuits.pub/2022/toy_model/index.html).
- Gao, X., Noller, Y., and Roychoudhury, A. Program repair, 2022. URL <https://arxiv.org/abs/2211.12787>.

- GitHub. Github copilot - your ai pair programmer, 2025. URL <https://github.com/features/copilot>.
- Google. Gemini cli. GitHub, 2025. URL <https://github.com/google-gemini/gemini-cli>.
- Google DeepMind. Gemini 3 pro: Best for complex tasks and bringing creative concepts to life, 2025. URL <https://deepmind.google/models/gemini/pro/>.
- Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al. Deepseek-rl incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025.
- He, Z., Yang, Q., Sheng, W., Zhong, X., Zhang, K., An, C., Shi, W., Cai, T., He, D., Chen, J., and Xu, J. Swe-swiss: A multi-task fine-tuning and rl recipe for high-performance issue resolution. <https://www.notion.so/SWE-Swiss-A-Multi-Task-Fine-Tuning-and-RL-Recipe-for-High-Performance-Issue-Resolution-21e174dedd4880ea829ed4c861c44f88>, 2025. Notion Blog.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- Hu, R., Peng, C., Ren, J., Jiang, B., Meng, X., Wu, Q., Gao, P., Wang, X., and Gao, C. Coderepoqa: A large-scale benchmark for software engineering question answering. *arXiv preprint arXiv:2412.14764*, 2024.
- Hu, R., Peng, C., Ren, J., Jiang, B., Meng, X., Wu, Q., Gao, P., Wang, X., and Gao, C. Understanding large language model performance in software engineering: A large-scale question answering benchmark. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3025–3029, 2025.
- Huang, D., Zhang, J. M., Harman, M., Zhang, Q., Du, M., and Ng, S.-K. Benchmarking llms for unit test generation from real-world functions. *arXiv preprint arXiv:2508.00408*, 2025.
- Jain, K. and Goues, C. L. Testforge: Feedback-driven, agentic test suite generation. *arXiv preprint arXiv:2503.14713*, 2025.
- Jimenez, C. E., Yang, J., Wettig, A., Yao, S., Pei, K., Press, O., and Narasimhan, K. R. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=VTF8yNQM66>.
- Kang, S., Yoon, J., and Yoo, S. Large language models are few-shot testers: Exploring llm-based general bug reproduction. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pp. 2312–2323. IEEE, 2023.
- Khatib, L., Mathews, N. S., and Nagappan, M. Assertflip: Reproducing bugs via inversion of llm-generated passing tests. *arXiv preprint arXiv:2507.17542*, 2025.
- Li, F., Jiang, J., Sun, J., and Zhang, H. Hybrid automated program repair by combining large language models and program analysis. *ACM Transactions on Software Engineering and Methodology*, 2024a.
- Li, H., Shi, Y., Lin, S., Gu, X., Lian, H., Wang, X., Jia, Y., Huang, T., and Wang, Q. Swe-debate: Competitive multi-agent debate for software issue resolution. *arXiv preprint arXiv:2507.23348*, 2025a.
- Li, H., Tang, Y., Wang, S., and Guo, W. Patchpilot: A cost-efficient software engineering agent with early attempts on formal verification. In *Forty-second International Conference on Machine Learning*, 2025b. URL <https://openreview.net/forum?id=ybODpT8ydV>.
- Li, W., Zhang, X., Guo, Z., Mao, S., Luo, W., Peng, G., Huang, Y., Wang, H., and Li, S. FEA-bench: A benchmark for evaluating repository-level code generation for feature implementation. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 17160–17176, Vienna, Austria, July 2025c. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.839. URL <https://aclanthology.org/2025.acl-long.839/>.
- Li, Z., Zhang, J., Yin, C., Ouyang, Y., and Rong, W. Procqa: a large-scale community-based programming question answering dataset for code search. *arXiv preprint arXiv:2403.16702*, 2024b.
- Lin, J., Berges, V.-P., Chen, X., Yih, W.-T., Ghosh, G., and Oğuz, B. Learning facts at scale with active reading, 2025. URL <https://arxiv.org/abs/2508.09494>.
- Liu, C. and Wan, X. Codeqa: A question answering dataset for source code comprehension. *arXiv preprint arXiv:2109.08365*, 2021.
- Lu, K. and Lab, T. M. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. <https://thinkingmachines.ai/blog/on-policy-distillation>.

- Luo, M., Jain, N., Singh, J., Tan, S., Patel, A., Wu, Q., Ariyak, A., Cai, C., Tarun Venkat, S. Z., Athiwaratkun, B., Roongta, M., Zhang, C., Li, L. E., Popa, R. A., Sen, K., and Stoica, I. DeepSWE: Training a state-of-the-art coding agent from scratch by scaling rl. <https://pretty-radio-b75.notion.site/DeepSWE-Training-a-Fully-Open-sourced-State-of-the-Art-Coding-Agent-by-Scaling-RL-22281902c1468193aabb9a8c59bbe33>, 2025. Notion Blog.
- Lv, W., Xia, X., and Huang, S.-J. Codeact: Code adaptive compute-efficient tuning framework for code llms. *arXiv preprint arXiv:2408.02193*, 2024.
- Ma, Y., Cao, R., Cao, Y., Zhang, Y., Chen, J., Liu, Y., Liu, Y., Li, B., Huang, F., and Li, Y. Lingma swe-gpt: An open development-process-centric language model for automated software improvement. *arXiv preprint arXiv:2411.00622*, 2024.
- Ma, Z., Peng, C., Gao, P., Meng, X., Zou, Y., and Xie, B. Sorft: Issue resolving with subtask-oriented reinforced fine-tuning. *arXiv preprint arXiv:2502.20127*, 2025.
- Meta FAIR CodeGen. Cwm: An open-weights llm for research on code generation with world models, 2025. URL <https://arxiv.org/abs/2510.02387>.
- Mini-SWE-agent. mini-swe-agent: The 100 line ai agent that solves github issues. <https://github.com/SWE-agent/mini-swe-agent>, 2025.
- Mündler, N., Müller, M., He, J., and Vechev, M. Swt-bench: Testing and validating real-world bug-fixes with code agents. *Advances in Neural Information Processing Systems*, 37:81857–81887, 2024.
- Nagarajan, V., Wu, C. H., Ding, C., and Raghunathan, A. Roll the dice & look before you leap: Going beyond the creative limits of next-token prediction. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=Hi0SyHMmkd>.
- Nan, Z., Guo, Z., Liu, K., and Xia, X. Test intention guided llm-based unit test generation. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*, pp. 779–779. IEEE Computer Society, 2025.
- Nashid, N., Bouzenia, I., Pradel, M., and Mesbah, A. Issue2test: Generating reproducing test cases from issue reports. *arXiv preprint arXiv:2503.16320*, 2025.
- Olmedo, R. mini-coder, September 2025. URL <https://ricardodominguez.github.io/blogs/minicoder.html>. [Online; accessed 2026-01-27].
- OpenAI. Gpt-5 is here, 2025. URL <https://openai.com/gpt-5/>.
- Pabba, A., Chen, S., Mathai, A., Chakraborty, A., and Ray, B. Refine: Enhancing program repair agents through context-aware patch refinement. *arXiv preprint arXiv:2510.03588*, 2025.
- Pan, J., Wang, X., Neubig, G., Jaitly, N., Ji, H., Suhr, A., and Zhang, Y. Training software engineering agents and verifiers with swe-gym. *arXiv preprint arXiv:2412.21139*, 2024.
- Peng, W., Shi, Y., Wang, Y., Zhang, X., Shen, B., and Gu, X. Swe-qa: Can language models answer repository-level code questions? *arXiv preprint arXiv:2509.14635*, 2025.
- Pipis, C., Garg, S., Kontonis, V., Shrivastava, V., Krishnamurthy, A., and Papailiopoulos, D. Wait, wait, wait... why do reasoning models loop?, 2025. URL <https://arxiv.org/abs/2512.12895>.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- Silver, D. and Sutton, R. S. Welcome to the era of experience. *Google AI*, 1, 2025.
- Tandon, A., Dalal, K., Li, X., Kocejka, D., Rød, M., Buchanan, S., Wang, X., Leskovec, J., Koyejo, S., Hashimoto, T., Guestrin, C., McCaleb, J., Choi, Y., and Sun, Y. End-to-end test-time training for long context, 2025. URL <https://arxiv.org/abs/2512.23675>.
- Tang, X., Kim, K., Song, Y., Lothritz, C., Li, B., Ezzini, S., Tian, H., Klein, J., and Bissyandé, T. F. Codeagent: Autonomous communicative agents for code review. *arXiv preprint arXiv:2402.02172*, 2024.
- Tang, Y., Li, H., Zhu, K., Yang, M., Ding, Y., and Guo, W. Co-patcher: Collaborative software patching with component-specific small reasoning models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=zhFEO67s5w>.
- Team, Q. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- Wang, L., Ramalho, L., Celestino, A., Pham, P. A., Liu, Y., Sinha, U. K., Portillo, A., Osunwa, O., and Maduekwe, G. Swe-bench++: A framework for the scalable generation of software engineering benchmarks from open-source repositories. *arXiv preprint arXiv:2512.17419*, 2025.

- Wang, S., Li, Z., Qian, H., Yang, C., Wang, Z., Shang, M., Kumar, V., Tan, S., Ray, B., Bhatia, P., Nallapati, R., Ramanathan, M. K., Roth, D., and Xiang, B. Re-Code: Robustness evaluation of code generation models. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13818–13843, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.773. URL <https://aclanthology.org/2023.acl-long.773/>.
- Wang, X., Li, B., Song, Y., Xu, F. F., Tang, X., Zhuge, M., Pan, J., Song, Y., Li, B., Singh, J., et al. Openhands: An open platform for ai software developers as generalist agents. In *The Thirteenth International Conference on Learning Representations*, 2024a.
- Wang, Z., Liu, K., Li, G., and Jin, Z. Hits: High-coverage llm-based unit test generation via method slicing. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, pp. 1258–1268, 2024b.
- Wei, Y., Duchenne, O., Copet, J., Carbonneaux, Q., ZHANG, L., Fried, D., Synnaeve, G., Singh, R., and Wang, S. SWE-RL: Advancing LLM reasoning via reinforcement learning on open software evolution. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=ULbl061XZ0>.
- Xia, C. S., Deng, Y., Dunn, S., and Zhang, L. Agentless: Demystifying llm-based software engineering agents. *arXiv preprint arXiv:2407.01489*, 2024.
- Xie, C., Li, B., Gao, C., Du, H., Lam, W., Zou, D., and Chen, K. SWE-fixer: Training open-source LLMs for effective and efficient GitHub issue resolution. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 1123–1139, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.62. URL <https://aclanthology.org/2025.findings-acl.62/>.
- Xue, Z., Li, L., Tian, S., Chen, X., Li, P., Chen, L., Jiang, T., and Zhang, M. Llm4fin: Fully automating llm-powered test case generation for fintech software acceptance testing. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 1643–1655, 2024.
- Yang, J., Jimenez, C. E., Wettig, A., Lieret, K., Yao, S., Narasimhan, K. R., and Press, O. SWE-agent: Agent-computer interfaces enable automated software engineering. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a. URL <https://openreview.net/forum?id=mXpq6ut8J3>.
- Yang, J., Lieret, K., Jimenez, C. E., Wettig, A., Khandpur, K., Zhang, Y., Hui, B., Press, O., Schmidt, L., and Yang, D. SWE-smith: Scaling data for software engineering agents. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2025. URL <https://openreview.net/forum?id=63iVrXc8cC>.
- Yang, L., Yang, C., Gao, S., Wang, W., Wang, B., Zhu, Q., Chu, X., Zhou, J., Liang, G., Wang, Q., et al. On the evaluation of large language models in unit test generation. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, pp. 1607–1619, 2024b.
- Zan, D., Huang, Z., Liu, W., Chen, H., Zhang, L., Xin, S., Chen, L., Liu, Q., Zhong, X., Li, A., Liu, S., Xiao, Y., Chen, L., Zhang, Y., Su, J., Liu, T., Long, R., Shen, K., and Xiang, L. Multi-swe-bench: A multilingual benchmark for issue resolving, 2025a. URL <https://arxiv.org/abs/2504.02605>.
- Zan, D., Huang, Z., Liu, W., Chen, H., Zhang, L., Xin, S., Chen, L., Liu, Q., Zhong, X., Li, A., et al. Multi-swe-bench: A multilingual benchmark for issue resolving. *arXiv preprint arXiv:2504.02605*, 2025b.
- Zhang, L., He, S., Zhang, C., Kang, Y., Li, B., Xie, C., Wang, J., Wang, M., Huang, Y., Fu, S., et al. Swe-bench goes live! *arXiv preprint arXiv:2505.23419*, 2025.
- Zhang, Y., Ruan, H., Fan, Z., and Roychoudhury, A. Autocoderover: Autonomous program improvement, 2024a. URL <https://arxiv.org/abs/2404.05427>.
- Zhang, Z., Liu, X., Lin, Y., Gao, X., Sun, H., and Yuan, Y. Llm-based unit test generation via property retrieval. *arXiv preprint arXiv:2410.13542*, 2024b.
- Zhao, Y., Huang, J., Hu, J., Wang, X., Mao, Y., Zhang, D., Jiang, Z., Wu, Z., Ai, B., Wang, A., Zhou, W., and Chen, Y. Swift: a scalable lightweight infrastructure for fine-tuning, 2024. URL <https://arxiv.org/abs/2408.05517>.

## Appendix

### A Related Work

**Software Engineering Tasks and Benchmarks.** Large language model (LLM) agents are increasingly integrated into a wide range of software engineering (SE) tasks. In this work, we focus on three areas most relevant to our study and introduce representative benchmarks: *program repair* (Yang et al., 2024a; Gao et al., 2022; Li et al., 2024a; Zhang et al., 2025; Zan et al., 2025b; Deng et al., 2025; Aleithan et al., 2024; Wang et al., 2025), *test generation* (Mündler et al., 2024; Huang et al., 2025; Nashid et al., 2025; Huang et al., 2025; Jain & Goues, 2025), and *code question answering* (Liu & Wan, 2021; Hu et al., 2024; Li et al., 2024b; Chen et al., 2025a; Hu et al., 2025). *Program repair* has emerged as a prominent application area, in which an LLM agent analyzes a code repository together with a natural-language issue report, identifies the underlying defects, and automatically proposes candidate patches. These fixes are expected to preserve intended program behavior and pass hidden test suites. A notable benchmark in this space is SWE-BENCH (Yang et al., 2024a), which formalizes real-world, repository-level repair tasks, later MULT-SWE-BENCH (Zan et al., 2025b) extends SWE-BENCH to multilingual settings. *Automated test generation* is another key application, where LLM agents synthesize context-aware and comprehensive test cases to improve coverage and enhance validation efficiency (Zhang et al., 2024b; Nan et al., 2025; Wang et al., 2024b; Xue et al., 2024; Yang et al., 2024b). A representative benchmark is SWT-BENCH (Mündler et al., 2024), a test-generation counterpart derived from SWE-BENCH. Finally, *code question answering* supports interactive development workflows by explaining code functionality, identifying potential problems, recommending improvements, and enforcing project-specific conventions. For example, CODEQA (Liu & Wan, 2021) provides free-form code comprehension questions, while SWE-QA (Peng et al., 2025) extends beyond CODEQA to operate in realistic, repository-scale environments.

**Software Engineering Agents** To address these real-world SE tasks, a growing number of LLM-based software engineering agents have been proposed. These agents typically rely on general-purpose LLMs as backend models and build specialized workflows or tool-use strategies to tackle specific SE problems. For *program repair agents* (Zhang et al., 2024a; Lv et al., 2024; Wang et al., 2024a; Li et al., 2025a), representative examples include SWE-AGENT (Yang et al., 2024a) and MINI-SWE-AGENT (Mini-SWE-agent, 2025), which adopt a “reason-then-act” paradigm with tool invocation, and AGENTLESS (Xia et al., 2024), which follows a static, predefined workflow rather than dynamic reasoning. However, these agents may produce overfitted patches during repair. Building on them, REFINE (Pabba et al., 2025) introduces a general enhancement agent designed to mitigate overfitting and strengthen existing program repair pipelines. For *test generation agents* (Khatib et al., 2025; Kang et al., 2023; Ahmed et al., 2025b), some approaches adapt program-repair workflows by modifying prompts and reusing their task structure. Beyond these, E-OTTER++ and ISSUE2TEST (Nashid et al., 2025) leverage execution feedback to iteratively produce higher-quality test cases. For *code question answering agents*, SWE-QA-Agent performs reasoning and tool-based actions to automatically locate answers, CODEAGENT (Tang et al., 2024) introduces autonomous communicative agents for code review, and CODEQA-AGENT (Ahmed et al., 2024) combines retrieval-augmented generation (RAG) with LLMs to further improve response accuracy.

**Training Specified Models for SE Agents** Although existing agentic workflows demonstrate strong effectiveness in solving SE tasks, most rely on proprietary, commercial LLMs, such as Claude-Sonnet or GPT-4.1, as their backbone. This dependence introduces high computational cost and raises concerns around scalability, accessibility, and on-device deployment. To alleviate these limitations, recent work has begun developing smaller, specialized models to replace heavyweight commercial LLMs (Meta FAIR CodeGen, 2025; Yang et al., 2024a; Ma et al., 2024; Pan et al., 2024; Xie et al., 2025; Ma et al., 2025). For example, CO-PATCHER (Tang et al., 2025) trains three lightweight models dedicated to bug localization, patch generation, and patch validation for program repair. SWE-EXP (Chen et al., 2025b) distills concise, transferable experience from prior agent trajectories, enabling continuous learning across issue-solving episodes. Similarly, SWE-MIRROR (Wang et al., 2025) extracts the semantic core of real-world issues, mirrors them into a different repository configured with a Gym-style environment, and replays them as verifiable issue-resolution tasks.

### B Details on Repository-Centric Experience Collection

Here, we provide more implementation details of our collection of Repository-Centric Experience (RCX).

#### B.1 Active Code Analysis for Software Design

We sample modules, files, and code chunks in a target repository as objects at varying granularity for active code analysis, sharing the underlying idea with seed-conditioned data synthesis (Nagarajan et al., 2025). We prioritize objects with richer

commit history, as these hotter regions are more likely to be central to the repository’s evolution and relevant to future tasks.

## B.2 Agentic FIM for Contextual Implementation

To operationalize the Contextual Implementation unit of RCX, we transform the traditional Fill-in-the-Middle (FIM) objective into an agentic task. Unlike standard FIM, which relies on passive context window stuffing, our Agentic FIM requires the model to actively explore the repository to discover cross-file dependencies.

We utilize static analysis to assist in constructing instances that require repository-level reasoning:

- **AST-Aligned Hole Sampling:** We parse source files into Abstract Syntax Trees (AST) to select candidate holes aligned with syntactic structures (complete functions or classes) rather than arbitrary text spans.
- **Dependency Verification:** To ensure the task requires cross-file reasoning, we employ a Language Server Protocol (LSP) backend (e.g., Pyright) to resolve symbols within the candidate hole to their definitions in other files.
- **Hole Classification & Weighted Sampling:** Candidates are classified as Positive (containing resolved cross-file calls) or Negative (local logic only). To maximize learning efficiency, we heavily downsample negative holes and prioritize positive holes that cover a diverse set of repository dependencies.

## B.3 PR Mirror for Evolutionary Replay

To learn from the software evolution of a target repository, we rely on its history of pull requests (PRs) to create issue-resolving tasks which encourage the agent to learn from previous pitfalls in the repository. We adapt the PR Mirror bug generation method from Yang et al. (2025), whereby changes from previous PRs are undone by an agent in the current (or target) commit of the repository. These bugs introduced by the agent are then validated against a set of tests from the repository’s test suite. While Yang et al. (2025) are cautious in their PR selection, our method strays with the goal of maximizing number of PRs used for bug creation, optimizing gain from this limited but vital resource.

Specifically, Yang et al. (2025) require PRs used for bug generation to have a linked issue with a problem statement, a test patch, and strict limits on number of lines changed. On the other hand, our method utilizes LLM agents for problem statement generation for PRs with no linked issue text (which is the majority); a deterministic test finding algorithm to find the relevant subset of the test suite to validate the bug against instead of relying on a test patch; and have loose restrictions on number of lines changed, leveraging the larger context windows of modern models. As a result, our method results in a roughly *18-fold* increase in task instances per repository as compared with the original PR Mirror method with little to no degradation in task quality.

## B.4 Test Reproduction for Semantic-Runtime Alignment

As mentioned, we use the same bugs generated by our adapted PR Mirror method for test generation trajectories. The difference being that the prompt used to feed in the task instance to the agent directly explains the workflow for writing reproduction tests. The agent is instructed to either edit existing testing files in the repository or place new ones within the existing testing framework. This way the agent must explore and understand the testing framework for the repository in order to succeed in its task.

## C The “Looping” Phenomenon

We observe a discrepancy between training data and inference behavior across all SFT models. As shown in Table 3, both RCL and TCL models produce significantly longer inference-time trajectories than the base model *and* their respective training data. Qualitative analysis reveals that SFT models get stuck more frequently in repetition loops (e.g., repeatedly issuing the same action without meaningful progress)—a pathology not present in the base model or the training data. This phenomenon aligns with recent findings by Pipis et al. (2025) regarding loops in reasoning models, and we leave addressing the root cause of this issue as an open challenge for future work.