IQuest

From Repositories to Training Environments: Scaling Long-Horizon Coding Agents with Environment Evolution Loop

Yang Tian[2,3] **Tianyu Zheng**[3,*] **Xiyuan Yang**[2,3] **Yaxin Du**[2,3]
Chengran Yang[3] **Bo Zhao**[2] **Jian Yang**[1,†] **Bryan Dai**[3]

[1]Beihang University [2]School of AI, Shanghai Jiao Tong University [3]IQuest Research

[*] *Project Leader* [†] *Corresponding Author*

GitHub Hugging Face

arXiv:2611.21152v2 [cs.CL] 17 Nov 2026

Abstract

Training coding agents for long-horizon tasks requires environments with verifiable feedback. Existing pipelines often follow a linear process that ends with validation, establishing usability but providing little guidance on improving training value. Recursive self-improvement (RSI) often requires model rollouts or model training followed by downstream evaluation, with costs and feedback delays amplified on long-horizon tasks. How can we use timely, reliable proxy signals to improve training value and synthesize environments at scale? We present **RepoFactory**, an automated pipeline that converts real repositories into verifiable environments where agents build complete projects from scratch given a language specification. Its Environment Evolution Loop optimizes a multiplicative objective combining validity, structural coupling, and implementation scope through proxy feedback. RepoFactory produces 10,159 environments across 17 software domains. Training Qwen3.5-35B-A3B on these environments yields RepoHorizen-35B-A3B, which achieves gains of 22.6 percentage points on NL2Repo-Bench and 8.5 points on average across five benchmarks covering long-horizon coding and agentic tasks in real-world production settings, compared with 17.5 and 6.4 points without environment evolution. These results highlight the training value of RepoFactory environments and the additional gains from the Environment Evolution Loop.

1 Introduction

Large language models (LLMs) are increasingly being developed into agents that can carry out consequential work in real-world environments (Team et al., 2026a; Xu et al., 2026; Team, 2026; Yang et al., 2026a; DeepSeek-AI et al.). Unlike static language tasks, agentic work unfolds through sustained interaction with environments. This transition changes the objective of model training: a capable agent must learn not only to produce correct answers, but also to make coherent progress toward completing long-horizon tasks (Li et al., 2026a; Tian et al., 2026; Bajoria et al., 2026).

Software engineering has become a major application of coding agents. Pipelines built around bug fixing and feature development have provided verifiable tasks for training agents on real codebases (Pan et al., 2025; Yang et al., 2025). A successful local change, however, is only part of the work involved in delivering software. Coding agents are increasingly expected to handle more complex development processes and produce complete, working systems (Yang et al., 2026c; Team et al., 2026b). Training on short-horizon tasks alone cannot cover this broader process, making it increasingly important to build repository-level environments for agent training.

Constructing environments for long-horizon tasks at the repository level remains challenging. **(i) Rigid construction pipelines limit the usable repository pool.** DeNovoSWE (Zhao et al., 2026) and MindForge (Chen et al., 2026) follow linear construction pipelines. They refine documentation or resolve build failures while keeping the reference implementation fixed. Repositories that remain unsuitable are discarded, even when further revision could make them usable. Scaling therefore depends on finding more repositories that meet the same fixed requirements. **(ii) Feedback focused on reliability overlooks training value.** RefactorForge (Ebrahimi et al., 2026) and TrajLong (Peng et al., 2026a) design data for specific capabilities but do not explore how environment feedback can improve task content. ProgramGym (Yang et al., 2026b) uses execution diagnostics and coverage feedback to repair builds and strengthen tests around a fixed reference implementation. These signals improve reliability without explicitly optimizing training value, which may remain limited when a task requires little imple-

mentation work or coordination across components. **(iii) Feedback from model execution is costly and time-consuming.** BenchEvolver (Wu et al., 2026) and SETA (Shen et al., 2026) use model performance to guide task revision. Obtaining this feedback requires model rollouts on the candidate tasks. For long-horizon coding tasks, each rollout can involve extensive implementation, increasing computational costs and delaying subsequent revisions. How to obtain timely, reliable feedback at low cost to improve training value and synthesize long-horizon coding environments at scale remains an open research question.

We present RepoFactory, an automated pipeline that transforms real-world repositories into verifiable environments. We design these environments around full repository generation task, in which an agent must build a complete, functioning repository from scratch given a language specification. We define a multiplicative objective combining validity, structural coupling, and implementation scope. It provides timely feedback at low cost without requiring a model to solve each candidate task. Our Environment Evolution Loop uses this objective to guide joint revisions to the oracle implementation, task specification, and tests. Validity feedback helps the loop repair initially unsuitable environments and recover repositories that would otherwise be discarded. Structural coupling and implementation scope serve as proxies for training value, directing revisions toward the implementation and integration work required by the task. A construction agent chooses how to modify each candidate and receives updated measurements after each revision, iterating within a fixed construction budget. Our main contributions are threefold:

- We introduce RepoFactory, an automated pipeline that converts real-world projects into 10,159 verifiable full-repository generation environments across 17 software domains.
- We develop an Environment Evolution Loop that optimizes a multiplicative objective combining validity, structural coupling, and implementation scope. Timely, low-cost proxy feedback guides improvements to both environment reliability and training value.
- Training Qwen3.5-35B-A3B on RepoFactory yields RepoHorizen-35B-A3B, with gains of 22.6 percentage points on NL2Repo-Bench and 8.5 points on average across five coding and agentic benchmarks. Environment evolution contributes 5.1 and 2.1 points to these gains, respectively.

2 Related Work

2.1 Scaling Verifiable Environments

Environment synthesis scales agent training by generating interactive tasks with executable verification (Kulikov et al., 2026; Ren et al., 2026; Zhu et al., 2025). EnvScaler (Song et al., 2026) and ScaleEnv (Tu et al., 2026) target tool use, while Endless Terminals (Gandhi et al., 2026), LiteCoder-Terminal (Peng et al., 2026b), and CLI-Universe (Hua et al., 2026) extend this approach to terminal-agent training. Broader task coverage, however, does not ensure that tasks preserve the interdependent requirements of real software or demand sustained development. RepoFactory derives full-repository generation tasks from real projects, grounding task behavior in existing implementations and tests. Reconstructing an entire repository requires sustained implementation and integration across these interdependent components.

2.2 Recursive Synthesis and Agent Self-Improvement

Self-training reuses successful solutions as supervision, while self-play generates new problems for agents to solve (Chen et al., 2024; Li et al., 2026b). BenchEvolver (Wu et al., 2026) evolves reference solutions to construct harder tasks, while SETA (Shen et al., 2026) adapts terminal environments using learner performance. DataMaster (Du et al., 2026) optimizes training data using downstream evaluation scores, adding model training to the feedback cycle. These approaches rely on solver rollouts or downstream training, increasing computational costs and delaying feedback. To our knowledge, RepoFactory is the first to evolve full-repository generation environments using an objective that combines validity checks with inexpensive code metrics. This objective provides timely feedback without requiring solver rollouts or downstream training.

3 RepoFactory

RepoFactory converts real-world repositories into verifiable environments where agents reconstruct complete projects from a language specifications. The Environment Evolution Loop uses timely, low-cost feedback from verification and code analysis

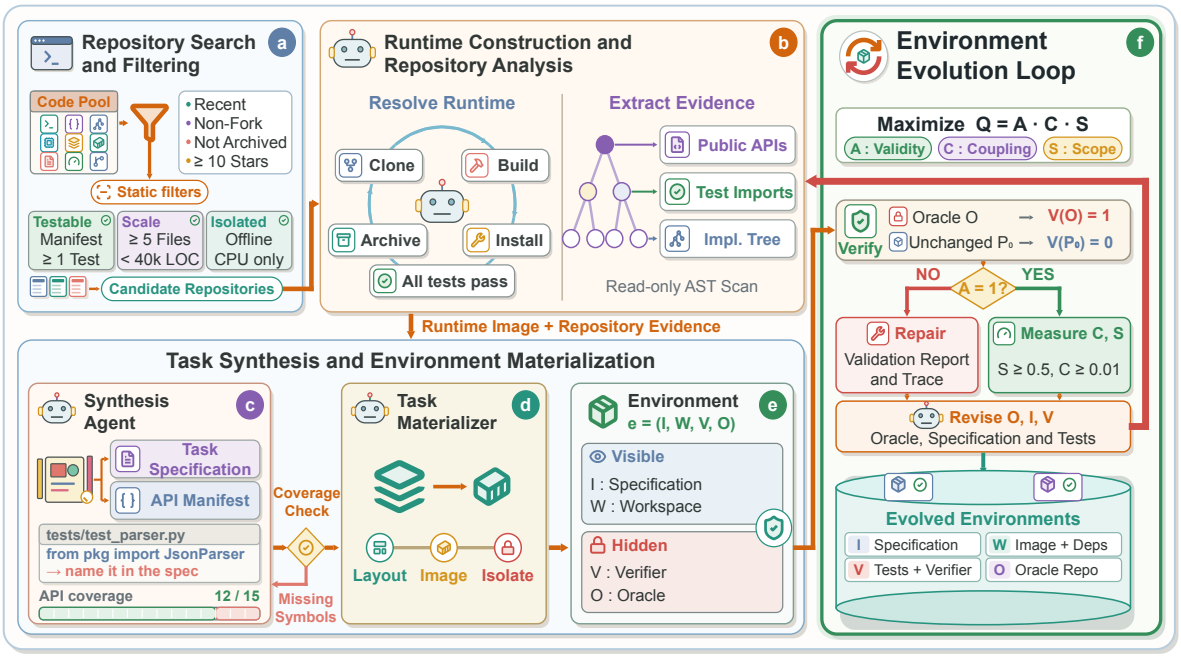


Figure 1: Overview of RepoFactory. The pipeline searches for repositories, reconstructs their runtimes, and synthesizes task specifications from repository evidence. The Environment Evolution Loop uses low-cost feedback to improve environment reliability and training value.

to jointly revise oracle implementations, task specifications, and tests, optimizing a multiplicative objective over validity, structural coupling, and implementation scope. All construction agents use Qwen3.8-Max. Figure 1 shows the overall pipeline.

3.1 Problem Formulation and Objective

Our goal is to transform a candidate repository r into a training environment $e = (I, W, V, O)$, where I is the language specification, W is an isolated workspace and runtime, V is the verifier, and O is the source repository's oracle implementation. During training, an agent π receives (I, W) and reconstructs a project $P = \pi(I, W)$, which is scored by $V(P) \in [0, 1]$.

We define the environment objective as

$$Q = A \cdot C \cdot S, \tag{1}$$

where A, C, and S measure validity, structural coupling, and implementation scope, respectively. The validity term requires full reward for the oracle and zero reward for the unchanged workspace P_0:

$$A = \begin{cases} 1, & V(O) = 1 \text{ and } V(P_0) = 0, \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

We measure structural coupling with normalized propagation cost:

$$C = \frac{R}{N^2}, \tag{3}$$

where N is the number of functions in the analyzed oracle code and R is the number of reachable ordered function pairs in its static internal call graph, including self-pairs.

Implementation scope combines code size and dependency breadth:

$$S = \log_{10}(\mathrm{LOC}) + \log_{10}(1 + D) - 2, \tag{4}$$

where LOC is the source code size and D is the number of direct runtime dependencies.

All three terms are evaluated on the current candidate. The Environment Evolution Loop iteratively revises the environment to maximize Q within a fixed construction budget.

3.2 Construction Pipeline

Repository Search and Filtering RepoFactory searches public code-hosting platforms for candidate repositories, excluding those included in established coding benchmarks to avoid evaluation overlap. Eligible projects must provide package metadata and a nonempty local test suite to support isolated installation and evaluation. To balance task complexity with reconstruction cost, we require at least five implementation files and fewer than 40,000 lines of Python code. We exclude projects that require external services, outbound network access, graphical interfaces, or specialized

hardware. Network-related projects remain eligible when their tests mock external calls or use only local connections.

Runtime Construction and Repository Analysis For each candidate repository, RepoFactory copies each candidate repository into an isolated task directory and builds a Docker image from its build files and dependency manifests. A agent installs dependencies, identifies test commands, and revises the image or execution configuration as needed. We retain a repository only when all tests pass in a fresh container with bounded resources and no network access. RepoFactory archives the validated image and records the installation and test commands, package files, and test results.

After establishing the runtime, RepoFactory analyzes the repository in a read only container to determine what the task specification must cover. A scanner based on abstract syntax trees extracts public classes, functions, and methods, together with their signatures, exceptions, and return information. It also records repository symbols imported by tests and builds an implementation tree excluding test files. These outputs guide task synthesis and, together with dependency metadata, provide the inputs for computing structural coupling and implementation scope.

Task Synthesis and Environment Materialization. RepoFactory prepares an isolated workspace with dependencies installed and disabled external network access. The coding agent must reconstruct the project from the task specification, without access to the original repository or evaluation code. The synthesis agent uses the repository analysis to write this specification and an accompanying API manifest. The specification describes the repository's functionality and structure, while the API manifest documents the public interfaces the agent must implement.

Functionally correct implementations can still fail if their interface names or import paths differ from those expected by the tests. RepoFactory therefore checks the specification and API manifest against the repository symbols imported by the original tests, returning missing symbols and their import records to the synthesis agent for revision. Once this check passes, RepoFactory combines the specification with the validated runtime image to create the environment. The original repository remains isolated as the oracle, and evaluation resources are available only during verification.

Environment Evolution Loop After materialization, RepoFactory first checks validity A. The oracle O must receive full reward, and the unchanged initial workspace must receive zero reward. If either condition fails, the construction agent uses the validation report and execution trace to identify the cause and repair the candidate. The revised environment is evaluated again, and optimization of scope and coupling begins only after both checks pass.

A valid environment may still require little implementation work. Implementation scope S therefore guides revisions toward more substantial reconstruction tasks. A larger project can still consist of independent components, so structural coupling C complements scope by directing evolution toward more interdependent functionality. The construction agent decides how to revise the oracle implementation, task specification, and tests, keeping the description and evaluation consistent with the revised functionality.

For each valid candidate, the agent compares S and C against thresholds of 0.5 and 0.01, respectively, and prioritizes any term below threshold. When both thresholds are met, the agent chooses which term to optimize under Q. Each revision is rechecked for validity before the scores are reassessed, with validity failures returned to repair. A revision targeting S or C is retained only if the selected score strictly increases; otherwise, RepoFactory restores the previous version. The loop runs for up to 12 rounds, assessing each revision through verifier execution and code analysis without solving-agent rollouts or downstream training and evaluation. RepoFactory uses this low-cost feedback to recover failed candidates and improve valid tasks, increasing both the yield and training value of repository-level environments.

4 Analysis of Synthesized Data

Effective training environments for long-horizon coding should expose agents to diverse software and require substantial repository reconstruction under reliable feedback across a broad range of difficulty. Our 10,159 RepoFactory environments span 17 domains and place agents in complete repository reconstruction tasks with executable validation. To collect SFT data, we run Qwen3.8-Max five times on each environment, filter out low reward trajectories, and retain 10,947 samples spanning 5,577 environments. We analyze the resulting

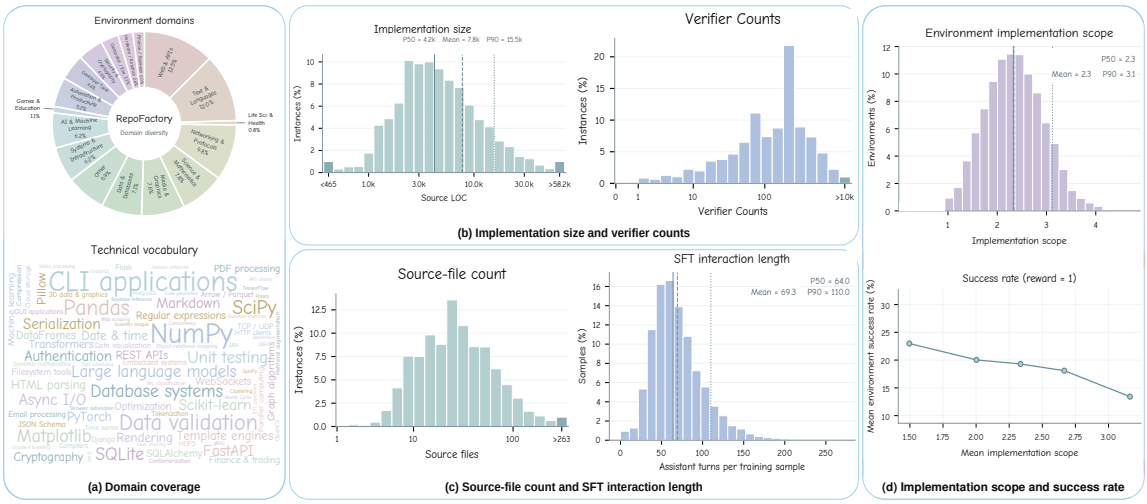


Figure 2: Statistics of RepoFactory environments and SFT data.

collection to characterize its diversity, scale, interaction length, and difficulty with the statistics summarized in Figure 2.

Diversity. The collection covers a broad range of application domains and technical ecosystems, without being dominated by any single area, as shown in Figure 2(a). The largest category accounts for only 12.5% of environments, and only 6.9% fall outside the named domains, indicating broad coverage rather than concentration. Its varied libraries, interfaces, and implementation conventions expose agents to the diversity of real software development. This breadth supports training agents that generalize across software projects instead of specializing in a narrow domain.

Scale and verification. These environments require agents to reconstruct substantial codebases rather than write isolated programs. As shown in Figure 2(b), the median oracle contains 4.2k lines of source code, with a substantial tail of larger systems. Across the 10,159 instances with successful test parsing, the verifier contains 2.09 million statically identified test definitions, providing dense coverage at the instance level. Because these tests are inherited from the source repositories, they exercise behavior across the public interface rather than checking only a few end to end outcomes. Success therefore depends on implementing a coherent repository whose components remain consistent under extensive verification.

Repository Scale and Interaction diversity. The environments and trajectories capture sustained development at the repository level. The median oracle repository contains 26 source files, and the median SFT trajectory contains 64 assistant turns, as shown in Figure 2(c). A successful trajectory must keep track of a multi-file project while successive edits are implemented and tested. Together, the long-horizon trajectories and varied operations provide rich supervision for maintaining project state as agents coordinate changes across files and refine their implementations through execution feedback.

Implementation Scope. Figure 2(d) shows the distribution of implementation scope S and its relationship with reconstruction performance. We measure each environment's success rate as the fraction of teacher trajectories receiving full reward. Environments with matched trajectories are ordered by S and divided into five equally sized groups, whose mean scope and success rates are plotted in the lower panel. Mean success falls from 23.0% in the group with the smallest scope to 13.4% in the group with the largest. This trend supports S as a proxy for the demands of repository reconstruction.

5 Experiments

5.1 Experimental Setup

Benchmarks. We evaluate our model on five benchmarks covering repository generation, software engineering, and terminal-based tasks. NL2Repo (Ding et al., 2025) is a long-horizon coding benchmark that requires agents to implement a complete repository from scratch based on a natural language description. Success requires sustained planning, tool use, and coordinated implementation

Model	Params	NL2Repo Bench	SWE-bench Verified	SWE-bench Pro	SWE-bench Multilingual	Terminal-Bench 2.1	Avg.
Frontier Closed-Source Models							
GPT-5.6-sol	-	56.8	96.2	64.6	-	88.8	76.6
GPT-5.5	-	45.1	82.6	58.6	77.8	73.8	67.6
Claude Opus 5	-	75.3	96.0	79.2	89.5	89.1	85.8
Gemini-3.1 Pro	-	33.4	80.6	54.2	44.0	70.3	56.5
Open-Weight Models (>200B)							
Qwen3.8-Max	2.4T-A95B	55.9	85.6	67.7	-	86.6	74.0
DeepSeek-V4-Pro	1.6T-A49B	38.5	80.6	55.4	76.2	78.7	65.9
GLM-5.2	744B-A40B	48.9	82.8	62.1	74.5	81.0	69.9
Kimi-K2.6	1T-A32B	-	80.2	58.6	76.7	67.2	70.7
Tencent-Hy3	295B-A21B	45.6	78.0	57.9	75.8	71.7	65.8
MiniMax-M3	428B-A23B	42.1	80.5	59.0	-	66.0	61.9
Open-Weight Models (<200B)							
Qwen3.5-122B-A10B	122B-A10B	27.7	73.4	49.5	67.2	44.9	52.5
Ornith-1.0-35B	35B-A3B	34.6	75.6	50.4	69.3	62.8	58.5
Qwen3.5-35B-A3B	35B-A3B	23.8	65.8	42.4	59.7	33.7	45.1
Ours							
RepoHorizen-35B-A3B	35B-A3B	46.4 (+22.6)	72.2 (+6.4)	47.6 (+5.2)	61.0 (+1.3)	40.5 (+6.8)	53.5 (+8.5)
w/o environment evolution	35B-A3B	41.3 (+17.5)	70.9 (+5.1)	45.1 (+2.7)	60.8 (+1.1)	39.3 (+5.6)	51.5 (+6.4)

Table 1: Main results on five agent benchmarks.

across multiple files and modules. SWE-bench Verified (Jimenez et al., 2024), SWE-bench Pro (Deng et al., 2025), and SWE-bench Multilingual (Zan et al., 2025) assess the ability to resolve issues in existing repositories, while Terminal-Bench 2.1 (Merrill et al., 2026) evaluates task completion in terminal environments. Together, these benchmarks measure coding-agent performance across both project construction and repository maintenance.

Evaluation Settings. We use Claude Code version 2.1.140 as the agent scaffold. The maximum context length is 262,144 tokens, and the maximum output length is 32,768 tokens. Context is automatically compacted when its usage reaches 80% of the maximum context length. For decoding, we set temperature to 1.0, top-p to 0.95, and top-k to 20.

5.2 Main Results

Table 1 compares RepoHorizen-35B-A3B and its base model with representative closed-source and open-weight models. On NL2Repo, which directly evaluates full-repository reconstruction, training raises the score from 23.8 to 46.4, a gain of 22.6 points. This is the strongest listed result among open-weight models below 200B parameters. RepoHorizen also surpasses GPT-5.5 and Gemini-3.1 Pro, which score 45.1 and 33.4, respectively, as well as larger open-weight models DeepSeek-V4-Pro and MiniMax-M3, which score 38.5 and 42.1. Improvements extend to all four remaining benchmarks, with gains of 1.3 to 6.8 points across issue resolution and cross-domain long-horizon tasks in Terminal-Bench 2.1. These results suggest that RepoFactory training transfers beyond repository generation to broader software engineering and terminal work at fixed model scale. Environment evolution further strengthens this performance across all five benchmarks, adding 5.1 points on NL2Repo and 2.1 points on average over the no-evolution variant. These additional gains support the effectiveness of our proxy metrics in guiding environment evolution that improve training value.

5.3 Environment Scaling

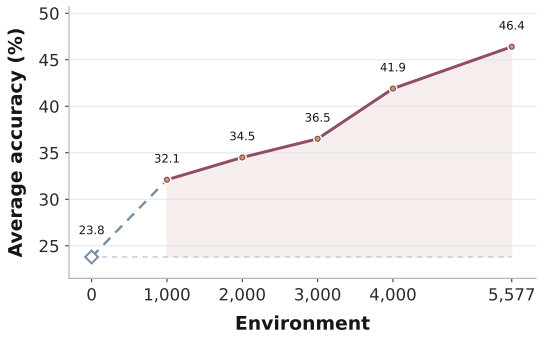


Figure 3: Average accuracy as a function of the number of RepoFactory training environments.

Figure 3 shows how NL2Repo performance changes as the training set grows. Performance increases steadily from 23.8% without RepoFactory

environments to 46.4% with 5,577 environments, consistent with the 22.6-point gain reported in Table 1. The improvement continues at the largest scale tested, with no evidence of saturation.

The gains become more pronounced beyond 3,000 environments. Adding the remaining 2,577 environments raises performance from 36.5% to 46.4%, suggesting that the larger collection exposes the model to repository structures and behavioral requirements that smaller subsets do not cover. Scaling the environment set therefore expands the supervision available for long-horizon repository reconstruction and continues to improve performance well beyond the earliest additions.

5.4 Scaling Across Model Sizes

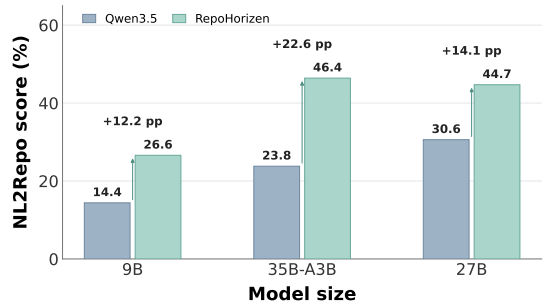


Figure 4: NL2Repo performance across model sizes.

We examine how the benefits of RepoFactory training vary with model capacity. Figure 4 reports NL2Repo gains of 12.2 to 22.6 percentage points across the 9B, 27B, and 35B-A3B models, showing that full-repository task remains useful across different parameter budgets. The 35B-A3B model gains the most despite the 27B model's higher baseline, suggesting that stronger initial performance does not necessarily translate into greater benefit from this training.

5.5 Effect of the Data Source

We compare RepoFactory with existing SWE datasets to assess the training value of its environments. We collect trajectories on six open datasets and apply the same cleaning procedure as for RepoFactory. Each baseline contains 11,000 training trajectories, slightly more than RepoFactory. We fine-tune Qwen3.5-35B-A3B separately on each collection and evaluate the models with the Claude Code harness. Table 2 includes both published results and these controlled comparisons. Among the published results, obtained with varying harnesses, the highest NL2Repo score is 28.9. RepoFactory reaches 46.4 and leads all six open-data baselines on all five benchmarks. On NL2Repo, its advantage over the strongest baseline, OpenSWE, is 21.2 points. The teacher is the same across our runs, so its strength alone cannot explain this advantage. The results support the value of the training experience provided by RepoFactory's environments.

Training Data	Open	N2R	Ver.	Pro	ML	TB
Reported by prior work (varying harnesses)						
DeNovoSWE	✗	27.1	–	–	–	–
MindForge	✗	23.5	73.8	51.3	67.8	–
Refact-92B	✗	28.9	–	–	–	–
TrajLong	✗	13.8	51.1	34.3	–	12.4
Ours (Claude Code harness)						
OpenSWE	✓	25.2	64.6	40.5	57.0	37.1
R2E	✓	22.6	62.0	40.5	59.7	29.2
ScaleSWE	✓	24.0	56.2	40.9	51.3	29.2
SWE-smith	✓	22.6	59.4	39.0	55.3	31.5
SWE-rebench	✓	23.1	68.8	42.0	59.0	29.2
SWE-rebench-v2	✓	24.0	62.4	36.8	54.0	33.7
RepoFactory	✓	**46.4**	**72.2**	**47.6**	**61.0**	**40.5**

Table 2: Performance of models trained on different SWE data sources, in percent. N2R = NL2Repo, Ver. = SWE-bench Verified, Pro = SWE-bench Pro, ML = SWE-bench Multilingual, TB = Terminal-Bench 2.1. Best result among our runs in **bold**.

5.6 Ablation of the Evolution Objective

The evolution objective must support both the recovery of usable environments and improvements in training value. We evaluate these effects separately, using oracle validation to measure candidate recovery and downstream training to assess S and C as optimization targets. For recovery, we uniformly sample 500 repositories from the filtered candidate set and compare runs allowing up to three repair steps and none. We compare the number of environments satisfying $A = 1$, with concurrency, timeouts, and infrastructure retry budgets held fixed.

Figure 6 shows that evolution increases the number of validated environments from 169 to 347, raising yield from 33.8% to 69.4% under the same acceptance criterion. These gains suggest that many initial failures stem from construction defects that can be corrected through execution feedback. Most of the improvement comes from the first repair, with smaller gains from subsequent steps. Recovering these otherwise discarded candidates expands the supply of valid environments from the same repository pool, supporting environment generation at scale.

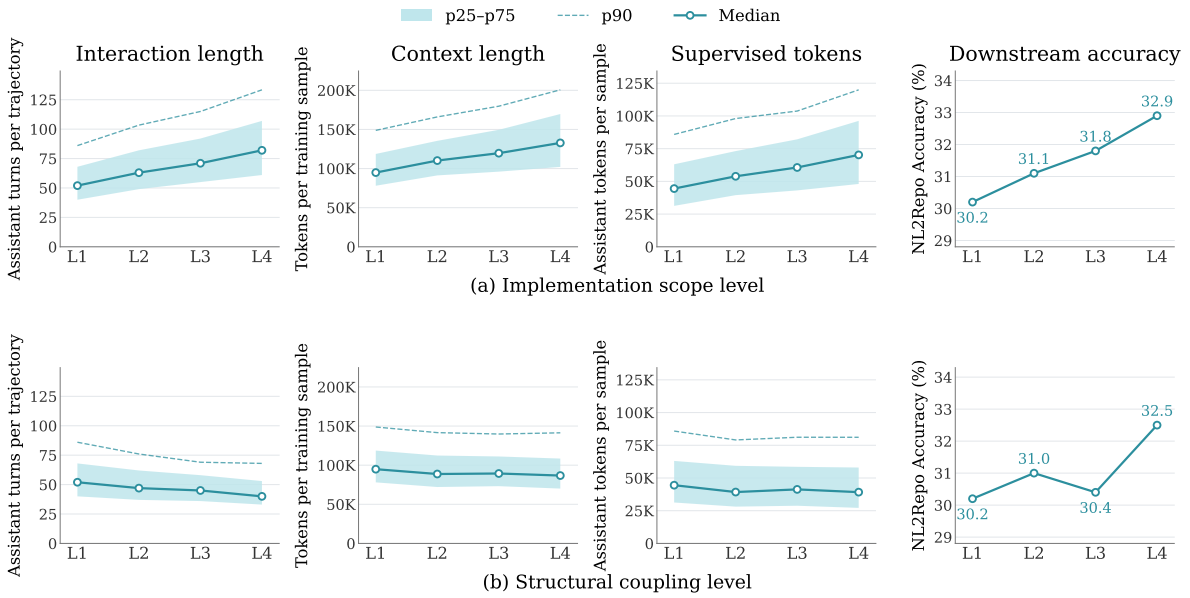


Figure 5: Statistics of the four difficulty levels and the corresponding NL2Repo results.

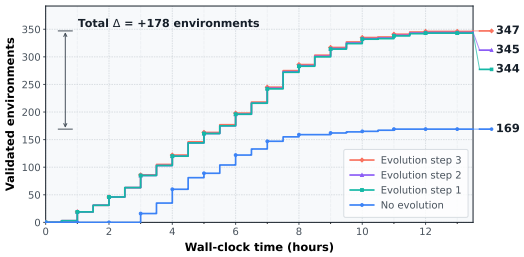


Figure 6: Cumulative number of validated environments over wall-clock time, sampled every 30 minutes.

To assess whether optimizing S and C improves training value, we run two independent evolution experiments from the same 1,000 randomly sampled environments. Each experiment targets one metric while retaining validity checks. We collect environments after rounds 3, 6, 9, and 12 as Levels 1 through 4, respectively. Each level supplies 2,000 trajectories for fine-tuning Qwen3.5-35B-A3B with the same initialization and optimization settings.

As shown in Figure 5, optimizing S produces longer interactions and contexts and more supervised tokens. NL2Repo accuracy increases steadily from 30.2% at Level 1 to 32.9% at Level 4. Expanding implementation scope thus provides more extensive development experience within a fixed number of trajectories, accompanied by improved reconstruction performance.

Optimizing C produces a different pattern. Accuracy fluctuates at intermediate levels but reaches 32.5% at Level 4, compared with 30.2% at Level 1. Interaction length decreases, while context length and supervised token counts decline only slightly. The improvement therefore does not depend on longer trajectories or additional supervision. This complements the scope ablation, suggesting that training value depends not only on the extent of implementation work but also on the relationships among the components being implemented.

6 Conclusion

We presented RepoFactory, an automated pipeline that constructs verifiable environments for full-repository generation from real-world projects. Its Environment Evolution Loop uses timely, low-cost feedback to optimize validity, implementation scope, and structural coupling, recovering failed conversions and improving task value. The resulting 10,159 environments span 17 software domains. Training Qwen3.5-35B-A3B on these environments yields RepoHorizen-35B-A3B, with gains of 22.6 percentage points on NL2Repo-Bench and 8.5 points on average across five coding and agentic benchmarks. Environment evolution contributes a further 5.1 and 2.1 points over training without evolution. These findings highlight the value of repository-level environments for coding-agent training and support low-cost proxy objectives as a practical means of improving their training value at scale.

7 Limitations

RepoFactory currently focuses on Python projects, with relatively limited coverage of other programming languages. Our proxy metrics provide useful, low-cost guidance for environment evolution, but characterize task content at a broad structural level. They do not yet explicitly distinguish the capabilities required by individual tasks. Extending language coverage and incorporating more detailed information about these capability requirements could support a wider range of development settings and more targeted training.

References

Sparsh Bajoria, Shreyanshu Ranjan, Adhitya M, Vineet Suri, Vikas Hassija, Vinay Chamola, and Amir Hussain. 2026. From language models to agentic ai: A survey of autonomous, action-enabled, and collaborative llm agents. *Cognitive Computation*, 18(1):103.

Yihao Chen, Shi Chang, Khaled Chawa, Feng Lin, Boyuan Chen, Shaowei Wang, and Ahmed E Hassan. 2026. Mindforge: Teaching small language models whole-life-cycle software engineering via source-free program synthesis. *arXiv preprint arXiv:2607.27146*.

Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pages 6621–6642. PMLR / OpenReview.net.

Anyi Xu DeepSeek-AI, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chenchen Ling, et al. Deepseek-v4: Towards highly efficient million-token context intelligence, 2026. *URL https://arxiv.org/abs/2606.19348*.

Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, Chetan Rane, Karmini Sampath, Maya Krishnan, Srivatsa Kundurthy, Sean Hendryx, Zifan Wang, Chen Bo Calvin Zhang, Noah Jacobson, Bing Liu, and Brad Kenstler. 2025. Swe-bench pro: Can AI agents solve long-horizon software engineering tasks? *CoRR*, abs/2509.16941.

Jingzhe Ding, Shengda Long, Changxin Pu, Huan Zhou, Hongwan Gao, Xiang Gao, Chao He, Yue Hou, Fei Hu, Zhaojian Li, Weiran Shi, Zaiyuan Wang, Daoguang Zan, Chenchen Zhang, Xiaoxu Zhang, Qizhi Chen, Xianfu Cheng, Bo Deng, Qingshui Gu, and 29 others. 2025. Nl2repo-bench: Towards long-horizon repository generation evaluation of coding agents. *CoRR*, abs/2512.12730.

Yaxin Du, Xiyuan Yang, Zhifan Zhou, Wanxu Liu, Zixing Lei, Zimeng Chen, Fenyi Liu, Haotian Wu, Yuzhu Cai, Zexi Liu, Xinyu Zhu, Wenhao Wang, Linfeng Zhang, Chen Qian, and Siheng Chen. 2026. Datamaster: Data-centric autonomous AI research. *CoRR*, abs/2605.10906.

Amir M Ebrahimi, Mohammed Mehedi Hasan, Seyed Shayan Daneshvar, Jiahuei Lin, Gopi Krishnan Rajbahadur, and Ahmed E Hassan. 2026. Software engineering llms: Why we need them and how software engineering helps build them. *arXiv preprint arXiv:2610.33851*.

Kanishk Gandhi, Shivam Garg, Noah D. Goodman, and Dimitris Papailiopoulos. 2026. Endless terminals: Scaling RL environments for terminal agents. *CoRR*, abs/2601.16443.

Zhanbo Hua, Yifan Yao, Weihao Xie, Yongchi Zhao, Minghao Liu, Ruizhi Qiu, Zhewei Huang, Zun Wang, Yiyan Ji, Yunhai Ye, Letian Zhu, Xinping Lei, Han Li, Zhiyuan Ma, Zili Wang, Zhaoxiang Zhang, and Jiaheng Liu. 2026. Cli-universe: Towards verifiable task synthesis engine for terminal agents. *CoRR*, abs/2606.22883.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R. Narasimhan. 2024. Swe-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Ilia Kulikov, Chenxi Whitehouse, Tianhao Wu, Yixin Nie, Swarnadeep Saha, Eryk Helenowski, Weizhe Yuan, Olga Golovneva, Jack Lanchantin, Yoram Bachrach, Jakob N. Foerster, Xian Li, Han Fang, Sainbayar Sukhbaatar, and Jason Weston. 2026. Autodata: An agentic data scientist to create high quality synthetic data. *CoRR*, abs/2606.25996.

Zhongzhi Li, Yucheng Shi, Zongxia Li, Ruhan Wang, Anhao Li, Zixun Huang, Junyao Yang, Lei Ke, Ninghao Liu, Haitao Mi, et al. 2026a. Recursive synthesis for long-horizon terminal tasks. *arXiv preprint arXiv:2608.05466*.

Zongxia Li, Hongyang Du, Chengsong Huang, Xiyang Wu, Lantao Yu, Yicheng He, Jing Xie, Xiaomin Wu, Zhichao Liu, Jiarui Zhang, and Fuxiao Liu. 2026b. Mm-zero: Self-evolving multi-model vision language models from zero data. *CoRR*, abs/2603.09206.

Mike Merrill, Alexander Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Shin, Thomas Walshe, E Kelly Buchanan, et al. 2026. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations*, volume 2026, pages 40903–40986.

Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. 2025.

Training software engineering agents and verifiers with swe-gym. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net.

Miao Peng, Qintong Zhang, Nuo Chen, Yuhan Li, Guochen Yan, Xinran Gu, Hongqiu Wu, Hai Wang, Lydell Huang, Wentao Zhang, et al. 2026a. Trajlong: Co-designing agentic and long-context supervision for mid-training. *arXiv preprint arXiv:2610.04973*.

Xiaoxuan Peng, Kaiqi Zhang, Xinyu Lu, Boxi Cao, Yaojie Lu, Hongyu Lin, Xianpei Han, and Le Sun. 2026b. Litecoder-terminal: Scaling long-horizon terminal environments for learning language agents. *CoRR*, abs/2605.29559.

Jincheng Ren, Siwei Wu, Yizhi Li, Kang Zhu, Shu Xu, Boyu Feng, Ruibin Yuan, Wei Zhang, Riza Batista-Navarro, Jian Yang, and Chenghua Lin. 2026. A self-evolving framework for efficient terminal agents via observational context compression. *CoRR*, abs/2604.19572.

Qijia Shen, Zhiqi Huang, Vamsidhar Kamanuru, Aznaur Aliev, Jay Rainton, Ahmed Awelkair, Zhichen Zeng, Jiajun Li, Shi Dong, Yueming Yuan, Boyuan Ma, Qizheng Zhang, Jiwei Fu, Yuzhen Mao, Wendong Fan, Ping Nie, Philip Torr, Bernard Ghanem, Changran Hu, and 3 others. 2026. SETA: scaling environments for terminal agents. *CoRR*, abs/2607.10891.

Xiaoshuai Song, Haofei Chang, Guanting Dong, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. 2026. Envscaler: Scaling tool-interactive environments for LLM agent via programmatic synthesis. In *Findings of the Association for Computational Linguistics, ACL 2026, San Diego, California, United States, July 2-7, 2026*, pages 8326–8357. Association for Computational Linguistics.

GLM Team et al. 2026a. Glm-5: from vibe coding to agentic engineering. arxiv preprint (2026). *arXiv preprint arXiv:2602.15763*.

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, Jianfeng Cai, Xinyuan Cai, Peizhou Cao, Yuxuan Cao, Ziwei Chai, Y Charles, et al. 2026b. Kimi k3: Open frontier intelligence. *arXiv preprint arXiv:2607.24653*.

XYZ Agentic Team. 2026. Ai4ai at scale: A full-pipeline system for enhancing llm agentic capabilities. Technical report, Technical report, XYZ AI Lab, July 2026. URL https://xyz-lab. ai/blogs/ai4ai ....

Yang Tian, Zhengpeng Shi, Yu Zhou, and Bo Zhao. 2026. Beyond function calling: Benchmarking tool-using agents under tool-environment unreliability. *arXiv preprint arXiv:2606.25819*.

Dunwei Tu, Hongyan Hao, Hansi Yang, Yihao Chen, Yi-Kai Zhang, Zhikang Xia, Yu Yang, Yueqing Sun, Xingchen Liu, Furao Shen, Qi Gu, Hui Su, and Xunliang Cai. 2026. Scaleenv: Scaling environment synthesis from scratch for generalist interactive tool-use agent training. *CoRR*, abs/2602.06820.

Yangzhen Wu, Aaron J. Li, Wenjie Ma, Li Cao, Ziheng Zhou, Mert Cemri, Shu Liu, Yuran Xiu, Chenxiao Yan, Haikun Zhao, Bin Yu, Ion Stoica, and Dawn Song. 2026. Benchevolver: Frontier task synthesis via solution-centric evolution. *CoRR*, abs/2606.01286.

Zhangchen Xu, Junda Chen, Yue Huang, Dongfu Jiang, Jiefeng Chen, Hang Hua, Zijian Wu, Zheyuan Liu, Zexue He, Lichi Li, et al. 2026. Autolab: Can frontier models solve long-horizon auto research and engineering tasks? *arXiv preprint arXiv:2606.05080*.

A Jian Yang, An Huang, Bo Dan, Chen Mu, Chengran Yang, Chuan Shi, Feng Chang, Hongxin Li, Hongyuan Tao, Jinxing Zhang, et al. 2026a. Iquest-q1: Advancing foundation capabilities for agentic cli systems. *arXiv preprint arXiv:2610.30949*.

Chengyi Yang, Renxu Bao, Jue Hou, Zongpei Teng, Ziheng Wang, Yujie Lin, Chengsong Huang, Junnan Zhu, Delai Qiu, Yining Wang, et al. 2026b. Programgym: Scaling training environments for long-horizon program reproduction agents. *arXiv preprint arXiv:2610.13395*.

Jian Yang, Wei Zhang, Jiajun Wu, Junhang Cheng, Shawn Guo, Haowen Wang, Weicheng Gu, Yaxin Du, Joseph Li, Fanglin Xu, Yizhi Li, Lin Jing, Yuanbo Wang, Yuhan Gao, Ruihao Gong, Chuan Hao, Ran Tao, Aishan Liu, Tuney Zheng, and 9 others. 2026c. Incoder-32b: Code foundation model for industrial scenarios. *CoRR*, abs/2603.16790.

John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. 2025. Swe-smith: Scaling data for software engineering agents. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*.

Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Shulin Xin, Linhao Zhang, Qi Liu, Aoyan Li, Lu Chen, Xiaojian Zhong, Siyao Liu, Yongsheng Xiao, Liangqiang Chen, Yuyu Zhang, Jing Su, Tianyu Liu, Rui Long, Ming Ding, and Liang Xiang. 2025. Multi-swe-bench: A multilingual benchmark for issue resolving. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*.

Jiale Zhao, Guoxin Chen, Fanzhe Meng, Wayne Xin Zhao, Ruihua Song, Ji-Rong Wen, and Kai Jia. 2026.

Denovoswe: Scaling long-horizon environments for generating entire repositories from scratch. *arXiv preprint arXiv:2606.10728*.

Yiqi Zhu, Apurva Gandhi, and Graham Neubig. 2025. Training versatile coding agents in synthetic environments. *CoRR*, abs/2512.12216.

A Implementation Details

This appendix describes the trajectory preparation and supervised fine-tuning configuration used for Qwen3.5-35B-A3B.

A.1 Trajectory Filtering and Cleaning

Before fine-tuning, we retain trajectories with rewards above 0.5 and apply offline quality checks. Deterministic checks examine record validity, dialogue structure, consistency between tool calls and results, response completeness, and the formatting of reasoning and code. They also screen for repetition, interaction loops, contamination, potentially sensitive or harmful content, and attempts to obtain information outside the intended task environment. We discard trajectories with clear violations and flag ambiguous cases for review.

A language-model judge reviews the flagged trajectories. These trajectories are included in the SFT data only when the judge determines that the apparent violation is benign. Confirmed violations and unresolved cases remain excluded.

A.2 SFT Hyperparameters

Hyperparameter	Value
Number of epochs	3
Rollout / global batch size	64 / 64
Per-device micro-batch size	1
Dynamic token batching	Enabled
Token budget per GPU	80,000
Optimizer	AdamW
Peak / minimum learning rate	1.0×10^{-6} / 1.0×10^{-7}
Learning-rate schedule / warmup	Cosine / 10%
Weight decay	0.1
Adam coefficients	$\beta_1 = 0.9$, $\beta_2 = 0.98$
Tensor parallel size	2
Pipeline parallel size	1
Context parallel size	4
Expert parallel size	8
Expert tensor parallel size	1
Sequence parallelism	Enabled

Table 3: Supervised fine-tuning hyperparameters for Qwen3.5-35B-A3B.

We fine-tune the model, using the settings in Table 3. The data are shuffled during training, with no online generation. Following the Qwen3.5 multi-turn objective, we compute the loss over assistant spans, including tool calls. System, user, and tool-result turns provide context but receive no loss.

The rollout and global batch sizes are both 64, and the per-device micro-batch size is 1. Dynamic token batching limits each GPU to 80,000 tokens, with no separate sequence-length cap. Tensor, context, and expert parallel sizes are 2, 4, and 8, respectively. Sequence parallelism is enabled, while pipeline and expert tensor parallel sizes are both 1.

We use AdamW with weight decay of 0.1 and coefficients $\beta_1 = 0.9$ and $\beta_2 = 0.98$. The learning rate warms up to 1.0×10^{-6} over the first 10% of training, then decays to 1.0×10^{-7} following a cosine schedule.

B Additional Dataset Analysis

B.1 Python Ecosystem Coverage

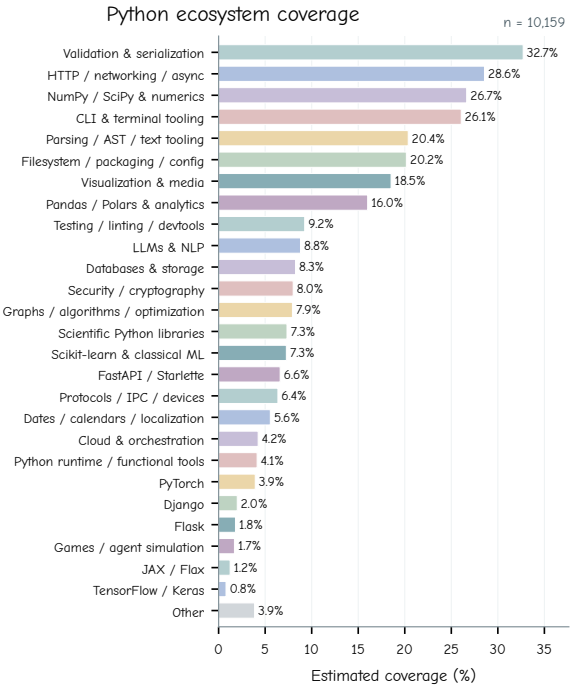


Figure 7: Estimated Python ecosystem coverage across RepoFactory's 10,159 environments. Categories may overlap. Other denotes environments without an ecosystem match.

Figure 7 summarizes the estimated coverage of Python ecosystems across RepoFactory's 10,159 environments. The estimates use the available Python repository records. We assign ecosystem labels by matching direct runtime dependency names

against predefined rules. For repositories without a dependency-based match, we use ecosystem terms in the first substantive paragraph of the project description. An environment may belong to multiple categories.

Validation and serialization are the most common category at 32.7%, followed by HTTP, networking, and asynchronous tooling at 28.6%. Numerical computing, command-line tools, text processing, and data analysis are also widely represented. The distribution extends from common application infrastructure to specialized scientific libraries, web frameworks, and machine-learning frameworks, covering a range of software stacks for repository reconstruction.

B.2 Development Operations in Training Trajectories

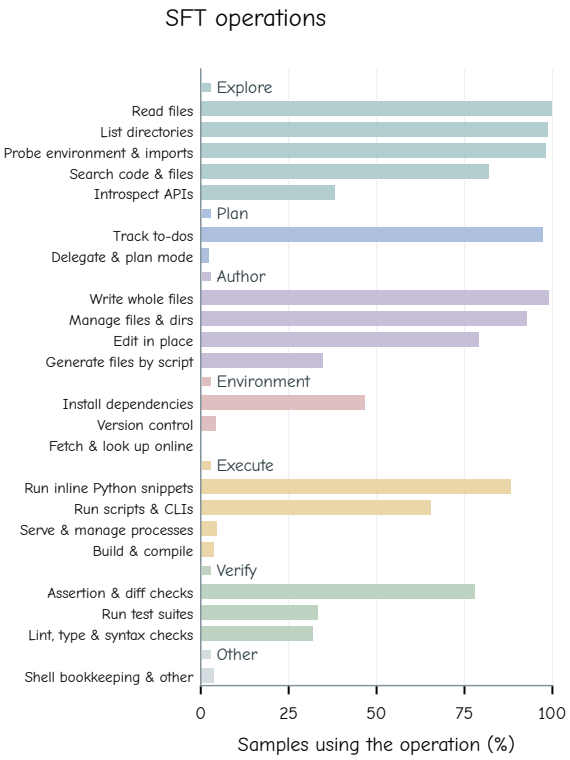


Figure 8: Development operations in the collected trajectories. Each bar shows the percentage of trajectories containing the operation. Categories may overlap.

We examine the development activities represented in the collected trajectories by classifying operations from tool names and shell-command contents. Figure 8 reports the percentage of trajectories containing each operation. A trajectory may contribute to several categories.

Reading and writing files occur in almost all trajectories, alongside frequent workspace exploration and code modification. Inline Python execution appears in 88.2% of trajectories, assertion or diff checks in 78.0%, and test-suite execution in 33.2%. These activities provide supervision for both implementing project code and examining its behavior through execution and local checks.