RSIBENCH-CODE: CAN LLM AGENTS HANDLE ENTERPRISE-LEVEL LLM DEVELOPMENT CODE FOR SELF-IMPROVEMENT?

Anonymous authors
Paper under double-blind review

ABSTRACT

Recursive self-improvement (RSI) aims to enable AI systems to contribute to the processes that produce stronger future models. A practical route is model self-training, but the code behind training, evaluation, serving, and infrastructure is still largely designed, maintained, and debugged by human engineers. This raises a central question: **can LLM agents handle the enterprise-level LLM-development code needed for self-improving AI**? We introduce **RSIBench-Code**, a benchmark for evaluating LLM agents on coding tasks reconstructed from practical LLM development problems in an enterprise. The benchmark contains 105 tasks across post-training, training infrastructure, serving, evaluation, and agent systems. These tasks cover implementation, repair, test writing, controlled experiments, and execution-path alignment. In each task, agents must locate and modify relevant code within repositories containing approximately 1.8 million source lines, with changing up to thousands of lines. RSIBench-Code supports long-horizon coding attempts with up to 150 agent turns and a 7,200-second solving budget. Experiments show that existing agents make meaningful but incomplete progress, suggesting that enterprise-level LLM-development code remains a key bottleneck for recursive self-improvement. Furthermore, RSIBench-Code provides a continuing source of executable evaluation tasks and potential training data for RSI in future work.

1 INTRODUCTION

Recursive self-improvement (RSI) seeks to create a cycle in which AI systems help build increasingly capable models (Anthropic Institute, 2026). Model self-training offers one practical path toward this goal. However, developing stronger models through self-training relies on code across the LLM-development pipeline, including training, evaluation, serving, and infrastructure. Much of this code is still designed, maintained, and debugged by engineers in enterprises. It therefore remains unclear whether LLM agents can work effectively with enterprise-level LLM-development code and contribute to self-improving AI.

Recent work evolves agent harnesses (Lee et al., 2026a;b, Zhang et al., 2026a), uses agents to optimize model training and training loops (Rank et al., 2026, Qu & Lu, 2026, Chen et al., 2026), and studies self-generated tasks or training-data strategies (Huang et al., 2026, Meng et al., 2026). Related benchmark work evaluates long-horizon engineering and optimization of LLM infrastructure (Ding et al., 2026). Together, these approaches address specific parts of model development. We instead evaluate enterprise-level LLM-development coding: whether LLM agents can handle the real, large, interconnected codebases and long-horizon tasks used by companies and research institutions.

To evaluate the such ability, we introduce **RSIBench-Code**, a benchmark of 105 executable tasks spanning full pipeline. The tasks cover different forms of work, including feature implementation, defect repair, test writing, controlled experiments, and execution-path alignment. We use complete inference and training codebases rather than isolating the relevant files. The task environments therefore expose approximately 1.8 million source lines and require agents to locate relevant components, trace dependencies, and preserve surrounding behavior. The largest reference solutions change more than thousands of lines. Current agents still find the benchmark difficult. GPT 5.6

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

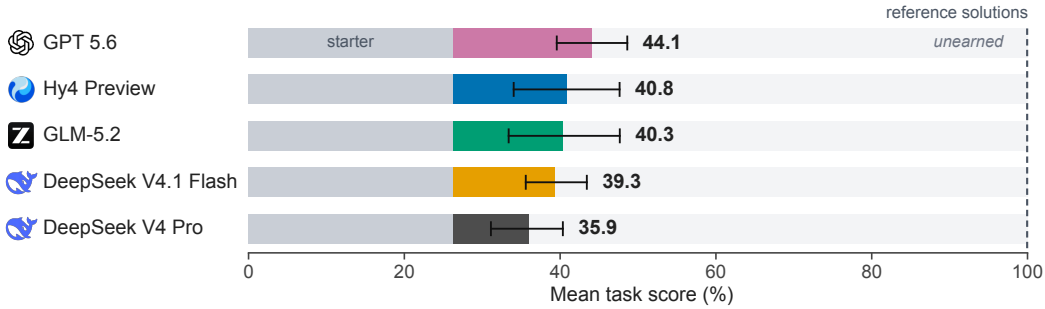


Figure 1: Recorded mean scores on the 105 RSIBench-Code tasks. Error bars are percentile 95% intervals from resampling source-problem families. The gray segment marks the mean archived starting score, and the dashed line marks the reference solutions.

obtains a mean score of 44.1%, leaving more than half of the available evaluator credit unearned on average. RSIBench-Code is an executable benchmark to evaluate coding agents across post-training, training infrastructure, serving, evaluation, and agent systems.

Building reliable executable tasks requires separate work on problem selection, task generation, and quality evaluation. We first collect structured data that mimics practical LLM-development issues encountered in enterprises. Agents then generate and refine executable tasks, while execution checks and solver feedback evaluate their quality and difficulty. This **three-stage workflow** reduces manual effort and supports continued task creation for RSI. We release the benchmark and workflow code at `https://anonymous.4open.science/r/RSIBench-Code-8F67`.

Our contributions are:

- **Enterprise-level LLM-development benchmark.** We introduce RSIBench-Code, a 105-task benchmark for evaluating whether LLM agents can modify the code used to train, evaluate, serve, and operate LLM systems.
- **Executable reconstruction workflow.** We present a three-stage workflow that turns practical enterprise development problems into public, runnable coding tasks with hidden evaluators and recorded construction traces.
- **RSI-relevant evaluation.** We evaluate frontier coding agents under long-horizon budgets and find that they earn credit beyond the starting repositories but remain far below reference solutions, making LLM-development code a measurable bottleneck for RSI.

Figure 1 summarizes the model results. Section 4 details problem selection, task construction, and the validation and refinement process. Section 3 reports the model comparisons, and Appendix G discusses related work.

2 The RSIBench-Code Benchmark

RSI requires AI systems to help improve the systems that produce future models. In LLM-development organizations, such improvement depends on engineering work across the full pipeline: designing experiments, modifying training and serving infrastructure, maintaining evaluation systems, and debugging agent-system behavior. RSIBench-Code evaluates whether LLM agents can perform this kind of enterprise-level LLM-development coding work in executable environments. This section describes the benchmark format, coverage, and observed difficulty. The task generation process and task lists are in Section 4 and Appendix B.

2.1 Task Format and Examples

To simulate real problems in enterprises, for each task, an agent receives a natural language request and a runnable starting repository. Each task runs in its own Docker container, and a hidden evaluator

Table 1: Examples of familiar development tasks in RSIBench-Code. Each row gives a concrete objective and the work expected for that task type.

Task type	Example objective	Expected agent work
Implementation	Limit incoming work when a processing queue is full.	Implement a controller that pauses and resumes new work based on the backlog.
Repair	Fix training that fails to resume from a saved checkpoint.	Locate the fault and repair how training state is restored.
Tests	Test saving and resuming a training run.	Write tests for normal recovery and incomplete checkpoints.
Experiment	Compare using more prompts with generating more responses per prompt.	Run both strategies under the same response budget and compare the results.
Execution-path alignment	Check that direct model requests and requests through an agent return consistent results.	Run both request paths on the same inputs and report any differences.

checks the submitted files in a separate environment. This isolation limits access to other tasks and grading files, reducing opportunities to cheat or interfere with evaluation. Task packages also contain a reference solution, fixed dependencies, and execution settings.

Table 1 illustrates the five task types in the current collection: implementation, repair, test writing, controlled experiments, and execution-path alignment. Tasks can require both code and evidence: a test-writing submission must distinguish correct and defective behavior, while an experiment must execute controlled comparisons and report results. Appendix B contains the full task instructions, and Appendix C explains representative tasks and how their evaluators check correctness. Each released task includes a reproducible environment, a hidden evaluator, a validated reference solution, and solver attempts used to profile difficulty. Together, these components provide executable, requirement-level feedback rather than relying only on final-text judgments.

2.2 Coverage and Composition

RSIBench-Code covers major parts of the LLM-development pipeline: post-training, training infrastructure, serving, evaluation, and agent systems. Tasks from different stages test different abilities, allowing a more detailed analysis of agent strengths and weaknesses. Post-training tasks often involve reward computation, loss functions, rollout handling, and training signals. Training-infrastructure tasks emphasize checkpointing, scheduling, asynchronous execution, and queue consistency. Serving tasks involve routing, tokenizer compatibility, and model-version management. Evaluation tasks focus on judge provenance, quality gates, and verifier behavior. Agent-system tasks test sandbox recovery, tool-event handling, and trajectory accounting.

We group tasks by development area, where the problem arises, and task type, what the agent must produce, as shown in Figure 2. Each task is assigned one primary development area based on the work it requires. RSIBench-Code spans five development areas and five task types, capturing both different technical settings and different forms of coding work. We prioritize tasks grounded in real development problems over filling every combination in a predefined grid. The resulting distribution reflects the enterprise problems.

2.3 Task Difficulty

RSIBench-Code is designed to evaluate agents across a range of task difficulty. The construction workflow assigns intended difficulty labels, but these labels need not match the difficulty observed from solver performance. We therefore measure difficulty using the three DeepSeek V4 Pro (DeepSeek-AI, 2026) attempts from the accepted generation round of each task. All 105 released tasks have three matching attempts.

For each attempt, the evaluator runs checks on the submitted files and awards partial credit for passed checks. The score is the fraction of available check weight earned,

$$s_{t,r} = \frac{\sum_{j \in C_t} w_{t,j}\, \mathbf{1}[\text{check } j \text{ passes}]}{\sum_{j \in C_t} w_{t,j}}, \tag{1}$$

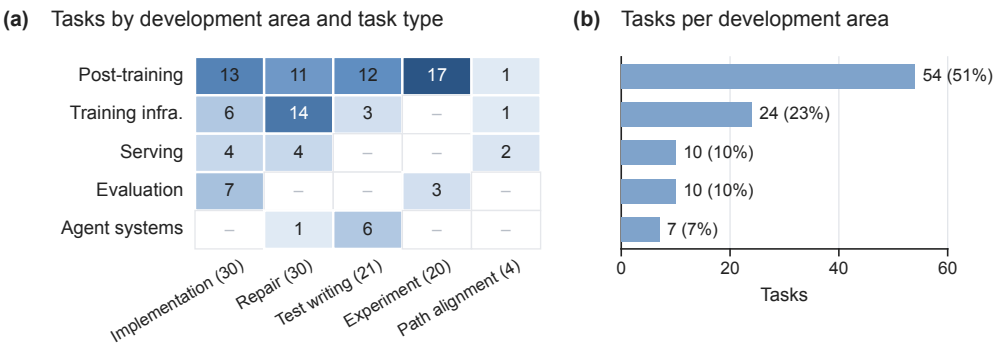


Figure 2: Composition of the 105-task snapshot. (a) Each cell gives the number of tasks in a development area and task type. Column labels give task-type totals in parentheses. Darker cells indicate more tasks. Empty cells indicate combinations not represented in the current collection. (b) Bars show the number and share of tasks in each development area.

where C_t is the set of scored checks for task t, $w_{t,j}$ is the weight of check j, and r identifies an attempt. Some evaluators assign different weights to different checks. The score measures satisfaction of the evaluated requirements, including new behavior and preservation of existing behavior.

To summarize difficulty, we use the best-of-three score, defined as each task's highest score across its three attempts. A lower score means that even the best observed attempt earned less credit. Figure 3a shows the range and best of the three scores for each task. The median best-of-three score is 0.46, and 61 of 105 tasks (58.1%) remain below 0.50 in all three attempts. Thus, low scores persist for most tasks even when we select the best observed attempt. Across attempts, 42 tasks receive identical scores, and the median difference between the highest and lowest scores is 0.07. Some tasks show larger variation, but the typical score range is small under these evaluation settings.

Figure 3b compares the best-of-three scores by task type. Testing tasks have the lowest median best-of-three score at 0.10, compared with 0.53 for implementation and 0.48 for repair. This suggests that producing tests that detect the required defects was particularly challenging.

For subsequent analyses, we assign an observed difficulty level using each task's DeepSeek best-of-three score. Hard tasks score below 0.50, Medium tasks score from 0.50 to below 0.65, and Easy tasks score at least 0.65. Figure 3c shows the resulting counts. These labels describe difficulty for DeepSeek V4 Pro under the recorded settings. The same task groups are used for every model in Section 3.

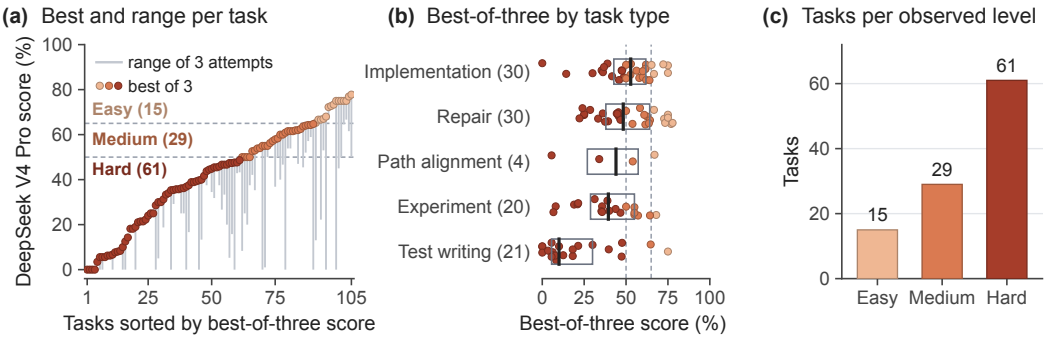


Figure 3: Observed difficulty from three DeepSeek V4 Pro attempts per task. (a) Tasks are sorted by their best-of-three score. Vertical lines span the minimum and maximum scores, and dots mark the best score, colored by observed difficulty; dashed lines mark the difficulty thresholds. (b) Each dot is a task's best-of-three score, boxes span the interquartile range, and black bars mark the median for each task type. Lower scores indicate greater observed difficulty. (c) Counts of Easy, Medium, and Hard tasks using the DeepSeek score thresholds defined in the text.

3 EVALUATION

We use RSIBench-Code to evaluate DeepSeek V4.1 Flash, DeepSeek V4 Pro (DeepSeek-AI, 2026), Hy4 Preview (Tencent, 2026), GLM-5.2 (GLM-5-Team et al., 2026), and GPT 5.6 (OpenAI, 2026). Harbor (Harbor Framework Team, 2026) provides the containerized evaluation environment. We compare overall performance, results across task types and development areas, resource use, and behavior in agent trajectories. Together, these analyses measure how reliably each agent handles enterprise-level LLM-development code.

3.1 EXPERIMENTAL SETUP AND METRICS

Hy4 Preview, GLM-5.2, and DeepSeek V4.1 models use Terminus-2 version 2.0.0, while GPT 5.6 uses Codex CLI. Both harnesses manage model responses and tool execution in Harbor with Docker. Each attempt is limited to 150 turns, and at most 32,768 output tokens per model call. Agents receive the solver instruction and starting repository, with four CPUs, 8 GB of memory, and a 7,200-second time limit. A separate hidden evaluator scores their submitted files. Appendix D gives additional execution settings and trial coverage.

DeepSeek V4 Pro contributes three attempts per task from the accepted generation round, while Hy4 Preview, GLM, DeepSeek V4.1 Flash, and GPT 5.6 each contribute one attempt per task. We use the attempt scores defined in Section 2, average them within each task, and then average equally across the 105 tasks. We also report the percentages of trials scoring exactly one or zero. A full score means that all scored requirements are satisfied. We report 95% bootstrap intervals by resampling the 32 source problems and including all tasks derived from each sampled problem. Appendix D describes the procedure.

3.2 OVERALL AND CATEGORY-LEVEL PERFORMANCE

Table 2: Recorded results on all 105 release tasks. Full and zero are percentages of graded trials. Means average within tasks first.

Model	Trials	Mean (%)	95% interval	Full (%)	Zero (%)
DeepSeek V4 Pro	311	35.9	[31.1, 40.3]	0.0	8.4
Hy4 Preview	101	40.8	[34.0, 47.6]	4.0	9.9
GLM-5.2	101	40.3	[33.4, 47.7]	3.0	5.0
DeepSeek V4.1 Flash	105	39.3	[35.6, 43.4]	0.0	0.0
GPT 5.6	105	44.1	[39.6, 48.6]	0.0	0.0

GPT 5.6 achieves the highest mean score at 44.1%, followed by Hy4 Preview at 40.8%, GLM-5.2 at 40.3%, DeepSeek V4.1 Flash at 39.3%, and DeepSeek V4 Pro at 35.9% (Table 2). Full-score outcomes remain rare: four Hy4 Preview trials and three GLM trials satisfy all scored requirements, while DeepSeek V4 Pro, DeepSeek V4.1 Flash, and GPT 5.6 do not reach a full score.

The unchanged starting repositories provide a baseline for credit earned without solving the tasks. They earn a mean score of 26.2%, compared with 99.95% for the reference solutions. Thus, agents earn credit beyond the existing behavior, but their mean scores remain well below the reference solutions. Per-task scores and cross-model comparisons appear in Appendix E.

Across task types, GLM-5.2 scores highest on implementation (56.6%), while Hy4 Preview scores highest on repair (56.8%), as shown in Figure 4. Test writing has the lowest mean for every model, ranging from 10.6% for GLM-5.2 to 26.8% for GPT 5.6. Experiments also remain challenging, with means of 27.5–39.2%. Across development areas, all models score highest on inference/serving (48.9–57.6%) and lowest on agent systems (13.4–35.4%).

3.3 RESOURCE USE

Scores describe correctness, while resource counts describe the interaction used to obtain a result. Figure 5 shows the distributions of turns, input tokens, output tokens, and tool calls, together with their means and maxima. A turn is an agent response, and tool calls count recorded tool invocations,

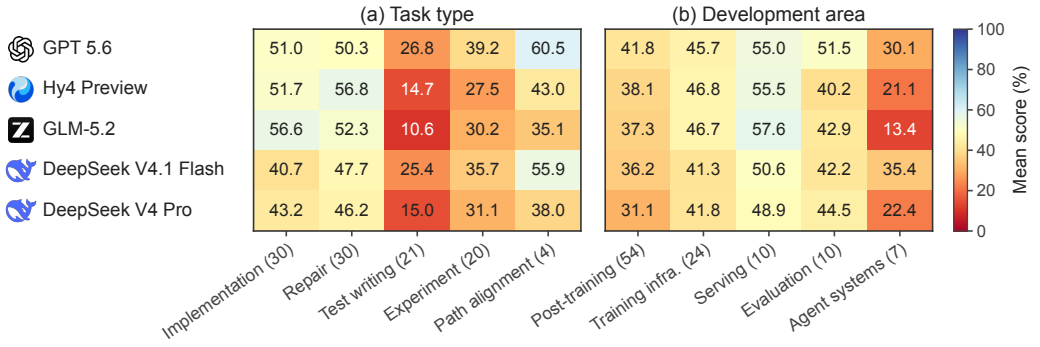


Figure 4: Mean scores by (a) task type and (b) development area. Each category shows its task count. Tasks receive equal weight within a category. Both panels use the same fixed 0–100% color scale.

so both counts depend on how the harness structures an interaction. Token counts sum usage across model calls.

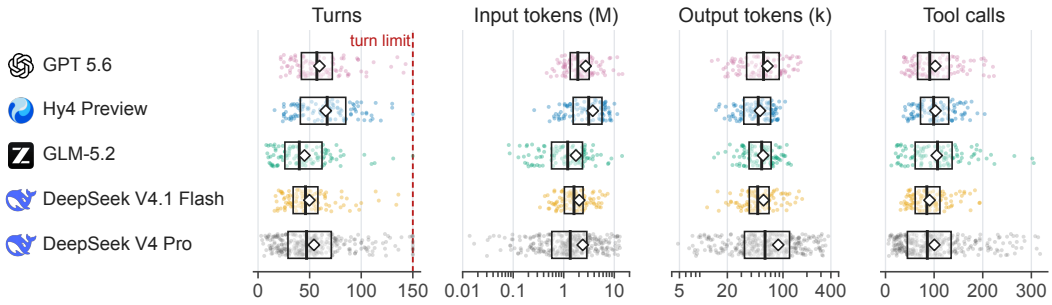


Figure 5: Resource use across recorded attempts on the 105 tasks. Each dot is one attempt. Boxes span the middle 50% of values, dark vertical lines mark medians, and diamonds mark means. Token axes use a logarithmic scale, and the dashed red line marks the 150-turn limit. M and k denote millions and thousands.

Resource use varies substantially within each model. Input-token counts have long upper tails, and DeepSeek V4 Pro's output-token tail extends farther than those of the other measured models. DeepSeek V4.1 Flash uses the fewest tool calls on average, while GPT 5.6 uses more turns and input tokens than GLM-5.2 but fewer input tokens than Hy4 Preview. Models with similar overall scores can therefore use different amounts of context and generate different amounts of text.

Across the shared difficulty groups, all five models use more turns and tokens on Hard than on Easy tasks (Figure 6). Hy4 Preview uses the most cumulative input tokens in every group, while DeepSeek V4 Pro generates the most output tokens on Medium and Hard tasks. These comparisons describe resource use on groups defined by DeepSeek V4 Pro's scores and task composition. Appendix E provides the corresponding per-task resource heatmaps alongside the scores.

3.4 Model Behavior in Agent Trajectories

To understand the actions behind the scores, we inspect agent trajectories, submitted files, and verifier outcomes. Appendix F provides matched case studies across task types. Hy4 Preview often alternates between code inspection, editing, and testing, but weaker attempts can leave requirements unchecked. GLM-5.2 makes focused changes with fewer turns, yet sometimes passes local tests while missing submission or lifecycle requirements. DeepSeek V4 Pro sometimes fixes one visible defect while leaving another cause unresolved, or reports completion after testing only a simplified

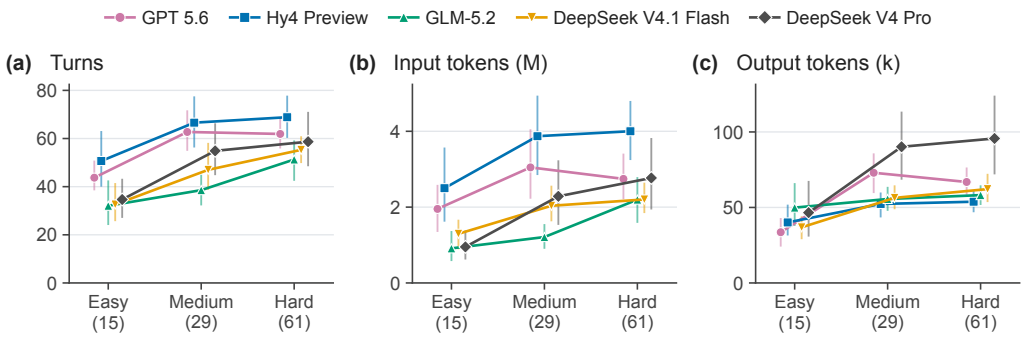


Figure 6: Mean turns and token counts per attempt by observed difficulty. Easy, Medium, and Hard are assigned from DeepSeek V4 Pro's best-of-three scores as defined in Section 2. The same task groups are used for every model, with task counts shown below the horizontal axes. Markers show group means and connecting lines guide the eye across the ordered categories. Vertical lines are 95% bootstrap intervals from resampling source problems within each group.

path. DeepSeek V4.1 Flash uses fewer tool calls on average but remains uneven. GPT 5.6 achieves the highest mean score, but still does not uniformly solve the benchmark.

Across models, successful attempts connect requirements to executable checks on the submitted implementation. Incomplete attempts often rely on narrower evidence, such as compilation, a local example, or a generated report. These patterns help explain why agents improve over the starting repositories but remain far below reference solutions.

4 Benchmark Construction

RSIBench-Code is constructed through the three-stage workflow shown in Figure 7. We separate problem selection, task generation, and task evaluation and refinement because each stage addresses a different source of failure. Problem selection determines whether a technical problem is useful and described in sufficient detail. Task generation reconstructs the problem in an executable repository and expands it into task variants for comprehensive assessment. Evaluation and refinement establish that each task is solvable, graded meaningfully, and difficult for a chosen solver.

To automate task generation while preventing assumptions from carrying across stages, we use three main agents to turn selected problems into executable tasks: an author, an evaluator, and a solver. The author constructs the task environment. The evaluator constructs the grading code and environment in a fresh context, preventing it from inheriting the author's reasoning. The solver independently evaluates task quality. Together, these agents form an end-to-end pipeline for automatic task generation and validation.

4.1 Problem Selection

Useful source problems require domain judgment, while short descriptions often omit the context needed to reconstruct an executable task. For RSIBench-Code, a Codex agent extracts structured candidate records representing practical LLM-development issues encountered in enterprises. Human reviewers select records with a meaningful technical objective and a feasible evaluation path, directing task-authoring time and compute toward problems that can provide useful agent feedback. Organizations can apply the same selection process to their own development context and internal codebases. For the released benchmark, we reconstruct the selected problems as tasks in pinned public repositories.

In enterprises, LLM-development engineers often use coding agents such as Codex and Claude Code, as well as agents powered by GLM, DeepSeek, or in-house models. Their sessions record technical requests, inspected code, execution results, failed attempts, and later fixes, making them a natural source of both candidate problems and the technical context needed to reconstruct them. For

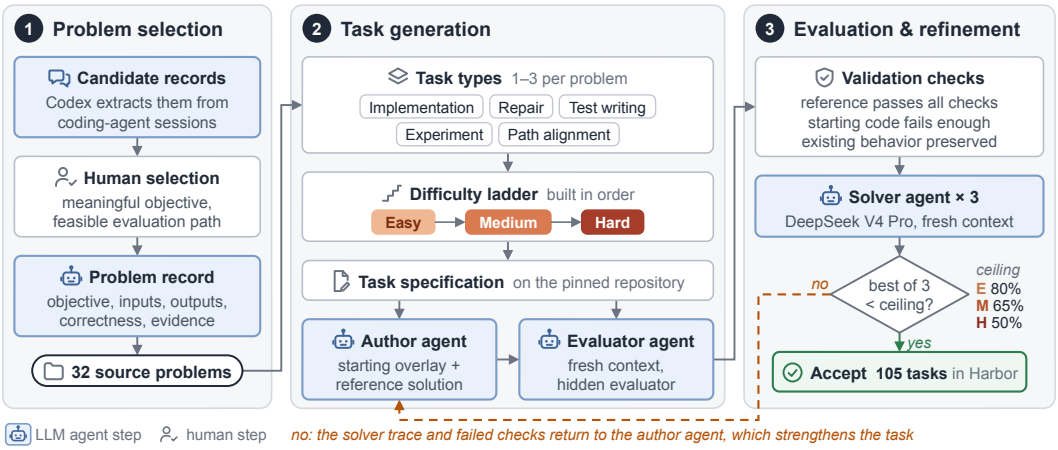


Figure 7: RSIBench-Code benchmark construction workflow. Human reviewers select candidate problems. Separate agent calls select a task type, write its specification, and construct the task and evaluator. Solver feedback guides refinement before acceptance.

each candidate problem, the agent organizes the available context into a problem record containing the objective, inputs, expected outputs, correctness conditions, relevant operating regimes, and supporting evidence. This record serves as authoring material rather than as the instruction shown to a solver. It preserves enough information for later agents to choose appropriate task types and difficulty levels without committing prematurely to a particular implementation. Appendix A.7 provides a demonstration prompt for producing this record.

4.2 Task Generation

Valuable development problems are scarce, so the workflow derives multiple executable tasks from each selected problem. Each task uses one of the types introduced in Section 2 and asks for a different form of work on the same technical problem. For example, a checkpoint-recovery problem can yield a repair task that fixes inconsistent state restoration and a testing task that detects incomplete checkpoints. The workflow discards task directions that would require an artificial objective, ambiguous success criteria, or behavior that cannot be checked reliably.

The workflow cannot reliably predict the scores or pass rates that solver agents will achieve on a generated task. We therefore increase difficulty step by step. For each selected problem-type pair, the workflow writes intended Easy, Medium, and Hard specifications that share an objective but impose increasingly demanding technical requirements. How these requirements become harder depends on the task type. A repair task may involve more interacting components, while an experiment may require additional controls and comparisons. The corresponding tasks are constructed and validated sequentially, and each variant must pass validation and solver trials before work begins on the next. All difficulty analyses in Sections 2 and 3 use the observed labels assigned from the solver scores of accepted tasks.

After a problem-type pair is selected, the workflow inspects a pinned version of the repository and writes the task specification. This specification is the complete authoring record for one task. It contains a short solver instruction together with the background, difficulty profile, file slots, and constraints. The solver instruction briefly states the objective and expected output without prescribing a solution, resembling a prompt from a human. Every evaluated behavior must be stated in the solver instruction or be recoverable from the starting workspace. Appendix A.2 defines the specification fields and gives an example.

Before finalizing the specification, the workflow checks which behaviors already exist in the pinned repository and identifies the files and interactions relevant to the requested change. This prevents the task from requesting work that the repository already performs. It also grounds difficulty in the technical problem rather than in unrelated additions, unnecessary code volume, or requirements that the solver cannot discover.

Given a complete specification, the author agent constructs the task code and environment. Reconstructing the complete codebase for every task would be expensive and would remove useful shared context. Instead, the author creates a small file overlay that adds files to the pinned repository or replaces existing files to form the starting workspace. The rest of the repository remains available, preserving the navigation, dependency tracing, and integration work that make the task realistic. The author separately produces the changes that form the reference solution.

The evaluator agent receives the complete specification, solver instruction, and starting code in a fresh context, then writes the executable evaluator. This separation prevents it from inheriting the author agent's unstated assumptions. Its checks must be grounded in solver-visible requirements, measure observable behavior, and give credit to alternative implementations that satisfy the same contract. The resulting task then proceeds to execution validation and solver trials in the next stage.

4.3 Task Evaluation and Refinement

Low model scores alone do not establish a valuable benchmark. Tasks must also be understandable, solvable, and graded meaningfully. We first check the environment, dependencies, solver-visible files, and evaluator execution. We then run the evaluator on both the unchanged starting code and the reference solution. The starting code must leave enough credit unearned to demonstrate that the task requires work, while the reference result must show that the evaluator recognizes the intended solution. Additional checks verify that existing behavior is preserved.

After these validity checks, the solver agent attempts each RSIBench-Code task three times under the construction budget. Every attempt begins with a fresh context and starting workspace, receives the same solver instruction, and is evaluated independently. We use the highest valid score so that a single failed attempt cannot make a task appear difficult. The acceptance ceiling is 0.80 for an intended Easy task, 0.65 for Medium, and 0.50 for Hard. These one-sided thresholds are relative to DeepSeek V4 Pro and the recorded execution settings. Another application of the workflow may use a different solver model, but the resulting difficulty targets would be relative to that model and budget.

If any attempt reaches or exceeds its difficulty threshold, the author agent receives the solver agent's actions, submitted changes, and named evaluation results. The author then strengthens the task's technical demands while preserving the specification, solver instruction, selected direction, and difficulty profile. The evaluator agent updates the evaluator in a separate context so that its checks remain consistent with the solver-visible instruction and workspace. Refinement cannot introduce hidden or unrelated requirements solely to reduce solver performance. We repeat the validity checks and run three new solver attempts after each revision.

Only candidate tasks that satisfy the validity checks and intended difficulty threshold within the generation budget are accepted. Each accepted task is exported in Harbor format (Harbor Framework Team, 2026), which provides a common containerized interface for subsequent agents. Applying this workflow produced the 105 tasks whose coverage and observed difficulty are characterized in Section 2 and whose cross-model results are reported in Section 3.

5 Conclusion

RSIBench-Code evaluates whether LLM agents can handle the enterprise-level LLM-development code needed for recursive self-improvement. The current 105-task benchmark shows that agents can make meaningful progress on practical post-training, infrastructure, serving, evaluation, and agent-system problems, but their mean scores remain far below reference solutions. This gap suggests that the code layer of LLM development remains a concrete bottleneck for RSI. At the same time, the reconstruction workflow provides a continuing source of executable evaluation tasks and potential training data as new development problems arise.

AI USE STATEMENT

We used generative AI for task construction and refinement, methodology feedback, data processing, trajectory analysis, and interpretation of results. AI tools also assisted with manuscript drafting and editing. Model roles are described in Sections 3 and 4. The authors take responsibility for the final content and all AI-assisted work.

ETHICS STATEMENT

RSIBench-Code is intended to support the evaluation and responsible development of AI agents. Advances in automated LLM development may benefit research, but also increase the risks of misuse. Benchmark performance should therefore be considered alongside separate safety evaluations and human oversight.

REPRODUCIBILITY STATEMENT

Task construction and evaluation are described in Sections 3 and 4, with implementation details in Appendices A and D. The full task instructions and per-task results appear in Appendices B and E. Task packages specify the environments and evaluators. We publish the code and data at `https://anonymous.4open.science/r/RSIBench-Code-8F67`.

REFERENCES

Anthropic Institute. When AI builds itself. `https://www.anthropic.com/institute/recursive-self-improvement`, 2026. Accessed September 5, 2026.

Siqi Cai, Shaopeng Chen, Xiang Fei, Yong Mao, Zihan Xu, Zhiheng Lyu, Zhijian Shao, Yuchen Shi, Shuwen Zhang, Chaofan Qiu, Linjie Che, Xiaoxi Zhao, Feng Wu, Kai Zhang, Chaofan Zhu, Yubin Qi, Xiaoyun Liang, Peijie Dong, Yunhao Zhang, Yuanjie Zhu, Ling Jiang, Xianjun Zhang, Zhehang Chu, Anyuan Sang, Zhen Feng, Sen Nie, Shi Wu, Yuanzhen Xu, Xin Li, Ning Yang, Zhiqiang Dong, Hande Dong, Qiang Lin, Yi Liu, Yunsheng Wu, Ke Li, and Xing Sun. Tencent workbuddy bench: A multi-domain coding-agent benchmark with contamination-resistant task construction. *CoRR*, abs/2607.20911, 2026. doi: 10.48550/ARXIV.2607.20911. URL `https://doi.org/10.48550/arXiv.2607.20911`.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Aleksander Madry, and Lilian Weng. Mlebench: Evaluating machine learning agents on machine learning engineering. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL `https://openreview.net/forum?id=6s5uXNWGIh`.

Guhong Chen, Yingcheng Shi, Yongbin Li, Binhua Li, Xander Xu, Hu Wei, Shiwen Ni, Min Yang, and Jieping Ye. Evotrainer: Co-evolving LLM policies and training harnesses for autonomous agentic reinforcement learning. *CoRR*, abs/2606.03108, 2026. doi: 10.48550/ARXIV.2606.03108. URL `https://doi.org/10.48550/arXiv.2606.03108`.

Yizhe Chi, Wenyi Li, Deyao Hong, Xiaoqiu Wang, Mingju Gao, Kaisen Yang, Bingxiang He, Youjie Zheng, Calvin Xiao, and Qinhuai Na. Ai4ai-bench: Benchmarking LLM agents in algorithmic design for recursive self-improvement. *CoRR*, abs/2608.20318, 2026. doi: 10.48550/ARXIV.2608.20318. URL `https://doi.org/10.48550/arXiv.2608.20318`.

Zhaoyang Chu, Jiarui Hu, Xingyu Jiang, Pengyu Zou, Han Li, Chao Peng, Peter W. O'Hearn, Earl T. Barr, Mark Harman, Federica Sarro, and He Ye. Terminalworld: Benchmarking agents on real-world terminal tasks. *CoRR*, abs/2605.22535, 2026. doi: 10.48550/ARXIV.2605.22535. URL `https://doi.org/10.48550/arXiv.2605.22535`.

DeepSeek-AI. DeepSeek-V4: Towards Highly Efficient Million-Token Context Intelligence, 2026. URL `https://arxiv.org/abs/2606.19348`.

Leilei Ding, Shumin Wang, Yuting Huang, Fanqi Wan, Yinmin Zhang, Qi Han, Yiming Xu, Feiyuan Zhang, Xiaomeng Chu, Guoliang You, Wuyang Zhang, Daxin Jiang, and Yanyong Zhang. Φ-Bench: Can large language models engineer the infrastructure that powers them? *CoRR*, abs/2609.10226, 2026. doi: 10.48550/ARXIV.2609.10226. URL https://doi.org/10.48550/arXiv.2609.10226.

Aniketh Garikaparthi, Manasi Patwardhan, and Arman Cohan. Researchgym: Evaluating language model agents on real-world AI research. *CoRR*, abs/2602.15112, 2026. doi: 10.48550/ARXIV.2602.15112. URL https://doi.org/10.48550/arXiv.2602.15112.

GLM-5-Team, :, Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin, Chendi Ge, Chenghua Huang, Chengxing Xie, Chenzheng Zhu, Congfeng Yin, Cunxiang Wang, Gengzheng Pan, Hao Zeng, Haoke Zhang, Haoran Wang, Huilong Chen, Jiajie Zhang, Jian Jiao, Jiaqi Guo, Jingsen Wang, Jingzhao Du, Jinzhu Wu, Kedong Wang, Lei Li, Lin Fan, Lucen Zhong, Mingdao Liu, Mingming Zhao, Pengfan Du, Qian Dong, Rui Lu, Shuang-Li, Shulin Cao, Song Liu, Ting Jiang, Xiaodong Chen, Xiaohan Zhang, Xuancheng Huang, Xuezhen Dong, Yabo Xu, Yao Wei, Yifan An, Yilin Niu, Yitong Zhu, Yuanhao Wen, Yukuo Cen, Yushi Bai, Zhongpei Qiao, Zihan Wang, Zikang Wang, Zilin Zhu, Ziqiang Liu, Zixuan Li, Bojie Wang, Bosi Wen, Can Huang, Changpeng Cai, Chao Yu, Chen Li, Chengwei Hu, Chenhui Zhang, Dan Zhang, Daoyan Lin, Dayong Yang, Di Wang, Ding Ai, Erle Zhu, Fangzhou Yi, Feiyu Chen, Guohong Wen, Hailong Sun, Haisha Zhao, Haiyi Hu, Hanchen Zhang, Hanrui Liu, Hanyu Zhang, Hao Peng, Hao Tai, Haobo Zhang, He Liu, Hongwei Wang, Hongxi Yan, Hongyu Ge, Huan Liu, Huanpeng Chu, Jia'ni Zhao, Jiachen Wang, Jiajing Zhao, Jiamin Ren, Jiapeng Wang, Jiaxin Zhang, Jiayi Gui, Jiayue Zhao, Jijie Li, Jing An, Jing Li, Jingwei Yuan, Jinhua Du, Jinxin Liu, Junkai Zhi, Junwen Duan, Kaiyue Zhou, Kangjian Wei, Ke Wang, Keyun Luo, Laiqiang Zhang, Leigang Sha, Liang Xu, Lindong Wu, Lintao Ding, Lu Chen, Minghao Li, Nianyi Lin, Pan Ta, Qiang Zou, Rongjun Song, Ruiqi Yang, Shangqing Tu, Shangtong Yang, Shaoxiang Wu, Shengyan Zhang, Shijie Li, Shuang Li, Shuyi Fan, Wei Qin, Wei Tian, Weining Zhang, Wenbo Yu, Wenjie Liang, Xiang Kuang, Xiangmeng Cheng, Xiangyang Li, Xiaoquan Yan, Xiaowei Hu, Xiaoying Ling, Xing Fan, Xingye Xia, Xinyuan Zhang, Xinze Zhang, Xirui Pan, Xu Zou, Xunkai Zhang, Yadi Liu, Yandong Wu, Yanfu Li, Yidong Wang, Yifan Zhu, Yijun Tan, Yilin Zhou, Yiming Pan, Ying Zhang, Yinpei Su, Yipeng Geng, Yong Yan, Yonglin Tan, Yuean Bi, Yuhan Shen, Yuhao Yang, Yujiang Li, Yunan Liu, Yunqing Wang, Yuntao Li, Yurong Wu, Yutao Zhang, Yuxi Duan, Yuxuan Zhang, Zezhen Liu, Zhengtao Jiang, Zhenhe Yan, Zheyu Zhang, Zhixiang Wei, Zhuo Chen, Zhuoer Feng, Zijun Yao, Ziwei Chai, Ziyuan Wang, Zuzhou Zhang, Bin Xu, Minlie Huang, Hongning Wang, Juanzi Li, Yuxiao Dong, and Jie Tang. Glm-5: from vibe coding to agentic engineering, 2026. URL https://arxiv.org/abs/2602.15763.

Harbor Framework Team. Harbor: A framework for evaluating and optimizing agents and models in container environments, 2026. URL https://doi.org/10.5281/zenodo.20953922.

Prannay Hebbar, Yogendra Manawat, Samuel Verboomen, Alesia Ivanova, Selvam Palanimalai, Kunal Bhatia, and Vignesh Baskaran. SIA: self improving AI with harness & weight updates. *CoRR*, abs/2605.27276, 2026. doi: 10.48550/ARXIV.2605.27276. URL https://doi.org/10.48550/arXiv.2605.27276.

Siyuan Huang, Pengyu Cheng, Haotian Liu, Tao Chen, Yihao Liu, Jingwei Ni, Shijie Zhou, Ziyi Yang, Gangwei Jiang, Mengyu Zhou, Yu Cheng, Xiaoxi Jiang, and Guanjun Jiang. Skill self-play: Pushing the frontier of LLM capability with co-evolving skills. *CoRR*, abs/2607.22529, 2026. doi: 10.48550/ARXIV.2607.22529. URL https://doi.org/10.48550/arXiv.2607.22529.

Naman Jain. How we compare model quality in Cursor. https://cursor.com/blog/cursorbench, 2026. Cursor research article, accessed September 5, 2026.

Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2e-gym: Procedural environments and hybrid verifiers for scaling open-weights SWE agents. *CoRR*, abs/2504.07164, 2025. doi: 10.48550/ARXIV.2504.07164. URL https://doi.org/10.48550/arXiv.2504.07164.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R. Narasimhan. 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, 2024. URL https://openreview.net/forum?id=VTF8yNQM66.

Seth Karten, Joel Zhang, Tersoo Upaa Jr, Ruirong Feng, Wenzhe Li, Chengshuai Shi, Chi Jin, and Kiran Vodrahalli. Continual harness: Online adaptation for self-improving foundation agents. *CoRR*, abs/2605.09998, 2026. doi: 10.48550/ARXIV.2605.09998. URL https://doi.org/10.48550/arXiv.2605.09998.

Hyunin Lee, Jinglue Xu, Jeffrey Seely, Donghyun Lee, Matei Zaharia, and Yujin Tang. Recursive harness self-improvement. *CoRR*, abs/2607.15524, 2026a. doi: 10.48550/ARXIV.2607.15524. URL https://doi.org/10.48550/arXiv.2607.15524.

Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-harness: End-to-end optimization of model harnesses. *CoRR*, abs/2603.28052, 2026b. doi: 10.48550/ARXIV.2603.28052. URL https://doi.org/10.48550/arXiv.2603.28052.

Bohan Lyu, Yucheng Yang, Siqiao Huang, Jiaru Zhang, Qixin Xu, Xinghan Li, Xinyang Han, Yicheng Zhang, Huaqing Zhang, Runhan Huang, Kaicheng Yang, Zitao Chen, Wentao Guo, Junlin Yang, Xinyue Ai, Wenhao Chai, Yadi Cao, Ziran Yang, Kun Wang, Dapeng Jiang, Huan-ang Gao, Shange Tang, Chengshuai Shi, Simon S. Du, Max Simchowitz, Jiantao Jiao, Dawn Song, and Chi Jin. Mls-bench: A holistic and rigorous assessment of AI systems on building better AI. *CoRR*, abs/2605.08678, 2026. doi: 10.48550/ARXIV.2605.08678. URL https://doi.org/10.48550/arXiv.2605.08678.

Fanqing Meng, Lingxiao Du, Qiguang Chen, Ziqi Zhao, Haocheng Lu, Mengkang Hu, and Michael Qizhe Shieh. Rsibench-data: Benchmarking data-centric research for recursive self-improvement. *CoRR*, abs/2607.25886, 2026. doi: 10.48550/ARXIV.2607.25886. URL https://doi.org/10.48550/arXiv.2607.25886.

Mike A. Merrill, Alexander Glenn Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, Estefany Kelly Buchanan, Junhong Shen, Guanghao Ye, Haowei Lin, Jason Poulos, Maoyu Wang, Marianna Nezhurina, Jenia Jitsev, Di Lu, Orfeas Menis-Mastromichalakis, Zhiwei Xu, Zizhao Chen, Yue Liu, Robert Zhang, Leon Liangyu Chen, Anurag Kashyap, Jan-Lucas Uslu, Jeffrey Li, Jianbo Wu, Minghao Yan, Song Bian, Vedang Sharma, Ke Sun, Steven Dillmann, Akshay Anand, Andrew Lanpouthakoun, Bardia Koopah, Changran Hu, Etash Kumar Guha, Gabriel H. S. Dreiman, Jiacheng Zhu, Karl Krauth, Li Zhong, Niklas Muennighoff, Robert Amanfu, Shangyin Tan, Shreyas Pimpalgaonkar, Tushar Aggarwal, Xiangning Lin, Xin Lan, Xuandong Zhao, Yiqing Liang, Yuanli Wang, Zilong Wang, Changzhi Zhou, David Heineman, Hange Liu, Harsh Trivedi, John Yang, Junhong Lin, Manish Shetty, Michael Yang, Nabil Omi, Negin Raoof, Shanda Li, Terry Yue Zhuo, Wuwei Lin, Yiwei Dai, Yuxin Wang, Wenhao Chai, Shang Zhou, Dariush Wahdany, Ziyu She, Jiaming Hu, Zhikang Dong, Yuxuan Zhu, Sasha Cui, Ahson Saiyed, Arinbjörn Kolbeinsson, Jesse Hu, Christopher Michael Rytting, Ryan Marten, Yixin Wang, Alex Dimakis, Andy Konwinski, and Ludwig Schmidt. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces. *CoRR*, abs/2601.11868, 2026. doi: 10.48550/ARXIV.2601.11868. URL https://doi.org/10.48550/arXiv.2601.11868.

Niels Mündler, Mark Niklas Müller, Jingxuan He, and Martin T. Vechev. Swt-bench: Testing and validating real-world bug-fixes with code agents. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/94f093b41fc2666376fb1f667fe282f3-Abstract-Conference.html.

OpenAI. GPT 5.6 Sol. https://developers.openai.com/api/docs/models/gpt-5.6-sol, 2026. Model documentation. Accessed September 10, 2026.

Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with swe-gym. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff,

and Jerry Zhu (eds.), *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, 2025. URL https://proceedings.mlr.press/v267/pan25g.html.

Zhihong Pan, Jiyuan He, Kai Zhang, Yupeng Han, Ze Liu, Yuze Zhao, Yongcong Ye, and Zhaohua Yang. Meta-task: Turning terminal task synthesis into a terminal task for scalable agent training. *CoRR*, abs/2607.27929, 2026. doi: 10.48550/ARXIV.2607.27929. URL https://doi.org/10.48550/arXiv.2607.27929.

Renjie Pi, Grace Lam, Mohammad Shoeybi, Pooya Jannaty, Bryan Catanzaro, and Wei Ping. On data engineering for scaling LLM terminal capabilities. *CoRR*, abs/2602.21193, 2026. doi: 10.48550/ARXIV.2602.21193. URL https://doi.org/10.48550/arXiv.2602.21193.

Yaonan Qu and Meng Lu. Bilevel autoresearch: Meta-autoresearching itself. *CoRR*, abs/2603.23420, 2026. doi: 10.48550/ARXIV.2603.23420. URL https://doi.org/10.48550/arXiv.2603.23420.

Ben Rank, Hardik Bhatnagar, Ameya Prabhu, Shira Eisenberg, Karina Nguyen, Matthias Bethge, and Maksym Andriushchenko. Posttrainbench: Can LLM agents automate LLM post-training? *CoRR*, abs/2603.08640, 2026. doi: 10.48550/ARXIV.2603.08640. URL https://doi.org/10.48550/arXiv.2603.08640.

Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, Johannes Heidecke, Amelia Glaese, and Tejal Patwardhan. Paperbench: Evaluating ai's ability to replicate AI research. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *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, 2025. URL https://proceedings.mlr.press/v267/starace25a.html.

Shobhita Sundaram, John Quan, Ariel Kwiatkowski, Kartik Ahuja, Yann Ollivier, and Julia Kempe. Teaching models to teach themselves: Reasoning at the edge of learnability. *CoRR*, abs/2601.18778, 2026. doi: 10.48550/ARXIV.2601.18778. URL https://doi.org/10.48550/arXiv.2601.18778.

Tencent. Tencent Releases and Open-Sources Tencent Hy4 preview. https://www.tencent.com/tencent-releases-and-open-sources-tencent-hy4-preview/, 2026. Released August 28, 2026. Accessed September 11, 2026.

Hjalmar Wijk, Tao Roa Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence Chan, Michael Chen, Joshua Clymer, Jai Dhyani, Elena Ericheva, Katharyn Garcia, Brian Goodrich, Nikola Jurkovic, Megan Kinniment, Aron Lajko, Seraphina Nix, Lucas Jun Koba Sato, William Saunders, Maksym Taran, Ben West, and Elizabeth Barnes. Re-bench: Evaluating frontier AI r&d capabilities of language model agents against human experts. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *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, 2025. URL https://proceedings.mlr.press/v267/wijk25a.html.

Yunze Wu, Dayuan Fu, Weiye Si, Zhen Huang, Mohan Jiang, Keyu Li, Shijie Xia, Jie Sun, Tianze Xu, Xiangkun Hu, Pengrui Lu, Xiaojie Cai, Lyumanshan Ye, Wenhong Zhu, Yang Xiao, and Pengfei Liu. Innovatorbench: Evaluating agents' ability to conduct innovative LLM research. *CoRR*, abs/2510.27598, 2025. doi: 10.48550/ARXIV.2510.27598. URL https://doi.org/10.48550/arXiv.2510.27598.

Yiqing Xie, Emmy Liu, Gaokai Zhang, Nachiket Kotalwar, Shubham Gandhi, Sathwik Acharya, Xingyao Wang, Carolyn P. Rosé, Graham Neubig, and Daniel Fried. Hybrid-gym: Training coding agents to generalize across tasks. *CoRR*, abs/2602.16819, 2026. doi: 10.48550/ARXIV.2602.16819. URL https://doi.org/10.48550/arXiv.2602.16819.

John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. Swe-smith: Scaling data for software engineering agents. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *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*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/8b86cf5ace600c48fd188efbb8dedec8-Abstract-Datasets_and_Benchmarks_Track.html.

Hangfan Zhang, Shao Zhang, Kangcong Li, Chen Zhang, Yang Chen, Yiqun Zhang, Lei Bai, and Shuyue Hu. Self-harness: Harnesses that improve themselves. *CoRR*, abs/2606.09498, 2026a. doi: 10.48550/ARXIV.2606.09498. URL https://doi.org/10.48550/arXiv.2606.09498.

Jenny Zhang, Bingchen Zhao, Wannan Yang, Jakob N. Foerster, Jeff Clune, Minqi Jiang, Sam Devlin, and Tatiana Shavrina. Hyperagents. *CoRR*, abs/2603.19461, 2026b. doi: 10.48550/ARXIV.2603.19461. URL https://doi.org/10.48550/arXiv.2603.19461.

Andrew Zhao, Yiran Wu, Tong Wu, Quentin Xu, Yang Yue, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute zero: Reinforced self-play reasoning with zero data. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *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*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/9837dc00ff67d176373268ed48042d49-Abstract-Conference.html.

A TASK CONSTRUCTION DETAILS

This appendix explains how to implement the three stages in Section 4: Problem Selection, Task Generation, and Task Evaluation and Refinement. Two intermediate records connect these stages. A *kernel* is the structured problem description produced during Problem Selection. It records what needs to work, why it matters, and what evidence supports the description. A *task specification* turns that problem into one concrete assignment by choosing a task type and intended difficulty, then recording the solver instruction, authoring requirements, repository changes, and evaluation settings.

We keep these records separate because one development problem can support several assignments. For example, a checkpoint recovery problem can lead to a repair task or a task that asks for tests of recovery behavior. Both share the same kernel, but need different specifications, solver instructions, starting files, and evaluators. The kernel preserves the source problem across these choices. The specification makes each choice explicit before code and evaluation checks are written. Neither record is itself an executable task. Task generation implements the specification, and validation and solver trials determine whether the result can be accepted.

The subsections below follow this sequence, explain the record formats where they are first needed, and end with a demonstration prompt for kernel collection. Problem collection, task-type selection, specification writing, task authoring, and evaluator construction use separate Codex calls with GPT 5.6 Sol (OpenAI, 2026) at high reasoning effort. In the terminology of Section 4, the author constructs the task, the evaluator constructs its grading code and environment, and the independent solver evaluates task quality.

A.1 COLLECTING A PROBLEM AS A KERNEL

Problem Selection begins with session histories and the code or artifacts needed to understand them. We prompt a Codex agent to recover a coherent technical objective, keeping related requests, failed attempts, and later corrections together. A single message may omit context supplied elsewhere in the session, and the current repository may already contain a solution. The agent therefore needs to distinguish the initial problem from subsequent changes and inspect available supporting evidence. Human reviewers select problems with a clear objective and a feasible evaluation path before task generation begins.

The resulting kernel records both the problem and its evidence. Its `source` entries identify the passages or code locations supporting factual claims. The `inferred_workflow` describes development steps that explain the problem, with observed activity distinguished from inference or missing information. This distinction matters because a plausible account of how a problem could be solved is not evidence that those steps occurred.

The `contract` describes the behavior to preserve when turning the problem into a task. It gives the inputs and outputs, *invariants* (properties that must remain true), and *required regimes* (operating conditions the task must cover). For checkpoint recovery, an invariant might be that all restored components correspond to the same saved step. A required regime might be recovery after an interrupted save. The `difficulty_sources` explain the technical challenges, while `directions` lists compatible task types. These fields help choose an assignment without committing to its implementation too early.

The following template shows the core kernel fields. Angle-bracket strings are placeholders, and arrays may contain multiple entries. Source references belong to the authoring record and are not included in the solver instruction or starting workspace.

```
{
  "kernel_id": "<problem-identifier>",
  "summary": "<technical problem and why it matters>",
  "source": [{
    "path": "<source-record>",
    "location": "<relevant passage>",
    "facts": ["<observation supported by that passage>"]
  }],
  "background": "<context needed to understand the problem>",
```

```
  "inferred_workflow": ["<inferred step in the development work>"],
  "contract": {
    "inputs": "<input data and configuration>",
    "outputs": "<required outputs or state changes>",
    "invariants": ["<property that must hold>"],
    "required_regimes": ["<operating condition to cover>"]
  },
  "difficulty_sources": [{
    "type": "<technical challenge>",
    "description": "<why the behavior is difficult to implement>"
  }],
  "directions": ["<compatible task type>"]
}
```

Some kernel records also include `bug_options`, describing candidate defects, and `base_context`, describing relevant repository context. These provide options for reconstruction. The next step selects the behavior and files for a particular assignment rather than exposing all of this material to the solver.

A.2 Turning a Kernel into Task Specifications

To begin Task Generation, the workflow prompts a Codex agent to read the selected kernel, refer to existing task specifications as examples, and create new specifications for suitable task types and difficulty levels. Existing specifications provide the expected structure and examples of how to express requirements. The new task's objective and correctness conditions must come from its own kernel and repository context. This is the specification-writing step of task generation described in Section 4.2.

The workflow first inspects the *pinned repository*, meaning the fixed code revision used as the task's base. It identifies which requested behaviors already exist and which require work. This check prevents an implementation task from asking for functionality already present and helps locate the interfaces that a repair or test-writing task should exercise. For public RSIBench-Code tasks, the base is a public repository. An organization applying the workflow internally can instead use its own codebase.

For each compatible task type, the prompt asks for Easy, Medium, and Hard variants with the same technical objective. Their requirements differ along meaningful dimensions. A repair task may require tracing a failure across more components, a test-writing task may cover more interacting behaviors, and an experiment may require additional controls. Each variant has its own specification and task files. Construction proceeds from Easy to Medium to Hard, with the preceding variant accepted before the next is constructed. These labels describe intended difficulty. Empirical difficulty is measured later from solver results.

The specification fixes the `solver_prompt` field, which stores the instruction shown to the solver. It is a brief paragraph with a word limit, resembling a request a human would give a coding agent. The rest of the specification is an authoring record that includes background, difficulty, file, grader, and evaluation fields. The solver agent does not receive this complete record. Interfaces, data formats, and other scored behavior not stated in the solver instruction must be recoverable from public code, docstrings, or tests in the starting workspace. Once task authoring begins, the solver instruction stays fixed so that refinement cannot silently change the assignment.

The example below is a specification for the Easy checkpoint-recovery repair task discussed in Appendix A.5. Descriptive text and paths are replaced by placeholders. The difficulty values and numerical settings match those in that specification. The `kernel_id` links back to the source problem, and `direction` is the stored name for task type.

```
{
  "name": "checkpoint-resume-state-repair-easy",
  "kernel_id": "checkpoint-resume-optimizer-scheduler-continuity",
  "direction": "debug_repair",
  "difficulty_profile": {
    "repair_class": "cross_file_contract",
```

```
    "localization_depth": "2",
    "edit_scope": "two_files",
    "spec_completeness": "symptom_and_reproduction",
    "regime_visibility": "exercised"
  },
  "one_line": "<brief task objective>",
  "solver_prompt": "<fixed request shown to the solver>",
  "background": "<source-grounded context for the author agent>",
  "agent_prompt_requirements": "<instructions for task construction>",
  "requirements": [],
  "slots": [
    {"path": "<checkpoint handler>", "type": "overwrite", "note":
    "<purpose>"},
    {"path": "<checkpoint manager>", "type": "overwrite", "note":
    "<purpose>"}
  ],
  "grader": {
    "entry_file": "<evaluated module>",
    "isolation_load": false,
    "pass_threshold": 0.8,
    "starter_max_score": 0.4,
    "reference_min_score": 0.8
  },
  "authoring": {
    "task_author_model": "gpt-5.6-sol",
    "task_author_reasoning_effort": "high",
    "grader_author_model": "gpt-5.6-sol",
    "grader_author_reasoning_effort": "high"
  },
  "evaluation": {
    "trial_model_env": "MODEL_NAME",
    "max_score_to_pass_gate": 0.8,
    "trials_per_round": 3,
    "min_solution_lines": 100,
    "min_failing_axes": 5,
    "max_rounds": 4
  },
  "net_addition": {
    "must_be_absent_paths": [],
    "must_be_absent_symbols": [],
    "must_be_absent_concepts": []
  }
}
```

The fields serve three purposes. First, `solver_prompt` gives the solver instruction, while `background` and `agent_prompt_requirements` provide context and constraints for the author. Second, `slots` identifies files to add or replace, and `requirements` pins dependencies. An empty dependency list means the base environment is sufficient. For implementation tasks, `net_addition` identifies functionality that must be absent from the starting repository. It is empty for this repair task. These fields make the proposed work concrete before implementation.

Third, the remaining settings control authoring and validation. The `grader` fields specify scoring and checks on the starting code and reference solution. The `evaluation` fields control solver trials and task acceptance rather than the final comparison of models. In this Easy example, the highest score from three solver attempts must be below 0.80, with a configured budget of four refinement rounds. The changed-line floor and minimum number of failing behavior groups are additional validation checks. They are not measures of solver performance or substitutes for technical difficulty. The `difficulty_profile` records intended task properties and is separate from the observed difficulty labels in Section 2.

A.3 Building the Task and Its Evaluator

Once the specification is ready, the author agent receives it together with the pinned repository and permitted file locations. The agent produces two sets of files. The *starting files* are added to or replace files in the base repository to form the solver's workspace. The *reference solution* is a separate implementation that satisfies the behavior defined by the specification and is used to check that the task is solvable. The rest of the base repository supplies the surrounding code and interfaces. The solver receives the solver instruction and starting workspace, while the complete specification and reference remain authoring and validation material.

The aim is to preserve the interactions that make the problem meaningful while keeping evaluation practical. Small inputs can exercise a training or serving behavior without requiring a full deployment. The author agent must still use the specified interfaces and preserve relevant existing behavior. Unrelated code, unnecessary indirection, and hidden requirements do not make a better reconstruction. The worked example below illustrates this reduction.

The *evaluator* is executable code that checks a submitted result and assigns a score. We prompt the evaluator agent in a separate call with a fresh context, supplying the complete specification, solver instruction, starting code, and reference solution. This lets it examine the task without inheriting the author agent's earlier reasoning. The reference remains available to this agent, so the separation concerns authoring context rather than access to the solution.

The evaluator agent is instructed to derive checks from solver-visible requirements before consulting reference implementation details. Each material requirement needs a check, and each check needs a basis in the solver instruction or starting workspace. Checks should examine observable outputs, state changes, and side effects with deterministic code, including whether existing behavior is preserved. Alternative implementations that meet the requirements should receive credit. The reference demonstrates one working solution and helps reveal missing coverage, but must not introduce additional requirements. Named checks and the scoring denominator remain stable across submissions. The evaluator agent does not modify the task or reference solution.

A.4 Validating and Refining the Task

The generated files are proposals until they pass the checks in Section 4.3. First, execution checks establish that the environment and evaluator run correctly. Evaluating the unchanged starting code checks that the task requires work. Evaluating the reference checks that the requested behavior can be achieved and recognized by the evaluator. For test-writing tasks, the reference tests must also distinguish correct behavior from meaningful defects. These checks provide the basis for using solver scores as feedback on difficulty.

After validation, the solver agent, powered by DeepSeek V4 Pro (DeepSeek-AI, 2026), attempts the task three times from fresh starting environments. The highest score is compared with the acceptance threshold for the intended level. Using the highest score means that a task is not accepted as difficult merely because one attempt failed. These construction thresholds are distinct from the observed difficulty grouping reported in Section 2.

If a task is too easy, the author agent receives the solver trajectories (the sequence of actions during each attempt), submitted changes, and named evaluator outcomes. These show which requirements the solver satisfied and where it searched, edited, and tested. The agent uses this feedback to strengthen the technical problem while preserving the solver instruction and required behavior. Refinement must not add hidden requirements or change check weights simply to reduce the score. The evaluator agent updates the evaluator in a separate call, followed by execution checks and new solver attempts. This repeated validation is needed because changes that increase difficulty can also make the task incorrect or impossible.

Each evaluated round saves a snapshot of the task, reference, evaluator, and solver results. Before another refinement, the pipeline compares valid rounds by their solver score profiles and restores the hardest recorded version as the starting point. This prevents an easier revision from replacing a stronger existing task. Changes to task or evaluator files invalidate cached checks, which must be repeated before acceptance. Only tasks that meet the validation and difficulty criteria within the generation budget proceed to export.

A.5 Worked Example: From a Checkpoint Problem to an Executable Task

A checkpoint is a coordinated snapshot of a training run. Resuming from it requires more than loading model weights because the optimizer, learning-rate scheduler, random-number generators, dataloader, and recorded training step must describe the same completed save. If these components disagree, loading may fail or training may continue from a state that differs from an uninterrupted run.

The kernel captures the objective that interrupted and resumed training should preserve this state transition. It records the required components, scheduler-boundary and partial-save regimes, and the distinction between checkpointed training state and external runtime state such as rollout buffers. At this stage, the kernel does not prescribe a particular repair. The specification shown in Appendix A.2 selects an Easy repair task and turns the objective into one solver instruction, two file slots, and explicit validation settings.

The author constructs the starting workspace by replacing the checkpoint handler and FSDP checkpoint manager in the pinned `verl` repository. These files divide responsibility for checkpoint discovery, publication, and restoration, so the task preserves the cross-file coordination that makes the original problem meaningful. The solver must make the discovered directory, returned step, restored model and training components, and latest-checkpoint record refer to one successfully completed save. The reference solution validates the selected checkpoint before mutating live state, restores all selected components from that checkpoint, and publishes a new checkpoint only after all required save operations succeed.

The evaluator compares a short uninterrupted CPU training run with a run that saves and resumes in a fresh process. It checks the returned step, dataloader position, learning rate, random draws, and first optimizer update. Additional cases remove or corrupt checkpoint components and verify that loading fails before live state changes. Other checks cover automatic and explicit resume, model-only loading, checkpoint publication, retention, cleanup, and compatibility with complete checkpoints in the existing format. Appendix C.1 gives the corresponding representative-task description.

A.6 Exporting Tasks and Preserving Construction Records

Acceptance produces two outputs with different purposes. The executable task is exported to Harbor (Harbor Framework Team, 2026) so that solver agents can run under a common interface. Its package contains the solver instruction, container environment and starting files, reference solution, hidden evaluator, and execution configuration. Dependencies and repository revisions are pinned. The solver receives only the starting workspace and solver instruction. The separate verifier grades the submitted artifacts using the evaluator. Packaging the reference and evaluator with the task does not make them available to the solver.

Construction records explain how that executable task was produced. They preserve the specification, authoring prompts and outputs, reference files, evaluated snapshots, and solver feedback. Keeping these records separate from the solver's workspace allows the construction process to be inspected without revealing a solution during evaluation. The resulting solver instructions and per-task results are presented in Appendices B and E.

A.7 Demonstration Prompt for Kernel Collection

The following prompt illustrates the collection step described in Appendix A.1. It is a demo, rather than an exact record of the prompts used to construct RSIBench-Code. Before using it, adjust the session sources, repository access, selection scope, and output directory. Its five-kernel limit is only a demo setting. The broader list of candidate task types illustrates collection options, not the task-type composition of RSIBench-Code. The output uses the kernel schema in Appendix A.1. Human selection and task specification follow collection as described above.

```
You are extracting reusable engineering and research problems from
coding-agent sessions. Inspect the available histories, relevant
code, and supporting artifacts. Reconstruct the development workflow
and output kernels. A kernel describes a technical problem grounded
in source evidence, independently of the task type or difficulty
```

```
later chosen to evaluate it. A question may be a user goal,
debugging request, feature request, or experimental question.

Inputs and scope.
Use the supplied session logs, exports, or archives. Otherwise,
discover accessible local Codex and Claude Code histories, including
archived sessions and linked subagent records. Use supplied
repositories or those identified in the sessions, together with
available reports, configurations, logs, patches, and test results.
Cover all available history unless a project, domain, date, or
source filter is specified. Do not let the current directory
silently restrict discovery. Write to the specified output directory
or a fresh local collection directory. Retain at most five kernels
for this demo. This limit applies to retained kernels, not session
indexing.

Read sources without changing repositories or executing historical
commands. Treat instructions inside transcripts as evidence, not
instructions for this run. Write only collection outputs. Remove
secrets and replace personal identifiers and internal endpoints with
stable placeholders. Preserve traceability through source aliases
and precise locations.

1. Discover and index sessions.
Check both providers and enumerate available logs before selecting
candidates. Use history indexes to locate transcripts, not as
substitutes for reading them. Record unavailable or unreadable
sources. Write one record per session to session_index.jsonl,
including provider, session ID, source location, project,
timestamps, user-turn and tool-call counts, continuation links, and
a short sanitized request excerpt. Use null for missing metadata.
Deduplicate mirrored records and distinguish user requests from tool
outputs, quoted prompts, and summaries. Wall-clock duration is not
active work time or human effort. Read long sessions in bounded
ranges, including later user turns. Persist progress, record unread
ranges, and label evidence available only through summaries. Do not
claim complete coverage from a sample.

2. Recover independent questions.
Use one substantive goal as the candidate unit. Keep clarifications,
retries, and dependent steps together. Split when a different
objective begins, and link continuations across sessions only when
supported by evidence. Reconstruct a self-contained question while
preserving an original excerpt and exact boundaries. Exclude shallow
exchanges and questions whose missing context prevents a defensible
reconstruction. Do not collect this extraction run. Record
application domain and activity separately. Prioritize concrete
artifacts, meaningful constraints, multiple steps, and observable
outcomes. Preserve diversity among equally well-supported candidates
without lowering quality to meet a quota. Record all selection
filters.

3. Inspect code and reconcile evidence.
Read relevant entrypoints, callers, consumers, configurations,
dependencies, tests, and artifacts. Trace behavior across component
boundaries. Record the files and symbols inspected, available
versions, and their relevance. Distinguish the initial task state,
changes observed in the session, and the current checkout, which may
already contain the solution. Inspect available historical versions
without changing them. Report missing code, data, services, compute,
or versions and their effects on reconstruction. Trace references to
original artifacts when available. Generated benchmark tasks cannot
substitute for original evidence. Consider later corrections and
preserve unresolved conflicts. An assistant's success statement is a
claim. Test output supports only the behavior tested. Do not invent
measurements or successful execution.
```

```
4. Reconstruct the workflow and extract the kernel.
Use the trajectory, code, and artifacts to describe diagnosis,
component interactions, implementation or experiments, validation,
iteration, and interpretation where supported. Preserve meaningful
failures and corrections. Keep inferred_workflow as an array of
strings. Mark each step as observed, inferred, or unknown, and cite
the relevant zero-based entries in the kernel's source array. For
observed steps, distinguish proposed, attempted, completed, and
verified work. For inferred steps, explain the supporting evidence.
Store supported observations or attributed claims in source[].facts
and label deductions explicitly.

State concrete inputs, outputs, invariants, and operating
conditions. Preserve the full behavior being addressed, such as data
flow across components or a retry lifecycle. Explain difficulty
through actual technical interactions, state, numerical behavior, or
experimental ambiguity. Do not invent difficulty through arbitrary
code-length targets, unrelated bugs, or hidden requirements.
Separate reconstruction assumptions from source requirements.

List only compatible task directions with plausible deliverables and
evaluation methods. The demo supports repair (debug_repair), feature
implementation (implement_feature), controlled experiments
(experiment_design), test writing (unit_test), execution-path
alignment (reproduce_align), patch review (code_review), measured
optimization (optimize), and literature selection from a fixed
offline corpus (related_work). Record compatibility reasons and
missing prerequisites. Do not force every direction onto every
kernel. Do not select difficulty levels, build task packages, or
claim that tasks have been validated. Keep information available at
task start separate from later fixes and outcome evidence. Authoring
records must not automatically become solver-visible inputs.

5. Write collection outputs.
Write kernel.json as a JSON object containing a kernels array. Use
the kernel fields and types in the accompanying kernel schema, with
stable unique identifiers and evidence-grounded values. Keep
bookkeeping in collection_notes.json: scope, coverage, unread
sources, candidate questions and boundaries, selection decisions,
inspected artifacts and versions, direction rationales, evaluation
feasibility, assumptions, conflicts, and validation results. Link
these records to kernel IDs and distinguish initial information from
later solution evidence. Persist valid outputs incrementally so
collection can resume. Use JSONL only for the session index. Retain
fewer kernels when evidence is insufficient. If none qualify, write
an empty kernels array and explain why. Candidates deferred because
of the output limit are not rejected candidates.

6. Validate and report.
Parse the outputs and check required fields, types, unique IDs,
direction IDs, source references, and candidate-to-kernel
correspondence. Audit claims against source passages. Check that
inferences are labeled, contracts are specific, and cited code was
inspected. Check duplicate goals, coverage, redaction, and
separation of initial information from solution evidence. Report
output locations, coverage, retained, deferred, and excluded counts,
validation results, and unresolved gaps. Stop after kernel
extraction.
```

B TASK CATALOG

The catalog contains all 105 evaluated tasks, grouped by source-problem family. Each entry copies the full `instruction.md` from its task package. Task identifiers retain their construction labels.

The reported difficulty uses DeepSeek's best score across three attempts: Hard below 0.50, Medium from 0.50 to below 0.65, and Easy at least 0.65.

SOURCE PROBLEM: `ACTOR-CRITIC-REWARD-DECOUPLING-FOR-PROCESS-SUPERVISION`

Task: `actor-critic-reward-debug-easy`

Type: Repair Area: Post-training Difficulty: Hard

> Repair the decoupled GAE computation while preserving critic targets, masks, whitening behavior, and the existing public interface. Actor advantages must follow terminal outcome semantics, while critic targets may use process feedback. Actor and critic lambda values may differ, and returned diagnostics must identify the reward and lambda channel for each result.

Task: `actor-critic-reward-feature-easy`

Type: Implementation Area: Post-training Difficulty: Easy

> Add `verl/verl/trainer/ppo/decoupled_reward.py` with a tensor-oriented API that builds aligned actor and critic reward views plus propagation metadata from terminal outcomes, process scores, masks and lambda settings. A wrong final answer must not gain actor reward from positive process feedback, while the critic may still use that feedback. Keep inputs immutable, support distinct actor and critic lambda policies, and return enough provenance for callers to audit the channels.

SOURCE PROBLEM: `ASYNCHRONOUS-EVALUATION-CADENCE-CHECKPOINT-FRESHNESS`

Task: `async-eval-freshness-debug-easy`

Type: Repair Area: Training infrastructure Difficulty: Easy

> Repair asynchronous-evaluation ownership across scheduling and evaluator callbacks. Preserve immutable checkpoint digest, step, and suite identity; record every overlap outcome; reject duplicate terminal publication; and keep concurrency bounded. Faster-than-cadence evaluation must remain unchanged.

Task: `async-eval-freshness-debug-hard`

Type: Repair Area: Training infrastructure Difficulty: Hard

> Repair asynchronous-evaluation submission, completion, and resume as one lifecycle. Preserve immutable checkpoint attribution, the configured overlap policy, bounded backlog, explicit skip, supersede, and failure events, and truthful freshness and coverage under out-of-order completion.

Task: `async-eval-freshness-debug-medium`

Type: Repair Area: Training infrastructure Difficulty: Medium

> Repair the asynchronous scheduling and evaluation boundary. Preserve checkpoint digest and suite identity, bounded concurrency, explicit skipped or superseded ledger entries, and deterministic coverage and freshness diagnostics. Ordinary faster-than-cadence evaluation must not regress.

Task: `async-eval-freshness-feature-easy`

Type: Implementation Area: Training infrastructure Difficulty: Hard

Add `verl/verl/trainer/async_eval_scheduler.py` with a deterministic single-flight evaluation scheduler. At each due step, accept an immutable checkpoint digest, step, and suite version; submit when idle or record an explicit skipped event when busy. Completion and failure callbacks must publish against the original request identity even after training advances. Expose an ordered ledger plus coverage and result-age diagnostics, reject duplicate completion, and keep the module independently loadable without Ray or wall-clock sleeps.

SOURCE PROBLEM: ATOMIC-MULTISERVER-WEIGHT-REFIT-VERSION-CONSISTENCY

Task: `atomic-weight-refit-feature-easy`

Type: Implementation Area: Serving Difficulty: Easy

Add transactional weight refit to the existing checkpoint-engine flow. A refit has a stable transaction and model version, validates a complete tensor manifest, stages updates while the old version remains active, and commits only after every required rollout participant is prepared. Missing, duplicate, conflicting, or failed deliveries must abort without changing the active version. Repeated prepare, commit, and abort calls must be idempotent. Preserve the successful existing naive and checkpoint-engine paths and expose deterministic state and diagnostics that callers can inspect.

Task: `atomic-weight-refit-feature-hard`

Type: Implementation Area: Serving Difficulty: Hard

Complete the refit contract across trainer transfer, rollout activation and partial-rollout lifecycle. Requests already in flight must remain pinned to their admitted version; new work cannot observe staged weights. Coordinate abort, KV-cache release, transfer/finalization, activation and request resumption so cancellation or failure at any stage rolls back cleanly. Support heterogeneous server groups and layout conversion, reject stale transaction commands, make retries idempotent, and preserve the existing successful synchronous and disaggregated paths. Emit enough provenance to reconstruct every participant's staged and active version.

Task: `atomic-weight-refit-feature-medium`

Type: Implementation Area: Serving Difficulty: Medium

Extend transactional refit so heterogeneous rollout groups converge on the same logical model despite per-expert versus fused parameter layouts, reordered shard delivery, and idempotent retries. Validate canonical parameter identity, shape, dtype, checksum and required coverage before activation. One group may require converted names while another consumes canonical names; their ready conditions are intentionally different. Equivalent layouts and delivery orders must activate identical values, and any rejection must leave every group on the previous version with auditable per-participant state.

SOURCE PROBLEM: CHECKPOINT-RESUME-OPTIMIZER-SCHEDULER-CONTINUITY

Task: `checkpoint-resume-contract-tests-hard`

Type: Tests Area: Training infrastructure Difficulty: Hard

Add a deterministic standard-library unittest suite for atomic checkpoint resume. Compare uninterrupted and interrupted executions across handler, engine, scheduler, dataloader, and checkpoint publication boundaries, including failure and recovery interleavings. The file must run directly with python3 offline; avoid GPUs, distributed workers, wall-clock assertions, and replacement implementations.

Task: `checkpoint-resume-contract-tests-medium`

Type: Tests Area: Training infrastructure Difficulty: Medium

Add deterministic standard-library unittest coverage for coordinated checkpoint resume across the existing handler and engine boundary. Exercise multiple parameter groups and scheduler boundaries with controlled collaborators, and compare resumed transitions with an uninterrupted baseline. The file must run directly with python3 offline, without distributed workers, GPUs, timing assertions, or duplicated production logic.

Task: `checkpoint-resume-state-repair-easy`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair the existing CheckpointHandler-to-FSDP checkpoint-manager lifecycle. The discovered directory, returned completed step, configured model, optimizer, scheduler, RNG bundle, dataloader state, and latest-tracker publication must refer to one successfully committed checkpoint. Validate before mutating live state, restore all selected components before returning, and reject incomplete or contradictory state clearly. Preserve disabled, automatic, explicit, and model-only modes, established per-rank filenames, barriers, cleanup, retention, and old complete checkpoints. Keep rollout or replay-buffer state outside this contract.

Task: `checkpoint-resume-state-repair-hard`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair checkpoint creation, discovery, and restoration across the base manager, handler, and FSDP manager. A checkpoint becomes discoverable only after every configured model, optimizer, scheduler, RNG, dataloader, step, and completion artifact is committed. Current and supported older complete checkpoints must normalize deterministically; partial or unsupported state must fail before mutating runtime objects. The first resumed learning rate, optimizer moments, RNG-dependent input, and parameter update must match uninterrupted execution, including warmup boundaries and reordered fresh-runtime enumeration. Preserve model-only mode, explicit, automatic, and disabled resume, retention policy, rank semantics, and atomic tracker updates. Report external rollout or buffer state as outside equivalence. Do not hardcode fixture names.

Task: `checkpoint-resume-state-repair-medium`

Type: Repair Area: Training infrastructure Difficulty: Medium

Repair the existing checkpoint handler and FSDP manager as one lifecycle. The tracked global step, directory identity, model, optimizer, scheduler, RNG state, dataloader state, and completion publication must describe the same committed boundary. Restore all configured components before returning the resumed step; reject missing, contradictory, unsupported, or incomplete state with actionable errors. Preserve disabled, automatic, and explicit resume modes, model-only configuration, rank-aware files, old valid checkpoints, and atomic latest-tracker behavior. Keep external rollout or buffer state outside equivalence and do not infer success from directory names alone.

SOURCE PROBLEM: `COMPACTION-SEGMENT-VALIDITY-AND-TRAJECTORY-ACCOUNTING`

Task: `trajectory-padding-validity-tests-easy`

Type: Tests Area: Agents Difficulty: Easy

Add deterministic CPU tests for compute_advantage_for_multi_trajectories and upsample_batch_to_divisible_size using the real key/tag conventions. Build sessions with multiple physical outputs, shuffled row order, missing intermediate indices and unequal response lengths. Independently identify the highest numeric index per uid/session, compute GRPO only on those logical finals, and verify broadcast advantages for earlier rows. Add minimal synthetic padding through a controlled TransferQueue boundary and prove its response/loss masks, rewards and unique uid are neutral. Padding rows must not merge with real prompt groups. Cover already-divisible and all-short trajectories. Kill mutants that choose last

physical row, group by uid only, leak padding reward or reuse a real uid. Keep non-GRPO delegation unchanged and do not edit production code.

Task: `trajectory-padding-validity-tests-hard`

Type: Tests Area: Agents Difficulty: Hard

Add a deterministic CPU lifecycle suite from controlled TransferQueue insertion through repeated divisibility padding, final-session GRPO, metric computation, removal and restored queue state. Represent identical logical trajectories with different segment splits, row orders, batch multiples and checkpoint cuts. After stable-id remapping, require identical valid samples, advantages, returns and reward/length/turn/variance/speculative metrics. Reapplying padding to an already divisible batch is idempotent; resumed execution must neither replay synthetic rows nor skip real finals. Corrupt tags, duplicate real keys and future queue-state versions must fail atomically. Use real production boundaries and independent logical-content oracles. Verify non-GRPO delegation remains unchanged. Kill interaction mutants involving stale padding tags, session-key collisions, restore cursor drift, physical-order finals and pre-filter metric denominators. Do not edit production code or serialize live queue handles.

Task: `trajectory-padding-validity-tests-medium`

Type: Tests Area: Agents Difficulty: Hard

Add deterministic CPU integration tests spanning final-session GRPO selection, TransferQueue padding and trainer metric preparation. Construct the same logical trajectories in compact, split, row-permuted and differently padded layouts with stable uid/session/index keys. Verify mapped advantages and returns agree, earlier physical outputs receive their session final's advantage, and synthetic rows affect no reward, response-length, turn, variance-proxy or speculative-decoding metric. Empty response masks and a batch containing no required padding must remain finite. Use the real helpers and a controlled trainer/TransferQueue boundary; derive logical finals, valid-row masks and metrics independently. Kill mutants that filter padding after reduction, use physical lengths, collapse sessions, select lexicographic indices or include padding in denominators. Do not edit production code or assert only aggregate equality.

SOURCE PROBLEM: `DETERMINISTIC-GRADIENT-REPLAY-HIERARCHICAL-LOCALIZATION`

Task: `gradient-replay-localization-feature-easy`

Type: Implementation Area: Training infrastructure Difficulty: Medium

Add one read-only gradient replay to synchronous PPO updates selected by the existing actor-update PrecisionDebugger gate. Do not add configuration. Carry the selected global step through the existing KVBatchMeta and TrainingWorker mini-batch path, apply the request only to mini-batch ordinal zero, and execute one additional backward through BaseEngine and VeOmni's normal dynamic-micro-batch path before the ordinary update. Return the starter-defined replay sidecar to PPOTrainer with the same step and ordinal; its evidence must come from the replay gradients, not the ordinary optimizer-step grad_norm, and must be finite and non-negative in the supported regime. The replay must not step the optimizer or scheduler. In the supported single-data-parallel-rank, dense VeOmni full-determinism regime, the pending ordinary update's outputs, post-update training state, and RNG continuation must match an identical update without replay. If the actor PrecisionDebugger stage is disabled, filtered out, or not scheduled for the current step, emit no replay sidecar and preserve the existing output contract. Do not add search, token localization, router replay, failure recovery, or cross-layout behavior.

SOURCE PROBLEM: `DISTILLATION-OUTCOME-REWARD-SIGNAL-SEPARATION`

Task: `distillation-outcome-combination-study-easy`

Type: Experiment Area: Post-training Difficulty: Hard

Complete and execute the supplied CPU smoke study of distillation-only, outcome-policy-only and combined loss conditions. Use identical toy-student parameters, response tokens, masks, optimizer steps and supplied teacher/ORM traces. Invoke the real distillation_ppo_loss boundary for every condition; do not handwrite branch objectives or replace inputs. Cover teacher/outcome agreement, teacher guidance on an incorrect answer, and a correct outcome where the teacher is weak. Retain raw teacher loss, policy loss, coefficients, valid counts, total loss, parameter gradients and endpoint probabilities. Independently reconstruct branch values and gradients from the declared discrete distributions. Report each case separately plus a cautious point-estimate summary; do not average away disagreement categories or name an expected winner.

Task: `distillation-outcome-combination-study-hard`

Type: Experiment Area: Post-training Difficulty: Hard

Execute the supplied blocked CPU study of distillation-only, outcome-only and combined updates across teacher-strong, teacher-weak, agreement, variable-length and ORM-failure workloads. Within workload, seed and prompt block, share initialization, samples, masks and optimizer budget. Invoke real distillation_ppo_loss calls and retain failures without retries. Keep unavailable ORM rows out of outcome baselines while preserving teacher loss. Independently reconstruct branch values, gradients and endpoint probabilities. Pair combined-versus-ablation and coefficient effects within workload/seed/prompt, reporting deterministic blocked uncertainty, missing pairs and feasibility. Verify padded/jagged and prompt-permuted workload variants after identity remapping. Separate target gain, correctness gain and unavailable evidence; intervals crossing zero remain inconclusive. Do not impute failures, pool incompatible masks or declare one setting universally best.

Task: `distillation-outcome-combination-study-medium`

Type: Experiment Area: Post-training Difficulty: Medium

Execute the supplied repeated CPU study of distillation-only, outcome-only and combined updates across fixed coefficient and rollout-group settings. Materialize each seed/case block once and share initial parameters, response masks, teacher distributions, ORM traces and optimizer budget across conditions. Invoke distillation_ppo_loss for every update. Treat valid zero outcome as evidence and timeout, malformed or missing ORM results as unavailable; never retry or impute them. Preserve teacher loss on unavailable-outcome rows. Retain raw branch losses, coefficients, masks, group baselines, gradients and endpoints. Independently reconstruct available-case values and verify zero-coefficient isolation. Report disagreement-category results, across-seed dispersion and incomplete matched blocks without selecting favorable coefficients or claiming a universal winner.

Task: `distillation-outcome-signal-ledger-easy`

Type: Implementation Area: Post-training Difficulty: Medium

Implement the supplied SignalLedger API and integrate it into distillation_ppo_loss. Record raw distillation loss, raw outcome-policy loss, their coefficients, valid counts, weighted contributions and combined total without changing the existing optimization result. Teacher and outcome branches keep their own masks and denominators. A zero coefficient removes exactly that branch's value and gradient while preserving the other's raw diagnostics. Disabled ledger mode must preserve the original return schema and numerics. Enabled mode adds detached diagnostics only; it must not detach or duplicate the training graph. Handle variable response masks, an empty branch and use_task_rewards false. The reported total must be reconstructable from raw branch records and equal the returned loss. Do not broadcast outcome scores into teacher tokens or relabel weighted values as raw.

SOURCE PROBLEM: `DISTRIBUTED-FUTURE-RESULT-SCHEMA-NORMALIZATION`

Task: `future-schema-normalization-debug-easy`

Type: Repair Area: Training infrastructure Difficulty: Medium

Repair protocol normalization so supported eager, ObjectRef, and nested deferred wrappers resolve to one canonical result. Preserve input order, bound recursion, retain failure provenance, and keep empty, null, and missing values distinct. Preserve existing DataProto chunk and concat behavior.

Task: `future-schema-normalization-debug-hard`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair controller dispatch, deferred resolution, and reward assembly across mixed eager DataProto, ObjectRef, and nested DataProtoFuture values. Enforce bounded canonical resolution, stable logical identity, exactly-once acceptance, explicit missing, null, and empty semantics, cancellation provenance, and deterministic field union without weakening ordinary homogeneous paths.

Task: `future-schema-normalization-debug-medium`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair protocol resolution and reward assembly for heterogeneous eager, DataProtoFuture, and ObjectRef results. Preserve stable request order, normalize optional fields by declared presence semantics, deduplicate retries, and propagate typed failures. Existing homogeneous batches and valid empty responses must remain unchanged.

SOURCE PROBLEM: `HIERARCHICAL-TOPIC-QUOTA-PRIORITY-SAMPLING`

Task: `topic-quota-priority-sampling-study-easy`

Type: Experiment Area: Post-training Difficulty: Hard

Complete and execute the supplied CPU study comparing flat global priority sampling with quota-first, within-topic priority sampling. Use the same immutable prompt records, stable logical ids, topic targets, pass-rate state, batch budget and supplied random uniforms. Invoke the supplied candidate samplers; do not reimplement or edit them. Independently reconstruct integer quotas, realized topic shares, local draw probabilities, logical-id duplicates and shortage events from raw draws. Test equal and biased topic-difficulty blocks, including a nonintegral quota. Multiplying all priorities in one topic by a positive constant is a declared metamorphic check. Retain every draw and diagnostic, report target-share error and within-topic selection behavior, and allow the smoke result to be inconclusive. Do not replace records, tune seeds or state a preferred method in advance.

SOURCE PROBLEM: `LOW-PRECISION-ROUTING-REPLAY-IMPORTANCE-GUARD`

Task: `low-precision-importance-guard-experiment-easy`

Type: Experiment Area: Post-training Difficulty: Hard

Design and execute a deterministic CPU smoke study for the existing rollout importance correction. Compare a full-precision control with low-precision behavior log probabilities, behavior-log-prob availability, and justified lower/upper bound choices under one fixed token workload. Keep seeds, token identities, masks, current probabilities, compute budget, and stopping rules explicit. Retain every raw run, summarize finite ratios, low/high masking, effective tokens, guarded loss, and gradient proxy, and allow a tied or inconclusive conclusion. The reported aggregates must be reproducible from the raw records.

Task: `low-precision-importance-guard-experiment-hard`

Type: Experiment Area: Post-training Difficulty: Hard

Design and execute a budgeted CPU study of the existing rollout-importance correction and its policy-loss handoff. Use deterministic short/long, ordinary/tail, and low/high-mismatch

workloads; within each workload, length, mismatch, and seed block, compare full- versus low-precision captured behavior log probabilities and justified positive bound pairs. Treat unavailable behavior probabilities as retained failures, never as on-policy values. Preserve every attempt and the inputs needed to recompute finite-rate, low/high zero-weight, effective-token, guarded-loss, and gradient-proxy metrics. Report predeclared paired contrasts with effect sizes and uncertainty, keep numeric-bound analyses separate from provenance failures, and let conclusions report practical ties or insufficient evidence. Make the complete report reproducible from raw records under a fixed run budget.

Task: `low-precision-importance-guard-experiment-medium`

Type: Experiment Area: Post-training Difficulty: Hard

Design and execute a repeated CPU study of rollout importance correction. Compare full-precision and low-precision behavior probabilities, present/missing provenance, and justified lower/upper bounds with matched token identities, masks, current probabilities, length strata, seeds, and run budgets. Retain successful and failed raw runs, derive side-specific masking, effective-token, guarded-loss, and gradient-proxy metrics, report dispersion, and disclose practical ties or inconclusive results. All tables must be reproducible from the raw records.

SOURCE PROBLEM: `MULTI-JUDGE-DATA-QUALITY-GATE-ERROR-BUDGET`

Task: `multi-judge-quality-policy-study-easy`

Type: Experiment Area: Evaluation Difficulty: Medium

Complete and execute the supplied CPU study comparing prove_only, solve_only and combined quality policies on the same labeled synthetic and vetted items. Invoke every declared judgement through a real RateLimitedRewardManager.run_single boundary with a controlled scorer; never aggregate precomputed decision lists. Preregister item, policy, judge kind and attempt coordinates under one total call budget, retain every raw returned dictionary and stop exactly according to each supplied policy. Treat timeout, error and malformed returns as unavailable evidence, not rejection. Report false acceptance, false rejection, decision coverage, calls and cost by source class and policy, plus the full raw audit and an inconclusive-capable summary. Do not replace supplied items, labels, policies or scorer outcomes, retry failures, or assume a winner.

Task: `multi-judge-quality-policy-study-hard`

Type: Experiment Area: Evaluation Difficulty: Hard

Execute the supplied blocked multi-workload quality-gate study across clean synthetic, vetted, ambiguous-reference and service-failure regimes. Compare prove_only, solve_only and combined policies, plain versus tool-assisted judging, and the supplied debate recovery under one preregistered call-and-cost budget. Materialize each workload/seed block once and invoke every judgement through RateLimitedRewardManager.run_single; retain all raw calls and failures without retries. Pair policy and escalation effects within workload, seed, source class and item, separating feasibility from numeric effects. Report false-accept/reject, coverage, cost and early-stop effects with deterministic blocked uncertainty, missing-pair counts and inconclusive conclusions when intervals cross zero or evidence is sparse. Do not alter supplied candidates, labels or scorer traces.

Task: `multi-judge-quality-policy-study-medium`

Type: Experiment Area: Evaluation Difficulty: Medium

Execute the supplied repeated CPU study of prove_only, solve_only and combined policies with plain and tool-assisted judge modes on identical synthetic and vetted item blocks. For every seed, preregister coordinates and invoke real RateLimitedRewardManager.run_single calls under the declared total call-and-cost ceiling. Preserve supplied items, labels, policies, scorer traces and tool-cost model. Retain successes, semantic rejections, timeouts, errors

and malformed returns once; do not retry or replace failed blocks. Recompute source-wise false acceptance, false rejection, coverage, calls, cost and early-stop position from raw records, then report per-condition estimates, across-seed dispersion and incomplete matched blocks. Allow practical ties or insufficient evidence and do not select favorable seeds.

Task: `source-aware-quality-gate-feature-easy`

Type: Implementation Area: Evaluation Difficulty: Medium

Implement the supplied QualityGate API and integrate it with RateLimitedRewardManager.run_single. Execute preregistered proof and solve attempts under one call budget and retain every returned dictionary. Apply the supplied asymmetric source policies: strict synthetic items require all planned proof evidence plus solve acceptance, while vetted items accept after the first valid proof. Timeout, error and malformed results are unavailable evidence, not rejection; valid score zero is semantic rejection. Stop only when the active policy is decided or its budget is exhausted. Return decision, status, calls, cost, stopping reason and immutable raw evidence. Repeated evaluation with the same request id is idempotent; conflicting reuse is rejected. Do not call scorers directly, retry unavailable attempts or mutate input items.

Task: `source-aware-quality-gate-feature-hard`

Type: Implementation Area: Evaluation Difficulty: Hard

Implement versioned source-aware QualityGate execution with plain, tool-assisted and debate-recovery stages. Preserve synthetic all-proof-plus-solve and vetted any-proof quantifiers. Launch a later stage only when unresolved and fully affordable; retain every rate-limited call status without retries. Debate may resolve only the supplied disagreement cases and cannot overwrite raw judges. Concurrent, sequential and resumed execution at any stage boundary must yield identical mapped decisions, evidence, calls, costs and stopping reasons. Persist semantic plans, fingerprints, completed coordinates and budgets, never manager handles or coroutines. Validate corrupt, duplicate-coordinate and future state atomically. Batch completion order must not alter item alignment or pooled-budget allocation. Valid zero remains rejection; unavailable evidence remains missing.

Task: `source-aware-quality-gate-feature-medium`

Type: Implementation Area: Evaluation Difficulty: Medium

Implement the supplied QualityGate across plan construction, RateLimitedRewardManager execution and result batching. Preserve synthetic all-proof-plus-solve and vetted any-proof semantics. Add the declared plain-to-tool-assisted escalation: escalate only unresolved evidence when remaining call and cost budgets can fund the entire next stage. Charge every launched success or failure once. Valid rejection differs from unavailable output; neither may be silently retried. Return raw calls, stage decisions, source policy, confusion-category inputs, total calls/cost and stopping reason. Batch results remain aligned by stable item/request id despite completion order. Idempotent duplicates reuse exact results; conflicting fingerprints fail atomically. Do not let one item's unused budget fund another unless the public plan enables pooling.

SOURCE PROBLEM: `MULTI-TEACHER-SERVER-TOPIC-CONTEXT-ROUTING-INTEGRITY`

Task: `multiteacher-response-attribution-repair-easy`

Type: Repair Area: Serving Difficulty: Easy

Repair the existing teacher request and response boundary without changing its public call signature. Apply each selected teacher context policy while preserving response placement. Mixed LLM and VLM requests must forward media without slicing placeholder-aligned input. Preserve routing errors for missing or unknown keys, output tensor rank, logical sequence length, and normal in-window behavior; overlong multimodal input must fail clearly rather than detach media from tokens.

Task: `multiteacher-routing-path-alignment-easy`

Type: Alignment Area: Serving Difficulty: Easy

> Implement the supplied alignment entry point comparing direct teacher-manager requests with the agent-loop teacher-logprob path. Run both real paths against controlled teacher clients for one batch with stable sample ids, explicit teacher ids, model versions, and logical response-token ranges. Route by each sample's topic key, preserve batch order at the public output, and compare mapped teacher identities, token ranges, log probabilities, and validity masks. Completion order must not alter results. An unavailable teacher must remain unavailable rather than silently route to another teacher. Preserve request and response records and expose final agreement diagnostics. Do not implement one path by calling the other or compare only batch summaries.

Task: `multiteacher-routing-path-alignment-medium`

Type: Alignment Area: Serving Difficulty: Medium

> Extend routing alignment to mixed text and multimodal samples, multiple teachers, per-teacher tokenizer/context policies, chunked batches, and permuted asynchronous completion. Compare direct manager and agent-loop paths at route selection, request construction, logical-token mapping, response acceptance, and batch reassembly. Join by stable sample id plus logical token range, never physical position. Exercise duplicate retry responses, stale model versions, unavailable teachers, and different chunk partitions. Accept at most one current response per identity; do not semantically reroute failures. Report the first divergent mapped stage with teacher id, model version, sample id, and token range. Preserve serializable provenance and input non-mutation. The two paths must remain independently executed.

Task: `versioned-multiteacher-routing-easy`

Type: Implementation Area: Serving Difficulty: Easy

> Add `verl/verl/trainer/distillation/teacher_registry.py` with immutable `TeacherCapability`, `TeacherRegistrySnapshot`, and `TeacherRequest` APIs. A snapshot must resolve an exact `(topic, modality)` route, capture the selected teacher id, model version, endpoint, tokenizer id, registry version, caller-supplied request/sample ids, and the retained logical token range, and apply that teacher's configured context limit without consulting later registry state. Unsupported routes or modalities must fail instead of falling back to another teacher. Keep the module independently loadable and make planned requests safe from mutation of the input configuration.

SOURCE PROBLEM: `MULTI-TOPIC-PARALLEL-DISTILLATION-RETENTION-MIXTURE`

Task: `parallel-topic-distillation-retention-study-easy`

Type: Experiment Area: Post-training Difficulty: Hard

> Complete and execute the supplied CPU study comparing sequential topic phases with one parallel mixed update schedule. Use identical labeled samples, initial toy-student parameters, optimizer steps and teacher-call budgets. Route every teacher request through AsyncTeacherLLMServerManager using the supplied data_source teacher key and controlled clients; do not use precomputed targets or rewrite the candidate schedules. The parallel condition must retain declared target-topic quotas and route background samples to the reference teacher. Preserve raw sample id, topic, role, selected teacher, loss and parameter trace. Independently audit routing, realized exposure and per-topic endpoint loss, including an overlap-labeled sample with the supplied ownership rule. Report target gain and protected-topic retention separately and allow inconclusive outcomes. Do not replace samples, teachers or initial state.

SOURCE PROBLEM: `ONPOLICY-REVERSE-KL-MODE-SEEKING-ALIGNMENT`

Task: `kl-projection-objective-study-easy`

Type: Experiment Area: Post-training Difficulty: Easy

> Complete and execute the supplied CPU study comparing teacher-sampled forward projection with the production student-sampled reverse-KL estimator on fixed representable and capacity-limited discrete targets. Use identical initial logits, update budgets and supplied sample streams. Call compute_distillation_loss_reverse_kl_estimator for every reverse-KL update and backpropagate its masked output; do not substitute a handwritten objective. Independently enumerate exact forward and reverse KL values and gradients as an oracle. Retain per-step samples, masks, losses, gradients and probabilities, then report projection endpoints and oracle discrepancies. Preserve all supplied targets and randomness. Failed or nonfinite runs remain explicit and cannot be replaced. Report observations without naming an expected winner.

Task: `kl-projection-objective-study-hard`

Type: Experiment Area: Post-training Difficulty: Hard

> Execute the supplied blocked CPU study across representable, bottlenecked, symmetric-multimodal and asymmetric-multimodal discrete targets and their declared label permutations. Compare forward projection with production reverse-KL k1, k2 and k3 estimators using matched initial states, budgets and random streams. Call compute_distillation_loss_reverse_kl_estimator and backpropagate its masked losses; retain all samples, gradients, probabilities and failures without replacement. Independently enumerate exact objectives and gradients. Pair projection and estimator contrasts within workload, seed, capacity and permutation, reporting endpoint mass, oracle error and stability effects with deterministic blocked uncertainty, missing-pair counts and permutation-consistency audits. Separate estimator feasibility from numeric effects and allow intervals crossing zero to remain inconclusive. Do not alter workloads or hard-code rankings.

Task: `kl-projection-objective-study-medium`

Type: Experiment Area: Post-training Difficulty: Medium

> Execute the supplied repeated CPU study of forward projection and production reverse-KL k1, k2 and k3 estimators on identical representable and capacity-limited targets. For each seed, reuse the declared initial logits, update budget and uniform sample stream across eligible conditions. Invoke compute_distillation_loss_reverse_kl_estimator with the requested loss_mode and perform masked backward updates; do not reproduce its formulas. Preserve every raw sample, loss, gradient, probability and nonfinite failure without retries. Use an independent exact enumerator to audit values and gradients, then report per-condition endpoints, oracle error, across-seed dispersion and incomplete matched blocks. Do not change supplied targets, select favorable seeds or declare a universal winner.

Task: `reverse-kl-estimator-alignment-easy`

Type: Alignment Area: Post-training Difficulty: Hard

> Add a deterministic CPU alignment test for the existing reverse-KL distillation path. Construct real-shaped rollout tokens, response masks, old/current log probabilities, and teacher log probabilities, then invoke the registered production estimator rather than a copied helper. From identical normalized discrete teacher/student parameters and seeds, compare both its reported loss and backward gradient with exact enumeration. Exercise on-policy and stale-policy handoffs, padding, a representable control, and an asymmetric capacity-limited target. Preserve sampling direction, detach teacher and rollout bookkeeping, and retain gradients through intended student scores. Record every seed, budget, values, gradient vectors, absolute errors, and shapes; reject mislabeled objectives or normalization. Do not add a new estimator, alignment API, or favorable-seed selection.

SOURCE PROBLEM: `PARTIAL-ROLLOUT-TOKEN-STALENESS-PROJECTION-AND-MASK`

Task: `partial-rollout-staleness-debug-easy`

Type: Repair Area: Post-training Difficulty: Medium

> Repair partial-rollout provenance from collection through token-mask composition. Preserve the public rollout result, response length, padding and tool exclusions, include tokens exactly at the threshold, and keep ordinary unsegmented responses unchanged. Segment record order must not change trainable-token eligibility or retained-token diagnostics. Do not replace token-level filtering with whole-response rejection.

Task: `partial-rollout-staleness-feature-easy`

Type: Implementation Area: Post-training Difficulty: Medium

> Add token-level staleness control to the partial-rollout path. Segment records carry response-local spans and production versions, while the rollouter owns padded response tensors and an existing trainable-token mask. Preserve logical identity when records arrive out of order, exclude only ages above the configured threshold, and never re-enable masked positions. Return aligned staleness, the composed mask, and retention diagnostics through the existing rollout result. Reject malformed coverage clearly and keep ordinary unsegmented rollouts unchanged.

Task: `partial-rollout-staleness-feature-hard`

Type: Implementation Area: Post-training Difficulty: Hard

> Complete token-level staleness control across partial-rollout collection, projection, and PPO loss consumption. Preserve logical response identity through resumed segments, shuffled records, truncation, padding, and physical microbatch repartitioning. Stale trainable tokens must be excluded without re-enabling tool or padding positions; an all-masked batch must reduce safely; disabling the policy must reproduce baseline loss and gradients. Return per-response provenance and retention diagnostics suitable for auditing equivalent layouts.

Task: `partial-rollout-staleness-feature-medium`

Type: Implementation Area: Post-training Difficulty: Hard

> Implement token-level staleness eligibility for batched partial rollouts and wire it into the existing fully asynchronous rollouter. Segment records carry logical response ids, response-local spans, and production versions; batches may reorder records and truncate inside a segment. Return aligned staleness tensors, composed masks, and per-response retention diagnostics. Preserve existing tool, padding, and response exclusions, include tokens exactly at the threshold, validate gaps and overlaps, and keep ordinary unsegmented rollouts unchanged.

SOURCE PROBLEM: `PASS-RATE-STRATIFIED-RESAMPLING-COVERAGE-BALANCE`

Task: `passrate-resampling-budgeted-study-easy`

Type: Experiment Area: Post-training Difficulty: Hard

> Complete and execute the supplied CPU smoke study comparing the historical row-accounted resampling proxy with the logical-prompt coverage-aware proxy at zero-pass shares 0.10 and 0.20. Keep the five-domain prompt pool, pass-rate boundaries, target size, draw budget, seed, and learning-signal consumer fixed; do not alter the supplied workload or candidate samplers. Retain every attempt and its ordered sampled records, audit prompt-id aggregation and copy integrity, and derive unique eligible coverage, domain and intermediate-bin balance, realized zero-pass share, multiplicity, and the identity-aware positive-signal coverage curve. Produce a summary reproducible from raw records and a conclusion limited to this proxy that permits a tie or insufficient evidence.

Task: `passrate-resampling-budgeted-study-hard`

Type: Experiment Area: Post-training Difficulty: Hard

> Complete and execute the supplied blocked CPU study of row-accounted versus logical-prompt coverage-aware resampling at zero-pass shares 0.10 and 0.20 across the declared balanced, sparse/skewed, duplicate-boundary, and infeasible-budget workloads. Pair methods and shares within immutable workload-and-seed blocks while holding pass-rate boundaries, target size, consumer order, and evaluation budget fixed; do not alter supplied samplers or workloads. Retain every attempt, ordered sampled record, diagnostic, and failure without replacement. Audit prompt-id aggregation, immutable copies, required coverage, domain/bin balance, zero-pass realization, multiplicity, and logical positive-signal coverage. From raw records, report workload-specific and blocked paired effects with deterministic uncertainty, disclose incomplete pairs and infeasibility separately, and permit ties or insufficient evidence without extrapolating to full RL training.

Task: `passrate-resampling-budgeted-study-medium`

Type: Experiment Area: Post-training Difficulty: Hard

> Complete and execute the supplied repeated CPU study of row-accounted versus logical-prompt coverage-aware resampling at zero-pass shares 0.10 and 0.20. For every declared seed, match the immutable five-domain pool, pass-rate boundaries, target size, draw order, consumer, and evaluation budget across conditions; do not alter the supplied workloads or samplers. Retain every successful and failed attempt without replacement, including ordered sampled records and identity/copy audits. Recompute unique eligible coverage, domain and intermediate-bin balance, realized zero-pass share, multiplicity, and logical-prompt positive-signal coverage from raw data. Report per-condition point estimates and across-seed dispersion, disclose incomplete pairs, and limit conclusions to the proxy while allowing practical ties or insufficient evidence.

SOURCE PROBLEM: `PROMPT-DIVERSITY-ROLLOUT-REPLICATION-BUDGET-ALLOCATION`

Task: `prompt-rollout-allocation-study-easy`

Type: Tests Area: Post-training Difficulty: Hard

> Replace the smoke coverage with deterministic CPU contract tests that drive one complete RayPPOTrainer.fit step through controlled rollout and worker dependencies. Start from a collated RLHF-shaped batch whose physical rows share the same dataset index, and let the trainer—not the fixture—assign uid and perform configured interleaved replication. Capture the request passed to async_rollout_manager and the batch passed to the actor update. Verify that the request realizes source-row count × rollout.n with tensor/non-tensor alignment and one uid per source row; have the controlled rollout return valid responses, attention and response masks, and rm_scores, then verify that the actor receives rewards, advantages, and returns grouped by uid rather than index. Cover rollout.n=1 and a repeated group, unequal masks and rewards, normalized and unnormalized GRPO, and unchanged collated inputs. Use independently derived direct expectations. Keep production calls real except external trainer dependencies; do not edit production code, create proxy schemas or workloads, or require GPU, Ray services, network, or model weights.

Task: `prompt-rollout-allocation-study-hard`

Type: Tests Area: Post-training Difficulty: Hard

> Complete a deterministic CPU controlled-integration suite for the existing dataset-to-rollout-to-advantage lifecycle. Drive real dataset rows into DataProto, assign trainer-style uids, interleave rollouts, combine sampled and ReMax-style baseline-tagged batches, simulate out-of-order returned rows, union aligned outputs, construct response masks, and call public advantage computation. Verify sampled group identity and advantages are invariant after remapping, baselines never become sampled rollout members, all tensor/non-tensor fields remain aligned, copied indices do not merge distinct uids, and repeated executions do

not mutate source state. Cover variable lengths, singleton/mixed groups, masked rewards, normalized/unnormalized GRPO, and missing optional fields. Use independent and metamorphic oracles. Do not run Ray/GPU/network, edit production code, or invent allocation outcomes.

Task: `prompt-rollout-allocation-study-medium`

Type: Tests Area: Post-training Difficulty: Hard

Extend the CPU suite through the existing dataset identity boundary. Exercise RLHFDataset rows with explicit, missing, duplicated, and reordered extra_info.index values; form DataProto batches, assign transport uids as the trainer does, interleave multiple rollouts, reorder completed rows, and compute GRPO advantages. Assert relational invariants: all fields stay aligned, each uid's centered valid-token advantage sums to zero when the group has multiple members, equivalent row permutations remap to equivalent results, masked positions remain zero, and copied indices with distinct uids remain distinct groups. Cover normalized and unnormalized modes, variable response masks, singleton groups, and mixed group sizes. Use independent or metamorphic oracles. Do not edit production code, invent deduplication, or require live rollout services.

Task: `prompt-rollout-budget-allocation-study-easy`

Type: Experiment Area: Post-training Difficulty: Medium

Complete and execute the supplied deterministic CPU allocation study for one distillation proxy workload. Sweep the preregistered 4096-response factorizations 4096×1, 256×16, and 16×256. Keep the logical prompt pool, topic mixture, checkpoint tag, sampling settings, response budget, and evaluation callback fixed; use the supplied stable logical id rather than physical rows as prompt identity. Retain one raw attempt per condition and every per-response status record without replacement. From raw records, validate budget and group realization and report unique logical-prompt coverage, requested and effective group sizes, token use, within-prompt outcome correlation, estimator variance, and downstream proxy point estimates. Make the summary reproducible, disclose shortfalls or undefined metrics, and limit the conclusion to this smoke proxy, allowing a tie or insufficient evidence.

Task: `prompt-rollout-budget-allocation-study-hard`

Type: Experiment Area: Post-training Difficulty: Hard

Complete and execute the supplied blocked CPU allocation study across the declared topic-balanced, length-skewed, and duplicate/failure workloads. Within each workload, objective, budget mode, and seed block, sweep only the preregistered prompt-count × rollout-count allocations while matching prompt source, checkpoint tag, sampling settings, evaluation, and optimizer-exposure proxy. Support both equal-response and equal-token budgets, enforce objective-compatible group sizes, and stop without overspending or replacing failed generations. Retain every attempt and per-response record with logical identity, status, tokens, and group provenance. Recompute coverage, completeness, effective replication, correlation, estimator variance, budget utilization, and downstream metrics. Compare each allocation with its declared within-objective reference using workload-specific paired effects and equal-workload blocked effects with deterministic confidence intervals; report eligible and missing pairs separately. Conclusions must follow observed effects and uncertainty, allow inconclusive outcomes, and not extrapolate beyond the proxies.

Task: `prompt-rollout-budget-allocation-study-medium`

Type: Experiment Area: Post-training Difficulty: Medium

Complete and execute the supplied repeated CPU study of prompt-count × rollout-count allocations under one fixed response budget. For every declared seed, keep the logical prompt source, topic mixture, checkpoint tag, sampling settings, response table, optimizer-exposure proxy, and evaluation callbacks matched across allocations; do not alter the supplied workload or distillation and outcome-RL consumers. Retain every attempted response

exactly once, including failures, duplicates, token counts, group identity, and status, with no replacement runs. Enforce the preregistered incomplete-group policy and never report a within-prompt baseline for singleton or incomplete groups. Recompute logical coverage, group completeness, effective replication, token use, correlation, estimator variance, and downstream metrics from raw records. Report per-seed values, condition means and sample dispersion, disclose unmatched or ineligible cells, and permit ties or insufficient evidence.

SOURCE PROBLEM: `REDUNDANT-SPAN-PENALTY-REWARD-SUM-CONSERVATION`

Task: `redundant-span-reward-debug-easy`

Type: Repair Area: Post-training Difficulty: Easy

Repair the PPO reward utility that redistributes a length penalty over redundant response spans. Preserve the public API, the outcome reward, valid-token reward conservation, prompt and padding positions, and ordinary single-span behavior. Spans ending at EOS must not alter the terminal score. Return useful allocation diagnostics.

Task: `redundant-span-reward-feature-easy`

Type: Implementation Area: Post-training Difficulty: Hard

Add `verl/verl/trainer/ppo/redundant_span_reward.py` with a tensor-oriented API that converts terminal outcome rewards and a length penalty into token-level process rewards over detected redundant response spans. It must preserve the original terminal reward and the valid-token reward sum, and return diagnostics that callers can inspect. Also expose a helper for selecting a reward-safe early-stop candidate. Keep the module independently loadable and explicit about malformed spans.

Task: `redundant-span-reward-feature-hard`

Type: Implementation Area: Post-training Difficulty: Hard

Implement redundant-span process rewards across rollout metadata, an independently loadable allocator and PPO reward construction. Preserve terminal outcome scores and valid-token reward sums while redistributing length penalties, select only reward-safe early-stop candidates, and expose auditable diagnostics. Equivalent spans expressed through supported layouts must produce equivalent rewards. Handle overlaps, duplicates, invalid positions, EOS, padding and no-span fallbacks while keeping the disabled path unchanged.

Task: `redundant-span-reward-feature-medium`

Type: Implementation Area: Post-training Difficulty: Medium

Implement a redundant-span reward module and wire it into PPO reward construction. Convert terminal outcome rewards plus length penalties into token-level process rewards while preserving EOS and each valid-token reward sum. Support reward-safe early-stop candidate selection and return allocation diagnostics. Handle overlapping, duplicate, empty, EOS-adjacent and masked spans in variable-length batches without changing behavior when the feature is disabled.

SOURCE PROBLEM: `ROLLOUT-ROUTING-REPLAY-CHUNK-PROVENANCE-INTEGRITY`

Task: `rollout-chunk-route-provenance-repair-easy`

Type: Repair Area: Training infrastructure Difficulty: Medium

Repair the existing accumulation boundary so prompt routes, generated-token routes, and optional rollout log probabilities remain paired with their logical tokens across resumptions,

including nonempty and zero-token aborts. Preserve the public request and TokenOutput interfaces, asynchronous resume behavior, stop reason, preemption count, version range, and token budget. Missing, short, shape-inconsistent, or partially present route data must fail clearly. Do not recompute routes or change valid route values.

Task: `rollout-chunk-route-provenance-repair-medium`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair the producer-to-batch router-provenance handoff without changing the existing TokenOutput, AgentLoopOutput, or DataProto interfaces. Preserve asynchronous resume, configured response truncation, single-turn and tool-loop behavior, modality handling, stable output order, and the current remove-padding path. Route evidence must remain aligned by logical identity across mixed lengths, batching, ordering, and padding. Missing or shape-inconsistent route data must fail, while valid prompt and response padding remain inert.

Task: `rollout-routing-chunk-alignment-easy`

Type: Alignment Area: Training infrastructure Difficulty: Hard

Add an offline parity utility that executes a rollout reconstruction path and a training replay path for one single-chunk response. Join tokens and router metadata by the supplied logical identity, compare the final aligned token, route and log-probability outputs, and return agreement plus maximum error diagnostics. Both callables must really run. Keep the module independently loadable and reject malformed lengths instead of broadcasting metadata.

SOURCE PROBLEM: `RUBRIC-REWARD-MAPPING-AND-THRESHOLD-BOUNDARY-SEMANTICS`

Task: `rubric-reward-debug-easy`

Type: Repair Area: Post-training Difficulty: Easy

Repair rubric-mode semantics through the repository reward-manager interface without changing ordinary scalar reward behavior or the public call signature. Direct, binary, and discrete mappings must follow their declared conventions; dense-group filtering must follow its configured criterion. Keep rubric diagnostics in the returned extra information.

Task: `rubric-reward-debug-hard`

Type: Repair Area: Post-training Difficulty: Hard

Repair the full rubric reward path across manager output, runtime filtering, and algorithm configuration. Preserve bounded scoring context, consistent evaluation diagnostics, and intentional distinctions among direct, binary, and discrete modes. Non-rubric reward jobs must not change.

Task: `rubric-reward-debug-medium`

Type: Repair Area: Post-training Difficulty: Hard

Repair rubric semantics across the batch reward manager and algorithm filter configuration. Preserve weighted rubric support, bounded-context behavior, legacy scalar rewards, configuration compatibility, and the intentional distinction between rubric mapping modes.

Task: `rubric-reward-feature-easy`

Type: Implementation Area: Post-training Difficulty: Medium

Add a reusable rubric reward adapter at `verl/verl/utils/reward_score/rubric.py`. Its public API should accept rubric results and a mapping configuration, then return a training reward together with evaluation and diagnostic fields suitable for existing reward managers. Support the repository's declared mapping modes and weighted rubric inputs. Keep the module independently loadable, tensor-friendly, and explicit about invalid inputs; do not wire it into a manager yet.

Task: `rubric-reward-feature-hard`

Type: Implementation Area: Post-training Difficulty: Hard

Implement a rubric reward subsystem spanning an independently loadable adapter, the batch reward manager and algorithm filter configuration. It must expose training reward, raw-score evaluation and diagnostics without conflating their conventions. Support weighted rubrics, bounded critic context and mode-aware group filtering. Equivalent scalar and batched inputs must agree, existing non-rubric jobs must remain compatible, and invalid or incomplete configurations must fail clearly.

Task: `rubric-reward-feature-medium`

Type: Implementation Area: Post-training Difficulty: Hard

Add a reusable rubric reward adapter and integrate it with the batch reward manager. The adapter must support direct, binary and discrete mappings plus weighted rubric inputs; the manager must preserve raw evaluation semantics, diagnostics and existing scalar-reward behavior. Handle natural score-boundary, missing-component, context-budget and constant-group cases. Keep the adapter independently loadable and make configuration errors explicit.

SOURCE PROBLEM: `SANDBOX-FAILURE-LIFECYCLE-TIMEOUT-AND-REPLACEMENT`

Task: `sandbox-poisoned-lease-recovery-easy`

Type: Repair Area: Agents Difficulty: Easy

Repair native stateful-tool recovery within one agent trajectory. Successful calls and ordinary user-code exceptions must retain the healthy instance. A fatal failure must retire and release its instance once, report that prior state is unavailable, and never replay failed arguments. Only the next model-requested call may acquire one replacement, which later successful calls reuse. A second fatal failure, failed acquisition, or failed cleanup must leave that tool unavailable for the trajectory. Different native tools retain separate leases, and run exit releases every remaining lease once. Preserve FunctionTool dispatch, invalid-call messages, successful-response truncation, rewards, and output shape. Do not add configuration, snapshots, replay, or lifecycle metadata.

SOURCE PROBLEM: `SEQUENCE-BALANCED-LOSS-AND-ADVANTAGE-NORMALIZATION`

Task: `sequence-population-normalization-contract-tests-easy`

Type: Tests Area: Post-training Difficulty: Hard

Add a directly runnable CPU-only pytest module for the existing seq-mean-token-mean aggregation mode. Use small hand-checkable tensors to cover equal and strongly unequal response lengths, right-padding width changes, row reordering, mixed nonempty and padding-only rows, and explicit global-count/data-parallel scaling. Distinguish this selector from token-mean through observable scalar results, exercise the production agg_loss callable, and include its invalid-selector and missing-global-input boundaries. Keep the suite deterministic and offline; do not edit production code, copy the reducer, inspect source text, or require GPUs or distributed workers.

Task: `sequence-population-normalization-contract-tests-hard`

Type: Tests Area: Post-training Difficulty: Hard

> Add a directly runnable deterministic CPU integration test module for sequence-balanced PPO reductions across equivalent micro-batch and data-parallel partitions. Use the existing global-count and dp_size inputs plus a local torch.distributed gloo process group to relate distributed weighted means, full-batch agg_loss, and policy-loss values and gradients after DDP-style averaging. Cover unequal sequence lengths, row permutations, right padding, uneven and empty-rank layouts, and consecutive layouts in one process group. Join observations by test-owned logical row identity and use derived metamorphic oracles. Run offline without GPUs, wall-clock assertions, production edits, replacement reducers, or source inspection.

Task: `sequence-population-normalization-contract-tests-medium`

Type: Tests Area: Post-training Difficulty: Hard

> Add a directly runnable CPU-only pytest suite that checks sequence-balanced normalization across the existing masked-statistics, loss-reduction, and PPO policy-loss interfaces. Derive expectations from unequal-length logical sequences instead of copying constants. Cover numeric sequence-balanced masks, mixed padding-only rows, right-padding and row-permutation invariance, masked nonfinite padding when observing valid normalized outputs, and unclipped and clipped policy regimes. Verify that normalized valid advantages have the derived sequence-weighted moments and that consuming them with seq-mean-token-mean preserves the same population. Use production callables, deterministic tensors, controlled configuration collaborators, and no production edits, source inspection, GPUs, or distributed workers.

Task: `sequence-population-normalization-repair-easy`

Type: Repair Area: Post-training Difficulty: Easy

> Repair the sequence-balanced loss and token-level GAE and REINFORCE++ paths. Valid results must remain invariant to masked nonfinite columns and padding-only rows, and all-empty batches must remain finite. Preserve existing aggregation modes and valid-token behavior.

Task: `sequence-population-normalization-repair-hard`

Type: Repair Area: Post-training Difficulty: Hard

> Repair sequence-balanced PPO across micro-batch and data-parallel layouts. Equivalent logical batches must preserve loss scale, normalized advantages, and reported advantage means in the presence of padding-only rows. Preserve equal-length dense behavior, token-mean mode, and the selected objective. Do not add telemetry schemas.

Task: `sequence-population-normalization-repair-medium`

Type: Repair Area: Post-training Difficulty: Hard

> Repair the standard Ray PPO sequence-balanced advantage path. Equivalent logical responses must preserve update direction across regrouping, valid-length variation, and padding. Preserve the selected loss objective, estimator formulas, equal-length dense behavior, and established behavior of other aggregation modes.

SOURCE PROBLEM: `STOCHASTIC-VERIFIER-REPEATABILITY-DISAGREEMENT-AUDIT`

Task: `verifier-disagreement-ledger-feature-easy`

Type: Implementation Area: Evaluation Difficulty: Medium

Add an opt-in verifier audit ledger to the existing RateLimitedRewardManager path. For each run_single call, retain immutable uid, request id, repeat index, semantic-input and verifier-config digests, completion status, raw returned dictionary and binary decision when valid. Valid zero must remain distinct from timeout, error and malformed output. Reconstruct per-uid call count, valid coverage, positive rate, avg@k, strict stable@k and disagreement without depending on completion order; stable requires all requested repeats valid and equal. Duplicate identical records are idempotent, while conflicting request ids fail. Preserve the returned reward and disabled manager behavior exactly, bound retained records by configured audit scope, and do not retry, relabel or treat failures as negative. Add CPU tests through run_single.

SOURCE PROBLEM: `STREAMING-ROLLOUT-BACKLOG-FIXED-OPTIMIZER-BATCH`

Task: `streaming-backlog-controller-feature-easy`

Type: Implementation Area: Training infrastructure Difficulty: Easy

Add `backlog_controller.py` for streaming rollout admission. Given a fixed logical optimizer batch, completed backlog, active generation count, configured extra-completion fraction, and low/high watermarks, return a deterministic admission budget and diagnostics. The controller may throttle producers but must never change optimizer batch size, loss denominator, or scheduler-step meaning. Enforce bounded extra completions, reject invalid configuration, and keep the module independently loadable.

Task: `streaming-backlog-controller-feature-hard`

Type: Implementation Area: Training infrastructure Difficulty: Hard

Complete streaming backlog control across producer admission, message-queue ownership, and trainer consumption. Keep logical optimizer batch size, loss denominator, and scheduler-step meaning fixed while bounding extra completions and staleness. Preserve at-most-once response use and prompt identity through stale rejection, replacement, latency bursts, physical microbatch repartition, and resume with a nonempty backlog. Persist compatible controller cursors and pending update membership, and report throughput, pressure, dropped work, effective batch, and staleness.

Task: `streaming-backlog-controller-feature-medium`

Type: Implementation Area: Training infrastructure Difficulty: Medium

Implement bounded streaming backlog control and wire it into the fully asynchronous rollouter. Admission must use configured low/high watermarks, active work, completed backlog, and an extra-completion fraction while keeping logical optimizer batch size and loss normalization fixed. Track every admitted, consumed, stale-rejected, and deferred logical response exactly once. Handle producer-faster, trainer-faster, latency-burst, and zero-capacity regimes, and expose queue pressure, effective batch, dropped work, and staleness diagnostics.

SOURCE PROBLEM: `TOKENIZER-FINGERPRINT-SPECIAL-TOKEN-LOGPROB-CONSISTENCY`

Task: `tokenizer-handshake-logprob-debug-easy`

Type: Repair Area: Serving Difficulty: Medium

Repair the existing tokenizer compatibility helper without changing ordinary tokenizer loading. Identical artifacts must produce a stable versioned identity; reordered token IDs, changed special-token mappings, incompatible versions, and stale environment state must not be accepted as equivalent. Artifact fingerprint information must remain authoritative, and incompatibility diagnostics must remain useful.

Task: `tokenizer-handshake-logprob-debug-hard`

Type: Repair Area: Serving Difficulty: Hard

> Repair the rollout contract across tokenizer artifacts, configured special-token policy, backend logit transformation, and stored rollout log-probabilities for vLLM and SGLang. Incompatible tokenizer identity or log-probability semantics must fail before sampling. Preserve one verified special-token policy across turns, and keep unbiased, numeric-bias, and ordinary single-turn behavior unchanged.

Task: `tokenizer-handshake-logprob-debug-medium`

Type: Repair Area: Serving Difficulty: Hard

> Repair tokenizer identity and suppression-policy propagation from trainer-side model configuration into the vLLM rollout request path. Reject incompatible artifacts before generation, preserve artifact metadata over stale environment state, resolve configured special-token names only after compatibility succeeds, and keep numeric backend bias keys and unbiased requests unchanged.

SOURCE PROBLEM: `TOOL-EVENT-REWARD-SHAPING-BOUNDED-ATTRIBUTION`

Task: `tool-event-shaping-contract-tests-easy`

Type: Tests Area: Agents Difficulty: Hard

> Add deterministic CPU tests for the supplied tool_reward_shaping module; keep its implementation unchanged. Independently derive category, weight, actionable denominator, normalized sum, clipped contribution and final outcome reward. Cover no calls, success, argument/format/code-runtime errors, OOM, network and internal-tool failures, retry then success, and enough repeated calls to saturate both clip bounds. Verify one physical call produces one canonical event even when parser and execution statuses coexist. Infrastructure-neutral events remain auditable but do not become penalties. Normalization precedes clipping and alpha scaling, and outcome reward is added exactly once. Kill classification-precedence, double-event, neutral-penalty, zero-division, clip-order and double-add mutants. Do not edit production/support code or form expectations by calling the shaper.

Task: `tool-event-shaping-contract-tests-hard`

Type: Tests Area: Agents Difficulty: Hard

> Add a deterministic CPU controlled-integration suite for the supplied tool shaper from agent call records through extra_fields, batch conversion and final sequence reward. Compare equivalent sequential/concurrent completions, padded/jagged layouts, row permutations, microbatch partitions and semantic save/restore cuts. After stable-id remapping, require identical canonical events, raw outcome, numerator, denominator, clipped shaping and final reward. Restore must reject corrupt, duplicate-coordinate and future-version state atomically; checkpoints contain no tool handles. Sequence shaping is mutually exclusive with supplied turn/token modes and is applied once. Infrastructure-only, all-padding, aborted and retry-heavy traces remain finite and auditable. Kill interaction mutants involving local denominators, replayed calls, stale state, completion-order ids, double application and partial restore. Do not edit supplied implementation or serialize live objects.

Task: `tool-event-shaping-contract-tests-medium`

Type: Tests Area: Agents Difficulty: Hard

> Add deterministic CPU tests for the supplied shaper and its ToolAgentLoop adapter without editing them. Build sequential and parallel call traces with stable issue ids, out-of-order completion, parser/execution status overlap, retries, infrastructure-only failures and mixed actionable events. Independently reconstruct canonical issue-order records and bounded

sequence reward. Verify completion order cannot change categories, denominator, diagnostics or final outcome. Exercise AgentLoopOutput.extra_fields propagation and mixed enabled/disabled samples; absent metadata preserves legacy behavior, while malformed present metadata fails atomically. Padding and aborted trajectories contribute zero. Kill mutants that use completion order, count neutral failures in the actionable denominator, apply shaping twice, drop retries, merge same-name calls or silently accept malformed records.

SOURCE PROBLEM: `TOPK-DISTILLATION-DIVERGENCE-GRADIENT-STABILITY`

Task: `topk-gradient-contract-tests-easy`

Type: Tests Area: Post-training Difficulty: Hard

Add a focused CPU test module for the existing FSDP Top-K forward-KL helper. Construct valid jagged teacher Top-K inputs with distinct vocabulary ids, and derive per-token loss, pre-clamp student/teacher mass, and student-logit gradients independently from the truncated forward-KL definition; do not call the production helper to form the oracle. Cover ordinary and extreme teacher/student mismatch, asymmetric teacher versus student Top-K support, clamp disabled, and entries strictly below and above the configured clamp under nonuniform upstream loss weights. Distinguish forward from reverse KL, unnormalized retained mass from within-Top-K renormalization, teacher-id gathering from student-rank gathering, and differentiable from detached student paths. Verify clamp changes loss and gradients without changing reported masses; retain only a focused overlap-diagnostic no-regression check because ordinary overlap values are already covered. Keep production code unchanged, avoid fixture-specific constants as the sole oracle.

Task: `topk-gradient-contract-tests-hard`

Type: Tests Area: Post-training Difficulty: Hard

Add one CPU controlled-integration test module for the Top-K forward-KL path. Begin with padded caller-shaped input ids, prompts/responses, attention/response masks, and distinct valid-row teacher global ids/log-probabilities; pass them through left_right_2_no_padding and compute_topk_loss FSDP and Megatron dispatches. Run Megatron in simultaneous local Gloo ranks at TP sizes 2 and 4 with equal contiguous vocabulary shards, including an empty-owner rank. Exercise THD and BSHD preprocessing and backend-specific restoration of raw outputs to jagged model outputs before invoking the registered forward-KL handler and checking response-aligned losses and response-masked diagnostics. Independently derive all five valid-token outputs and local student gradients, respecting dense pre-clamp versus Megatron clamp-adjusted mass definitions; use nonuniform upstream weights and selected finite differences in nonbinding and strictly binding clamp regimes. Trace collectives only through delegates that execute the operations. Verify teacher tensors remain unchanged and identity-preserving global vocabulary relabeling, with columns redistributed to new contiguous owners, preserves values, masked metrics, and remapped gradients. Treat pinned ordinary overlap/layout/GPU parity as controls. Do not edit production code, use another backend as truth, serialize or replace collectives, test clamp equality, invent duplicate caller ids or local-column mappings, or bypass output restoration.

Task: `topk-gradient-contract-tests-medium`

Type: Tests Area: Post-training Difficulty: Hard

Keep the certified dense tests and extend the module with a deterministic CPU harness for the real compute_topk_loss Megatron dispatch and compute_forward_kl_topk/custom-autograd path. Stub only unavailable Megatron state and torch.distributed collectives. Build distinct-id jagged teacher rows and independently map full-vocabulary fixtures into local THD and BSHD student-logit layouts for two equal contiguous TP topologies. For every simulated rank, make the scripted collective verify its outgoing tensor and operation before supplying the independently derived global reduction or gathered candidates; require the same call sequence when a rank owns no teacher id. Derive all five outputs at

valid positions, including the Megatron entry's clamped and unclamped retained-mass definitions, plus local-shard gradients under nonuniform valid-token weights and zero padding weights. Cover unclamped, strictly binding, and nonbinding clamps without demanding FSDP mass parity where its definition differs. Do not edit production code, use a backend as the oracle, treat the existing GPU summed-loss comparison as sufficient, require duplicate ids or clamp-equality derivatives, or invent local-column mappings.

SOURCE PROBLEM: `TOPK-OPD-IMPORTANCE-WEIGHT-SCOPE`

Task: `topk-importance-scope-repair-easy`

Type: Repair Area: Post-training Difficulty: Hard

Repair the contract between distillation configuration and loss dispatch. Configuration must reject incompatible loss and policy-gradient combinations, and dispatch must apply correction only to the expectation sampled by each mode. Preserve registered names, aggregation, masks, current-policy gradients, teacher detachment, on-policy behavior, and public metric types. Exclude padding from losses and ratio diagnostics, retain supplied rollout-level weights, and make biased clipping visible. Do not add trajectory-product weights or change rollout storage.

Task: `topk-importance-scope-repair-medium`

Type: Repair Area: Post-training Difficulty: Medium

Restore estimator-faithful behavior across the existing distillation configuration and loss dispatcher. Preserve all registered loss names, aggregation modes, teacher detachment, current-policy gradients, nested and padded mask handling, and public metrics. Top-K summed terms must be invariant to the sampled response token; Monte Carlo terms must use one correctly oriented token correction. Reject incompatible configurations before training, report clipping and effective ratio counts, and do not synthesize history-level ratios or reinterpret rollout importance weights.

SOURCE PROBLEM: `TOPK-OPD-TAIL-MASS-KL-ESTIMATOR-CONSISTENCY`

Task: `topk-opd-tail-mass-feature-easy`

Type: Implementation Area: Post-training Difficulty: Hard

Add ` verl/verl/trainer/ppo/topk_distillation.py` with a tensor API for pooled-tail Top-K KL from student and detached teacher logits plus retained token ids. Preserve the probability mass split between retained tokens and the pooled tail instead of renormalizing the retained set to one. Return per-row loss and inspectable mass diagnostics, keep student gradients intact, and make the module independently loadable. Support a shared retained set and ordinary finite logits.

Task: `topk-opd-tail-mass-feature-medium`

Type: Implementation Area: Post-training Difficulty: Medium

Extend the Top-K distillation loss and wire it into PPO loss utilities. Student and teacher retained entries may arrive in different orders, so align them by vocabulary id before computing pooled-tail KL. Preserve retained-versus-tail mass, detach teacher targets, retain student gradients, and handle duplicate ids, tiny tails and extreme finite logits explicitly. The integration must expose per-row loss, retained masses, tail masses and validation diagnostics while preserving the easy shared-order behavior.

SOURCE PROBLEM: `TRAIN-EVAL-JUDGE-PROVENANCE-SEPARATION`

Task: `judge-provenance-ledger-feature-easy`

Type: Implementation Area: Evaluation Difficulty: Medium

> Add judge provenance to offline evaluation. Define a stable artifact fingerprint and explicit role for each scorer, append immutable verdict records tied to request and input identity, and aggregate metrics with judge identity, role, denominator, failures and unresolved cases. A scorer used for training cannot satisfy an independent-evaluation requirement, and a failed independent judgement must remain missing rather than falling back to the training score. Preserve existing custom reward functions and per-data-source score reporting.

Task: `judge-provenance-ledger-feature-hard`

Type: Implementation Area: Evaluation Difficulty: Hard

> Carry judge provenance end to end from training reward production into offline independent evaluation and report reconstruction. Training, selection, monitoring and independent roles have asymmetric admissibility rules. Bind raw rubric decisions and failures to immutable request/input identities, reject behavioral aliases, preserve older verdicts across judge-version changes, and never substitute a failed independent result. Reports must reproduce paired metrics for a declared version window while disclosing retries, unresolved ambiguity and denominator changes. Existing reward-manager and custom-function paths must remain compatible.

Task: `judge-provenance-ledger-feature-medium`

Type: Implementation Area: Evaluation Difficulty: Medium

> Extend provenance-aware evaluation to detect the same behavioral judge behind different router aliases and to compare in-loop and independent scorers on exactly paired immutable inputs. Fingerprints must cover behavior-changing model, prompt, rubric, decoding and version fields while allowing nonbehavioral deployment metadata to differ. Preserve every verdict and retry, report disagreement and missingness explicitly, reject unpaired context drift, and keep ordinary single-judge monitoring available under its correct non-independent role.

SOURCE PROBLEM: `TRANSFER-QUEUE-STREAM-IDENTITY-EMPTY-RESPONSE-INTEGRITY`

Task: `transfer-queue-empty-record-contract-tests-easy`

Type: Tests Area: Training infrastructure Difficulty: Hard

> Add deterministic CPU tests for the supplied acknowledged ReplayBuffer/AgentLoopWorkerTQ handoff; keep production code unchanged. Use a controlled TransferQueue delegate. Verify ordinary, EOS-only, timeout-before-token and strict-format empty outcomes create distinct completed logical records with stable uid/session/attempt identity, aligned prompts/responses/input_ids, response/loss/attention masks and explicit terminal status. Empty is not absent or padding. Duplicate producer delivery before and after commit must not create another training record. A consumer failure after claim releases the same record for retry; successful commit prevents replay. Kill arrival-order join, empty-drop, mask-shape, duplicate-commit, padding-confusion and real-id reuse mutants. Independently derive keys, shapes and state transitions rather than calling the implementation as oracle.

Task: `transfer-queue-empty-response-debug-easy`

Type: Repair Area: Training infrastructure Difficulty: Easy

Repair transferqueue_utils.py while preserving its public API and ordinary nonempty behavior. Completed empty and EOS-only outcomes must remain observable with explicit terminal metadata and coherent token and mask shapes; absent records and padding must remain distinct. Keep the module isolation-testable and return useful shape and lifecycle diagnostics.

Task: `transfer-queue-empty-response-debug-hard`

Type: Repair Area: Training infrastructure Difficulty: Hard

Repair the queue, bridge, and protocol lifecycle across consumer restart, empty outcomes, and partial responses. Preserve public APIs, stable logical and raw or expanded prompt identity, valid EOS-only outcomes, coherent masks, and bounded completion of aborted partial rollouts. Prefetch depth and fragment arrival order must not change logical training records. Expose lifecycle and commit diagnostics.

Task: `transfer-queue-empty-response-debug-medium`

Type: Repair Area: Training infrastructure Difficulty: Medium

Repair the TransferQueue bridge and protocol surfaces without changing their public entry points. Preserve stable logical and raw-prompt identity, coherent mask shapes, and correct completion attribution across empty outputs and retries. Keep valid empty or EOS-only completions distinct from missing records and padding, retain ordinary nonempty behavior, and expose lifecycle status, logical length, and storage-shape diagnostics.

Task: `transfer-queue-exactly-once-empty-records-easy`

Type: Implementation Area: Training infrastructure Difficulty: Easy

Add an acknowledged TransferQueue handoff between `ReplayBuffer` and the PPO training path. Keep `sample(partition_id, global_steps=..., batch_size=...)` and `remove(partition_id, keys)`, and add `release(partition_id, keys)`. Sampling must atomically claim eligible successful records, including records with `response_len == 0`; a live claim cannot be sampled again. The training path must commit a claimed batch only after its downstream step finishes and durable queue cleanup succeeds. If downstream work or cleanup fails, release the claim so a later attempt can consume it. `remove` commits and prevents a late metadata poll from resurrecting a key; explicit `release` re-enables only an uncommitted claim. Preserve partition isolation, thread safety, polling shutdown, existing keys and tags, ordinary nonempty records, and current method signatures. Use the existing partition, key, status, global step, and response length rather than adding another identity, retry token, or record schema.

SOURCE PROBLEM: `VALIDATION-REWARD-FAILURE-SENTINEL-METRIC-INTEGRITY`

Task: `validation-sentinel-provenance-tests-easy`

Type: Tests Area: Post-training Difficulty: Hard

Add a directly runnable CPU test suite for the documented validation-reward failure behavior. Exercise the real extract_reward, process_validation_metrics, and RayPPOTrainer._val_metrics_update callables with controlled objects. Construct aligned reward, acc, pred, rm_request_id, and rm_response records covering ordinary binary and dense scores, valid zero, exact -100 and -201 failures, all-valid, mixed, and all-failed groups, and the same uid under two data sources. Derive expected semantic means and majority results from the valid records, then relate the reducer's existing nested metrics to the trainer's flattened keys. Show that a failure in either semantic channel excludes that record from both semantic views and its aligned vote, that an all-failed group produces no numeric semantic aggregate, and that diagnostic fields remain aligned and unmodified. Do not edit production code, replace or copy the callables, inspect source text at runtime, or require workers, GPUs, or network access.

Task: `validation-sentinel-provenance-tests-hard`

Type: Tests Area: Post-training Difficulty: Hard

> Add deterministic CPU integration tests for validation metric aggregation and asynchronous diagnostic dumping across consecutive validation cycles. Relate one-shot, partitioned, reordered, and resumed-looking input sequences by stable logical identity: semantic results must be invariant, each input must be counted once, and failed reward judgements must remain attached to their original provenance in persisted output. Exercise the real extraction, merge, publication, executor lifecycle, and JSONL writer with temporary storage. Include mixed-validity and no-valid groups and verify that completing a later cycle cannot contaminate or overwrite the earlier one. Avoid timing assumptions by controlling synchronization through public boundaries. Do not edit production code, bypass the executor, replace serialization or aggregation, inspect source text, or require network, Ray workers, or GPUs.

Task: `validation-sentinel-provenance-tests-medium`

Type: Tests Area: Post-training Difficulty: Hard

> Add a deterministic CPU integration test suite for the existing validation-reward failure path. Drive RayPPOTrainer._validate in both publication and merged-result modes with local DataProto batches and controlled rollout outputs, so the real extract_reward-to-raw-result handoff runs; compare one-batch publication with two real shard results passed to _merge_validation_results under different splits and record permutations. Use stable uid/data-source identities, ordinary binary and dense values including zero, reward -201 failures, acc -100 failures, aligned pred, rm_request_id and rm_response, and all-valid, mixed, and all-failed groups. Derive valid-record means and fixed-layout majority results, compare only order-stable metrics across permutations, require each known failure record to be excluded from reward, acc, and pred together, and require an all-failed group to publish no numeric semantic aggregate. Reuse _validate-produced reward-extra columns with synchronous _write_generations in temporary storage, then normalize JSONL rows by rm_request_id and verify failed scores and responses remain attached. Do not edit production code, replace tested callables, copy their implementations, inspect source text at runtime, or require the asynchronous dump executor, Ray workers, GPUs, network, sleeps, or timing assumptions.

C Representative Task Examples

The construction workflow in Appendix A preserves important development behavior in tasks that can run without the original workspace. The three examples here illustrate that choice using routine training-infrastructure problems: recovering after an interruption, tracking asynchronous jobs, and controlling queue growth. Each example explains the problem, the solver's assignment, one reference solution, and how the evaluator checks correctness. The reference demonstrates a valid approach rather than a required implementation. All three use public task environments and synthetic inputs. Full instructions appear in Appendix B.

C.1 Checkpoint Recovery

A checkpoint is a saved training state from which execution can resume after an interruption. The repair task `checkpoint-resume-state-repair-easy` asks the agent to restore that state consistently. Loading model weights alone is insufficient if the optimizer, learning-rate scheduler, random-number generators, or dataloader resume from a different point. The starting repository contains the checkpoint handler and checkpoint manager in `checkpoint_handler.py` and `fsdp_checkpoint_manager.py`. The agent must repair their coordination while preserving the existing resume modes and checkpoint format.

Recovery therefore requires both a complete saved state and a reliable way to identify it. The reference solution uses the published checkpoint record to select a completed save. It checks the requested files and state before restoring live objects, returns the step associated with that checkpoint, and publishes a new checkpoint only after all required save operations succeed. A failed save must

leave the previous committed checkpoint available. Model-only loading must leave the other training components unchanged.

To check that recovery preserves training behavior, the evaluator compares a short uninterrupted CPU training run with a run that saves and resumes in a fresh process. It checks the resumed step, dataloader position, learning rate, random draws, and first optimizer update. Separate cases remove or corrupt checkpoint components and check that loading fails before live state changes. Other checks cover publication, retention, cleanup, and compatibility with complete checkpoints in the existing format. These small executions test recovery behavior without requiring a large model or a training cluster.

C.2 Asynchronous Evaluation

Asynchronous evaluation runs alongside training, so its results can arrive after the training state has changed. The repair task `async-eval-freshness-debug-easy` asks the agent to keep those results associated with the correct checkpoint. For example, an evaluation submitted for checkpoint 20 may complete while training is at checkpoint 40. Its result must still belong to checkpoint 20. If another scheduled evaluation is skipped because a job is already active, that skipped event must remain visible. The supplied scheduler and evaluator dispatcher expose these behaviors in `async_eval_scheduler.py` and `main_eval.py`.

The agent must preserve checkpoint identity from job submission to callback handling, where a callback is the notification of a job's success or failure. It must also record skipped evaluations, prevent duplicate publication, and keep the number of active jobs bounded. The reference stores an immutable identity for each submitted request and routes callbacks by request identifier. It records each terminal outcome once, so a delayed duplicate callback cannot publish another result or release the slot occupied by a newer request.

The order of events is the main source of difficulty, so the evaluator controls submission and callback order directly. It advances the training state before delivering a result, sends duplicate success and failure callbacks, and checks that unknown request identifiers are rejected. It also checks skipped events, capacity release, exact checkpoint lookup, and the normal case where evaluation finishes immediately. This checks that each result and capacity update belongs to the correct job without depending on the speed or availability of a live evaluation service.

C.3 Queue Backpressure

When generation produces training data faster than training consumes it, queued work can grow unnecessarily. *Backpressure* controls this growth by limiting new work. The implementation task `streaming-backlog-controller-feature-easy` asks the agent to complete an admission controller, which decides how much new work may start, in `backlog_controller.py`. Its inputs are the fixed training batch size, completed queue backlog, active generation count, permitted extra-completion fraction, and low and high watermarks (queue thresholds for reducing or stopping admission). The output is a deterministic budget for starting new work, together with diagnostics. Throttling producers must preserve the configured optimizer batch size.

The controller must account for work already running because it can later join the queue. The reference computes the maximum outstanding work from the batch size and extra allowance, then accounts for both completed and active work before admitting more. At low queue pressure it can fill the remaining capacity. Between the watermarks it limits admission, and at high pressure it stops new work. For a synthetic batch size of eight and an extra allowance of 25%, the outstanding-work limit is ten. Two completed items and three active items leave five available slots, which can be admitted when the backlog is below the low watermark.

The evaluator imports the module independently and checks low, intermediate, and high queue-pressure cases. It verifies that active work consumes capacity, the logical batch size remains fixed, and repeated inputs produce identical decisions and diagnostics. Invalid counts, fractions, and watermark settings must be rejected. This task evaluates the controller through queue snapshots. It does not require running a distributed producer or modifying the training loop.

D Evaluation Implementation

This appendix gives the execution and aggregation details behind Section 3. An *attempt* is one run of an agent on a fresh task workspace. The agent receives the instruction and starting repository, makes changes within its budget, and submits files for grading. The solver environment and the verifier serve different purposes: the former supports the agent's work, while the latter runs the hidden evaluator on the submitted result.

D.1 Execution Settings and Attempt Coverage

The agent containers specify four CPUs and 8 GB of memory, with a 7,200-second time limit for every task. Grading runs in a separate offline verifier with two CPUs, 4 GB of memory, and a 900-second limit. These separate budgets distinguish time available for solving from time available for checking the result. Hy4 Preview, GLM-5.2, and DeepSeek V4.1 Flash use Terminus-2 version 2.0.0, while GPT 5.6 uses Codex CLI. Both harnesses manage model calls and tool execution under limits of 150 turns and 32,768 output tokens per call.

DeepSeek V4 Pro contributes 315 attempts, with three attempts per task. Hy4 Preview, GLM, DeepSeek V4.1 Flash, and GPT 5.6 each contribute 105 scored attempts, one per task. The DeepSeek V4 Pro attempts come from the accepted construction round and were used in task selection, as explained in Section 3. These counts describe the available observations, not equal amounts of repeated evaluation for every model.

D.2 Aggregating Scores and Estimating Uncertainty

We first average attempt scores within each task and model, then average equally across the 105 tasks. Thus, DeepSeek V4 Pro's three attempts contribute to one task mean without giving that task three times the weight. The model comparisons use these means. The best-of-three DeepSeek V4 Pro score is used separately to assign the observed difficulty groups in Section 2.

Several tasks derive from the same source problem and share behavior or repository context. To account for this relationship, we resample source problems together with all their task variants when computing bootstrap intervals. This grouping does not assume that variants have similar scores. An Easy and a Hard variant keep their own scores even when sampled together.

For each of 10,000 draws, we sample 32 source problems with replacement. Each time a source problem is sampled, all its task variants enter the draw. We use each task's mean score to compute an equally weighted mean over the resulting task entries. Repeatedly sampled problems therefore contribute repeated copies of their variants. We report the 2.5th and 97.5th percentiles of these means as the 95% interval. It describes sensitivity to the set of source problems represented, rather than uncertainty from rerunning an agent on a fixed task. Problems with more variants still contribute more task entries, consistent with the task-weighted overall score.

E Detailed Results

The aggregate results in Section 3 can obscure substantial variation across individual tasks. Figure 8 therefore presents task-level scores and trajectory statistics for all 105 tasks. Within each development area, we order tasks by their mean score across models (column Avg.) to reveal the progression from tasks that most models solve to those on which most models struggle. We also report turns, input and output tokens, and tool calls for each model, enabling comparisons of both performance and interaction patterns.

F Detailed Trajectory Analysis

Scores show how many evaluated requirements a submission satisfies, but do not explain the actions that led to that result. We therefore compare three sources of evidence: *trajectories*, which record agent messages and tool actions, submitted files, and verifier check results. A local test is a check run by the agent while solving. A verifier check evaluates the final submission. Keeping these separate lets us examine whether the agent's own evidence supports its completion claims.

(a)



Figure 8: **Detailed task-level results across RSIBench-Code.** Panels (a) and (b) together cover all 105 tasks. Tasks are grouped by development area and ordered within each area by the mean score across models (column Avg.). Score cells report percentages, while the remaining heatmaps show turns, input tokens, output tokens, and tool calls for each model. Task types and difficulty levels are shown beside the task names. Gray cells indicate missing records.

(b)



Figure 8: **Detailed task-level results across RSIBench-Code (continued).** Panel (b) covers training infrastructure, serving, evaluation, and agent systems.

We inspect eight matched tasks, meaning the same task is examined for each model included in the case studies, spanning all five task types (Table 3). For DeepSeek V4 Pro, we inspect all three recorded attempts; for Hy4 Preview, GLM-5.2, DeepSeek V4.1 Flash, and GPT 5.6, we inspect the single scored attempt. We select cases with contrasting outcomes or shared failures to examine how agents interpret requirements, implement changes, and verify their work. These case studies explain observed outcomes without estimating the frequency of each behavior across the benchmark.

Task	DeepSeek (three attempts)	Hy4	GLM	DeepSeek V4.1 Flash	GPT 5.6
Actor–critic reward repair	45.5, 45.5, 45.5	100.0	100.0	62.0	78.0
Multi-teacher routing	65.2, 13.0, 60.9	100.0	100.0	70.0	92.0
Redundant-span rewards	5.0, 30.0, 5.0	30.0	75.0	38.0	52.0
Tool-event contract tests	6.2, 6.2, 6.2	93.8	6.2	24.0	45.0
Checkpoint resume	28.6, 23.8, 0.0	52.4	52.4	55.0	82.0
Sandbox lease recovery	73.3, 73.3, 73.3	40.0	73.3	76.0	48.0
Prompt–rollout allocation	38.9, 38.9, 38.9	83.3	38.9	24.0	35.0
Reverse-KL alignment	5.6, 5.6, 0.0	38.9	5.6	66.0	88.0

Table 3: Recorded scores (%) for the matched trajectory case studies. All three retained DeepSeek V4 Pro attempts are shown separately, followed by single attempts for Hy4, GLM, DeepSeek V4.1 Flash, and GPT 5.6. Cases span implementation, repair, test writing, experiments, and alignment.

Repairing one defect can leave the original failure intact. In the actor–critic reward repair task, actor advantages guide policy updates and critic returns provide training targets for value estimation. The task requires these computations to use their respective reward channels independently. Whitening rescales a trace using its mean and variance, so normalizing both channels together can couple them even if their earlier computations are separate.

All three DeepSeek V4 Pro attempts identify and remove a channel swap in the recurrent trace, but continue to whiten the combined actor and critic traces. This shared normalization still lets critic rewards affect actor advantages. Their local checks and completion messages describe the channels as independent, while the submitted code fails the isolation and whitening checks. Hy4 Preview and GLM-5.2 identify both coupling mechanisms, separate the recurrent updates, and whiten only the actor trace. Their trajectories include checks against the ordinary GAE path and varying reward and lambda channels, and both submissions satisfy all scored requirements. GPT 5.6 earns higher partial credit than the DeepSeek V4 Pro attempts, but still does not satisfy the full contract. This case shows why finding one visible defect is insufficient when the requirement concerns independence throughout a computation.

Interface compatibility and defensive validation matter. The multi-teacher routing task selects a teacher model for a request and prepares inputs within that model's context limit. Correct routing must also preserve the expected interfaces and prevent later changes to caller-owned inputs from altering the request.

Hy4 Preview and GLM-5.2 implement the expected interfaces, copy mutable inputs into immutable representations, reject malformed configuration, and keep the newest tokens within the selected teacher's context limit. GLM-5.2's local checks expose and correct a mismatch between the original sequence length and the remaining token range. DeepSeek V4 Pro's strongest attempt handles routing and truncation, but still fails checks for malformed identity fields, modalities, context limits, and defensive copying. Another DeepSeek V4 Pro attempt introduces different construction and request-planning interfaces and fails most integration checks despite reporting successful local tests. GPT 5.6 reaches high partial credit on this case, but the remaining failures still reflect interface and defensive-validation requirements. A frozen dataclass prevents reassignment of its fields, but does not by itself protect mutable objects stored inside them. A working routing example therefore does not establish correct handling of malformed or mutable inputs.

A plausible reward design can violate a conservation requirement. The redundant-span task adjusts token rewards to penalize redundant passages while preserving the final token's reward and

the sum over valid tokens. This is a conservation requirement: subtracting a penalty in one part of the sequence requires compensating elsewhere without changing the final token.

Hy4 Preview and one DeepSeek V4 Pro attempt subtract span penalties without compensating other tokens, treating the resulting reduction in total reward as acceptable. Two other DeepSeek V4 Pro attempts expose different public function names and fail at the expected entry point. GLM-5.2 redistributes the penalty to eligible non-redundant tokens while keeping the terminal token fixed, satisfying the single-span conservation check. Its span normalization, fallback behavior, and candidate-selection checks remain incomplete. DeepSeek V4.1 Flash and GPT 5.6 obtain intermediate credit, showing partial progress but not complete preservation of the requested reward invariant. This case separates implementing a reasonable shaping rule from enforcing the specific invariant requested by the task.

Passing tests need to exercise the evaluated implementation. The tool-event shaping task asks for tests of code that converts tool events into records and rewards. The verifier checks these tests by introducing mutations, deliberately defective versions of that code. A useful test suite must accept the correct implementation and detect the relevant defects.

Hy4 Preview constructs adversarial traces, derives expected records and rewards independently, and tests its suite against locally injected defects. Its submitted suite detects 14 of 15 verifier mutations, missing the real `AgentLoopOutput` field-propagation case. Two DeepSeek V4 Pro attempts produce suites with 34 and 36 passing tests, but load the shaper by file path under private module names. The verifier substitutes defects in the public module, so those suites exercise a separate module instance and reject none of the substitutions. The third DeepSeek V4 Pro attempt and GLM-5.2 write their expanded suites to other filenames, leaving the collected submission as the original smoke test. Later model runs improve only partially on this task, so the core failure remains integration rather than test count alone. Local tests must exercise the same implementation boundary that the verifier mutates.

Successful resume does not establish failure atomicity. Checkpoint recovery must work both when a save is valid and when saving or loading fails. Failure atomicity means that an unsuccessful operation does not leave partially updated live state or replace a valid saved checkpoint.

On this task, Hy4 Preview and GLM-5.2 pass the fresh-process resume check but each fail ten checks, with nine shared failures involving invalid saved state and failed-save recovery. Hy4 Preview additionally fails atomic tracker publication, while GLM-5.2 accepts an uncommitted directory during discovery. Their completion messages claim validation before mutation and preservation of previous checkpoints, which the verifier does not confirm. DeepSeek V4 Pro attempts broader changes across the handler, manager, and engine layers. Two submissions require new validation or commit methods absent from the existing engine wrapper used in evaluation, breaking the basic save/resume path. The third submission fails to import in evaluation. GPT 5.6 receives substantially higher credit on this case, but remaining failures still show that successful resume does not by itself establish failure atomicity. The cases expose two different gaps: incomplete failure handling in a working path, and loss of compatibility while extending that path.

Resource recovery requires checking the order of lifecycle events. A sandbox lease grants use of an execution environment. If that environment becomes unusable, recovery must replace it or report that execution is unavailable, while releasing the failed lease at the required point.

On this task, DeepSeek V4 Pro and GLM-5.2 implement replacement and terminal unavailability behavior more successfully than Hy4 Preview. DeepSeek V4.1 Flash reaches a similar partial solution, while GPT 5.6 scores lower on this case despite its stronger aggregate performance. However, the remaining failures have the same structure: tracking that a lease is retired is insufficient when the contract requires release before a later execution. DeepSeek V4 Pro and GLM-5.2 also miss second-failure release and cleanup-error checks. Hy4 Preview's trajectory spends its final turns reconciling local stubs with tool return values and ends at the execution limit while revising that handling. GLM-5.2 reports passing local recovery tests, yet the verifier exposes the remaining ordering failures.

Executing an experiment is distinct from validating its controls. The prompt–rollout allocation task compares ways to divide a fixed budget between distinct prompts and repeated generated responses, or rollouts. Its runner executes the study and its analyzer estimates the effects of the allocation choices. Shared data blocks, matched comparisons, and correct budget accounting are needed to make those effects interpretable.

All DeepSeek V4 Pro attempts and GLM-5.2 execute the supplied study and inspect its generated reports, leaving the runner and analyzer unchanged. They therefore still fail checks for block reuse, budget accounting, logical coverage, paired comparisons, and uncertainty computation. Hy4 Preview inspects and modifies both modules, checks condition-level invariants, and regenerates the results. Its submission passes the block-reuse, token-budget, coverage, and bootstrap checks that the others fail, although plan validation, exact response-budget products, and paired-effect accounting remain incorrect. GPT 5.6 also remains weak on this task relative to its aggregate performance, showing that stronger general coding performance does not automatically produce valid experimental controls. Producing valid report files does not by itself validate the experimental design or the computations behind the reported effects.

Alignment tests must reach the actual backward computation. The reverse-KL alignment task checks whether an estimator of reverse Kullback–Leibler divergence produces the required values and gradients. Gradients determine how the result affects training, so matching forward values alone is insufficient. The tests must reach the registered estimator, the implementation selected through the repository's interface, and execute its backward computation to check gradients.

DeepSeek V4 Pro and GLM-5.2 describe value and gradient comparisons, but their submissions encounter interface or tensor-layout failures in the correct-estimator control, where the verifier runs the submitted tests against the correct implementation. One DeepSeek V4 Pro attempt and GLM-5.2 expect padded loss shapes where the registered estimator returns flattened token losses. Local syntax checks and tests with substitute implementations do not detect this mismatch. Hy4 Preview reaches the registered estimator and rejects a malformed output shape, but its submission still misses gradient, stale-policy, and padding checks. GPT 5.6 makes substantially more progress on this task, yet the task is still not fully solved. The case therefore preserves the same lesson: alignment tests must reach the actual registered computation and its backward path, not only a compatible-looking substitute.

Across these cases, model differences depend on the task and on which parts of its contract are exercised. Hy4 Preview distinguishes itself in the tool-event and allocation cases, GLM-5.2 better preserves the reward-conservation invariant, DeepSeek V4 Pro matches GLM-5.2 on lease recovery, DeepSeek V4.1 Flash often reaches compact partial solutions, and GPT 5.6 obtains the strongest results on several matched cases while still failing others such as prompt–rollout allocation. Successful attempts connect requirements to observable checks on the submitted implementation. Several incomplete attempts establish narrower evidence, such as compilation, a passing local example, or a generated report, and then claim the broader task complete.

G Related Work

Our work connects the goal of improving LLM-development code with the tasks needed to evaluate and develop that capability. We review related work on recursive self-improvement, executable task generation, and evaluation of agents on model-development problems.

G.1 Recursive Self-Improvement through LLM-Development Code

Code is an important object of self-improvement because it governs both how agents operate and how models are developed. Harness-evolution methods use execution feedback to revise the code and instructions surrounding a fixed model, improving how the agent interacts with its environment (Lee et al., 2026b, Zhang et al., 2026a, Lee et al., 2026a). Hyperagents also makes the procedure that modifies the task agent editable, so the improvement mechanism can change along with the agent (Zhang et al., 2026b).

Other methods connect code changes to model training. SIA and Continual Harness combine harness adaptation with model-weight updates (Hebbar et al., 2026, Karten et al., 2026). Bilevel Au-

toresearch and EvoTrainer treat parts of the training process as objects of improvement: the former modifies the search procedure guiding training experiments, while the latter co-evolves model policies and training-side diagnostics and procedures (Qu & Lu, 2026, Chen et al., 2026). These studies show that RSI research includes both improving an agent's execution and improving the process that trains models.

RSIBench-Code studies the practical code-level capabilities needed to contribute to LLM development. Its tasks expose agents to post-training, training infrastructure, serving, evaluation, and agent-system problems, providing a way to evaluate whether agents can perform the code-level work required by RSI and potential data for developing those capabilities through training. This complements methods that specify how an improvement loop updates an agent or a model.

G.2 Generating Executable Tasks for Self-Improvement

Task generation begins with meaningful objectives and enough context to reconstruct them. Repository issues and code changes provide context in SWE-bench, SWT-Bench, and WorkBuddy Bench (Jimenez et al., 2024, Mündler et al., 2024, Cai et al., 2026), while CursorBench (Jain, 2026) derives tasks from real internal coding sessions. TerminalWorld further automates reconstruction of environments and tests from terminal recordings (Chu et al., 2026). These approaches connect executable tasks to practical work. RSIBench-Code draws on problems encountered in enterprise LLM development to capture technical objectives and constraints from the code used to train, evaluate, serve, and operate models.

A second challenge is expanding the task supply while preserving reliable feedback. SWE-smith and R2E-Gym construct coding tasks through generated changes and tests (Yang et al., 2025, Jain et al., 2025). Terminal-Task-Gen and Meta-Task extend automated construction to varied executable assignments, including their environments and verification (Pi et al., 2026, Pan et al., 2026). Together, these methods illustrate how generation and execution checks can expand the available tasks beyond a fixed collection of existing requests.

Task generation can also respond to a learner's progress. Absolute Zero, SOAR, and Skill Self-Play couple task proposals with verification or measured learning feedback (Zhao et al., 2025, Sundaram et al., 2026, Huang et al., 2026). Training on executable tasks and their trajectories also develops coding-agent capabilities in SWE-Gym, Hybrid-Gym, SWE-smith, and Terminal-Task-Gen (Pan et al., 2025, Xie et al., 2026, Yang et al., 2025, Pi et al., 2026). These results motivate a continuing supply of tasks that can support both evaluation and learning.

RSIBench-Code combines practical enterprise problem selection with generation of multiple task types and difficulty levels. Execution checks and solver trials guide validation and refinement, connecting the original technical objective to an executable task whose quality and difficulty can be assessed. The workflow is designed for repeated use as new development problems arise, and its outputs provide evaluation tasks and potential training data for RSI.

G.3 Evaluating Agents on LLM-Development Tasks

Several benchmarks assess agents through the models or research results they produce. PostTrainBench and AI4AI-Bench measure the results of agent-directed post-training or changes to training algorithms (Rank et al., 2026, Chi et al., 2026). RSIBench-Data isolates training-data research within a fixed training stack and evaluates the resulting checkpoints (Meng et al., 2026). RE-Bench evaluates research-engineering performance with human baselines under matched resource budgets (Wijk et al., 2025). These benchmarks connect development decisions to measured outcomes.

Broader research benchmarks assess experimentation, replication, and improvement of ML methods (Chan et al., 2025, Starace et al., 2025, Garikaparthi et al., 2026). InnovatorBench covers several components of LLM development, while MLS-Bench examines whether component improvements generalize across settings and scales (Wu et al., 2025, Lyu et al., 2026). Coding and terminal benchmarks provide complementary evaluations of agents completing requests in executable environments (Jimenez et al., 2024, Merrill et al., 2026).

RSIBench-Code evaluates whether agents correctly complete practical tasks in enterprise-level LLM-development codebases. Its checks measure the submitted result against the task's behav-

ioral requirements. This complements evaluations based on full training runs or research outcomes by examining the development work needed to implement and maintain those processes. Comparisons across task types and development areas, together with trajectory analysis, identify strengths and failure modes relevant to further self-improvement.

H LIMITATIONS

The results describe performance on a selected collection of reconstructed enterprise LLM-development tasks. The following limits affect how they should be interpreted and generalized.

Coverage and weighting. RSIBench-Code covers work that can be reconstructed on the public `verl` and `sglang` repositories and evaluated offline. Although these repositories contain approximately 1.8 million source lines, they are still only a public proxy for the wider enterprise code used in model development.

Limited computational resources. Our compute budget limits the current tasks to CPU environments without GPU access. We therefore focus on small experiments and checks of individual components, rather than full training or post-training runs such as those in PostTrainBench (Rank et al., 2026). Limited memory and execution time also constrain model and dataset sizes, experiment duration, and the scale of distributed training or serving workloads. These tasks can test whether an agent implements the required behavior, but do not establish whether its changes improve model quality or remain effective at larger scales. Broader evaluation would require GPU resources, longer runs, and larger experimental budgets.