arXiv:2610.02405v2 [cs.AI] 16 Nov 2026

When Terminal-Agent Training Stalls: Demystifying Data Generation and Verification Challenges

Xi Qin
SAP Lab
Palo Alto, CA 94304
grace.qin@sap.com

Isabel Kurth*
SAP Lab
Palo Alto, CA 94304
isabel.kurth@t-online.de

Xin Cui
SAP Lab
Palo Alto, CA 94304
xin.cui@sap.com

Elin Park
SAP Lab
Palo Alto, CA 94304
elin.park@sap.com

Alexander Schaefer
SAP Lab
Palo Alto, CA 94304
alexander.schaefer01@sap.com

Yaad Oren
SAP Lab
Palo Alto, CA 94304
yaad.oren@sap.com

Abstract

Automated terminal-agent training relies on generated environments and verifiers to turn open-ended interaction into reward. Yet a runnable Docker image and executable test suite do not guarantee that the resulting reward is faithful, informative, or even attributable to agent behavior. We present a pipeline that generates containerized terminal tasks and their executable verifiers, and use it to demystify failure modes in task generation and live-environment evaluation. Our analysis separates four requirements that are often conflated: structural validity, verifier faithfulness, policy solvability, and infrastructure reliability. In a Qwen3.5-9B GRPO training, the model overfits on a relatively easy dataset in 20 steps. In contrast, in a controlled Qwen3.5-9B comparison, adding hard tasks reduces mean rollout $\mathrm{pass@2}$ from 81.3% to 20.6%, restoring informative failures without changing the training configuration. We further identify duplicate task generation, incomplete anti-shortcut safeguards, Dockerfile failures, asynchronous $\mathrm{pass@}k$ bias, and environment errors caused by concurrency constraints. All impact the precision in measurement. These findings should motivate the community for a terminal agent specific evaluation harness that pairs structural and semantic checks, and task solvability with extensive verifier audits.

1 Introduction

Autonomous terminal agents, as models that issue shell commands, read outputs, and iterate to complete engineering tasks, are a promising direction for automating software operations. Reinforcement learning is a natural training approach: either a binary pass-or-fail signal or multi-state checks from a test suite, provides a clean reward without human labeling. The bottleneck lies in environment harness. Writing executable terminal tasks with correct initial states, verifiable final states, and graded rewards is labor-intensive, and the resulting datasets are small.

Gandhi et al. Gandhi et al. [2025] address environment construction with Endless Terminals, a pipeline that procedurally generates tasks without human annotation. Their pipeline produces structurally valid tasks: the container boots, the test suite runs, and this process discards roughly half of all generated candidates by OpenAI O3 model, one of the top models at that time. This seemingly sparse reward would highly possibly make RL training difficult: even with 16 samples per task, most batches

*This work was done when Isabel was an intern at SAP.

Workshop on Who Verifies the Agents? Toward Reliable Agent Development, NeurIPS 2026.

return all-zero reward and produce no gradient. Future work is left to quantify the impacts of this sparsity issue on training instability.

In this work, we aim to bridge this gap by first analyzing the challenges of data generation and verification in this brittle pipeline.

Endless Terminals and Harbor

Endless Terminals Gandhi et al. [2025] generates terminal-agent tasks in five stages: task description and ground truth, initial-state pytest, final-state pytest, Apptainer container definition, and packaged task bundle. Each stage is a batched LLM call conditioned on earlier outputs. The pipeline requires no human annotation and scales with inference computation. Structure-wise, we preserve this five-stage skeleton and replace the pipelines of data generation and curation, evaluation harness, and LLM providers.

Harbor Harbor Framework Team [2026] is an open-source benchmark framework for terminal agents. It prescribes a fixed directory layout per task: an instruction file, a `task.toml` metadata file, a Docker environment, a reference solution, and a pytest verifier that writes binary reward. Harbor's CLI builds task images, attaches an agent, drives terminal interaction, and collects results. We adopt Harbor as our task format and evaluation harness.

We make three contributions.

1. **Terminal Task Dataset** We generate 18,185 tasks with 8192-token context windows from Claude Opus models, **the largest** to our best knowledge.
2. **Task Quality Improvements.** We design the generation prompts thoroughly around three dimensions: engineer voice (casual, storytelling with deliberate imperfections), anti-leakage (prohibiting restatement of verification criteria), and domain coverage (round-robin balancing across eight task categories). We replace syntactic complexity levels with three levels of cognitive difficulty and thread difficulty through every generation stage.
3. **Structural Validity Is Not Benchmark Validity.** We demystify three failure modes in automated terminal task generation and verification: task calibration, verifier loopholes, and live-environment infrastructure failures can corrupt the reward signal used to evaluate and train agents.

2 Acknowledgment

The authors gratefully acknowledge Kanishk Gandhi of Stanford University Human-Centered Artificial Intelligence Institution, for his participation in constructive discussions that helped sharpen the analysis and presentation of this work.

3 Data Generation

In our five-stage pipeline (Figure 1), stages 1–3 produce the task description, ground truth, initial-state tests, and final-state tests. Stage 4 produces a Dockerfile to build the sandbox for simulation and verification of a terminal task. Stage 5 writes a Harbor task directory instead of a flat bundle. There are three types of models involved across stages: 1. Generation model for task generation; 2. Generation model for solution generation; 3. The agent policy model. We use Claude Opus 4.6/4.8 for task generation. Following a distillation-inspired approach, we use a less capable model for solution generation than for task generation. For example, when Claude 4.8 Opus generates tasks, Claude 4.8 Sonnet generates these tasks' solutions. The agent policy model is then limited to smaller parameter scales (3B, 4B, or 9B) to reduce training and inference costs.

Here are the details in each stage.

3.1 From Apptainer to Docker and Harbor

The original pipeline targets Apptainer because SLURM clusters require daemonless unprivileged containers. Our setting infuses Harbor, which expects a Dockerfile per task. Stage 4's prompt was rewritten around specific constraints in Docker: heredoc handling, the absence of `systemd`, restrictions on Docker-in-Docker, non-root user execution, and wrapper `CMD` scripts for tasks that

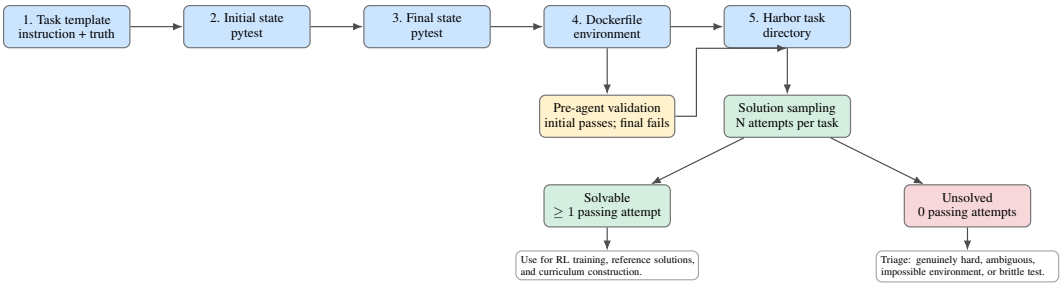


Figure 1: Overall pipeline. Stages 1–5 generate Harbor task directories. Pre-agent validation filters structural failures before rollout. Solution sampling assigns empirical solvability labels, which control RL training signal.

need a service running at agent connect time. The Harbor task directory layout is listed in Appendix B. During training and evaluation, the verifier writes reward 1 when the final-state pytest passes and 0 otherwise. We follow this convention throughout.

3.2 Task-Template Prompt Redesign

The original prompt is approximately thirty lines and specifies a rubric for the `<task>` and `<truth>` XML blocks. It produces two systematic failure modes.

Worksheet Voice. The model produces tasks that sound like graded assignments: numbered requirements, phrases such as "The output should…" and "Ensure that…," and openings like "Please complete the following task." These patterns signal to the agent that it is being evaluated, not helping.

Grader Leakage. The model restates verification criteria inside the task body. For example, asking the agent to "produce a log of every change" when the verifier checks exactly that log. Leakage collapses the task from problem solving to transcription.

We redesign the prompt around three dimensions.

Voice. The narrator is an engineer: casual, and whom could be possibly wrong. The prompt enumerates opening modes – trigger/debug, build, reverse engineering, recovery, migration, investigation, optimization – and requires the model to rotate across them. Each task must contain at least one deliberate imperfection: a hedge, a possibly wrong guess, or a register cue. Banned phrasings are listed verbatim so the model can self-check before emitting.

Anti-leakage. The prompt enumerates leakage behaviors: restating verification criteria, requesting artifacts whose only purpose is grading, enumerating sub-requirements, and explaining the root cause in the task body. The `<truth>` section requires four items in order: initial state, expected final state, invariants, and anti-shortcut guards.

Domain Coverage. To prevent the model from collapsing to a small attractor (in our case, "Python data bug"), the system message lists multiple domains and task types and instructs deliberate rotation. A self-check asks three questions: did the diagnosis leak, does the task contain at least one imperfection, and does the first sentence drop the agent into a situation rather than announcing a project.

The system message carries standing rules. The user message offers task variation: a domain drawn from balanced category buckets, a fictitious narrator persona, and a difficulty level with level-specific guidance. We share the core prompts in Appendix A.

3.3 Difficulty and Category Balancing

Previous work in Gandhi et al. [2025] defines six complexity levels by command count, sampled uniformly and inserted as a bare phrase. We think command count is a poor proxy for solve difficulty: `chmod +x foo` and a one-line `sed` bug fix are both one-command tasks. So we replace it with a three-level `difficulty` dimension defined in cognitive terms. Easy tasks have a single obvious cause and a direct fix. Medium tasks require investigation before action. Hard tasks satisfy at least

one of five criteria: misleading symptom, conflicting constraints, required domain knowledge, large state space, or reverse engineering. Each level includes examples rather than a bare phrase. Difficulty is threaded through every generation stage and written into `task.toml`, aligning with Harbor's metadata schema.

The core round-robin sampling strategy applies across ten high-level buckets enforces category balance: `file_ops`, `networking`, `services`, `security`, `debugging`, `data_processing`, `git_ops`, `system_admin`, `build_tools`, and `text_processing`. Without this, a single batched call produces correlated tasks: if the first task involves `git rebase`, several subsequent tasks do as well.

3.4 Pre-Agent Validation

After Stage 4 builds the Docker image, the pipeline runs the generated tests against the container before any agent rollout. A task survives only when the initial-state test passes and the final-state test fails. The initial pass requirement checks that the Dockerfile correctly seeds the world described in the ground truth. Missing input files, wrong port bindings, or incorrect permissions are caught here rather than during rollout.

This pre-agent validation step is absent from the pipeline in Gandhi et al. [2025], which discovers such mismatches only during solution sampling.

3.5 Solution Generation

After task generation, a solution generation stage samples agent trajectories using Harbor's `Terminus2` agent. The companion script `collect_harbor_results.py` scans Harbor job directories, groups trials by task, and writes `solution.json` with per task $\mathrm{pass@}k$ metrics.

Solution generation assigns an empirical solvability label to each task. Tasks with at least one passing attempt are candidates for RL training datasets. Tasks with zero passing attempts require triage: they may be genuinely hard, underspecified, impossible, or over strictly graded.

4 Dataset Analysis

4.1 Solvability of Generated Tasks

We generate three datasets as summarized in Table 1. Details of assigning categories are in Appendix D. We identify 8 attempts being a sweet point for most solution generation, which is proved by figure 8 in Appendix.

The $5.6\times$ improvement in solvability from $V1$ to $V2$ comes from two sources. Longer context allows more complete environment specifications, reducing cases where the agent cannot infer the missing state. The redesigned prompt reduces grader leakage and worksheet voice, producing tasks that an agent can approach as genuine engineering problems.

Table 1: Dataset overview. For solvability measured by $\mathrm{pass@1}$, Claude 4.6 Sonnet filters for $V1$ and $V2$, and Claude 4.7 Opus for $V3$ specifically.

Dataset	Model	Context	Total tasks	$\mathrm{pass@1}$	Category Assignment
$V1$ (Dataset-4k)	Opus 4.6	4096	5038	$\sim$10%	Round-robin
$V2$ (Dataset-8k)	Opus 4.6	8192	4929	56.5%	Difficulty and category coupling
$V3$ (Dataset-5k)	Opus 4.8	8192	8218	39.6%	Difficulty first, category next

4.2 Difficulty and Category Breakdown

Overall, to avoid heavy tails in easy and hard, we make medium tasks 63%, then easy and hard 18.5% separately. Within $V3$, difficulty is distributed as follows: easy (57%), medium (36%), and hard (7%), after solvability filtering. A subset of randomly sampled 1,000 tasks, has pass rate that monotonically decreases with difficulty (easy: 97%, medium: 92%, hard: 89%). This confirms the cognitive difficulty definitions are meaningful.

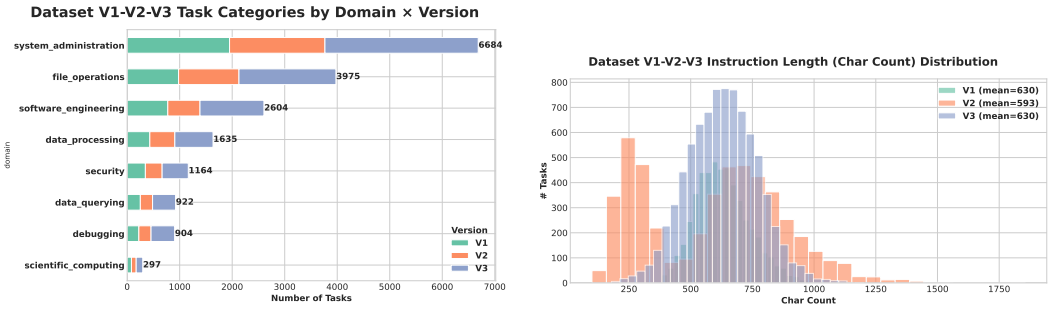


Figure 2: Left: task categories in all three datasets; no single category dominates. Right: instruction length distributions for all datasets.

Round-robin construction makes category coverage near uniform. The eight common task domains are shown in Figure 2 left graph; domain distribution remains in a universal pattern across all datasets.

4.3 Instruction Length Distribution

In Figure 2 right graph, $V2$ instructions has a broader coverage than $V1$ and $V3$, with some extremely long instructions with 1000 more characters, reflecting richer environment specifications enabled by the larger context window. But the caveat is the heavy left tail down to 200 characters, meaning a significant number of unusually short instructions that could detail the trajectory. $V3$ covers a broader range than $V1$ with a nice bell-shaped and tightly centered distribution, which shows a more balanced instruction design.

4.4 Verifier Quality

Assertion count improved steadily from $V1$ (mean=24.3) to $V3$ (mean=27.6), shifting noticeably rightward and having the least low-count tail, suggesting more thorough verifiers over time. However, the semantic type breakdown reveals a trade-off: $V1$ had the highest process output assertions (5.7 per task), which degraded to 2.7 in $V3$. File content assertions similarly collapsed from 5.8 in $V2$ to 1.7 in $V3$, as shown in 3. We show more details in graphs in Appendix C.

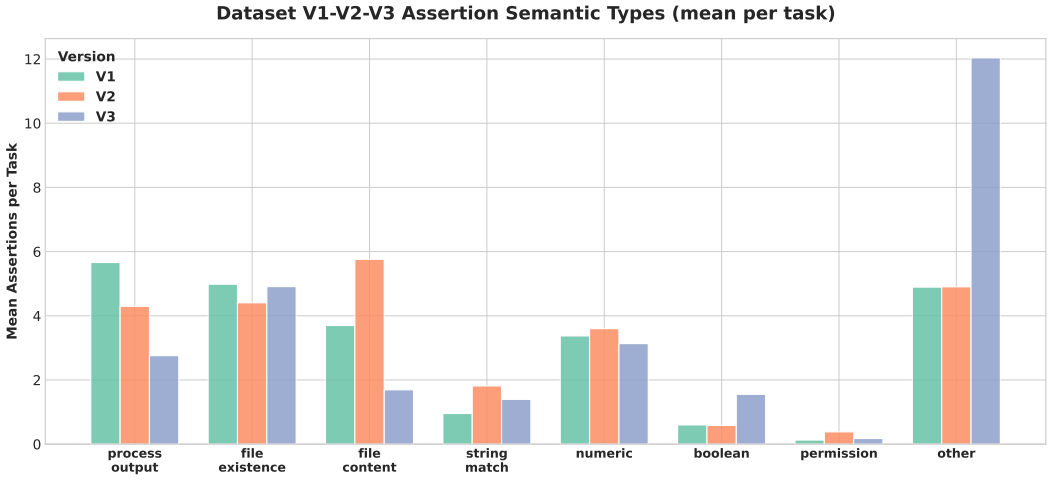


Figure 3: Assertion semantic types across $V1$, $V2$, and $V3$ datasets.

5 Training Experiments

We use training rollouts to test whether a generated task distribution supplies useful RL signal, rather than treating a runnable verifier as sufficient. All experiments use SkyRL Cao et al. [2025]Griggs et al. [2025], binary reward from Harbor's verifier, and the `Terminus2` agent with at most eight turns per rollout.

Before training started, we should assess the solvability baseline, i.e. the range of per-task pass rates where RL training can make progress. In extreme cases, too low rate (near 0%) causes the absence of gradient signals; too high rate (near 100%) causes marginal training gains. Table 2 shares these baselines. As a ceiling reference, Claude Opus 4.6's solve rate (pass@1) over three validation datasets (100, 151, 517 samples respectively from $V1$ to $V3$) from $E1$ to $E3$ are 94.1%, 98.0% and 11.2% respectively; similarly, Qwen3.5-9B has the solve rate in the same three datasets at 86.3%, 92.0%, 9.8% respectively.

Table 2: Training evidence for dataset calibration.

Experiments	Policy Model	Dataset	Baseline Solve Rate
$E1$	Qwen2.5-3B	$V1$	7.8% pass@1
$E2$	Qwen3.5-9B	$V1$+$V2$	81.3% mean pass@2
$E3$	Qwen3.5-9B	$V1$+$V2$+$V3$ hard	20.6% mean pass@2

5.1 $E1$: Sparse Rewards on $V1$

Our first experiment started before QWEN 3.5 was available. So we train `Qwen/Qwen2.5-3B-Instruct` with PPO on 457 tasks downsampled from $V1$ Dataset-4k for two epochs (228 gradient steps). Each step contains four tasks with four samples per task; the maximum sequence length is 4096 tokens. We evaluate on 51 held-out tasks every 20 steps.

The base model achieves 7.8% pass@1 across the whole $V1$ dataset. Training unfortunately does not improve it: validation pass@1 ranges from 3.9% to 5.9%, and the final checkpoint reaches 5.9%. Mean reward during training stays around 6.3% with a standard deviation of 0.11. With roughly 10% task solvability, most batches contain no successful trajectory. Training reward consequently remains zero for most of the run, with only isolated positive reward steps. This is the low end of the solvability band: a structurally valid task set can still provide too little signal for learning.

5.2 $E2$–$E3$: Hard Tasks Restore Headroom for a 9B Model

We next run GRPO with `Qwen/Qwen3.5-9B` on an H200 instance, holding the instance, batch size (eight tasks), two samples per prompt, maximum turns, and sequence length fixed while changing only the task distribution. Due to the limitation of computation resources and considering solvability ceiling, we still downsample from this mixture. The experiment uses 3,238 training tasks. Then $E3$ adds 3,291 hard tasks from $V3$, for 6,529 training tasks.

Table 3: Changing the training distribution changes rollout difficulty.

Experiments	Training set	Steps	Mean pass@2	Mean reward std.
$E2$	$V1$+$V2$	20	81.3% -> 100%	0.42
$E3$	$V1$+$V2$+$V3$ hard	190	9.8% -> 10.1%	0.37

On $V1$+$V2$, Qwen3.5 9B model solves 81.3% of attempts on average for evaluation set; four of the first 20 steps have pass@2 of 100%. Adding $V3$ hard tasks reduces the mean pass@2 to 20.6% and only lifts by 1pp after 190 steps. The mean reward standard deviation remains nonzero in both runs, but the harder mixture restores failures that can distinguish trajectories. Thus the same training configuration moves from a frequently saturated distribution to one with substantially more headroom solely by changing task difficulty.

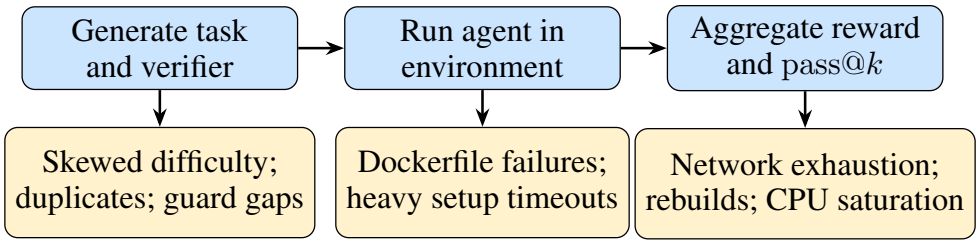


Figure 4: Failure modes enter at generation, environment execution, and operation stages; they require separate structural, solvability, and error rate checks.

Table 4: Observed failure modes from distinct runs.

Failure mode	**Observed scope**	**Implication**
Hard task under-representation	7.0% of a 3,310 task source batch	Insufficient hard cases can leave capable policies near reward saturation.
Dockerfile parse failures	$\sim$74% heredocs parsing errors	Generation defects can dominate apparent agent failure.
Docker network pool exhaustion	$\sim$9.1% in solution generation	Infrastructure errors bias aggregate evaluation until controlled.
Docker build failures	$\sim$9.7% in solution generation	Fragile setup steps consume trials without producing trajectories. This defect remains to be fixed.

This result complements $E1$. Tasks that are too difficult yield almost no positive reward; tasks that are too easy leave too few informative failures. Task generation must therefore be calibrated to the target policy rather than optimized for a single global solvability rate.

6 Failure Modes and Measurement Fidelity

Generated terminal tasks can fail before an agent acts, while an agent acts, or while the harness measures the outcome. These cases have different consequences: a malformed task is not evidence of limited agent capability, and a successful container build is not evidence that a task supplies useful training signal. Figure 4 organizes the observed failure modes by the stage at which they enter the pipeline.

6.1 Generation and Verifier Failures

The most frequent categories are predominantly mechanical workflows, and difficulty is sampled independently of category. A nominal difficulty label does not guarantee that the task creates discriminative reward variation. Therefore, the combination of difficulty level and category distribution should be enforced during round-robin generation.

Deduplication exposes a second generation failure. This process removed 1,753 of 6,682 raw candidate tasks (26.2%) at a similarity threshold of 0.85; 1,461 removed tasks were near duplicates with similarity between 0.85 and 0.95. Repetition inflates dataset size without expanding the actual novel behaviors for agents to learn. Finally, a baseline quality sample found anti-shortcut guards in roughly 60% of task truth blocks, but the generation pipeline did not validate their presence. Missing guards create a reward hacking risk: behavior tailored to the test can receive reward without meeting the intended engineering objectives.

6.2 Environment and Evaluation Failures

Environment construction is a direct confound in live evaluation. In one initial test generation of $V2$, we identified approximately 74% of errored trials were attributed to a Docker heredoc generation bug. In a later 8,218-task solution generation run, Docker network pool exhaustion affected 9.1% of trials and varied Docker build failures affected 9.7%. The former was reduced by periodic network pruning; the latter includes heterogeneous causes such as network dependent package installation,

Git operations, heredoc handling, and database initialization. These are harness failures, not policy outcomes.

Some limitations persist even after failed trials are removed. The same large run rebuilt each task image for every attempt, spending a mean 163 seconds in environment setup; retaining images would reduce rebuilds but was infeasible at the reported scale because of disk demand. Twelve concurrent Docker builds also saturated roughly 29.7 of 32 CPU cores. Separately, three locally evaluated tasks exceeded a 600-second startup limit because their initial states seeded large databases, repositories, or log exports. Together, these observations motivate a designated infrastructure and isolating heavy tasks rather than interpreting every non-pass as an agent mistake.

6.3 Evaluation Measurement Confounds

Two operational findings affect how we interpret evaluation results. First, when vLLM returns more than two samples asynchronously, completion order biases prefix $\mathrm{pass@}k$: short, early returning trajectories can dominate $\mathrm{pass@1}$ or $\mathrm{pass@2}$. We therefore report $\mathrm{pass@2}$ only for the two-sample runs above; for larger offline samples, $\mathrm{pass@}k$ should be estimated from the full set of successful trajectories.

Second, evaluation infrastructure can masquerade as model failure. On the same 151-task validation set, increasing evaluation concurrency to 40 containers caused 144 tasks to error and yielded 4.6% $\mathrm{pass@2}$. Reducing concurrency to 16 containers reduced errors to one task and yielded 76.8% $\mathrm{pass@2}$. We consequently treat task level error rate as a required companion metric to $\mathrm{pass@}k$ for evaluation of the live environment.

7 Discussion and Future Work

Task Solvability Our experiments identify two failure modes: $V1$ is too sparse for Qwen2.5-3B PPO, while the $V1$+$V2$ mixture is frequently near-ceiling for Qwen3.5-9B GRPO. If we want to build a full spectrum of different policy model's solvability, we miss more ablation choices due to the cost constraints. Still, from the solve rate jump from Qwen2.5-3B to Qwen3.5-9B, we observe a great advancement in model capability. Adding $V3$ hard tasks reduces Qwen3.5-9B's solve rate from 81.3% to 20.6% without changing the training configuration. The appropriate training distribution is therefore model dependent; a single global pass rate cannot certify that a benchmark will produce useful gradient signal.

Demystifying live-environment evaluation. A low score in a terminal environment does not, by itself, show that an agent failed at the intended task. It can instead reflect a malformed Dockerfile, a timeout in environmental initialization, exhausted Docker network address pools, or biased aggregation of asynchronously returned samples. Therefore, for researchers working on similar pipelines, we strongly suggest separating three questions from an end-to-end benchmarking:

whether an environment is structurally valid?

whether its verifier faithfully measures the stated task?

whether the target policy can solve it?

Our current pipeline incorporates practical mitigations for these confounds. Examples are like the Dockerfile generator prefixes BuildKit syntax to support heredocs. Evaluation isolates known heavyly build tasks with a longer timeout, limits container concurrency, and runs periodic Docker cleanup to reduce stale network and cache failures.

Finally, we pair $\mathrm{pass@}k$ with task-level error rates and avoid prefix $\mathrm{pass@}k$ for asynchronous runs with more than two samples. These controls do not make a generated verifier fully faithful, but they prevent common infrastructure failures from being reported as agent incapability.

Limitations and Future Work

In measurement, asynchronous vLLM completion order biases prefix $\mathrm{pass@}k$ when more than two samples are generated, so evaluating larger inference batches than existing hundreds of tasks requires an estimator over the full set of trajectories. Moreover, concurrent container failures can dominate an evaluation score; task level error rates must accompany $\mathrm{pass@}k$. The present experiments demonstrate

calibration effects, but do not yet show end-to-end improvement across a full training run in the middle of the solvability band. These constraints are also largely driven by the high costs of RL training and evaluation.

In data, 16.6% tasks show few assertions check the desired command outcomes, and overlap score between task description and test suite flags potential drift from the real task intent.

Therefore, we target the followings as our next steps.

- Scale up training experiments for longer steps and larger training and evaluation samples.
- Filter and improve tasks having low assertion matching score and description–test overlap score.
- Generate more tasks in scientific computing, data science, and security.
- Add intricate multi-stage tasks to broaden the curriculum for capable agents.

8 Related Work

Terminal agent benchmarks. SWE-bench Jimenez et al. [2024] and related work evaluate agents on real GitHub issues in Docker containers, but construction is human curated and the task set is static. Harbor Harbor Framework Team [2026] provides a harness for executing terminal agents against programmatically defined tasks. TMAX Ivison et al. [2026] generates terminal tasks using LLMs; our approach differs in the RL training loop, the solvability analysis, and the integration with Harbor's evaluation API. Gandhi et al. Gandhi et al. [2025] is the direct predecessor; we preserve their generation skeleton and extend it along the dimensions described above.

RL for code and tool-use agents. DeepSeekCoder Guo et al. [2024] and related work show that GRPO with binary reward drives measurable improvement on code generation when the base model already solves a fraction of training problems. Our experiments confirm this pattern and extend it to the terminal domain, where environment statefulness and open action spaces make the solvability requirement harder to satisfy.

Curriculum learning for RL. Several works on curriculum learning for RL Cai et al. [2026]Parashar et al. [2026] show that training on tasks calibrated to the model's current capability accelerates learning. Our solvability hypothesis is an instantiation of this principle: the dataset must be filtered to keep pass rates in the range where gradient signal is present but not saturated.

References

Pengxiang Cai, Tianchen Fang, Xiaohan Li, Qingyuan Zeng, Guocong Li, and Jintai Chen. Curriculum reinforcement learning can incentivize reasoning capacity in llms beyond the base model, 2026. URL `https://arxiv.org/abs/2606.22317`.

Shiyi Cao, Sumanth Hegde, Dacheng Li, Tyler Griggs, Shu Liu, Eric Tang, Jiayi Pan, Xingyao Wang, Akshay Malik, Graham Neubig, Kourosh Hakhamaneshi, Richard Liaw, Philipp Moritz, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. Skyrl-v0: Train real-world long-horizon agents via reinforcement learning, 2025.

Kanishk Gandhi, Shivam Garg, Noah D. Goodman, and Dimitris Papailiopoulos. Endless terminals: Scaling rl environments for terminal agents. *arXiv preprint arXiv:2601.16443*, 2025.

Tyler Griggs, Sumanth Hegde, Eric Tang, Shu Liu, Shiyi Cao, Dacheng Li, Charlie Ruan, Philipp Moritz, Kourosh Hakhamaneshi, Richard Liaw, Akshay Malik, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. Evolving skyrl into a highly-modular rl framework, 2025. Notion Blog.

Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y.K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming – the rise of code intelligence, 2024. URL `https://arxiv.org/abs/2401.14196`.

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`.

Hamish Ivison, Junjie Oscar Yin, Rulin Shao, Teng Xiao, Nathan Lambert, and Hannaneh Hajishirzi. Tmax: A simple recipe for terminal agents, 2026. URL https://arxiv.org/abs/2606.23321.

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*, 2024. URL https://openreview.net/forum?id=VTF8yNQM66.

Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li, Yu Zhang, James Caverlee, Dileep Kalathil, and Shuiwang Ji. Curriculum reinforcement learning from easy to hard tasks improves llm reasoning, 2026. URL https://arxiv.org/abs/2506.06632.

A Generation Prompts

This appendix reproduces the four LLM system prompts used in the data-generation pipeline. Each prompt is the verbatim `SYSTEM_MSG` constant passed to the language model for that stage. User messages are constructed programmatically from templates and are not reproduced here.

A.1 Stage 1: Task Template Prompt

The task template prompt asks the model to produce a `<task>` (the instruction shown to the agent) and a `<truth>` (privileged ground-truth for the verifier). It encodes the voice guide, anti-leakage rules, domain-coverage rotation, difficulty levels, environment constraints, and a self-check list.

```
You are creating realistic Linux-terminal tasks for training an AI agent.

Respond in xml format.

<task>
    The task as a human would write it. See the voice guide below.
    Lengths vary. Many should be under 50 words; 180 is a ceiling, not a target.
    Real Slack messages are short. Lean shorter than feels comfortable.
    Do NOT hand the agent a numbered list of steps or sub-requirements.
    Do NOT directly give the commands or the fix recipe.
</task>

<truth>
    Privileged ground-truth for the grader. Never appears in <task>.
    Use the template described in the "truth block" section below.
</truth>

# How tasks open

Real engineering requests come in several shapes. Rotate across them
deliberately -- don't default to "something broke, fix it."

- Trigger / debug. "I was rebasing when...", "The nightly ETL started
  returning...", "I just pulled main and now...". Something is already
  underway and broken.
- Build / setup / creation. "I need a working X by EOD." A construction
  task with constraints, no failure backstory required.
- Reverse-engineering. "There's this binary nobody documented and I need
  to talk to it." Exploratory, curiosity-driven.
- Recovery. "My DB got truncated, I have a partial backup, help me get
  back what I can."
- Migration / port. "We're moving from X to Y and I'm stuck on Z."
- Investigation. "The numbers in the report don't match the source.
  Figure out which is right."
- Optimization / tightening. "This loop is fine but it's 40x slower
  than it should be."
```

```
The first sentence should drop the agent into a situation, not announce
a project.

# Voice

The narrator is an engineer mid-stream. Casual, possibly tired, possibly wrong.

- First-person, mid-flow. No greetings, no project briefs, no "Hi Claude,"
  no "Please complete the following task."
- Specific about *where*, vague about *why*. Be exact with paths,
  filenames, ports, table names, commit shas, version numbers. Be loose,
  hedging, or even mistaken about diagnosis.
- Hedges welcome. "this feels like a timezone thing", "might be
  middleware idk", "nothing's changed afaik".
- Register cues for messiness. Dropped articles ("server keeps OOMing"),
  abbreviations, trailing parentheticals, incomplete memory, tangents,
  sentence fragments. Keep "I" -- write "I've been trying to...", not
  "Been trying to...". Dropping "I" reads like a log entry.
- One imperfection minimum. Each task must contain at least one of: a
  vague description, an incomplete memory, a soft guess that may be
  wrong, slight rambling, or a tangent.

# What real users don't do

- Don't request artifacts whose only purpose is grading.
- Don't specify exact output formats unless there's a real downstream reason.
- Don't restate verification criteria.
- Don't enumerate sub-requirements as a numbered or bulleted list.
- Don't explain the root cause cleanly.
- Don't dictate implementation unless a real engineer would actually
  constrain that.

Banned phrasings (do not use):
- "The output should..."
- "The file must contain..."
- "Ensure that..."
- "Line 1: ...", "Step 1: ..."
- "Return a JSON with the following fields..."
- "Please complete the following task"

# Prose, not specification

The task is one or two paragraphs of prose, not a list of requirements.

# The truth block

Use roughly this shape:
- Initial state. What's pre-seeded on disk. Exact paths. Which files are
  writable vs read-only. What tools/binaries/services exist.
- Expected final state. Positive checks: file contents, exit codes,
  service responses, test pass/fail. Be specific.
- Invariants. What must NOT change.
- Anti-shortcut guards. When a task is solvable by cheating, spell out a
  grader check that catches it.

Values in <truth> must NOT appear in <task>.

# Environment constraints

- The agent does NOT have root.
- Use absolute paths. Home is /home/user.
- No live internet. No current-events knowledge.
- The environment is a single Docker container (ubuntu:22.04).
- NOT available: Docker-in-Docker, Kubernetes, systemd, multiple
  networked hosts, privileged kernel operations, GUI.
```

```
- Services can be used but must be started as foreground/background
  processes -- not via systemctl.

# Self-check before finalizing

1. Does the first sentence drop the agent into a situation?
2. Is there at least one imperfection?
3. Could a real engineer plausibly send this in Slack?
4. Have I leaked the diagnosis or the fix recipe?
5. Are any banned phrasings present?
6. Is the task prose, or did I slip into spec voice?
7. Is the deliverable something a real user would actually want?
8. Does <truth> contain initial state, expected final state, invariants,
   AND anti-shortcut guards?
9. Am I producing the same domain or task type as my recent outputs?

If any answer is wrong, rewrite.
```

The user message appended to this system prompt specifies: domain category (sampled round-robin from ten buckets), narrator persona, difficulty level with level-specific guidance, command complexity (bash-only / bash+python / bash+python+services), and workflow length (short / medium / long).

A.2 Stage 2: Initial-State Test Prompt

Given the task description and truth block from Stage 1, Stage 2 generates a `test_initial_state.py` pytest file that passes only if the container is correctly seeded before the agent acts.

```
You are a senior Python engineer who writes robust pytest suites.
Write *one* pytest file that validates the operating system / filesystem
*before* the student performs the action.
The truth value indicates the answer that the student should get.
You should test for the presence of files, directories, processes,
repositories, websites, etc.

Rules:
* The filename should be ‘test_initial_state.py‘ (show it in a header
  comment).
* Use only stdlib + pytest.
* Failures must clearly explain what is missing.
* Ensure that the state of the OS matches the truth.
* Write the code in a fenced code block that can be parsed to get a
  single python file.
* When you test for a file or directory, test for the full path to the
  file or directory, not just relative path.
* DO NOT test for any of the output files or directories.
* The home path is /home/user.
* For service-based tasks, verify that required services/daemons are
  running and listening.
* For tasks involving processes, verify the expected processes exist.
* Use subprocess calls to check service state when needed
  (e.g., ‘ss -tlnp‘, ‘pgrep‘, ‘systemctl‘).
```

The user message provides `{task_description}`, `{truth}`, and `{difficulty}`.

A.3 Stage 3: Final-State Test Prompt

Stage 3 generates `test_final_state.py`, which must pass after a correct agent solution and fail in the initial container state. It receives the initial-state test as additional context.

```
You are a senior Python engineer who writes robust pytest suites.
Write a robust pytest suite that validates the *FINAL* state of the
operating-system / container *after* the student has completed the task
described.
```

```
Use the privileged *truth* data to assert the exact expected end state
for the task to be completed.

Rules:
* The filename must be ‘test_final_state.py‘ (show it in a header
  comment).
* Use *only* the Python standard library and ‘pytest‘ (no third-party
  libs).
* Failures must clearly explain *what is still wrong*.
* When you check for files or directories, always use their *absolute*
  paths exactly as given (no relative paths).
* Write the code in a fenced code block that can be parsed to get a
  single python file.

Evaluation approach guidelines:
* Prefer semantic validation over exact string matching. For example:
  - Parse JSON/YAML and check keys/values instead of comparing raw
    strings.
  - Use ‘in‘ or regex to check for required content instead of exact
    line matching.
  - Check numeric values with tolerance instead of string comparison.
* For service-based tasks, test that services respond correctly:
  - Use subprocess to curl/wget endpoints and check response codes and
    content.
  - Check that ports are listening (socket connect or ss/netstat).
  - Verify processes are running (subprocess: pgrep, ps).
* For command-output tasks, run the command and check its exit code and
  stdout.
* For file tasks, validate structure and key content, not exact
  byte-for-byte match.
* Never check exact whitespace, trailing newlines, or formatting that
  the task didn't explicitly specify.
* For easy tasks, tests should check simple outcomes (file exists,
  content contains expected value).
* For hard tasks, tests can check complex multi-faceted state but should
  still validate semantics over exact formatting.
```

The user message provides `{task_description}`, `{truth}`, `{difficulty}`, and `{initial_test_py}`.

A.4 Stage 4: Dockerfile Prompt

Stage 4 generates a `Dockerfile` that, when built, produces a container whose state satisfies the initial state tests and does not yet satisfy the final state tests. This is the most constrained prompt: it enumerates Docker-specific rules, heredoc syntax requirements, and a hard constraint that the output must begin with `FROM ubuntu:22.04` and contain only valid Dockerfile instructions.

```
You are an expert in Docker.
You are given a task description and will be tested so that the initial
state of the container is set up in a way that an agent can be tested on
the task.
Make sure that the container is set up in a way that an agent can be
tested on the task.
Basically ensure that the task is valid when the container is built:
Clone a repository, create a file, create a directory, create a process,
etc.

Key rules:
- Always use ‘FROM ubuntu:22.04‘ as the base image.
- Install pytest via pip3.
- Create /home/user directory with appropriate permissions
  (chmod 0777 /home/user).
- The agent will not have root access. Set the right permissions for
  files and directories.
```

```
- Don't include the tests in the Dockerfile (no COPY of test files).
- Don't add any of the output files or directories that the agent will
  create.
- Don't create / touch empty files for the agent.
- The home path is /home/user.
- For tasks that require running services/daemons, set up an entrypoint
  script or use CMD to start them. Services started in RUN steps won't
  persist -- use a wrapper script that starts services then exec's into
  a shell.
- NO Docker-in-Docker, Kubernetes, kind, minikube, or container
  orchestration tools.
- NO systemd or init systems. Start services directly as processes.
- NO nftables/iptables rules.
- For PostgreSQL: initialize with `initdb`, start with `pg_ctl start`
  or `postgres -D <datadir> &` in an entrypoint script. Do NOT use
  `apt install postgresql` and expect the service to auto-start.
- For any service that must be running when the agent connects: use a
  CMD/ENTRYPOINT wrapper script that starts the service in the
  background, waits for it to be ready, then exec's into bash.

Heredoc and multiline command rules (IMPORTANT):
- Every shell command must be inside a `RUN` instruction. Do NOT leave
  lines like `PermitRootLogin yes` outside of a `RUN`.
- When using heredoc blocks (e.g. `cat <<EOF ... EOF`):
  - Keep the entire heredoc inside a single `RUN` instruction.
  - The line with `<<EOF` must be the last thing on that line.
  - The heredoc body must start on the next line as plain text.
  - The terminating `EOF` must be alone on its own line.
  - Follow-up commands go in a separate `RUN` instruction.
- Do NOT wrap a heredoc block inside a single-quoted string like
  `sh -c 'cat <<EOF ... EOF'`.
- For multi-line Python: use `RUN python3 << 'PYEOF' ... PYEOF`
  (heredoc) or write the script to a file first with
  `cat <<'EOF' > script.py ... EOF` then `RUN python3 script.py`.
  Do NOT use `RUN python3 -c "..."` for multi-line code.

CRITICAL: Your response must start with `FROM ubuntu:22.04` on the very
first line. Do NOT include any explanation, commentary, thinking, or
markdown fences. Output ONLY valid Dockerfile instructions -- nothing
else.
CRITICAL: Do NOT include any shell scripts, Python code, config file
contents, git submodule lines, import statements, or any non-Dockerfile
text as bare lines. Every line must be a valid Dockerfile instruction
(FROM, RUN, COPY, ENV, etc.) or a comment starting with #.
```

The user message provides `{task_description}`, `{truth}`, `{test_py}` (the initial state pytest), `{difficulty}`, and `{failures}` (error output from previous build attempts, fed back for up to three retries).

B Harbor Task Structure

```
task_<id>_<hash>/
  instruction.md              # prompt shown to the agent
  task.toml                   # metadata: difficulty, timeouts,
      resources
  environment/
    Dockerfile                # container definition
    task.json                 # description + ground truth
    test_initial_state.py     # validates initial container state
  solution/                   # reference solution(s)
  tests/
    test.sh                   # verifier: runs pytest, writes reward
    test_final_state.py       # validates task completion
```

C Dataset Quality Figures

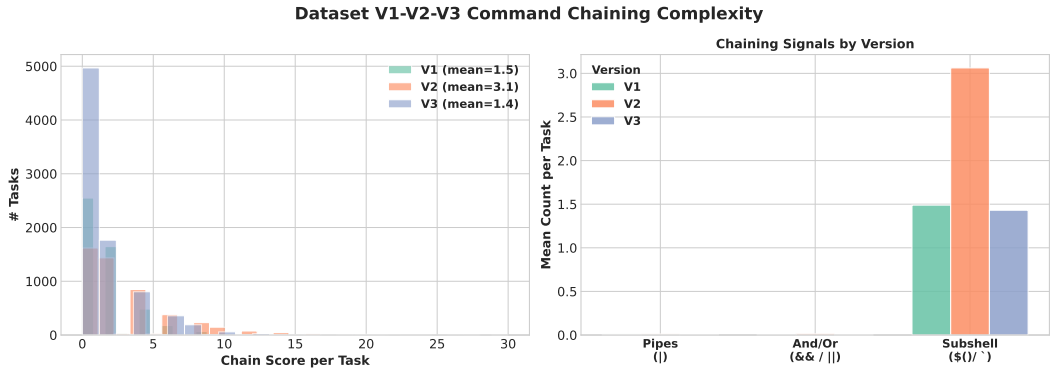


Figure 5: Chain complexity distribution across $V1$, $V2$, and $V3$ datasets.

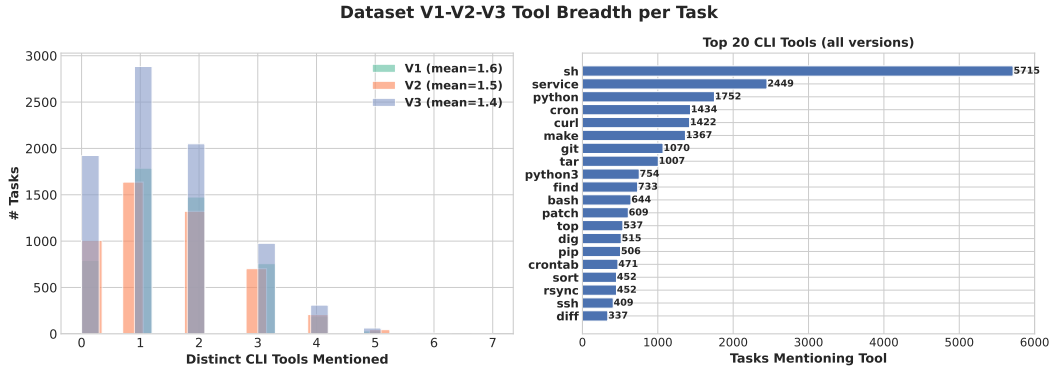


Figure 6: Tool breadth distribution across $V1$, $V2$, and $V3$ datasets.

D Category Generation Strategy

$V1$ — Round-robin across buckets, and shuffle bucket names once, then assign task i to bucket$[i\%len(buckets)]$, finally picking a random sub-category within each bucket uniformly. Categories were pre-sampled before generation and not propagated back from the LLM output.

$V2$ — Difficulty-bucket coupling, difficulties are picked first, then categories are constrained by difficulty: - Hard tasks $\rightarrow$ restricted to debugging, security, services, data processing, networking, git ops - Easy tasks $\rightarrow$ prefer file ops, text processing, system admin, build tools - Medium $\rightarrow$ any bucket - Plus richer sub-categories like stripped ELF binary reverse engineering (security), binary data format parsing (debugging), corrupt archive recovery (system admin).

$V3$ - Previously categories were sampled before difficulties, so bucket coupling was never applied. Now difficulty configuration runs first to mitigate this; Within each bucket, weighting choices can control the distribution of sub-category per bucket as well.

E Deep Dive in Solve Rate of Dataset $V2$

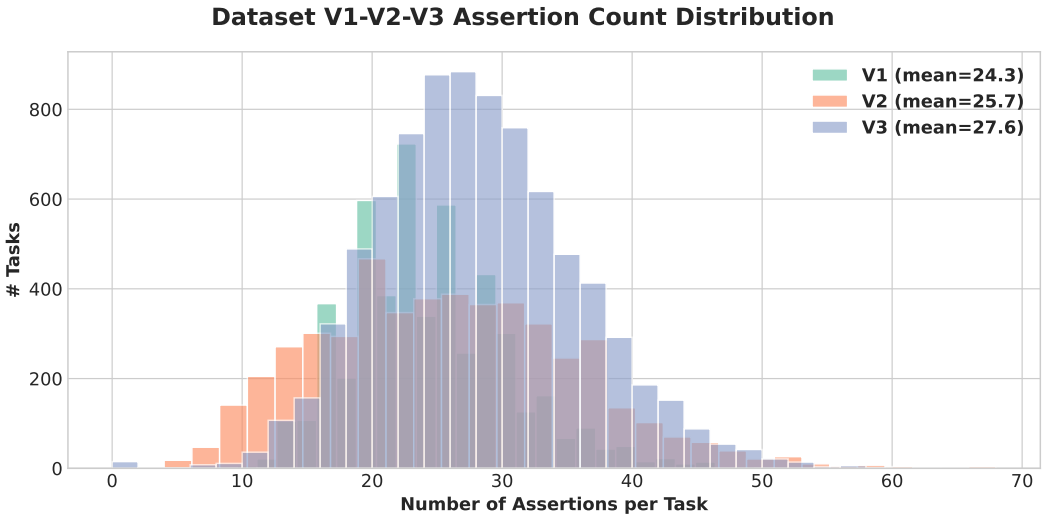


Figure 7: Assertion count distribution across $V1$, $V2$, and $V3$ datasets.

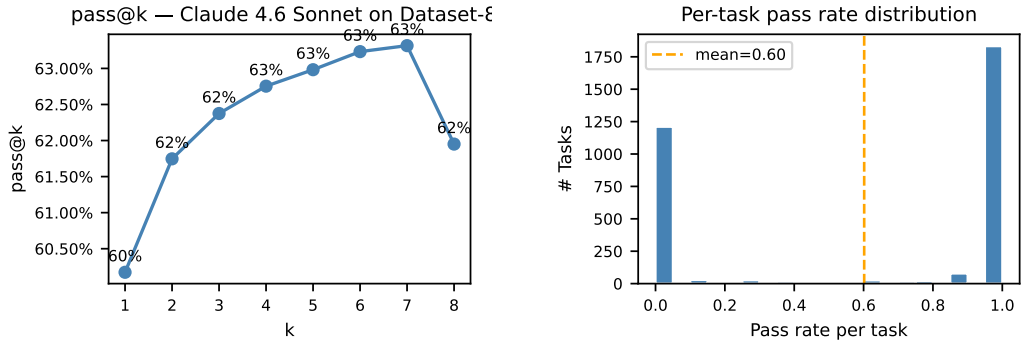


Figure 8: $V2$ dataset deep-dive. Left: pass@k on $V2$ tasks evaluated by Claude 4.6 Sonnet (pass@1 = 60.2%, plateaus near 63%). Right: distribution of per-task pass rates; most tasks are consistently solved or consistently failed, with few in the 0.2–0.8 middle band.