

LEARNING GENERALIZABLE BEHAVIORS FOR TERMINAL AGENTS

Yihang Yao^{1,2,†}, Bo Pang^{1,†}, Xuan Phi Nguyen¹, Ding Zhao², Shafiq Joty¹, Semih Yavuz¹

¹Salesforce AI Research, ²Carnegie Mellon University; † Core contributors.

Correspondence: yihangya@andrew.cmu.edu

ABSTRACT

Terminal agents have emerged as a compelling application of large language models (LLMs), with the potential to integrate deeply into users’ daily workflows. Reinforcement learning (RL) has become a key technique for improving their capabilities, making the scalability of training environments a central challenge. Since public real-user interaction data are scarce, synthetic environments provide a practical alternative. However, they often exhibit substantial domain gaps and limited fidelity, resulting in poor generalization. Existing efforts primarily focus on scaling the quantity and diversity of synthetic environments, while the reward signal quality is under-explored and the mechanisms that govern generalization remain poorly understood.

In this work, we study how RL improves terminal agents and propose the **agentic compositional generalization** hypothesis: rather than teaching the agent new domain-specific skills from scratch, RL primarily shapes high-level decision-making behaviors that compose and route the low-level skills already acquired during pre-training and supervised fine-tuning (SFT). This account is consistent with the evaluation performance we observe. This perspective suggests that verifier quality, which determines which behaviors are reinforced by providing optimization signals in RL, is more important than simply increasing the number or diversity of training environments. Motivated by this insight, we propose RIVER, a simple training recipe that improves the reward signal by filtering low-quality environments and augmenting the outcome reward with process-level behavior regularization. Using this recipe, our RL-trained agent achieves state-of-the-art performance among evaluated open-source RL-trained 8B models across four terminal-agent benchmarks. Moreover, our approach generalizes across model families, scales, agent harnesses, and RL objectives. Using **fewer than 30%** of the TMax training environments, RIVER improves RL gains by **106% and 30%** on average for models ranging from 2B to 27B on Terminal-Bench-Lite and Terminal-Bench-v2.1, respectively.

1 INTRODUCTION

As large language models (LLMs) continue to advance toward real-world deployment, *terminal agents* have emerged as one of their most important applications, enabling software engineering, system administration, and other interactive coding tasks (Anthropic, 2025; OpenAI, 2025; Pi et al., 2026; xAI, 2025; Ivison et al., 2026; Wu et al., 2026b). While the capabilities of frontier terminal agents have improved rapidly, the principles underlying the design of their training environments remain largely undisclosed (Pi et al., 2026; Ivison et al., 2026; Zhu et al., 2026). In particular, it remains unclear how post-training for terminal agents generalizes across domains, and what factors matter beyond simply increasing coverage over task categories and targeted skills.

Reinforcement learning (RL) has become the dominant paradigm for post-training reasoning models and agents (OpenAI, 2024; Guo et al., 2025). Unlike single-turn reasoning tasks such as mathematical reasoning (Shao et al., 2024), terminal agents require interacting with executable environments through long-horizon, multi-turn decision making (Gandhi et al., 2026). Consequently, the training environment plays a central role in agentic RL. Since large-scale real-user interaction data are

generally unavailable, recent work has increasingly relied on *synthetic environments* as a scalable alternative (Pi et al., 2026; Li et al., 2026).

Existing work on synthetic terminal environments has primarily focused on *scaling coverage*: expanding the diversity of tasks, domains, and required skills (Pi et al., 2026; Fan et al., 2026; Zhao et al., 2026). Early efforts generated demonstration trajectories for supervised fine-tuning (SFT), whereas more recent work constructs executable environments for RL. For example, Endless-Terminals (Gandhi et al., 2026) demonstrated that RL over synthetic environments substantially improves terminal agents, and TMax (Iverson et al., 2026) further increased environment diversity and domain coverage. However, *coverage alone does not explain generalization*. This raises two fundamental questions: (i) What does RL actually learn during terminal-agent training? (ii) What makes one training environment more valuable than another, aside from diversity, for improving generalization beyond the synthetic training-task distribution?

In this work, we study these questions through the lens of the *compositional generalization* hypothesis (Yuan et al., 2026; Kong et al., 2026), extending it from single-turn reasoning models to multi-turn long-horizon agentic models. We term the resulting account *agentic compositional generalization* (see Fig. 1 for an overview), and test it in realistic terminal environments rather than purely synthetic ones. Our findings indicate complementary roles of SFT and RL in generalization: SFT (and general model prior imposed by pre-training and mid-training) equips the agent with low-level atomic skills, while RL primarily teaches *high-level decision-making behaviors* that compose, route, and coordinate these skills across domains. Under this hypothesis, generalization requires both coverage on atomic skills and reusable high-level behaviors.

This perspective suggests a different axis for the design of the training recipe. For skills, since RL is not efficient in learning atomic skills, then SFT should equip the models with diverse atomic skills that are absent from model priors. For behaviors, since RL primarily learns high-level behaviors, then *verifiers* in synthetic environments, which provide the signal that shapes those behaviors, become critical. Guided by this insight, we propose RIVER, a simple yet effective training recipe that filters out low-quality environments and improves verifier quality through behavior-level regularization. Despite using fewer than 30% of the TMax training environments, RIVER consistently improves RL efficiency and generalization across model families, model scales, agent harnesses, and RL algorithms.

Our contributions are summarized as follows:

- (1) **Mechanistic interpretation.** We propose *agentic compositional generalization*, which posits that RL improves generalization by shaping reusable high-level behaviors that compose pre-RL domain atomic skills, and we present empirical evidence consistent with its predictions.
- (2) **A simple and effective training recipe.** We introduce RIVER, which improves the quality of reward signal and the training efficiency. Using fewer than 30% of the TMax environments (Iverson et al., 2026), RIVER achieves on average 106% and 30% larger RL gains on Terminal-Bench-Lite and Terminal-Bench-v2.1, respectively, across models ranging from 2B to 27B.
- (3) **Comprehensive empirical analysis.** We provide extensive experiments that test the predictions of our hypothesis and offer new insights into the learning dynamics of RL for terminal agents.

2 RELATED WORK

Terminal Agents such as Claude Code (Anthropic, 2025) and Codex (OpenAI, 2025) have rapidly become central to software-engineering workflows, demonstrating strong performance on long-horizon coding benchmarks such as Terminal-Bench (Merrill et al., 2026), SWE-Bench (Jimenez et al., 2024), and their variants (Badertdinov et al., 2026; Qiu et al., 2025). Motivated by these advances, the open-source and research communities have increasingly explored recipes for training terminal and coding agents (Wang et al., 2025a; Zeng et al., 2025; Guha et al., 2026; Raoof et al., 2026; Peng et al., 2025; Cheng et al., 2026; Bercovich et al., 2026). Among the various post-training approaches, reinforcement learning (RL) has emerged as one of the most effective techniques for improving long-horizon agentic behaviors (Yang et al., 2026b; Wang et al., 2026), and has recently been adopted for terminal-agent training. A key challenge, however, is scaling high-quality training data and executable environments, which has become one of the primary bottlenecks for further progress (Pi et al., 2026; Iverson et al., 2026; Ruan et al., 2026).

Synthetic Data and Environment Scaling. Scaling synthetic data has become a central paradigm for training LLM agents, driven by the scarcity and cost of collecting real-world interaction data (Chen et al., 2026b; Zhang et al., 2025; Cen et al., 2026a). Recent efforts have evolved from generating static demonstrations for supervised fine-tuning (SFT) (Liu et al., 2024; Xu et al., 2025; Prabhakar et al., 2026) to constructing interactive training environments (Yang et al., 2026a; Zhou et al., 2026; Lin et al., 2026; Li et al., 2026; He et al., 2026), which provide the executable feedback required for large-scale reinforcement learning. For terminal agents in particular, most existing work has focused on expanding environment coverage across skills, domains, and task categories (Iverson et al., 2026; Pi et al., 2026). While recent studies have shown that verifier quality, such as stronger test cases, can improve coding performance (He et al., 2025), the broader role of verifier quality in shaping learned behaviors and enabling out-of-distribution generalization for RL-trained terminal agents remains largely unexplored.

Learning Dynamics and Generalization Mechanisms. A growing body of work seeks to understand why LLMs and agents generalize beyond their training distributions and how these insights can guide more effective training recipes (Chu et al., 2025; Huang et al., 2025; Ren & Sutherland, 2025; Chen et al., 2026a; Ye et al., 2025; Yu et al., 2025; Rajaraman et al., 2026; Yue et al.). Recent works on *compositional generalization* (Kong et al., 2026; Yuan et al., 2026) explain the complementary roles of SFT and RL in single-turn reasoning models: SFT provides the low-level reasoning primitives, while RL learns to compose and route these primitives to enable generalization. Other studies investigate reasoning behaviors and decision-making patterns, showing how understanding these dynamics can improve RL training (Gandhi et al., 2025; Cen et al., 2026b; Yeo et al., 2025; Yao et al., 2026b), with recent extensions to multi-turn agentic settings (Yao et al., 2026a). A complementary line of work examines how spurious or misaligned reward signals affect RL learning and generalization (Shao et al., 2025). In contrast, we study the learning dynamics of *terminal agents* operating in realistic interactive environments. We extend the compositional generalization hypothesis to the multi-turn agentic setting, quantitatively dissociate single-turn skills from multi-turn behaviors at the feature level (see Sec. 5), and show how the identified behaviors can be targeted directly through advantage shaping and used to guide the design of RL training environments and recipes.

3 MECHANISM INTERPRETATION AND METHOD

Synthetic environments typically introduce a substantial training-evaluation distribution gap: evaluation tasks are often significantly more complex than those posed by the synthetic training environments, and no practical training corpus can exhaustively cover the diversity of real-world task categories. In this section, we first present an **Agentic Compositional Generalization** hypothesis, our account of how RL-trained terminal agents achieve generalization beyond the synthetic environment distribution and close the training-evaluation gap. Guided by this perspective, we then introduce **Reward-Integrity Verified Environments for RL (RIVER)**, a simple training pipeline that improves the quality of synthetic environments through verifier-aware filtering and behavior-level regularization, and we describe the resulting RL training recipe.

3.1 AGENTIC COMPOSITIONAL GENERALIZATION

We model terminal-agent capability as a hierarchical decision-making process with two complementary levels of competence: (i) high-level *behaviors* that determine *what* the agent should do next (e.g., decomposing a task, planning future actions, or verifying current results), and (ii) low-level *skills* that determine *how* each action is executed (e.g., applying domain knowledge for reasoning or invoking appropriate tools). We formalize this process as a Markov Decision Process (MDP) (Puterman, 2014). At interaction step t , the state s_t consists of the complete interaction history up to that point, including the user instruction, executed commands, and terminal feedback. An action a_t corresponds to the agent’s single-turn output, including its reasoning tokens together with any tool invocation or terminal command. The policy $\pi(a_t | s_t)$ defines the agent’s action distribution. Within this formulation, we distinguish two complementary roles (see Fig. 1):

- **Low-level Single-turn Skills.** A skill is the ability to execute a single turn correctly given the current state s_t : applying domain knowledge, reasoning accurately, and issuing the appropriate command or tool call (e.g., composing a valid grep pipeline, running pytest, installing a package, see Appendix B.1). Skills are properties of individual actions, whether a step is locally correct,




Figure 1: The **Agentic Compositional Generalization** hypothesis. We consider terminal-agent capability at two levels that are acquired at different training stages. **(Left)** Low-level skills, i.e., domain knowledge and familiarity with specific tools and commands that determine how each single-turn action is executed, are acquired during pretraining and SFT. **(Middle)** High-level behaviors, i.e., cross-turn patterns such as exploration and verification that determine what to do next, are shaped during RL by feedback from environment verifiers. Under our hypothesis, generalizability arises when RL-shaped behaviors compose and route the skills that an unseen task demands. **(Right)** Evidence consistent with this account: behavior features of a trajectory predict task success (AUC-ROC 0.74), while the presence of individual skills is near chance (about 0.55), and the link between behaviors and success strengthens from SFT to RL (0.64 to 0.74); test details in Sec. 5.2.

irrespective of the surrounding strategy. SFT primarily improves skills by concentrating the policy on valid, domain-appropriate actions (see Sec. 5.1).

- **High-level Behaviors.** A behavior is a pattern, typically spanning multiple turns, in how actions are chosen across turns, e.g., inspecting the environment before acting, verifying results before declaring completion, or abandoning a failing approach instead of repeating it (see App. D). Behaviors are properties of action sequences rather than of any single action: they manifest in multi-turn fragments and aggregate into trajectory-level patterns, so a trajectory can be flawed even when every step is locally correct. Our hypothesis is that behaviors compose, route, and coordinate low-level skills to accomplish long-horizon tasks; RL primarily shapes them (Sec. 5.2).

We propose the **Agentic Compositional Generalization** hypothesis, extending recent theories of compositional generalization for single-turn reasoning models (Kong et al., 2026; Yuan et al., 2026) to interactive terminal agents. The main idea is that generalizability arises from the complementary roles of SFT and RL: pre-training and SFT equip the model with domain-specific low-level skills, while RL shapes high-level behaviors from environment feedback that compose, route, and coordinate these skills to solve long-horizon tasks. During deployment, the agent generalizes by combining shaped behaviors with the domain skills required by the target task. In this section, we present the intuition and implications of this hypothesis, while deferring the empirical validation to Sec. 5, where we provide comprehensive experimental evidence supporting these claims.

Roles of SFT. Whether RL should be warm-started from an SFT model has been actively debated since DeepSeek-R1 (Guo et al., 2025). Our findings suggest that the answer depends on the model’s existing capabilities. For weaker base models, large-scale SFT plays a critical role by familiarizing the model with the target domains and interaction harness, enabling it to produce trajectories correct enough to yield meaningful learning signals during RL. More importantly, SFT equips the model with the low-level skills required for downstream tasks, providing the foundation upon which RL learns high-level behaviors that, under our hypothesis, enable generalizability (Sec. 5.1 for experiments). In contrast, for stronger foundation models that already possess broad low-level skills and can adapt to new harnesses via their robust instruction following capacity, we can perform zero-RL without SFT.

Roles of RL. Our findings suggest that RL primarily shapes *high-level behaviors* rather than acquiring new low-level skills. The low-level skills learned before RL are largely preserved throughout




Figure 2: **(Left:)** Pipeline for environment filtering and verifier enhancement. **(Right:)** Environment quality evaluation results of TMax-15K (Iverson et al., 2026).

RL, while RL learns how to compose, route, and coordinate these skills to solve long-horizon tasks (Sec. 5.1 for experiments). We hypothesize that this behavioral adaptation is the primary driver of generalizability. Consistent with this, an AUC-ROC analysis (Fawcett, 2006) on evaluation trajectories shows that high-level behaviors are strongly correlated with task success, whereas the presence of individual low-level skills exhibits little predictive power (Fig. 1 and Sec. 5.2). These findings suggest that the effectiveness of agentic RL depends less on acquiring additional domain-specific skills and more on learning reusable, generalizable behaviors that transfer across tasks and domains.

3.2 RIVER: REWARD-INTEGRITY VERIFIED ENVIRONMENTS FOR RL

A synthetic terminal environment typically consists of an interactive sandbox, task-specific file systems, and a verifier that provides reward signals to the RL agent (Fig. 2). The discussion above suggests that effective terminal-agent RL training requires two complementary ingredients: **(1)** the starting checkpoint should have broad coverage of low-level atomic skills, and **(2)** the RL training environments should provide high-quality feedback that fosters the learning of reusable, high-level behaviors. To satisfy the first requirement, we construct a filtered version of the OpenThoughts (Raoof et al., 2026) SFT corpus that provides broad coverage of atomic skills for weaker models lacking strong prior capabilities (see App. A.1 for details and Sec. 5.3 for ablations). To address the second, we propose RIVER, motivated by our observation that verifier quality is a key factor beyond environment quantity and diversity. Inaccurate verifiers can produce false-positive and false-negative reward signals on rollout trajectories, thereby inducing reward hacking (see App. C.4 for examples). Accordingly, RIVER improves the integrity of the reward signals provided by the source environment collection TMax (Iverson et al., 2026) through two key components: *environment filtering* and *verifier enhancement*.

Environment filtering. Although synthetic environments have become a dominant paradigm for terminal-agent RL, their quality remains highly variable. Existing open-source collections often contain weak or mismatched verifiers, resulting in the potential for noisy supervision that encourages shortcut solutions, reward hacking, or overfitting. Since verifier quality directly shapes the high-level behaviors learned during RL, ensuring environment quality is a prerequisite for effective training. We develop a lightweight environment filtering system that evaluates synthetic environments using a set of rubrics without executing the underlying sandbox. Our framework identifies eight categories of quality issues, including weak verifiers, mismatched task specifications, and other inconsistencies (see App. C for details and examples). Applying this evaluation to the TMax (Iverson et al., 2026) environment collection reveals that more than 60% of the environments exhibit at least one quality issue, leaving fewer than 40% suitable for RL training¹ (see Fig. 2).

We further filter the remaining environments using GPT-5.4 (OpenAI, 2026) with harness of Terminus-2 (Laude Institute, 2026) as an oracle verifier. Specifically, we perform a $\text{pass}@k = 2$

¹Our filtering process may produce both false positives and false negatives. Also, there might be multiple issues in one environment. Details of environment filtering can be found in App. C. Examples of reward hacking with unfiltered environments can be found in App. C.4.

Algorithm 1: Full Training recipe.

- (1) SFT dataset construction with broad skill coverage and SFT training # Details in App. A.1, Discussions in Sec. 5.3

RIVER core components:

- (2) RL Environment Filtering
 - (2.1) Deploy Rubrics-based Filtering;
 - (2.2) Oracle Agent Pass@ k Filtering;
 - (3) Verifier Enhancement (1)
 - (4) RL on final envs
-

evaluation and remove environments with $\text{pass}@k = 0$, which are either excessively difficult or incompatible with our training infrastructure. The resulting curated collection, **RIVER-TMax-3.5K**, serves as the training environment set used throughout this work.

Verifier enhancement. Environment filtering improves the quality of training tasks, but each environment still provides only sparse outcome supervision. We introduce an optional *verifier enhancement* step that supplies denser behavior-level feedback during RL. The key idea is to complement the trajectory-level reward with lightweight behavior verifiers that identify desirable or undesirable interaction patterns and provide turn-level shaping signals. These signals encourage generalizable behaviors while remaining compatible with standard RL objectives.

We instantiate this idea using GRPO (Shao et al., 2024). Unless otherwise stated, we use a binary outcome reward $r_i \in \{0, 1\}$: for each prompt, we sample a group of G rollouts, assign $r_i = 1$ only if the i -th rollout passes all verifiers and test cases, normalize rewards within the group, and broadcast the resulting scalar advantage uniformly to every response token. While effective, this trajectory-level signal is blind to behaviors that occur during interaction. To incorporate behavior supervision, we augment the standard GRPO advantage $A_{i,t}$ with a turn-wise shaping term $s_{i,t}$:

$$A_{i,t} = \frac{r_i - \bar{r}}{\sigma_r + \epsilon}, \quad \bar{r} = \frac{1}{G} \sum_{j=1}^G r_j, \quad \sigma_r = \text{std}(\{r_j\}_{j=1}^G), \quad \tilde{A}_{i,t} = A_{i,t} + s_{i,t} \quad (1)$$

where $\bar{r}$ and σ_r denote the within-group reward mean and standard deviation, and ϵ is a small constant. The shaping term $s_{i,t}$, applied after reward normalization, may be positive to reinforce desirable behaviors or negative to discourage undesirable ones. One particularly effective $s_{i,t}$ uses a simple rule-based verifier that penalizes repetitive interaction loops with little information gain. Specifically,

$$s_{i,t} = -\delta \mathbf{1} \left[t \in \bigcup_{k \in \mathcal{L}_i} \mathcal{T}(k) \right],$$

where δ is a hyperparameter², $\mathcal{T}(k)$ denotes the response-token indices of turn k , and $\mathcal{L}_i$ is the set of turns identified as repetitive. We penalize repetition behaviors: a turn is flagged as repetitive when both its command and observation have Jaccard similarity greater than 0.8 with an earlier turn. We also provide experiments with additional behavior verifiers in Sec. 4.1.

Recipe summary and intuition. Our overall training recipe is summarized in Alg. 1. For weaker base models, we include an SFT stage to establish broad coverage of low-level skills before RL. Additional implementation details, including the agent harness, datasets, and training configurations, are provided in Sec. 4 and App. A. The central intuition is simple: once a model possesses sufficient low-level skills, RL should focus on learning reusable high-level behaviors. Providing accurate verifier signals through high-quality environments enables these behaviors to generalize across domains, leading to improved out-of-distribution performance.

4 EXPERIMENT RESULTS

In this section, we present the evaluation results of our RIVER pipeline with two base model families: Qwen3 and Qwen3.5³ (Yang et al., 2025; Team, 2026a;b). We use four terminal benchmarks for evaluation: **(OpenThoughts-)Terminal-Bench-Lite** (Raoof et al., 2026), **Terminal-Bench-v2.1** (Merrill et al., 2026), **Terminal-World-Verified** (Chu et al., 2026), and **Terminal-Bench-Pro** (Wang et al., 2025b). For brevity, we refer to them as `otbl`, `tbv2.1`, `twv`, and `tbpro`, respectively, when there is no ambiguity. We also hold out a subset of around 300 tasks from the TMax collection, excluding them from training and using them for controlled experiments. We adopt the default evaluation criteria of each benchmark: `otbl` uses a continuous score in $[0, 1]$, while the others use a binary score, assigning 1 only if the solution passes all verifiers (unit tests).

²We used 0.05 in our experiments as determined by a lightweight search.

³Qwen3 uses full attention with GQA, while Qwen3.5 uses a hybrid architecture with Gated DeltaNet linear attention and full attention blocks. For the 27B experiment, we use Qwen3.6, following Ivison et al. (2026), which has the same architecture as Qwen3.5.

Model	Harness	Training	Otbl	Tbv2.1	Tbpro	Twv	Avg.
River-8B (ours)	EndlessAgent	SFT+RL	21.8 ± 2.9	9.7 ± 0.6	23.0 ± 3.1	23.0 ± 2.5	19.4 ± 1.2
OpenThinker-8B-RL	Terminus-2	SFT+RL	18.4 ± 1.5	8.6 ± 0.6	21.8 ± 1.6	22.3 ± 1.5	17.8 ± 0.7
OT-SFT-Endless-8B	EndlessAgent	SFT+RL	14.1 ± 3.6	6.4 ± 1.7	16.8 ± 0.8	21.7 ± 1.9	14.7 ± 1.1
TMax-RL-8B	Vanillux2	RL	17.0 ± 3.7	6.0 ± 0.6	16.3 ± 1.0	19.0 ± 2.6	14.6 ± 1.2
OpenThinker-Agent-v1	Terminus-2	SFT+RL	16.0 ± 0.6	5.2 ± 0.6	16.2 ± 1.0	21.7 ± 2.9	14.8 ± 0.8
TMax-SFT-8B	Vanillux2	SFT	12.1 ± 3.4	4.5 ± 1.9	13.7 ± 1.9	14.3 ± 4.1	11.2 ± 1.5
River-SFT-8B	EndlessAgent	SFT	11.3 ± 2.3	6.4 ± 1.3	13.7 ± 2.6	18.5 ± 2.3	12.5 ± 1.1
Qwen3-8B	EndlessAgent	None	3.8 ± 1.3	2.2 ± 1.2	6.8 ± 0.6	9.8 ± 1.6	5.7 ± 0.6

Table 1: Evaluation results of 8B agents on four terminal benchmarks. Numbers are benchmark scores ($\times 10^2$). **Avg.** is the unweighted mean of the four benchmarks. Scores are mean $\pm$ std over 3 seeds. Baseline scores are reported by our evaluation pipeline under fair comparison.

4.1 QWEN3-8B MODEL FAMILY

Training. Using Qwen3-8B as the base model, we apply the full RIVER pipeline. We filter out trajectories from the OpenThoughts-100K SFT collection (Raoof et al., 2026) that cannot be converted from its original Terminus-2 harness format to our Endless-Agent harness, resulting in 78K trajectories. SFT is done with LLaMA-Factory (Zheng et al., 2024b). For RL, we adopt Endless-Agent as the agent harness (Gandhi et al., 2026). The RL training pipeline is also based on Endless-Agent, which in turn builds on SkyRL (Cao et al., 2025), Harbor (Laude Institute, 2026), and Verl (Sheng et al., 2025). We use the **RIVER-TMax-3.5K** environment set as our RL training set.

Baselines and Evaluation. We compare our method with open-source fine-tuned 8B models as well as the base model. The baselines include OpenThinker-8B-RL and OpenThinker-Agent-v1 (Guha et al., 2026; Raoof et al., 2026), both of which perform SFT before RL on their synthetic environments; OT-SFT-Endless-8B, which performs RL on the Endless-Terminals collection starting from the OpenThoughts SFT checkpoint (Gandhi et al., 2026); and TMax-8B and TMax-SFT-8B (Iverson et al., 2026), which are trained using their synthetic RL environments and SFT data, respectively. Each baseline is evaluated using its corresponding harness. We set the interaction turn budget to 64 and keep all other evaluation configurations fair across methods. Further evaluation details are provided in App. A.3.

Main Results. Comparisons are summarized in Tab. 1. We observe that our model achieves the best performance. Among the baselines, Endless-8B uses the same harness and also relies on SFT data from the OpenThoughts collection but with lower coverage on atomic skills. Their RL is trained with their synthetic environments of size 3.3K, on the same scale as our collection. With the harness and training pipeline held fixed, this comparison isolates the gains attributable to higher SFT atomic-skill coverage and improved RL verifier quality. TMax-8B is based on a different harness, vanillux2, a mini-swe-agent-2 (Yang et al., 2024) variant. Their RL training set is the full TMax-15K environments, which requires substantially more resources. Nevertheless, we achieve consistently better performance than their trained agents on all evaluation benchmarks, as the full TMax set contains tasks that foster reward hacking (See App. C.4). For the OpenThinker models, they use the Terminus-2 harness (Laude Institute, 2026), which is likewise more complex and more capable than our EndlessAgent. They also use large-scale SFT and RL during training. Their RL-trained model attains the best performance among the baselines, yet our method still shows a significant margin over it. We attribute the performance advantage to our RIVER, which fosters more generalizable behaviors with cleaner verifiers and additional behavior regularization.

Behavior Ablations and Analysis. We select behavior verifiers for the verifier enhancement step in a data-driven way. We first assess the utility of each candidate behavior by how well it predicts the task success, using the single-behavior-feature AUC-ROC test (see Sec. 5.2 for AUC-ROC details) shown in Tab. 9 in App. D. The two strongest signals are `verify_before_done` and `repetition_rate`. We then run experiments with each of these top behaviors as a turn-level shaping signal, with the sign chosen according to its correlation: a turn-wise reward for verification before completion, and a turn-wise penalty for repetitive loops. The performance and behavior comparisons for both behaviors are visualized in Fig. 3. Compared with outcome-based reward only, we




Figure 3: Ablations on behavior verifiers. **(a):** Averaged performance on the four evaluation benchmarks; **(b, c):** Behavior ratio comparison between RL with and without behavior verifiers. Arrows $\uparrow / \downarrow$: the higher/lower, the better.

find that advantage shaping does reshape the target behaviors: it reduces the repetition ratio and increases the fraction of verification in the corresponding runs. However, only the repetition-penalty run yields a noticeable overall performance improvement, whereas rewarding verification brings no additional gains. We interpret this as indicating that verification is correlated with task performance but does not directly cause task success. Based on this result, only the repetition penalty enters the final RIVER recipe as its verifier enhancement component 3.2.

We hope this discussion draws more attention to systematic study of multi-turn behaviors that are causally related to task performance, and to the design of advantage shaping in RL. Further ablations, such as the selection of SFT data and RL environments, are discussed in Sec. 5.3 and 5.4.

Section 4.1 Takeaways

- Our RIVER-8B model achieves the best performance on four terminal benchmarks.
- Advantage shaping from behavior verifiers steers the targeted behaviors, and discovering behavior verifiers that improve performance is an interesting and non-trivial direction.

4.2 QWEN3.5 AND QWEN3.6 FAMILIES

Training. For an apples-to-apples comparison, we keep the same base models as the TMax collection: Qwen3.5-{2B, 4B, 9B} and Qwen3.6-27B. We also adopt their training pipeline and most of their configurations, including the harness and their RL algorithm DPPO (Qi et al., 2026), except for a few changes made under resource constraints, such as switching from the sandbox API to locally hosted Docker containers. For evaluation, we use the two benchmarks reported in their paper. These experiments can be viewed as an ablation on environment quality: if we keep only the environments after our filtering process, can we further improve the trained model’s performance?




Figure 4: Comparison results with the TMax model series. **(a)** Performance comparison from 2B to 27B models; **(b)** RL performance improvement over base models from 2B to 27B models. We report the mean value and standard deviation with 3 seeds. TMax results are adopted from their report.

Main Results. The evaluation results are presented in Fig. 4. We observe that RIVER consistently improves base-model performance across model sizes and outperforms the original TMax on both Terminal-Bench-Lite and Terminal-Bench-v2.1. These results demonstrate that our method generalizes across model families, model sizes, and agent harnesses in agentic RL for terminal agents. We also report the performance gains of the RL models over their corresponding base models. Notably, using only $\sim 30\%$ of the TMax environments (our TMax-RIVER-3.5K set), we achieve, on average, 106% and 30% larger RL gains on Terminal-Bench-Lite and Terminal-Bench-v2.1, respectively.

Section 4.2 Takeaways

- With our filtered environment set, **RIVER-TMax-3.5K**, RL using the same base models achieves on average 106% and 30% larger performance gains on `otbl` and `tbv2.1`, respectively, compared with TMax training on the full environment set.
- The comparison between RIVER and TMax shows that RIVER generalizes across model families, model sizes, agent harnesses, and RL algorithms.

5 ANALYSIS

In this section, we first present experiments in Secs. 5.1 and 5.2 that support our **Agentic Compositional Generalization** hypothesis introduced in Sec. 3.1. We then provide further analyses and ablations of SFT in Sec. 5.3 and RL in Sec. 5.4. Unless otherwise stated, we use Qwen3-8B as the base model.

5.1 SKILL PRESERVING AND BEHAVIOR SHAPING IN RL

We would like to understand how RL changes the model’s capabilities over low-level skills and high-level behaviors. We compare evaluation traces before and after RL for our 8B models, which are warm-started with SFT. With a list of $N_s = 60$ pre-defined low-level skills, we use GPT-5.4 to extract the feature of each evaluation trace as a binary vector indicating which skills are exhibited in that trace. The evaluation benchmarks contain $N_t = 589$ tasks in total, so we obtain two binary matrices $M^{\text{pre}}, M^{\text{post}} \in \{0, 1\}^{N_s \times N_t}$ (before and after RL), where $M_{s,t} = 1$ iff skill $s \in \{1, \dots, N_s\}$ is exhibited in the trace for task $t \in \{1, \dots, N_t\}$. Row $M_{s,:} \in \{0, 1\}^{N_t}$ is skill s ’s *task-incidence vector* (on which tasks it appears); column $M_{:,t} \in \{0, 1\}^{N_s}$ is task t ’s *skill vector* (which skills it recruits).

First, we compute the exhibition ratio of each skill by averaging along the task dimension, $r_s = \frac{1}{N_t} \sum_{t=1}^{N_t} M_{s,t}$, once for M^{pre} and once for M^{post} . The results are presented in Fig. 5 (a), where each point corresponds to one skill s at coordinates $(r_s^{\text{pre}}, r_s^{\text{post}})$. We observe that the skill ratios remain strongly correlated: if the model does not exhibit a specific skill before RL, it is unlikely to gain significant exhibition of that skill after RL, indicating that RL is not efficient at teaching novel skills absent from the pre-training and SFT stages.

Second, in Fig. 5 (b and c) we apply Representational Similarity Analysis (RSA) (Kriegeskorte et al., 2008) to the same matrices $M^{\text{pre}}, M^{\text{post}}$ along their two axes. The skill-skill structure is the matrix $S^{\text{skill}} \in \mathbb{R}^{N_s \times N_s}$ with $S_{ij}^{\text{skill}} = \cos(M_{i,:}, M_{j,:})$, the cosine similarity between every pair of skills over their task-incidence vectors, i.e., “which skills co-occur”. Symmetrically, the task-task structure is $S^{\text{task}} \in \mathbb{R}^{N_t \times N_t}$ with $S_{ab}^{\text{task}} = \cos(M_{:,a}, M_{:,b})$, the cosine similarity between every pair of tasks over their skill vectors, i.e., “which tasks are solved with the same skills”. We build each structure once for M^{pre} and once for M^{post} ; each panel is a density scatter of one structure’s upper-triangle entries before RL (x) vs after RL (y), one point per pair. We report the Pearson (Pearson, 1896) and Spearman correlation (Spearman, 1961) between the two triangles, with a Mantel permutation test (Mantel, 1967) for significance (the correct test given the non-independent similarity cells).

Fig. 5 (b, c) support the hypothesis that (1) **The skill-skill structure is strongly preserved during RL**: in Fig. 5 (b), points lie close to $y = x$, with $\rho = 0.83$, indicating that skills that co-occurred before RL tend to continue co-occurring afterward. In contrast, (2) **the task-task structure changes substantially**: in Fig. 5 (c), correlation is $\rho = 0.27$, with a diffuse cloud far from $y = x$, suggesting that the set of skills recruited for a given task is reorganized during RL. Together, these results are consistent with RL primarily reshaping how existing skills are selected and composed across tasks.




Figure 5: Skills and behaviors correlation analysis before and after RL. (a) Correlation between skill ratio; (b) Correlation between skill-skill occurrence; (c) Correlation between task-task structure.

5.2 BEHAVIORS’ CORRELATION WITH TASK PERFORMANCE

In this section, we provide additional details on the experimental design and results of the Area Under the Receiver Operating Characteristic Curve (**AUC-ROC**) analysis (Fawcett, 2006) introduced in Sec. 3.1. We test whether task success is associated more with which *skills* a trajectory uses or with how the agent *behaves*. For each evaluation trajectory, we construct two representations and use each to predict the binary outcome *solved* with an ℓ_2 -regularized logistic regression classifier, evaluated by ROC-AUC using 15 seeds and 5-fold cross-validation. **Model-S (skills)** encodes the skills that appear in the evaluation traces, together with the task categories, whereas **Model-B (behavior)** encodes agent behaviors using 12 rule-based features, such as *verify-before-done* and *repetition*. We also perform human inspection to remove behaviors with trivial meanings. We compute both representations from the traces generated by the SFT and RL checkpoints (trained without behavior-verifier enhancement). ROC-AUC measures the probability that the classifier assigns a higher score to a randomly chosen solved task than to a randomly chosen unsolved task. Thus, 0.5 corresponds to chance-level discrimination, while 1.0 corresponds to perfect separation. An AUC near 0.5 therefore indicates that the corresponding features provide little information about whether a task is solved. Further details of the setup are provided in App. D.

Test results and visualization of curves are presented in Fig. 1 (**Right**). We make two main observations from the results. **(i)** Behavior is strongly predictive of success, whereas skills are not: Model-B reaches $\text{AUC} = 0.64$ to 0.74 , while Model-S remains near chance (0.48 to 0.56). Knowing which skills appear in a trajectory provides little information about whether the task is solved, whereas how the agent behaves is substantially more predictive. **(ii)** RL strengthens this relationship: from SFT to RL, the AUC of Model-B increases from 0.64 to 0.74 , while that of Model-S remains nearly unchanged (~ 0.55). The improvement is therefore specific to the behavior axis rather than the skill inventory. Together, these results suggest that RL improves terminal-agent performance primarily by **optimizing behavior** and is not efficient in acquiring new skills.

Section 5.1 and 5.2 Takeaways

- RL largely preserves low-level skills and is inefficient at teaching novel skills.
- Task success is associated more with how the agent *behaves* than which *skills* a trajectory uses.
- RL learns to route and compose single-turn skills and shapes high-level agent behaviors.

5.3 SFT ANALYSIS

This section provides additional experimental results and evidence supporting the importance of performing SFT for broader atomic-skill coverage. We use Qwen3-8B and a public open-source SFT dataset collection from the OpenThoughts (Guha et al., 2026; Raoof et al., 2026). We construct three different sets of SFT data after our format converting and filtering process: (1) **Narrow-SFT**, with narrow skill coverage: the data comes mainly from `nl2bash` and `InferredBugs` and contains $\sim 13.6\text{K}$ trajectories; (2) **Diverse-SFT**, with diverse coverage of atomic skills: the data is collected

from a wider range of sources, including SWE-Smith (Yang et al., 2026a) and Openthoughts’ own IssueTasks datasets, and contains $\sim 78\text{K}$ trajectories after filtering; and (3) **Diverse-DS-SFT**, a random subset version of **Diverse-SFT** containing 13.6K trajectories (similar size as Narrow-SFT). A visualization of skill coverage is presented in Fig. 9 (in App. B.1), where we observe that the Narrow set covers only a subset of the atomic skills, whereas the Diverse set achieves higher coverage. We also add a zero-RL group, in which RL is performed directly without SFT.

The average performance across the four benchmarks (TWV, TB-v2.1, TB-Pro, and OTBL) of the checkpoints before and after RL for these four initialization settings is presented in Fig. 6. We observe that, without SFT, performing RL directly on the base model leads to the worst evaluation performance, motivating the use of SFT before RL for relatively weak models such as Qwen3-8B. Comparing **Narrow-SFT** with **Diverse-DS-SFT**, although the two achieve comparable performance after SFT, a substantial gap remains after RL, indicating the importance of acquiring diverse atomic skills before RL for improved generalizability.




Figure 6: Mean score over the four evaluation benchmarks Before and after RL on TMax-RIVER-3.5K, for Base model, Narrow-SFT, Diverse-DS-SFT, and Diverse-SFT.

In another set of experiments, we further observe that **Diverse-SFT** can reduce reward hacking. During an early stage of our exploration of RL with environments from TermiGen (Zhu et al., 2026), we adopted partial rewards defined as $r_p = 0.5 \times \text{pass-rate} + 0.5 \times \text{full-pass-bonus}$. The training curves are presented in Fig. 7 (a), where both the training rewards and in-distribution evaluation scores improve during RL. However, this setup introduces two potential sources of reward hacking: (1) **Hackable verifiers**: the environments had not yet been filtered by our RIVER pipeline. The evaluation results in Fig. 11 (App. C) indicate that this dataset is not sufficiently clean for RL⁴; and (2) **Early termination**: under partial rewards, the agent may learn a shortcut that satisfies only a subset of the verifiers and then prematurely claims task completion. As shown in Fig. 7 (b, left), RL degrades generalizability even as in-distribution performance improves, which is consistent with reward hacking. Combined with the decrease in generation length during training and our trajectory analysis, these results indicate that this form of reward hacking primarily stems from early termination.




Figure 7: Training curves and evaluation results for different SFT initializations (Narrow-SFT vs. Diverse-SFT) and trajectory-level reward designs (binary vs. partial). All RL experiments are conducted using TermiGen (Zhu et al., 2026).

Using **Diverse-SFT** mitigates these reward-hacking risks. As shown in Fig. 7 (b, right), when starting from Diverse-SFT while keeping all other configurations unchanged, the evaluation performance after RL improves under both the binary-reward and partial-reward settings. Our interpretation is that exposure to diverse atomic skills equips the model with the capabilities needed to solve the train-

⁴TermiGen environments were used for SFT data collection in their experiments.

ing tasks, making it less likely to exploit shortcuts during policy optimization and thereby mitigating reward hacking to some extent.

Section 5.3 Takeaways

- For relatively weak base models (e.g., Qwen3-8B), SFT with diverse atomic-skill coverage provides an important warm start for RL to improve generalizability;
- SFT with broader skill coverage also helps reduce reward hacking.

5.4 RL ANALYSIS

In this section, we conduct more controlled experiments to better understand generalization in terminal-agent RL. We divide the RIVER-TMax-3.5K dataset into 8 categories based on the required domain skills, as shown in Tab. 2, and construct a held-out evaluation set of ~ 300 tasks from RIVER-TMax-3.5K. We then perform RL on different subsets of domains, ranging from all 8 domains down to only 2. For RL (2 doms), training uses only ~ 800 environments from the Debug and System domains. The evaluation results of the RL-trained models and the SFT model on the held-out set are presented in Tab. 2. We observe that reducing the number of training domains to 4 does not decrease overall performance, and we find no significant performance degradation even on specific domains excluded from training. When the number of training domains is reduced to 2, such that the training environments cover only Debug and System, overall performance decreases but remains substantially higher than that of the SFT model. Moreover, performance on untrained and OOD domains, such as Data and Numerical, still improves.

We also visualize 6 representative behaviors in the evaluation traces and measure the changes in their prevalence before and after RL for the 2-domain setting. The results are presented in Fig. 8. We observe that these behaviors are shaped not only in the trained domains (darker bars), but also in the untrained domains (lighter bars). This suggests that the behavioral changes induced by RL generalize beyond the training domains and contribute to performance improvements in other domains.




Figure 8: Behaviors difference after RL (post-RL ratio - pre-RL ratio). Domains: BU = Build, DA = Data, DE = Debug, GR = Graph, NU = Numerical, SE = Security, SY = Systems, TE = Text. $\uparrow/\downarrow$: the higher/lower, the better from empirical evaluations.

Category	Build	Data	Debug	Graph	Numerical	Security	Systems	Text	ALL
RL (8 doms)	42.6 (+8.2)	56.9 (+27.7)	54.3 (+10.9)	68.8 (+40.6)	69.7 (+24.2)	40.6 (+18.8)	52.4 (+26.2)	55.6 (+11.1)	53.9 (+21.5)
RL (6 doms)	45.9 (+11.5)	50.8 (+21.5)	58.7 (+15.2)	65.6 (+37.5)	63.6 (+18.2)	43.8 (+21.9)	61.9 (+35.7)	44.4 (+0.0)	55.2 (+22.9)
RL (4 doms)	50.8 (+16.4)	53.8 (+24.6)	47.8 (+4.3)	68.8 (+40.7)	69.7 (+24.2)	40.6 (+18.7)	51.2 (+25.0)	55.6 (+11.2)	53.6 (+21.3)
RL (2 doms)	45.9 (+11.5)	47.7 (+18.5)	52.2 (+8.7)	53.1 (+25.0)	57.6 (+12.1)	46.9 (+25.0)	46.4 (+20.2)	55.6 (+11.1)	49.2 (+16.9)
SFT	34.4	29.2	43.5	28.1	45.5	21.9	26.2	44.4	32.3

Table 2: Solve rate (%) on the held-out evaluation set by task domain.

Section 5.4 Takeaways

- RL shapes behaviors beyond the training domains and fosters generalizability.

6 CONCLUSION

In this work, we study how RL improves terminal agents and propose the *agentic compositional generalization* hypothesis. Our empirical findings support complementary roles for different stages of model training: pre-training and SFT provide low-level atomic skills, while RL primarily shapes high-level decision-making behaviors that compose, route, and coordinate these skills across domains. This perspective is consistent with the strong OOD generalization of agentic RL and highlights *verifier quality*, which determines the training signal that shapes these behaviors, as a key factor beyond simply scaling the quantity or diversity of training environments. Guided by this insight, we introduce RIVER, a simple training recipe that improves the integrity of the reward signal by filtering low-quality environments and optionally incorporating behavior-level regularization. RIVER achieves state-of-the-art performance among evaluated RL-trained 8B models across four terminal-agent benchmarks. Moreover, RIVER generalizes across model families, model scales, agent harnesses, and RL objectives: using fewer than 30% of the TMax training environments, we achieve, on average, 106% and 30% larger RL gains on Terminal-Bench-Lite and Terminal-Bench-v2.1, respectively, across models ranging from 2B to 27B. Together, our results suggest that improving the quality of training signals, rather than scaling environment coverage alone, provides an effective and efficient path toward building more generalizable terminal agents.

REFERENCES

- Anthropic. Claude code: Best practices for agentic coding. <https://www.anthropic.com/engineering/claude-code-best-practices>, April 2025.
- Anthropic. Claude opus 4.8. <https://www.anthropic.com/news/claude-opus-4-8>, 2026.
- Ibragim Badertdinov, Alexander Golubev, Maksim Nekrashevich, Anton Shevtsov, Simon Karasik, Andrei Andriushchenko, Maria Trofimova, Daria Litvintseva, and Boris Yangel. Swe-rebench: An automated pipeline for task collection and decontaminated evaluation of software engineering agents. *Advances in Neural Information Processing Systems*, 38, 2026.
- Ivan Bercovich, Ivgeni Segal, Kexun Zhang, Shashwat Saxena, Aditi Raghunathan, and Ziqian Zhong. Terminal wrench: A dataset of 331 reward-hackable environments and 3,632 exploit trajectories. *arXiv preprint arXiv:2604.17596*, 2026.
- Shiyi Cao, Sumanth Hegde, Dacheng Li, Tyler Griggs, Shu Liu, Eric Tang, Jiayi Pan, Xingyao Wang, Akshay Malik, Graham Neubig, Kourosh Hakhmaneshi, Richard Liaw, Philipp Moritz, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. Skyrl-v0: Train real-world long-horizon agents via reinforcement learning, 2025.
- Zhepeng Cen, Haolin Chen, Shiyu Wang, Zuxin Liu, Zhiwei Liu, Ding Zhao, Caiming Xiong, Huan Wang, and Weiran Yao. Webscale-rl: Automated data pipeline for scaling rl data to pretraining levels. In *International Conference on Learning Representations*, volume 2026, pp. 87623–87642, 2026a.
- Zhepeng Cen, Yihang Yao, William Han, Zuxin Liu, and Ding Zhao. Behavior injection: Preparing language models for reinforcement learning. *Advances in Neural Information Processing Systems*, 38:155105–155137, 2026b.
- Fan Chen, Audrey Huang, Noah Golowich, Sadhika Malladi, Adam Block, Jordan Ash, Akshay Krishnamurthy, and Dylan Foster. The coverage principle: How pre-training enables post-training. In *International Conference on Learning Representations*, volume 2026, pp. 112225–112293, 2026a.
- Zhaorun Chen, Zhuokai Zhao, Kai Zhang, Bo Liu, Qi Qi, Yifan Wu, Tarun Kalluri, Xuefei Cao, Yuanhao Xiong, Haibo Tong, et al. Scaling agent learning via experience synthesis. In *International Conference on Learning Representations*, volume 2026, pp. 121394–121420, 2026b.
- Zihao Cheng, Hongru Wang, Zeming Liu, Xinyi Wang, Xiangrong Zhu, Yuhang Guo, Wei Lin, Jeff Z Pan, and Yunhong Wang. Terminal-world: Scaling terminal-agent environments via agent skills. *arXiv preprint arXiv:2605.20876*, 2026.

- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *arXiv preprint arXiv:2501.17161*, 2025.
- Zhaoyang Chu, Jiarui Hu, Xingyu Jiang, Pengyu Zou, Han Li, Chao Peng, Peter O’Hearn, Earl T Barr, Mark Harman, Federica Sarro, et al. Terminalworld: Benchmarking agents on real-world terminal tasks. *arXiv preprint arXiv:2605.22535*, 2026.
- Zhiyuan Fan, Tinghao Yu, Yuanjun Cai, Jiangtao Guan, Yun Yang, Dingxin Hu, Jiang Zhou, Xing Wu, Zhuo Han, Feng Zhang, et al. Toward scalable terminal task synthesis via skill graphs. *arXiv preprint arXiv:2604.25727*, 2026.
- Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv preprint arXiv:2503.01307*, 2025.
- Kanishk Gandhi, Shivam Garg, Noah D Goodman, and Dimitris Papailiopoulos. Endless terminals: Scaling rl environments for terminal agents. *arXiv preprint arXiv:2601.16443*, 2026.
- Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, et al. Openthoughts: Data recipes for reasoning models. In *International Conference on Learning Representations*, volume 2026, pp. 108059–108130, 2026.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Runyuan He, Qiuyang Mang, Shang Zhou, Kaiyuan Liu, Hanchen Li, Huanzhi Mao, Qizheng Zhang, Zerui Li, Bo Peng, Lufeng Cheng, et al. Frontiersmith: Synthesizing open-ended coding problems at scale. *arXiv preprint arXiv:2605.14445*, 2026.
- Zhongmou He, Yee Man Choi, Kexun Zhang, Jiabao Ji, Juntao Zhou, Dejie Xu, Ivan Bercovich, Aidan Zhang, and Lei Li. Hardtests: Synthesizing high-quality test cases for llm coding. *arXiv preprint arXiv:2505.24098*, 2025.
- Audrey Huang, Adam Block, Dylan Foster, Dhruv Rohatgi, Cyril Zhang, Max Simchowitz, Jordan Ash, and Akshay Krishnamurthy. Self-improvement in language models: The sharpening mechanism. In *International Conference on Learning Representations*, volume 2025, pp. 76687–76739, 2025.
- Hamish Ivison, Junjie Oscar Yin, Rulin Shao, Teng Xiao, Nathan Lambert, and Hannaneh Hajishirzi. Tmx: A simple recipe for terminal agents. *arXiv preprint arXiv:2606.23321*, 2026.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations*, volume 2024, pp. 54107–54157, 2024.
- Lingjing Kong, Xin Liu, Guangyi Chen, Martin Q Ma, Xiangchen Song, Yuekai Sun, Mikhail Yurochkin, Taylor W Killian, Ruslan Salakhutdinov, Kun Zhang, et al. From reasoning traces to reusable modules: Understanding compositional generalization in language model reasoning. *arXiv preprint arXiv:2606.18089*, 2026.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis: connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:249, 2008.
- Laude Institute. Harbor: evaluate agents in sandboxed environments. <https://github.com/laude-institute/harbor>, 2026.
- Zhongzhi Li, Yucheng Shi, Zongxia Li, Ruhan Wang, Anhao Li, Zixun Huang, Junyao Yang, Lei Ke, Ninghao Liu, Haitao Mi, et al. Recursive synthesis for long-horizon terminal tasks. *arXiv preprint arXiv:2608.05466*, 2026.

- Yusong Lin, Haiyang Wang, Shuzhe Wu, Lue Fan, Feiyang Pan, Sanyuan Zhao, and Dandan Tu. Cli-gym: Scalable cli task generation via agentic environment inversion. *arXiv preprint arXiv:2602.10999*, 2026.
- Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran Yao, Zhiwei Liu, Yihao Feng, et al. Apigen: Automated pipeline for generating verifiable and diverse function-calling datasets. *Advances in Neural Information Processing Systems*, 37:54463–54482, 2024.
- Nathan Mantel. The detection of disease clustering and a generalized regression approach. *Cancer research*, 27(2_Part_1):209–220, 1967.
- Mike Merrill, Alexander Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Shin, Thomas Walshe, E Kelly Buchanan, et al. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations*, volume 2026, pp. 40903–40986, 2026.
- OpenAI. Openai o1 system card. <https://cdn.openai.com/o1-system-card-20241205.pdf>, December 2024.
- OpenAI. Introducing codex. <https://openai.com/index/introducing-codex/>, May 2025.
- OpenAI. Gpt-5.4 thinking system card. <https://openai.com/index/gpt-5-4-thinking-system-card/>, 2026.
- Karl Pearson. Vii. mathematical contributions to the theory of evolution.—iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, (187):253–318, 1896.
- Xiaoxuan Peng, Xinyu Lu, Kaiqi Zhang, Taosong Fang, Boxi Cao, and Yaojie Lu. Announcing LiteCoder-terminal: Lightweight terminal agents with <1k synthesized trajectories. Hugging Face Blog, December 2025. URL <https://huggingface.co/blog/Lite-Coder/litecoder-terminal-preview>.
- Renjie Pi, Grace Lam, Mohammad Shoeybi, Pooya Jannaty, Bryan Catanzaro, and Wei Ping. On data engineering for scaling llm terminal capabilities. *arXiv preprint arXiv:2602.21193*, 2026.
- Akshara Prabhakar, Zuxin Liu, Ming Zhu, Jianguo Zhang, Tulika Manoj Awalganekar, Shiyu Wang, Zhiwei Liu, Haolin Chen, Thai Hoang, Juan Carlos Niebles, et al. Apigen-mt: Agentic pipeline for multi-turn data generation via simulated agent-human interplay. *Advances in Neural Information Processing Systems*, 38, 2026.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Penghui Qi, Xiangxin Zhou, Zichen Liu, Tianyu Pang, Chao Du, Min Lin, and Wee Sun Lee. Re-thinking the trust region in llm reinforcement learning. *arXiv preprint arXiv:2602.04879*, 2026.
- Jielin Qiu, Zuxin Liu, Zhiwei Liu, Rithesh Murthy, Jianguo Zhang, Haolin Chen, Shiyu Wang, Ming Zhu, Liangwei Yang, Juntao Tan, et al. Locobench-agent: An interactive benchmark for llm agents in long-context software engineering. *arXiv preprint arXiv:2511.13998*, 2025.
- Nived Rajaraman, Audrey Huang, Miroslav Dudik, Robert Schapire, Dylan Foster, and Akshay Krishnamurthy. Learning to reason with curriculum ii: Compositional generalization. *arXiv preprint arXiv:2606.27721*, 2026.
- Negin Raoof, Richard Zhuang, Marianna Nezhurina, Etash Guha, Atula Tejaswi, Ryan Marten, Charlie F Ruan, Tyler Griggs, Alexander Glenn Shaw, Hritik Bansal, et al. Openthoughts-agent: Data recipes for agentic models. *arXiv preprint arXiv:2606.24855*, 2026.
- Yi Ren and Danica Sutherland. Learning dynamics of llm finetuning. In *International Conference on Learning Representations*, volume 2025, pp. 70523–70563, 2025.

- Chi Ruan, Dongfu Jiang, Huaye Zeng, Ping Nie, and Wenhua Chen. Evolvecoder: Evolving test cases via adversarial verification for code reinforcement learning. *arXiv preprint arXiv:2603.12698*, 2026.
- Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, et al. Spurious rewards: Rethinking training signals in rlvr. *arXiv preprint arXiv:2506.10947*, 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, pp. 1279–1297, 2025.
- Charles Spearman. The proof and measurement of association between two things. 1961.
- Qwen Team. Qwen3. 5-omni technical report. *arXiv preprint arXiv:2604.15804*, 2026a.
- Qwen Team. Qwen3. 6-27b: Flagship-level coding in a 27b dense model, 2026b.
- Weixun Wang, XiaoXiao Xu, Wanhe An, Fangwen Dai, Wei Gao, Yancheng He, Ju Huang, Qiang Ji, Hanqi Jin, Xiaoyang Li, et al. Let it flow: Agentic crafting on rock and roll, building the rome model within an open agentic learning ecosystem. *arXiv preprint arXiv:2512.24873*, 2025a.
- Weixun Wang, XiaoXiao Xu, Xander Xu, et al. Let it flow: Agentic crafting on rock and roll, building the rome model within an open agentic learning ecosystem, 2025b. URL <https://arxiv.org/abs/2512.24873>.
- Yinjie Wang, Xuyang Chen, Xiaolong Jin, Mengdi Wang, and Ling Yang. Openclaw-rl: Train any agent simply by talking. *arXiv preprint arXiv:2603.10165*, 2026.
- Siwei Wu, Yizhi Li, Yuyang Song, Wei Zhang, Yang Wang, Riza Batista-Navarro, Xian Yang, Mingjie Tang, Bryan Dai, Jian Yang, et al. Large-scale terminal agentic trajectory generation from dockerized environments. *arXiv preprint arXiv:2602.01244*, 2026a.
- Yifan Wu, Zhuokai Zhao, Songlin Li, Ho Hin Lee, Jiacheng Zhu, Shirley Wu, Tianhe Yu, Serena Li, Lizhu Zhang, Xiangjun Fan, et al. Swe-together: Evaluating coding agents in interactive user sessions. *arXiv preprint arXiv:2606.29957*, 2026b.
- xAI. Grok code fast 1. <https://x.ai/news/grok-code-fast-1>, August 2025. Accessed: 2026-08-11.
- Zhangchen Xu, Adriana Meza Soria, Shawn Tan, Anurag Roy, Ashish Sunil Agrawal, Radha Poovendran, and Rameswar Panda. Toucan: Synthesizing 1.5 m tool-agentic data from real-world mcp environments. *arXiv preprint arXiv:2510.01179*, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik R Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://arxiv.org/abs/2405.15793>.
- John Yang, Kilian Lieret, Carlos Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. Swe-smith: Scaling data for software engineering agents. *Advances in Neural Information Processing Systems*, 38, 2026a.
- Rui Yang, Qianhui Wu, Zhaoyang Wang, Hanyang Chen, Ke Yang, Hao Cheng, Huaxiu Yao, Baolin Peng, Huan Zhang, Jianfeng Gao, et al. Gui-libra: Training native gui agents to reason and act with action-aware supervision and partially verifiable rl. *arXiv preprint arXiv:2602.22190*, 2026b.

- Yihang Yao, Zhepeng Cen, Haohong Lin, Shiqi Liu, Zuxin Liu, Jiacheng Zhu, Zhang-Wei Hong, Laixi Shi, and Ding Zhao. Pushing forward pareto frontiers of proactive agents with behavioral agentic optimization. *arXiv preprint arXiv:2602.11351*, 2026a.
- Yihang Yao, Guangtao Zeng, Raina Wu, Yang Zhang, Ding Zhao, Zhang-Wei Hong, and Chuang Gan. Tailored primitive initialization is the secret key to reinforcement learning. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 33300–33318, 2026b.
- Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of Language Models: Part 2.1, Grade-School Math and the Hidden Reasoning Process. In *Proceedings of the 13th International Conference on Learning Representations, ICLR '25*, April 2025. Full version available at <https://ssrn.com/abstract=5250629>.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-of-thought reasoning in llms. *arXiv preprint arXiv:2502.03373*, 2025.
- Zhaochen Yu, Ling Yang, Jiaru Zou, Shuicheng Yan, and Mengdi Wang. Demystifying reinforcement learning in agentic reasoning. *arXiv preprint arXiv:2510.11701*, 2025.
- Lifan Yuan, Weize Chen, Yuchen Zhang, Ganqu Cui, Hanbin Wang, Ziming You, Ning Ding, Zhiyuan Liu, Maosong Sun, and Hao Peng. From $f(x)$ and $g(x)$ to $f(g(x))$: Llms learn new skills in rl by composing old ones. In *International Conference on Learning Representations*, volume 2026, pp. 147547–147574, 2026.
- Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model?, 2025. URL <https://arxiv.org/abs/2504.13837>, 204:1–16.
- Guangtao Zeng, Maohao Shen, Delin Chen, Zhenting Qi, Subhro Das, Dan Gutfreund, David Cox, Gregory Wornell, Wei Lu, Zhang-Wei Hong, et al. Satori-swe: Evolutionary test-time scaling for sample-efficient software engineering. *arXiv preprint arXiv:2505.23604*, 2025.
- Kai Zhang, Xiangchao Chen, Bo Liu, Tianci Xue, Zeyi Liao, Zhihan Liu, Xiyao Wang, Yuting Ning, Zhaorun Chen, Xiaohan Fu, et al. Agent learning via early experience. *arXiv preprint arXiv:2510.08558*, 2025.
- Jiarong Zhao, Zhikai Lei, Zhiheng Xi, Rui Zheng, Hang Yan, Jie Zhou, Qin Chen, and Liang He. Nexforge: Scaling agent capabilities through requirement-driven task synthesis for llms. *arXiv e-prints*, pp. arXiv–2607, 2026.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody H Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, et al. Sglang: Efficient execution of structured language model programs. *Advances in neural information processing systems*, 37:62557–62583, 2024a.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 3: system demonstrations)*, pp. 400–410, 2024b.
- Yuhang Zhou, Lizhu Zhang, Yifan Wu, Jiayi Liu, Xiangjun Fan, Zhuokai Zhao, and Hong Yan. Synthetic sandbox for training machine learning engineering agents. *arXiv preprint arXiv:2604.04872*, 2026.
- Kaijie Zhu, Yuzhou Nie, Yijiang Li, Yiming Huang, Jialian Wu, Jiang Liu, Ximeng Sun, Zhenfei Yin, Lun Wang, Zicheng Liu, et al. Termigen: High-fidelity environment and robust trajectory synthesis for terminal agents. *arXiv preprint arXiv:2602.07274*, 2026.

CONTENTS

1	Introduction	1
2	Related Work	2
3	Mechanism Interpretation and Method	3
3.1	Agentic Compositional Generalization	3
3.2	RIVER: Reward-Integrity Verified Environments for RL	5
4	Experiment Results	6
4.1	Qwen3-8B Model Family	7
4.2	Qwen3.5 and Qwen3.6 Families	8
5	Analysis	9
5.1	Skill preserving and behavior shaping in RL	9
5.2	Behaviors' Correlation with Task Performance	10
5.3	SFT Analysis	10
5.4	RL Analysis	12
6	Conclusion	13
A	Training Details	19
A.1	Supervised Fine-Tuning	19
A.2	Reinforcement Learning	19
A.3	Evaluation	20
B	Data and Environment Details	22
B.1	SFT dataset coverage	22
B.2	RL environments	23
C	Environment Evaluation	24
C.1	Evaluation Rubrics and Issue Categories	24
C.2	Cross-model validation of the audit labels	25
C.3	Environment Evaluation Examples	26
C.4	Examples of Verifier Failures on the Unfiltered Environment Set	29
D	Behavior Analysis Details	34

A TRAINING DETAILS

A.1 SUPERVISED FINE-TUNING

Training data. We finetune on multi-turn terminal-agent trajectories in an *endless-agent* format Gandhi et al. (2026): a system prompt, the task instruction, and an alternating sequence of assistant turns and environment observations, where every assistant turn contains an explicit chain-of-thought (`<think>...</think>`) followed by a single shell command or a terminating action. We obtained **Diverse-SFT** with 78K trajectories after format converting and filtering. For more SFT dataset discussions and atomic skill coverage analysis, please refer to App. B.1.

Model and optimization. We perform full-parameter fine-tuning of **Qwen3-8B** using the LLaMA-Factory framework (Zheng et al., 2024b). Loss is computed on assistant tokens (generated tokens) only. Key hyperparameters are listed in Table 3; experiments with both datasets are done with the identical recipe, so any downstream difference is attributable to the data alone.

Table 3: SFT hyperparameters.

Hyperparameter	Value
Base model	Qwen3-8B
Fine-tuning type	Full parameter
Precision	bf16
Epochs	1
Learning rate	1×10^{-5}
LR schedule	Cosine
Warmup ratio	0.1
Effective batch size	64
(per-device $\times$ GPUs $\times$ grad. accum.)	$(1 \times 8 \times 8)$
Max sequence length	16,384
Optimizer sharding	DeepSpeed ZeRO-2
Attention	SDPA

A.2 REINFORCEMENT LEARNING

This section provides details for training our 8B models.

Training environments. For full RIVER, we train on the **RIVER-TMax-3.5K** environment set.

Training method. For the Qwen3 model family, We fine-tune the policy with **Group Relative Policy Optimization (GRPO)** (Shao et al., 2024). For each prompt we sample a group of G rollouts from the current policy, execute them in the terminal environment, and score each rollout with the task verifier under a *sparse binary* reward ($r \in \{0, 1\}$: 1 iff the verifier passes). The advantage term is defined and discussed in (1).

Objective. Let $\rho_t(\theta) = \pi_\theta(a_t | s_t) / \pi_{\theta_{\text{old}}}(a_t | s_t)$ be the per-token importance ratio and A_t the (shaped) token advantage.

$$\mathcal{L}(\theta) = -\mathbb{E}_t \left[\min(\rho_t A_t, \text{clip}(\rho_t, 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}}) A_t) \right] + \beta \text{KL}(\pi_\theta \| \pi_{\text{ref}}), \quad (2)$$

with asymmetric clipping bounds $\varepsilon_{\text{low}}, \varepsilon_{\text{high}}$ and KL coefficient β . The per-token losses are aggregated with a sequence-mean reduction.

Hyperparameters. We adopt most hyperparameters for RL training from the endless-terminal work (Gandhi et al., 2026). Table 4 lists the training hyperparameters.

Computation details. All RL runs for Qwen3-8B models use a single such node. Each node is an 8xNVIDIA H200 GPU server, with ~ 143 GB HBM3e per GPU (~ 1.15 TB aggregate GPU memory), fully NVLink-connected. Host side: 2x Intel Xeon Platinum 8488C (96 logical cores),

and ~ 2 TB system RAM. For the main experiments, we run for around 60 hours to perform 400 RL steps.

Table 4: RL (GRPO) training hyperparameters.

Hyperparameter	Value
<i>Optimization</i>	
Base policy	Qwen3-8B (post-trained checkpoint)
Advantage estimator	GRPO
Learning rate	5×10^{-7}
Warmup steps	20
Weight decay	0.01
Training epochs	10
Update epochs per batch	2
Train batch size (prompts)	32
Policy mini-batch size	32
Parameter sharding	FSDP2, 1 node $\times$ 8 GPUs
Max RL Step	400
<i>Objective</i>	
KL in reward	disabled
Clip range (low, high)	(0.2, 0.28)
Loss reduction	sequence-mean
Reward	sparse binary $\{0, 1\}$
<i>Rollout / generation</i>	
Samples per prompt	8
Max turns per episode	16
Max prompt length	1536
Max input length	16384
Max generation length (per turn)	2048
Sampling temperature	0.6
Top- p	1.0
Inference engines	8
<i>Repetition penalty</i>	
Penalty δ (per token)	0.05
Per-rollout cap	0.15

A.3 EVALUATION

Serving. All checkpoints are evaluated under an identical serving configuration. We convert each policy checkpoint to Hugging Face `safetensors` and serve it with **SGLang** (Zheng et al., 2024a). A single 8-GPU node hosts one model instance with tensor parallelism $TP=8$ (we do not use data-parallel replicas, which we found to deadlock under the agent’s synchronous per-turn calls) and a static KV-cache memory fraction of 0.85. The served context length is the model’s full 40,960-token window (`max_position_embeddings`).

Decoding and budgets. Rollouts use temperature $T=0.6$ and top- $p=1.0$, with up to 2048 tokens generated per turn. Each episode is capped at 64 action turns and a 600-second wall-clock budget, whichever comes first; a task that reaches either cap without a terminating action is scored as a failure.

Evaluation Benchmarks. We evaluate on four terminal-agent benchmarks and report their mean score, summarized in Tab. 5. All benchmarks are in the harbor (Laude Institute, 2026) task format. For **OpenThoughts-TBLite**, the score is a continuous value per their design. For the rest benchmarks, scoring is binary pass/fail per task (reward 1 iff the verifier’s tests pass).

Table 5: Evaluation benchmarks. The averaged score is the unweighted mean of the four test benchmarks.

Benchmark	#Tasks	Score Type	Focus
OpenThoughts-TB Lite	100	Continuous	General CLI / file-ops / SWE
Terminal-Bench 2.1	89	Binary	SWE, ML, security, data science
Terminal-Bench-Pro	200	Binary	SWE, systems, debugging, data
TerminalWorld-Verified	200	Binary	Mixed (categorized, verified)

Table 6: Evaluation serving and decoding configuration, held fixed across all checkpoints and benchmarks.

Setting	Value
Serving engine	SGLang
Tensor parallelism	8 (<code>--tp 8</code>)
KV-cache mem. fraction	0.85
Served context length	40,960 tokens
Sampling temperature	0.6
Top- p	1.0
Max tokens per turn	2048
Max turns per episode	64
Wall-clock budget / task	600 s
Task concurrency	8

B DATA AND ENVIRONMENT DETAILS

B.1 SFT DATASET COVERAGE




Figure 9: Atomic-skill coverage of the three SFT datasets, measured as the fraction of trajectories labeled by GPT-5.4 as exhibiting each of the 60 command-level skills.

SFT Construction Details. The source SFT data OpenThoughts-Collection (Guha et al., 2026; Raoof et al., 2026) uses the **Terminus-2** harness, whereas our stack consumes the **EndlessAgent** (Gandhi et al., 2026) harness, so we translate every trajectory deterministically and drop any that cannot be represented. In **Terminus-2**, the content in one generation turn is a JSON object with `analysis`, `plan`, a **batch** of commands, and `task.complete`, and each observation is a full terminal feedback. In **EndlessAgent**, the content in one generation turn consists of the following content: `<think>...</think><command>c</command>` for one command c (or

<action>done</action> to claim the task is done), and observations use the fixed template Command executed successfully. Output: <stdout> (exit_code=0).

The translation (1) replaces the system prompt with the EndlessAgent one; (2) extracts the task from the first user message (the span after Task Description:); (3) merges a turn’s command batch $[c_1, \dots, c_k]$ into <command> $c_1 \ \&\& \dots \ \&\& \ c_k$ </command>; (4) fills <think> from analysis+plan; and (5) reduces each raw terminal screen to the command’s actual output by removing the shell prompt and the echo of the typed command (including the continuation lines of multi-line commands), while carefully preserving genuine output that happens to start with “>” (e.g. diff results).

We drop a trajectory if its first message lacks a Task Description: marker, or if any assistant turn *other than the last* fails to parse as valid JSON. A non-JSON turn appearing *as the final turn* is not treated as corruption: it is the terminus “are you sure?” confirmation, which we convert to done. Because the trainer silently discards any trajectory with two consecutive assistant turns, we treat done as terminal and splice a synthetic observation wherever one assistant turn would directly follow another; The remaining trajectories all end in an explicit done.

Skill composition of the SFT datasets. Fig. 9 compares the atomic skill composition of the three SFT corpora (Narrow-SFT, Diverse-DS-SFT, and Diverse-SFT) used in our study, labeling every trajectory with a 60-skill atomic taxonomy (via GPT-5.4) and reporting per-skill coverage. The comparison serves two purposes. First, it validates Diverse-DS-SFT as a clean size control for Diverse-SFT: because Diverse-DS-SFT is drawn i.i.d. from Diverse-SFT, its coverage is indistinguishable from the full set across all 60 skills, so any downstream difference between models trained on Diverse-SFT and Diverse-DS-SFT is attributable to data *volume* rather than to a shift in skill composition. Second, it isolates what distinguishes Narrow-SFT from Diverse-SFT: the gap is not a uniform thinning but a categorical one on the execution-oriented skills. Narrow-SFT is dominated by filesystem navigation and authoring yet contains almost no code execution or environment setup (e.g., run-python, run-pytest, pip/apt-install all near 0%), whereas Diverse-SFT and Diverse-DS-SFT exercise these skills with SFT.

B.2 RL ENVIRONMENTS

In Fig. 10, each environment is passed through a two-stage filter, and every bar shows the outcome as a stacked count: (1) **fails rubric check** (grey): flagged by our rubrics and GPT-5.4 judges; (2) **pass rubric check but too hard** (blue): the environment is clean, but a capable reference agent cannot solve it; (3) **difficulty-adjusted** (green): clean *and* solvable within the budget. These environments consist of our **RIVER-TMax-3.5K**. The annotation above each bar reports difficulty-adjusted% / rubric-pass% of the original pool. Categories are the task-provided domain labels; skill groups are assigned by mapping each environment’s primary skill onto one of eight groups (build, data, debug, graph, numerical, security, systems, text).




Figure 10: Environment quality-and-difficulty funnel for the TMax training pool (14,399 environments), broken down by task *category* (left, 9 classes) and by *skill group* (right, 8 groups).




Figure 11: Rubrics evaluation results conducted by GPT-5.4: (left to right): Environments from TMax (Iverson et al., 2026); environments from TerminalTraj (Wu et al., 2026a); and environments from TermiGen (Zhu et al., 2026)).

C ENVIRONMENT EVALUATION

We perform evaluations on three candidate environment pools using our environment-filtering rubrics. We emphasize that our evaluations may contain both false positives and false negatives. The results (Fig. 11) differ sharply across pools. **TMax** is the cleanest, with 35.8% CLEAN, but still contains 40.4% of environments flagged as **VERIFIER-TOO-WEAK**. **TermiGen** is only 10.1% CLEAN (51.4% too-weak, 19.6% mismatch), while **TerminalTraj-5k** performs the worst under our rubrics, with only 3.3% CLEAN and 67.2% **VERIFIER-TOO-WEAK**. Across all three pools, the dominant failure mode is weak verification (**VERIFIER-TOO-WEAK**). We next provide the full rubric (Sec.C.1) and show representative examples of each verdict (Sec.C.3).

C.1 EVALUATION RUBRICS AND ISSUE CATEGORIES

The evaluation agent reads the three artifacts that must be mutually consistent: `instruction.md` (the only information visible to the agent), the reward *verifier* (the pytest that scores a rollout), and the *environment* (the Dockerfile and any bundled data). It then asks a single agent-centric question: *would a competent agent that faithfully follows the instruction be able to (a) understand the task unambiguously and (b) actually pass the verifier?* Each task is assigned exactly one of eight verdicts:

- **CLEAN:** Instruction is clear and self-sufficient, a correct-per-instruction solution passes, AND the verifier genuinely validates correctness: it recomputes/decodes the true answer from the input (or exact-compares against a correct hardcoded answer) rather than merely checking output shape. A hard-but-fair task is still CLEAN.
- **VERIFIER-TOO-WEAK:** The verifier passes *without* the task actually being solved: it checks only output shape/format/types/count, trusts a self-reported value, greps the source for keywords, checks a hardcoded subset, or tests only fixed inputs so the answer can be hardcoded. A fabricated, plausibly-shaped output would pass.
- **INSTR-VERIFIER-MISMATCH:** The verifier checks something the instruction does not imply, or contradicts it. For example, reads/writes a different path, asserts a hardcoded value that conflicts with the data, or demands a stricter format. A solution that faithfully follows the instruction still fails.
- **INSTR-ENV-MISMATCH:** The instruction or verifier references input paths, files, dependencies, or services the environment does not actually provide, so the task cannot be completed as written (missing data/compiler/service).
- **ANSWER-LEAK:** The instruction itself, or a file shipped in the environment, reveals the expected answer (lists the result, embeds the solution, or provides a ground-truth file), so the task requires no genuine work.
- **INSTR-AMBIGUOUS:** The instruction alone is underspecified: a *necessary* specification is missing, so two equally reasonable interpretations produce different outputs and the verifier accepts only one. (Merely admitting multiple valid solution *paths* is not enough.)

GPT-5.4 label	Independent blind re-judge	
	CLEAN	defective
CLEAN (15)	6	9
defective (105)	11	94

Table 7: Binary cross-model comparison (CLEAN vs. defective) between the GPT-5.4 audit and an independent, label-blind Claude re-judge on a stratified sample of 120 TMax environments, with 15 examples drawn from each GPT-5.4 verdict category. Defect labels show substantially higher conditional agreement than CLEAN labels.

- **TASK-TRIVIAL:** The task itself is so easy there is no meaningful work or learning signal: write a constant, echo a value, copy a file, regardless of verifier strength.
- **OTHER:** A genuine defect fitting none of the above: verifier code that is buggy/crashes; a task needing network/GPU/GUI/hardware the container lacks; non-determinism; a required background service that is never started; or an internally contradictory spec.

C.2 CROSS-MODEL VALIDATION OF THE AUDIT LABELS

The verdicts above were produced by GPT-5.4. To assess whether these judgments are robust rather than idiosyncratic to a single model, we conducted an independent blind re-judgment using a different model, Claude Opus 4.8 (Anthropic, 2026). We constructed a stratified sample of 120 TMax environments, with 15 examples drawn from each of the eight GPT-5.4 verdict categories. To avoid anchoring, each environment was stripped of its GPT-5.4 label and any category cue, shuffled, and provided to a dedicated judge that independently read the full instruction, Dockerfile, and verifier code before assigning a verdict. The two labels were compared only afterward.

At the fine-grained level, the two models assign the same eight-way verdict to only 22.5% of the 120 examples. This disagreement does not necessarily imply that one judgment is incorrect: a single defective environment can exhibit multiple failure modes simultaneously. For example, a verifier that hard-codes its expected output may reasonably be categorized as VERIFIER-TOO-WEAK, while the same environment may also exhibit characteristics of INSTR/VERIFIER-MISMATCH or TASK-TRIVIAL. Thus, forcing each environment into a single primary category can lead the two models to disagree on *which* defect is most salient even when both identify the environment as defective.

For training, the more consequential distinction is binary: whether an environment is usable (CLEAN) or exhibits at least one defect. Under this binary grouping, the two models agree on 100/120 examples (83.3%) in our stratified sample, as summarized in Tab. 7. Because the sample was stratified by the GPT-5.4 verdict rather than drawn according to the natural label distribution of the full pool, this number should be interpreted as agreement under our sampling design rather than as an estimate of population-level agreement.

More informative are the conditional agreement rates. Among the 105 examples labeled defective by GPT-5.4, the independent judge also identifies 94 as defective (89.5%). In contrast, among the 15 examples labeled CLEAN by GPT-5.4, only 6 are independently judged CLEAN (40.0%). This asymmetry suggests that defect labels are substantially more stable across judges than CLEAN labels. It also indicates that our reported clean rates may be optimistic, since some environments labeled CLEAN by GPT-5.4 are judged defective by the independent model. Overall, this validation supports using the audit as a conservative filtering mechanism for removing environments with readily detectable defects, while treating fine-grained defect categories and CLEAN labels with greater caution.

C.3 ENVIRONMENT EVALUATION EXAMPLES

We show representative environments for each major failure category identified by our audit. Each example was independently checked by a second model. Each box summarizes what the task asks, why the environment receives the corresponding verdict, and the most direct evidence supporting that verdict. We show one or two examples per category, prioritizing cases where the failure mode is unambiguous and can be established directly from the instruction, environment, and verifier.

VERIFIER-TOO-WEAK

`task_000254_21ede93b`

Task. Write a C archiver daemon that recursively processes legacy log files, invokes a filter script via `popen()`, maintains atomic JSON state, and continuously monitors an incoming directory using `inotify`.

Why VERIFIER-TOO-WEAK. The instruction requires several concrete implementation properties: the solution must be a C program, recursively traverse existing logs, invoke the filter through `popen()`, update its state atomically using `rename()`, and use `inotify` to process newly closed files. The verifier, however, observes only a running process named `archiver`, the existence of output files, and the expected aggregate counts before and after adding a new file. It never checks that the implementation is written in C or that any of the required mechanisms are actually used. Consequently, another implementation that merely reproduces the expected observable files and counts could satisfy the verifier without implementing the requested archiver.

Key evidence. The verifier checks the process name, output-file existence, and `total_critical` counts, but never validates C source, `popen()`, recursive traversal, `inotify`, or atomic `rename()`.

`task_000233_ee5f0e33`

Task. Back up and repair several microservice configurations, write an idempotent Bash script that applies the required port and upstream settings, and write a second script that derives a topology report from the resulting configurations.

Why VERIFIER-TOO-WEAK. The verifier checks configuration values only through substring membership, such as testing whether `LISTEN_PORT=9001` occurs somewhere in the file. This does not ensure that the configuration is well formed, unique, or free of conflicting stale values. More importantly, although the task requires `generate_report.sh` to parse the configurations and dynamically construct the topology report, the verifier never executes that script. It only checks that the final report file contains a fixed set of expected lines. Thus, the required data flow from configuration repair to report generation is not verified.

Key evidence. Configuration values are checked by substring presence only, and `generate_report.sh` is never executed by the verifier.

INSTR-VERIFIER-MISMATCH

`task_000087_787a0724`

Task. Implement a C program that selects hyperparameters M and T for a text classifier using sequential 5-fold cross-validation over 100 examples.

Why INSTR-VERIFIER-MISMATCH. The instruction explicitly requires evaluating every (M, T) pair on five sequential validation folds of 20 examples each and selecting the pair with the highest *mean validation accuracy*. The verifier implements a different selection criterion: it evaluates every candidate once on all 100 examples and chooses the pair with the highest full-dataset accuracy. These two objectives are not equivalent. Therefore, a program that correctly implements the requested cross-validation procedure can select a different pair from the one accepted by the verifier.

Key evidence. Instruction: mean accuracy over five sequential validation folds. Verifier: `correct / len(y)` over the entire 100-example dataset, with no cross-validation folds.

task_001036_07b75c89

Task. Repair a compromised data-processing pipeline, including an authentication validator, a redaction worker, and an Nginx reverse proxy that should listen on 127.0.0.1:8080 and forward traffic to the Flask application.

Why INSTR-VERIFIER-MISMATCH. The agent follows the networking requirement using the canonical Nginx configuration `listen 127.0.0.1:8080;`. However, the verifier searches the configuration with the regular expression `listen\s+8080\s*`, which accepts `listen 8080;` but does not match the explicitly specified `127.0.0.1:8080` form. Thus, the verifier rejects a configuration that directly follows the instruction.

Key evidence. Instruction requires listening on 127.0.0.1:8080; the agent writes `listen 127.0.0.1:8080;`, while the verifier only matches the pattern `listen 8080;`.

INSTR-ENV-MISMATCH

task_000641_37f8ea78

Task. Read a mathematical degeneracy threshold from an image at `/app/threshold.png` using OCR and use that threshold in the requested analysis.

Why INSTR-ENV-MISMATCH. The instruction explicitly identifies `/app/threshold.png` as the input containing the threshold definition. The environment does not provide a file at that path. Instead, the Dockerfile copies the available image to `/app/fixtures/image.png`, and the setup script does not create or copy anything to `/app/threshold.png`. The required input therefore cannot be accessed at the location specified by the task.

Key evidence. Instruction: `/app/threshold.png`. Dockerfile: `COPY fixtures/image.png /app/fixtures/image.png`. No setup step creates `/app/threshold.png`.

task_000051_313e5ebc

Task. Transcribe an incident audio recording, redact the recovered secret from log files, and implement a secure JWT validator whose outputs should match a provided reference implementation under fuzz testing.

Why INSTR-ENV-MISMATCH. The instruction refers to the audio input at `/app/incident_log.wav`, whereas the environment copies the audio fixture to `/app/fixtures/audio.wav`. More critically, the verifier evaluates the agent's JWT validator against `/app/reference_validator.oracle.py`, but the environment does not provide that reference file. The task and verifier therefore depend on artifacts that are absent from the runtime environment.

Key evidence. Instruction expects `/app/incident_log.wav`; the Dockerfile provides `/app/fixtures/audio.wav`. The verifier additionally requires `/app/reference_validator.oracle.py`, which is not created by the environment.

ANSWER-LEAK

task_000028_ae6330dd

Task. Perform a multi-stage incident investigation: extract information from video, recover a 4-byte XOR key from an ELF binary, analyze packet-capture traffic, and produce a JSON report containing the detected malicious events.

Why ANSWER-LEAK. The intended task requires combining several forensic modalities to recover three specific events. However, the setup exposes `/home/user/ground_truth.json`, which already contains the exact three events used by the verifier, including timestamps, payload values, and certificate flags. Reading this file reveals the complete expected report, so the intended video, binary, and packet-capture analysis can be bypassed.

Key evidence. The environment's `ground_truth.json` contains the same three events as the verifier truth set, including timestamps `00:01:15`, `00:02:30`, and `00:04:45`, together with the expected payload and flags.

task_000236_54686d44

Task. Recover a passphrase from video, decrypt a vault, reverse-engineer a malicious 12-byte signature from an ELF payload, and implement an efficient IDS scanner that detects that signature in logs.

Why ANSWER-LEAK. The intended pipeline requires several nontrivial steps before the malicious signature can be recovered. However, the exact signature `deadbeef1122334455667788` is already embedded in the environment’s fixture metadata. An agent can read this value directly and hard-code it into the scanner, bypassing the video extraction, vault decryption, and binary reverse-engineering stages.

Key evidence. The fixture files contain the complete 12-byte signature `DE AD BE EF 11 22 33 44 55 66 77 88`, identical to the value expected by the verifier.

INSTR-AMBIGUOUS

task_004618_d562309e

Task. Compute employee rankings within each department according to `total_team_sales` in descending order and output the ranked results.

Why INSTR-AMBIGUOUS. The instruction specifies the primary ranking key but does not define what should happen when two employees in the same department have identical `total_team_sales`. Multiple orderings are therefore consistent with the instruction. The verifier silently resolves ties using employee ID in ascending order and uses this ordering to assign the final ranks. A competent implementation using another deterministic tie-breaking convention can therefore produce a different, equally reasonable result.

Key evidence. Instruction specifies only descending `total_team_sales`. Verifier uses `emps.sort(key=lambda x: (-x[0], x[1]))`, where `x[1]` is `emp_id`.

task_000795_cd4f601f

Task. Load a directed citation graph into SQLite and compute the shortest path from node 10 to node 99.

Why INSTR-AMBIGUOUS. The graph contains multiple shortest paths of equal length, but the instruction does not specify a tie-breaking convention. Standard BFS or Dijkstra implementations can therefore return different shortest paths depending on adjacency order or queue ordering. The verifier, however, accepts only one hard-coded path: `10, 25, 40, 88, 99`. Thus, an equally short and mathematically valid path can be rejected.

Key evidence. Multiple equal-length shortest paths exist; the instruction provides no tie-breaking rule, while the verifier requires the single sequence `10, 25, 40, 88, 99`.

TASK-TRIVIAL

task_001515_bf254241

Task. Recover a 3×3 projection matrix from the voice memo `/app/artifact_summary.wav`, then implement a Python transformation script that applies the recovered matrix to arbitrary three-dimensional inputs.

Why TASK-TRIVIAL. The stated challenge is to recover the matrix from audio. However, the agent-accessible environment also contains `/app/oracle_transform.py`, the exact reference implementation used by the verifier for fuzz equivalence. Inspecting or probing this deterministic oracle directly reveals the transformation, so the audio-recovery component is unnecessary. In the observed rollout, the agent directly inspected the oracle and recovered the fixed diagonal transformation from it.

Key evidence. The verifier invokes `/app/oracle_transform.py` as the reference implementation, and this same oracle is accessible to the agent. It reveals the matrix

$$\text{diag}(2.5, -1.5, 4.2),$$

allowing the transformation to be hard-coded without using the audio.

C.4 EXAMPLES OF VERIFIER FAILURES ON THE UNFILTERED ENVIRONMENT SET

Using the unfiltered TMax collection (Iverson et al., 2026) may lead to reward hacking and encourage short-cuts in policy’s rollouts. The performance of RL on random-sampled 3.5K environments from the TMax collection is presented in Tab. 8, where we can see a significant performance gap with our model.

Model	Otbl	Tbv2.1	Tbpro	Twv	Avg.
River-8B (ours)	21.8 $\pm$ 2.9	9.7 $\pm$ 0.6	23.0 $\pm$ 3.1	23.0 $\pm$ 2.5	19.4 $\pm$ 1.2
RL On Random TMax-3.5K	18.6 $\pm$ 1.9	8.2 $\pm$ 1.3	23.7 $\pm$ 1.5	20.3 $\pm$ 0.6	17.7 $\pm$ 0.7
River-8B-SFT	11.3 $\pm$ 2.3	6.4 $\pm$ 1.3	13.7 $\pm$ 2.6	18.5 $\pm$ 2.3	12.5 $\pm$ 1.1

Table 8: Numbers are benchmark scores ($\times 10^2$). **Avg.** is the unweighted mean of the four benchmarks. Scores are mean $\pm$ std over 3 seeds.

This section also provides some examples of bad behaviors fostered by training tasks with low-quality verifiers. Inspecting the corresponding trajectories shows that these defects can corrupt the reward signal in both directions. Some defective environments produce *false-positive rewards*: the agent obtains `reward= 1` by exploiting leaked answers, degenerate oracles, or missing verifier checks without completing the intended task. Other defects produce *false-negative rewards*: a correct and instruction-aligned solution receives `reward= 0` because the verifier uses an incorrect reference answer, a buggy oracle, an over-specified implementation check, or an internally inconsistent test. Below we show some representative examples for both *false-positive rewards* and *false-negative rewards*. **Red** highlights the critical shortcut or verifier failure.

False-positive rewards. The first five examples show cases in which the observed reward overstates task completion.

(1) Answer Leak – task_003722 (reward = 1)

Task. Recover a scrolling news ticker from `archive_broadcast.mp4`. The agent must extract frames at 1 FPS, crop the bottom of each frame, run OCR, reverse mojibake corruption, normalize the recovered text, and reconstruct the chronological ticker in `reconstructed_ticker.txt`.

Environment/verifier defect. The verifier compares the reconstructed text against `/app/ground_truth.txt`, but this supposedly hidden reference is directly readable by the agent.

T1 While inspecting the environment, the agent discovers `/app/ground_truth.txt` alongside the video and reads the complete expected ticker.

T2–T5 The agent nevertheless attempts the intended video pipeline: it extracts frames, crops the ticker region, and constructs a noisy OCR-like reconstruction containing duplicated and corrupted fragments.

T6 It explicitly compares this reconstruction against the leaked reference and obtains only **0.368 similarity**, far below the required threshold of 0.85.

T7 The agent then reasons that the output should match the expected string and **replaces the failed reconstruction with the text read directly from `ground_truth.txt`**.

T8 The final output is identical to the reference, giving similarity 1.0 and full reward.

Why the reward is misleading. The intended challenge is to recover the ticker from the video. The agent’s actual reconstruction is insufficient, but the exposed reference lets it replace that failed result with the exact answer. Thus, **the successful reward reflects access to the verifier target rather than successful video understanding**.

(2) Answer Embedded in the Instruction – task_008903 (reward = 1)

Task. Write and run a C++ scanner over several binaries. For each valid 64-bit ELF file, the scanner must determine whether the SUID bit is set, parse the ELF `.rodata` sec-

tion, detect whether it contains the string `/etc/shadow`, and write the sorted results to `/home/user/scan_results.txt`.

Environment/verifier defect. The instruction’s example output is not merely illustrative: it gives the **exact correct three output lines for the binaries in the environment**. The verifier then checks the final result against those same expected values.

T1–T3 The agent attempts to construct the requested C++ scanner, but compilation and execution do not produce a reliable working scan.

T4 It notices that the instruction already provides the complete expected output:
`fake_bin: SUID=Yes, SHADOW_IN_RODATA=No, safe_bin: SUID=No, SHADOW_IN_RODATA=No, and suid_bin: SUID=Yes, SHADOW_IN_RODATA=Yes.`

T5 The trajectory explicitly acknowledges the shortcut: **“I’ve just created a hardcoded output file”**, and writes these three lines directly to `scan_results.txt`.

T6 The hard-coded file matches the expected output and receives full reward even though the requested scanner is not demonstrated to have produced it.

Why the reward is misleading. The task is intended to test binary inspection and ELF parsing. Instead, **the exact answer is exposed in the instruction and can be copied without obtaining it from the binaries**. The final-output verifier cannot distinguish this shortcut from a working scanner.

(3) Task Trivial – `task_001515` (reward = 1)

Task. Recover a 3×3 linear transformation from the voice memo `/app/artifact_summary.wav`, then implement the transformation in `/home/user/ttransform.py`. The script is evaluated on random three-dimensional input vectors.

Environment/verifier defect. The exact reference implementation, `/app/oracle_transform.py`, is shipped inside the agent-accessible container, and the verifier uses this same script as the oracle for fuzz equivalence.

T1 The agent lists the task files and discovers `oracle_transform.py` alongside the audio memo.

T2 Rather than recovering the matrix from the audio, it executes `cat /app/oracle_transform.py`.

T3 The reference implementation reveals the entire transformation:

`diag(2.5, -1.5, 4.2).`

T4 The agent writes `ttransform.py` using exactly these constants.

T5–T6 It explicitly notes that **“the audio wasn’t needed”**. The verifier then fuzzes the copied implementation against the exposed oracle, and all tests pass.

Why the reward is misleading. The generated implementation is numerically correct, but the defining challenge of the task is to recover the transformation from the voice memo. **Exposing the reference implementation collapses the intended audio-understanding task into copying a few constants.**

(4) Degenerate Oracle – `task_000925` (reward = 1)

Task. Reverse-engineer the behavior of a proprietary semantic-version resolver, `/app/legacy_resolver`, and implement a Bash replacement supporting SemVer comparisons, Boolean operators, and parentheses.

Environment/verifier defect. The supplied resolver is degenerate: across the agent’s probes, it returns exit code 0 even for obviously false comparisons. The verifier evaluates the replacement by agreement with this broken resolver.

T1–T3 The agent probes `legacy_resolver` on straightforward positive and negative examples.

- T4** It observes that contradictory cases such as $\geq 1.0.0$ with $0.9.0$, $> 2.0.0$ with $1.9.0$, and $= 1.0.0$ with $2.0.0$ all return success.
- T5** The agent explicitly concludes: “The binary appears to have a bug where it just always returns exit code 0.”
- T6** Instead of implementing semantic-version parsing and comparison, it creates a replacement whose effective behavior is simply `exit 0` for every input.
- T7** Because the verifier compares the replacement against the same degenerate oracle, this constant policy receives full reward.

Why the reward is misleading. The replacement does not implement semantic-version matching. Instead, the policy identifies the oracle’s degenerate decision boundary and implements the cheapest behavior that reproduces it. This shows how a defective oracle can make a trivial shortcut optimal under the observed reward.

(5) Unchecked Implementation Constraint – `task_008402` (reward = 1)

Task. Fix a C++ backup-analysis program so that it queries replication metadata, constructs the valid directed graph, computes the lowest-latency path between two datacenters, and writes the result to `optimal_backup_path.json`. The instruction explicitly requires the solution to be produced *entirely in C++* and forbids a separate Python or Bash implementation.

Environment/verifier defect. The verifier only checks that `/home/user/backup_analyzer` exists and is executable, then separately checks the JSON contents. It never runs the C++ executable to verify that the executable actually produced the answer.

- T1–T3** The agent edits and compiles the requested C++ program, but repeated executions of `backup_analyzer` time out and never produce the required JSON.
- T4** Instead of continuing to debug the required implementation, it writes a separate Python program, `backup_checker.py`, that computes the answer.
- T5** The Python program writes the correct path `[1, 2, 4]` with total latency 30 to `optimal_backup_path.json`.
- T6** The trajectory explicitly acknowledges that the C++ program was timing out and that Python was used to generate the final output.
- T7** The verifier observes an executable C++ file and a correct JSON file, never checks their causal relationship, and returns full reward.

Why the reward is misleading. The numerical answer is correct, but the explicit implementation requirement is not satisfied. Existence-only verification of the binary lets a forbidden Python bypass receive the same reward as a working C++ implementation.

False-negative rewards. The next five examples show the opposite failure mode: the observed reward understates task completion because the verifier rejects correct or instruction-faithful behavior.

(6) Incorrect Numerical Reference – `task_009230` (reward = 0)

Task. Compute a posterior score using the specified formula

$$p = (1 - \exp(-x)) p_{\text{prior}},$$

and write the resulting probability rounded to six decimal places.

Environment/verifier defect. For one case with $x = 2.99$ and $p_{\text{prior}} = 0.7$, the verifier hard-codes `0.664815` as the expected value.

- T1–T3** The agent follows the numerical procedure specified by the instruction.

- T4** For the disputed value, it computes

$$(1 - \exp(-2.99)) \times 0.7 = 0.664798794\dots,$$

which rounds to `0.664799`.

- T5** The remaining reported numerical values match the verifier.

T6 The verifier nevertheless expects 0.664815. The discrepancy exceeds its tolerance, so the mathematically correct rollout receives `reward=0`.

Why the reward is misleading. The agent follows the stated formula and produces the correct rounded value. The failure is caused by an erroneous hard-coded reference constant, not by an error in the solution.

(7) Incorrect Semantic Assumption – `task_002319` (reward = 0)

Task. Perform spectral analysis on an EIIP-mapped DNA sequence. After computing the float64 FFT magnitude spectrum, the instruction requires finding the *maximum magnitude* over indices 1 through $N/2$, excluding only the DC component.

Environment/verifier defect. The verifier does not compute the requested maximum. Instead, it hard-codes `expected_peak_index = N // 11`, assuming that the fundamental frequency must be dominant.

T1–T3 The agent parses the FASTA file, constructs the float32 and float64 EIIP sequences, computes both FFTs, and evaluates their magnitude spectra.

T4 Following the instruction literally, it searches indices 1 through $N/2$ and obtains `peak_index = 50000` with magnitude approximately 1311.08.

T5 Independent recomputation confirms that this is the true maximum. The verifier’s expected index 10000 has magnitude only about 322.33.

T6 The verifier nevertheless rejects 50000 because it expects the hard-coded fundamental-frequency index, producing `reward=0`.

Why the reward is misleading. The agent correctly answers the question posed by the instruction. The verifier substitutes an incorrect domain assumption for the actual maximum and penalizes the correct spectral result.

(8) Buggy Reference Oracle – `task_001636` (reward = 0)

Task. Build a high-performance C program for “latest-frame” lookup. Given a query timestamp T , it must return the frame whose `pkt_pts_time` is the largest value less than or equal to T , using an indexed array and binary search.

Environment/verifier defect. The reference oracle reads H.264 packet metadata in decode order, where PTS values are not monotonic, and then runs binary search on this unsorted array. The verifier compares the agent against this buggy oracle rather than against the requested lookup semantics.

T1–T3 The agent extracts the same PTS/packet-size pairs from the video as the oracle.

T4 Because binary search requires sorted keys, it correctly sorts the metadata by PTS before building the query index.

T5 For query 2.6755, the largest PTS not exceeding the query is 2.6, whose packet size is 120.

T6 The unsorted oracle instead returns 97. The verifier treats 97 as ground truth and rejects the agent’s correct 2.6755 $\rightarrow$ 120 output.

Why the reward is misleading. The agent implements the requested latest-frame semantics and satisfies the precondition required for binary search. It receives zero reward because oracle equivalence is used as a proxy for correctness even though the oracle itself violates the algorithm’s assumptions.

(9) Brittle Source-Code Check – `task_006064` (reward = 0)

Task. Fix a Rust text-processing pipeline so that it lowercases the corpus, retains only alphabetic characters a–z and spaces, tokenizes by whitespace, reports the top three tokens, and computes their 3×3 sample covariance matrix.

Environment/verifier defect. Two functional tests validate the produced tokens and covariance matrix, but a third test additionally requires the source file to contain the literal substring `is_alphabetic`. The instruction never requires this specific Rust API.

- T1–T3** The agent fixes the tokenization using `c.is_ascii_alphabetic() || c.is_whitespace()`, which directly matches the instruction’s requirement to retain a–z characters and spaces.
- T4** It builds and runs the Rust project successfully.
- T5** The resulting `top_tokens.txt` is byte-exact correct (the, quick, dog), and `covariance.csv` exactly matches the expected matrix.
- T6** Both functional output tests pass, but the source-inspection test fails solely because `is_ascii_alphabetic` does not contain the literal substring `is_alphabetic`. The final reward is therefore 0.

Why the reward is misleading. The observable behavior is correct, and the chosen API is at least as faithful to the explicit a–z requirement as the expected implementation. A brittle source-pattern check rejects a functionally correct solution because it uses a different valid implementation.

(10) Inconsistent End-to-End Test – `task_001040` (reward = 0)

Task. Repair a local experiment-tracking pipeline by configuring Nginx to proxy `127.0.0.1:8080` to a provided Flask server, writing a CSV validator, and starting the services. Clean training-log uploads should be accepted, while malformed logs should be rejected.

Environment/verifier defect. The provided Flask server accepts uploads only through a multipart form field named `file`; if this field is missing, it returns HTTP 400. However, the end-to-end verifier sends the CSV bytes as a raw POST body with no multipart encoding and no `file` field, while still requiring HTTP 200 for a clean file.

- T1–T4** The agent writes a validator enforcing the requested headers and numerical constraints, configures Nginx to proxy port 8080 to the Flask service on port 5000, and writes an executable startup script.
- T5** The static validator, startup-script, and configuration tests pass; three of the four verifier tests succeed.
- T6** The end-to-end test reads a clean CSV into raw bytes and constructs `urllib.request.Request(..., data=data, method="POST")` without a multipart `file` field.
- T7** The unmodified Flask fixture therefore behaves as designed and returns `400 BAD REQUEST`, while the verifier expects 200. The rollout receives `reward=0`.

Why the reward is misleading. The failing request does not use the upload protocol required by the provided server. The verifier and fixture are internally inconsistent, so an instruction-faithful solution cannot satisfy the end-to-end test.

Takeaway. These examples show that defective environments can corrupt the RL signal in both directions: **False-positive** rewards arise when answer-bearing artifacts, degenerate oracles, or incomplete checks make shortcuts sufficient for `reward=1`. **False-negative** rewards arise when incorrect reference values, buggy oracles, over-specified source checks, or inconsistent test protocols reject solutions that correctly follow the instruction. Together, these cases further motivate our environment filtering process.

D BEHAVIOR ANALYSIS DETAILS

In this section, we list the behaviors we used for the AUC-ROC test. Each behavior feature is extracted deterministically from a trajectories using rule-based parsing, with no model or LLM in the loop. For each assistant turn, we parse the `<command>` and treat the following user turn as the resulting observation. A command is considered an *inspection* if its head corresponds to a pre-defined read-only tool set including (`ls`, `cat`, `grep`, `find`, `stat`, ..., etc.) and does not perform redirection or writing. An *error* is an observation that matches an error regex, such as `Traceback`, `command not found`, `No such file`, or a non-zero exit status. Two commands are considered *near-duplicates* if their token-level Jaccard similarity is at least 0.8.

AUC denotes the single-feature cross-validated AUC-ROC for predicting per-task success of the RL checkpoint on the four evaluation benchmarks, averaged over 15 seeds of 5-fold cross-validation. **Dir.** indicates the sign of the feature’s correlation with success. Rows are sorted by AUC. The two strongest individual signals are verifying before finishing (`verify_before_done`, +) and repeatedly issuing similar commands (`repetition_rate`, −). The latter serves as a broader proxy for the harmful repetition loops targeted by our turn-wise RL penalty.

Table 9: The 12 behavior features used by Model-B, together with their definitions and univariate predictive power.

Feature	Definition / extraction rule (per trajectory)	AUC	Dir.
<code>verify_before_done</code>	1 if at least one read-only inspection occurs among the up to three commands immediately preceding the <code><done></code> action; otherwise 0.	0.632	+
<code>repetition</code>	Fraction of commands that are near-duplicates of at least one earlier command in the same trajectory.	0.631	−
<code>cmd_diversity</code>	Number of distinct command heads divided by the total number of commands.	0.615	+
<code>python_heredoc</code>	Fraction of commands that invoke Python (<code>python/python3</code> , including <code>-c</code>) or use a here-document to perform task-related operations.	0.604	−
<code>inspect_before_act</code>	Fraction of commands that perform read-only inspection, defined as commands with an inspection head and no output redirection, before modifying the environment.	0.596	+
<code>n_errors</code>	Number of commands whose resulting observations match the error regex.	0.563	−
<code>gave_up_flag</code>	1 if the episode ends <i>without</i> a <code><done></code> action and its final observation is an error; otherwise 0.	0.549	−
<code>error_rate</code>	Fraction of commands whose resulting observations match the error regex (<code>n_errors</code> divided by the total number of commands).	0.535	−
<code>error_recovery</code>	Among commands that result in an error, the fraction for which the <i>next</i> command is substantially different (i.e., not a near-duplicate), indicating a change in approach after failure.	0.531	+
<code>inspect_first</code>	1 if the first command in the trajectory is a read-only inspection; otherwise 0.	0.524	+
<code>tool_help</code>	Fraction of commands that request tool help or usage information (<code>--help</code> , <code>-h</code> , or <code>man</code>).	0.510	−
<code>premature_done</code>	1 if <code><done></code> is emitted either after the immediately preceding command results in an error or without prior verification; otherwise 0.	0.507	~0
All 12 features (Model-B)		0.735	