# TMAX: A SIMPLE RECIPE FOR TERMINAL AGENTS

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Terminal-using agents have quickly become a popular and important downstream application of language models (LMs). Despite their prevalence, relatively little academic work has examined RL-based training of these models, due to difficult benchmarks, a lack of data, and a lack of simple baseline recipes. We present TMAX, an open RL recipe for terminal agents that brings small open models closer to the frontier. While simple, our recipe achieves **33% on Terminal-Bench 2.0 with only 9B parameters**, outperforming much larger models from prior work. Our recipe first involves generating data using a novel taxonomy, combining difficulty control, personas, and verifier diversification, which allows us to cheaply generate large amounts of terminal environments for RL and SFT training. We then train open-weight models using RL with our data, using a simple, outcome-only recipe, with small adjustments to improve stability. Our recipe works for models from 2B to 27B and across model families, dominating the Pareto curve for open terminal-agent performance at smaller sizes. We will publicly release our data, models, and code as a strong baseline for future open academic work on terminal agents.

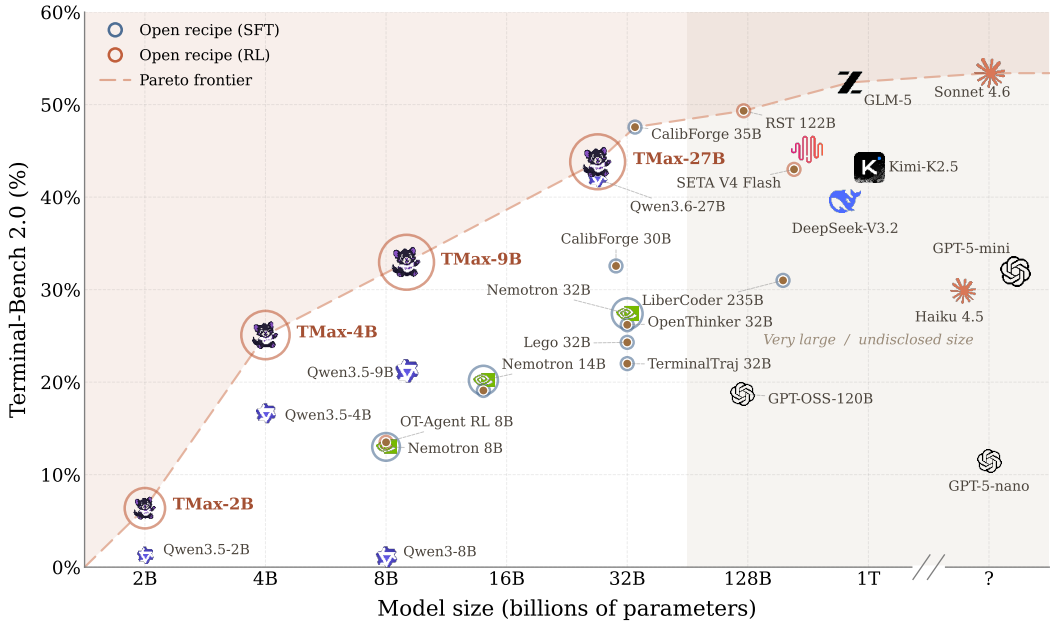


Figure 1: Terminal-Bench 2.0 performance of TMAX alongside prior and concurrent models. TMAX and Qwen use our harness; other scores are self-reported or from the leaderboard. **TMAX-9B dominates the Pareto curve for models under 10B parameters.** See Table 15 for full results.

## 1 INTRODUCTION

Terminal-based agentic coding products have quickly become incredibly popular, with models expected to perform increasingly complex and long-running tasks through the terminal (Anthropic,



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

2025; Cursor Research et al., 2026). However, existing academic work has largely focused on bug-fixing (Jimenez et al., 2024) or relatively simple terminal tasks (OpenThoughts-Agent Team, 2025; Gandhi et al., 2025), rather than the complex, long-horizon tasks exemplified by Terminal-Bench (Merrill et al., 2026). We aim to close this gap by providing a large dataset of complex terminal tasks and a recipe for training small yet capable open-weight terminal agents.

We first introduce TMAX-13K, a dataset of RL environments spanning diverse domains, skills, and difficulty levels. We explicitly control task difficulty and use continuous rewards where appropriate, while keeping the generation process simple and scalable. Compared with many prior datasets, TMAX-13K is substantially more challenging, as measured by lower pass rates from Gemini-3-Flash-Preview. While we get the strongest results using stronger models as data generators, we also show that training on data generated by the same model can still provide strong improvements, showing that our pipeline is effective in the absence of a stronger teacher model. Secondly, while recent work has focused primarily on generating terminal-agent data (Wu et al., 2026b; Zhu et al., 2026), RL training for terminal agents remains underexplored. We address this gap by developing a simple RL recipe for training open-weight terminal agents on TMAX-13K. Our recipe builds on fully asynchronous RL infrastructure implemented in open-instruct (Lambert et al., 2025) and uses DPPO (Qi et al., 2026) for policy optimization. To improve training stability, we use an FP32 LM head and increase the group size. **Our best 9B model achieves 33% on Terminal-Bench 2.0**, outperforming similarly sized open models and selected closed models, including Claude Haiku 4.5.

We hope that our overall recipe and data prove useful for future academic work exploring terminal agents, as we believe this setting provides a number of unique and interesting challenges. Concretely, our contributions are:

- We introduce TMAX-13K, a dataset of 13,565 RL environment instances with explicit control over difficulty and diversity. Our data is generated using a simple and scalable procedure, avoiding complex multi-agent loops.
- We train open-weight terminal agents and achieve state-of-the-art performance among reported models under 10B parameters under default Terminal-Bench settings. Our best 9B model reaches 33% on Terminal-Bench 2.0.
- We provide a simple, reproducible open RL recipe for achieving this performance and will publicly release all elements required for training our models. Our recipe outperforms prior open RL recipes such as Endless Terminals and OpenThinker-Agent.
- We show that terminal-based RL training can generalise non-trivially across harnesses and tasks, providing strong evidence our training improves general model capabilities.

We will publicly release our code, checkpoints and data for future use and reproduction.

## 2 Background and Related Work

**Data for terminal agents.** Early terminal-agent datasets paired short instructions with shell commands (Lin et al., 2018). More recent work either constructs environments from existing repositories or questions (Yang et al., 2025; Wu et al., 2026b; Lin et al., 2026; Yang et al., 2026; Shen et al., 2026b; Shi et al., 2026), or synthesises them using seed tasks and combinatorial prompting (Gandhi et al., 2025; Zhu et al., 2026; Peng et al., 2026; Li et al., 2026; Meng et al., 2026). We follow the latter approach. We find that repository-based construction often produces imbalanced domain coverage, whereas synthesis provides greater control over domain balance. Compared with prior work on task synthesis, we vary tasks along more axes, including task difficulty, and release substantially more data. We also incorporate a repair stage and show that strong performance can be achieved without a commonly used but expensive validation step. We discuss prior and concurrent work on terminal RL data in more detail in § B.

**RL for terminal agents.** Closest to our work is Endless Terminals (Gandhi et al., 2025), which performs RL training on a Qwen 3 8B model using its generated data, but with a significantly simpler harness and setup. OpenThoughts-Agent (Raoof et al., 2026) explores RL training for terminal agents, but focuses on comparing existing datasets and the effect of varied SFT, finding that RL dataset quality matters significantly. More recently, SETA (Shen et al., 2026b) and ROME (Wang et al., 2026b) explore terminal RL training, but show limited gains or use fairly complex setups

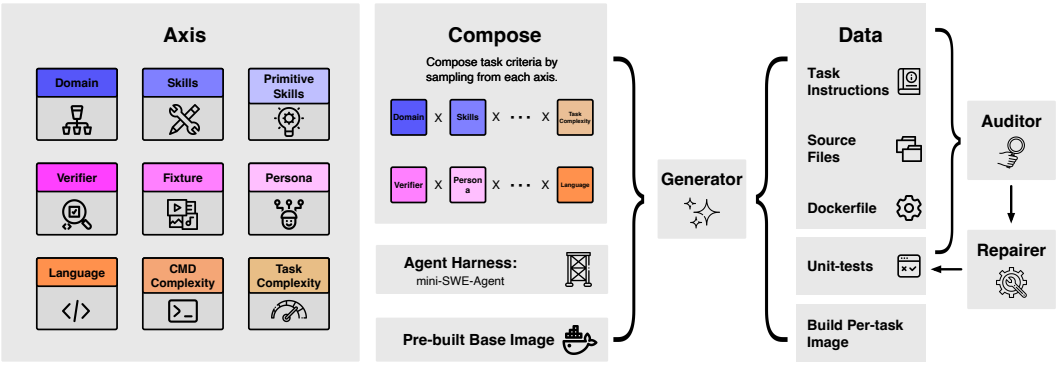


Figure 2: **TMAX Data Pipeline**. Each task is composed by hierarchically sampling from 9 structured axes, after which a *generator* instantiates a task as a Dockerfile, unit-test verifier, source files, and task instructions. An *auditor* then labels whether each verifier actually validates its instruction—flagging answer leakage, instruction–verifier mismatch, and overly weak checks—and a *repairer* rewrites the verifier where fixable. Unrepairable tasks are dropped. Tasks are built atop a pre-built per-domain base image and served through a mini-SWE-agent harness. Composing axes yields combinatorially many task signatures with explicit, per-axis control over difficulty and diversity. The single build step, with audit and repair as a one-shot post-pass, keeps generation cheap to scale.

to train their models. In contrast, we show large gains and use a simple, straightforward recipe. River (Yao et al., 2026) examines RL training data, finding that rubric-based repair and filtering of existing datasets can improve performance. It also investigates process-level behaviour regularisation. In contrast, we generate our own data, and use fairly straightforward outcome rewards.

## 3 Terminal Data Generation

Training autonomous terminal agents requires large-scale and diverse coding tasks. With real-world terminal data often scarce or proprietary (Merrill et al., 2026), synthetic generation is used to bridge the data gap (Pi et al., 2026; Zhu et al., 2026). However, existing pipelines fall short along three axes. First, they rely on complex multi-stage generation procedures that are hard to scale. Second, they generate homogeneous task suites with limited coverage against real-world tasks and Terminal-Bench (Merrill et al., 2026). Third, they offer little control over difficulty, often producing bi-modal task pools that are either trivially solved or unsolvable by current models. Hence, we propose a simple terminal environment generation pipeline that is scalable, diverse, and difficulty-aware. We use Gemini-3-Pro (Pichai et al., 2025) as our generation model, as it requires a strong enough model to generate coherent tasks and code fixtures.

### 3.1 Generation Pipeline

We adopt a simple compositional generation framework: each synthetic task is sampled as a product of structured axes (Table 7). We follow Pi et al. (2026) to seed the first three axes – domains, skills, primitive skills – and introduce six orthogonal axes explicitly targeting diversity and difficulty, unlike prior work.

**Diversity via hierarchical sampling.** The compositional sampler is itself our main diversity mechanism, and drawing from each axis yields combinatorially many distinct task signatures. Further, we add two specific diversity dimensions:

*Persona diversification.* We augment generation with a set of user personas, motivated by prior work on persona-conditioned data synthesis (Ge et al., 2025; Lambert et al., 2025), to further diversify generated tasks. Personas are domain-specific (6–18 per domain, e.g. "red-team operator crafting an evasion payload" for security), so each is always reasonable for its sampled skill set.

*Multi-modal fixtures.* Prior work on RL environments largely focuses on *text-in / text-out* problems. Here we extend the task inputs beyond plain text by shipping a concrete artifact for some tasks: a

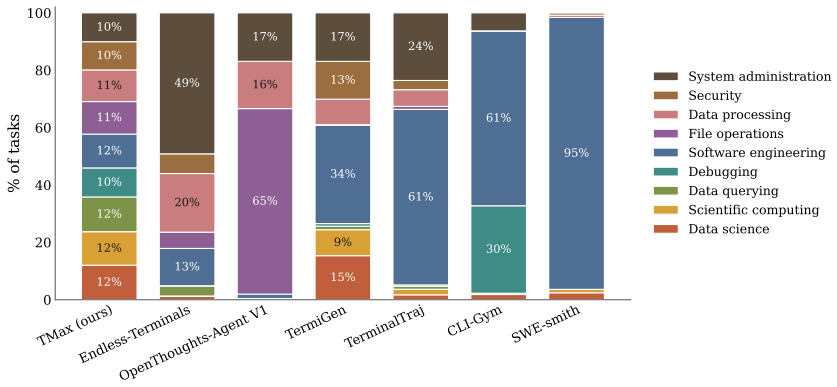


Figure 3: **Data Composition**. Domain distribution of tasks across terminal datasets. Prior datasets skew heavily toward one or two domains, whereas our compositional sampler yields balanced coverage across all nine domains. In-bar percentages are rounded to the nearest whole number. OpenThoughts Agent refers to the original OpenThoughts data (OpenThoughts-Agent Team, 2025).

PNG image, audio file, video, stripped binary, vendored package, or multi-service compose stack (Table 9). The policy itself remains a text-only language model and training is unchanged: rather than perceiving these artifacts natively, the agent inspects them through standard terminal tooling (e.g., OCR, audio transcription, `ffmpeg`), so no multi-modal model is required.

**Difficulty via explicit calibration.** Unlike previous works, we make an explicit effort to calibrate task difficulty, avoiding the bi-modal distribution where tasks are either trivially solved or unsolvable by current models (Table 1). We address this through two mechanisms:

*Fine-grained complexity.* We introduce two complexity axes, *task complexity* (from a few shell commands to intricate workflows of 30–60 commands) and *command complexity* (from bash-only to bash + code + system services), for granular control over how hard each task is. We sample uniformly across complexity buckets by default, with optional per-bucket up-weights to match specific model capabilities or induce a curriculum.

*Graded verifiers.* Previous RL tasks often rely on exactly matching a gold output. Here we extend with richer verifier kinds (Table 8): metric-threshold (e.g., accuracy $\geq 0.95$), adversarial-corpus (accept clean, reject malicious), fuzz-equivalence (bit-exact match against an oracle), or multi-protocol (protocol-level requests against a service). Thresholds give us a continuous difficulty knob, while multi-condition variants naturally extend task length.

**Verification via LM Judge.** Finally, inspired by Yao et al. (2026), we apply an LM judge (GPT 5.6 Sol) to classify our data and use it to identify weak, ambiguous, or contradictory verifiers and task descriptions. The vast majority (76.4%) of tasks are marked as having verifiers too weak to enforce the correct solution, and 10.9% of tasks are judged as leaking the answer in the initial environment state, highlighting that a repair pass is much-needed to improve task validity. We then provide the original environment definition along with the label to GPT 5.6 Sol and construct a repaired instance. We then do a final spot-check to ensure that images have the required dependencies and can be built (see §C.2 for the pre-repair rubric counts). We retain tasks originally judged as correct and repaired tasks.

**Do we need to filter?** Unlike prior works that validate task quality and correctness through expensive teacher generation (Zhu et al., 2026; Wu et al., 2026b), we deliberately skip this step entirely, as our RL training applies effective soft filtering via active sampling and zero-std filtering (§4.1). We find the all-zero rate is roughly 40% (see §E.4).

Using this pipeline, we generate a dataset of 13,565 tasks, which we name TMAX-13K. We measure overlap between dataset task descriptions and the Terminal-Bench and TB Lite tasks using a 13-gram sliding window, and find no sign of overlap (§C.5).

| Data | Dataset size | Pass@1 Gemini | Pass@4 Gemini | Pass@8 Gemini | Mean turns | Mean tokens / run | Domain balance | Skill-type balance |
|---|---|---|---|---|---|---|---|---|
| **TMAX-13K (Ours)** | 13k | 49% | 57% | 60% | 31.5 | 446K | 0.997 | 0.738 |
| Endless Terminals | 2.4k | 92% | 94% | 95% | 15.7 | 117K | 0.481 | 0.284 |
| Terminal Gen | 3k | 57% | 64% | 66% | 20.9 | 305K | 0.646 | 0.477 |
| Terminal Traj | 5.5k | 54% | 63% | 65% | 15.1 | 141K | 0.363 | 0.374 |
| CLI-Gym | 1.5k | 41% | 52% | 55% | 43.8 | 769K | 0.283 | 0.061 |
| SWE-Smith | 59k | 54% | 69% | 72% | 41.2 | 582K | 0.146 | 0.042 |

Table 1: **Terminal datasets and difficulty-related statistics** with Gemini-3-Flash-Preview, evaluated on a fixed subsample of 250 tasks per dataset with 8 rollouts each. Pass@$k$ is mean pass@$k$ across tasks. Mean tokens/run is the sum of tokens over turns, in thousands. For domain and skill-type balance, higher indicates more even category coverage with a maximum of 1.

### 3.2 COMPARING TERMINAL DATASETS

We compare our generated data against past work, both against a subset of prior terminal agent datasets, and a prior SWE-bench-targeting dataset, SWE-Smith (Yang et al., 2025).

**Composition** A good training dataset should have broad coverage and a balanced composition across target domains. By design, our generation framework exposes domain as an explicit sampling axis, so we can directly calibrate per-domain mass at generation time. To compare across datasets, we use Gemini-3-Pro to annotate every task for our and six comparison datasets, with one of these nine domain labels. As shown in Figure 3, prior terminal datasets concentrate 34–95% of their mass on a single domain (e.g., `software_engineering` accounts for 95% of SWE-Smith and over 60% of TerminalTraj and CLI-Gym), while ours spreads roughly uniformly across all nine. This balanced coverage avoids biasing the policy toward any narrow domain and better covers real-world usage.

**Difficulty & Balance** To compare difficulty across datasets, we evaluate Gemini-3-Flash-Preview on a fixed subsample of 250 tasks per dataset (8 rollouts each) and report mean pass@$k$ in Table 1. We additionally compute a *balance score* for each categorical axis (domain, skill-type, task complexity, command complexity) as the fraction-of-uniform diversity of its empirical distribution:

$$\text{Balance} = \frac{\exp(H)}{N}, \qquad H = -\sum_{i=1}^{N} p_i \log p_i, \tag{1}$$

where $p_i$ is the proportion of tasks in bucket $i$ and $N$ is the number of buckets. The score lies in $[1/N, 1]$ with 1 being perfectly uniform; see §C.4 for derivation. This highlights that our data requires significantly more turns than other datasets apart from CLI-Gym and SWE-Smith, but is significantly more balanced than them. Difficulty-wise, we observe that it has the lowest pass@8 apart from CLI-Gym, suggesting genuine difficulty and not high variance in rollout outcomes. While some of these tasks may be unfeasible for the model, we do not attempt to filter them out ahead of time as our RL infrastructure already filters out zero pass rate instances.

## 4 TRAINING TERMINAL AGENTS

We now validate our data by practically applying it to RL training. We aim to validate our data with a simple recipe, providing a testbed in which we believe there is significant room for improvement.

### 4.1 EXPERIMENTAL SETUP

**Algorithm** We train models using DPPO (Qi et al., 2026), a variant of GRPO (Shao et al., 2024) that masks tokens when inference and training logprobs deviate. Specifically, we mask tokens based on a binary approximation of the total variation (TV) divergence. We use a token level loss (Yu et al., 2025) and run training fully asynchronously. We filter groups with zero standard deviation and use active sampling to ensure full batches following Olmo 3 (Olmo et al., 2026).

**Infrastructure** We extend open-instruct (Lambert et al., 2025) for terminal agent training. For rollouts, we use vLLM (Kwon et al., 2023), and we use the same mini-SWE-agent-based harness as used for verification and SFT data generation. We build a custom sandboxing infrastructure utilising slurm for requesting additional compute resources, and Apptainer for setting up and isolating environments. We additionally ensure that the language model head of models is kept in FP32 precision to minimise training-inference mismatch (MiniMax et al., 2025) (see §E.3). We train on a cluster of H100 nodes, and use 2 nodes for training and 6 for inference. Training typically takes 1–2 days. We remove a small number of training data samples and eval samples that are unable to run correctly in our infrastructure (see §H).

**Harness.** Unless otherwise stated, we use a simple harness based on mini-SWE-agent (Yang et al., 2024) with persistent shell. We find that Terminus-2 harness (Merrill et al., 2026) is more brittle with small models, as it requires agents to send raw keystrokes to interact with the terminal (see §D).

**Evaluation** We primarily evaluate on Terminal-Bench 2.1 (Merrill et al., 2026) and Terminal-Bench Lite (OpenThoughts-Agent Team, 2026). We use Harbor (Harbor Framework Team, 2026) and our custom sandboxing infrastructure for evaluation. Due to compatibility issues, we filter a small number of samples from Terminal-Bench Lite and Terminal-Bench 2.1 which are unsupported by our sandbox (see §H). We run 3 trials for all evaluations and report mean and standard error. We run all evals using a single H200 and 4 concurrent trials. We do not override timeout defaults.

For the evaluations in Fig. 1, we use Terminal-Bench 2.0 (the earlier version) to compare directly with past work. We additionally use Daytona as the backend as recommended by Terminal-Bench authors to make results more reproducible. We find that the choice of backend can non-trivially affect performance: for instance, Daytona sandboxes often perform installs faster than our locally-hosted runs, resulting in fewer timeouts. However, the cost of these services makes them expensive to develop against.[1]

**Hyperparameters** Unless otherwise stated, we run RL training for 500 steps and evaluate every 100 steps. For all runs, we report scores on the checkpoint scoring the highest on Terminal-Bench Lite to avoid unfairly biasing against some datasets that collapse later during training in our setup (e.g., CLI-Gym). We use a group size of 32 and 8 prompts per batch, with a maximum sequence length of 65536 tokens and 64 maximum tool calls. When training Qwen 3.5 and 3.6 models, we do not perform any initial SFT. We edit chat templates such that thinking is preserved on intermediate turns (often called 'interleaved thinking'), which has been shown to improve performance in agentic settings (MiniMax, 2025). We show further hyperparameters in §E.1.

### 4.2 Core Results

**TMAX-13K outperforms other terminal datasets.** We first apply RL training to Qwen 3.5 9B (Qwen Team, 2026) using a variety of datasets from prior work and show results in Table 2. We include prior datasets focusing on Terminal-Bench alongside SWE-Smith, as a representative sample of a dataset focusing on SWE-Bench tasks. We note that in many cases, we are the first to openly apply these datasets for RL training (e.g., in the case of TermiGen and TerminalTraj). We call the model trained on our data TMAX-9B. We find that training on TMAX-13K matches or exceeds training on other datasets, and achieves the highest mean performance across both datasets.

To further investigate this, we plot the average number of steps taken by the model during training for each dataset in Fig. 4. We find that some lower scoring datasets (OpenThoughts) use fewer turns during training or even collapse in turn count over the course of training, indicating lower difficulty. In contrast, TMAX-13K remains stable at requiring 25-30 steps to solve, maintaining high difficulty throughout training. While easier datasets like Endless Terminals also maintain higher step counts, they do so at the cost of filtering increasingly more fully-solved samples during training (see §E.4), requiring up to seven times as many rollouts as the minimum needed to produce a single batch.

We additionally plot reasoning and tool-call token usage in Fig. 5. We observe token usage increasing over the initial 150 training steps before plateauing (and then becoming somewhat unstable

[1] One RL training run of TMAX-9B would cost roughly $3,150 on Daytona, using average runtimes from the run and pricing from `https://www.daytona.io/pricing`.

| RL Dataset / Model | TB Lite | TB 2.1 |
|---|---|---|
| None (Qwen3.5-9B) | $49.4 \pm 1.9$ | $20.5 \pm 1.7$ |
| TerminalTraj (Wu et al., 2026b) | $52.5 \pm 1.9$ | $15.9 \pm 1.1$ |
| TermiGen (Zhu et al., 2026) | $54.7 \pm 2.0$ | $21.2 \pm 1.2$ |
| LiteCoder (Peng et al., 2026) | $56.5 \pm 1.8$ | $22.0 \pm 1.4$ |
| OpenThinker-Agent (Raoof et al., 2026) | $59.4 \pm 2.1$ | $22.7 \pm 1.6$ |
| CLI-Gym (Lin et al., 2026) | $61.6 \pm 1.9$ | $24.6 \pm 1.6$ |
| SETA (Shen et al., 2026b) | $59.2 \pm 1.8$ | $26.5 \pm 1.8$ |
| CalibForge (Meng et al., 2026) | $59.9 \pm 2.1$ | $26.5 \pm 1.7$ |
| Terminal-Lego (Yang et al., 2026) | $63.2 \pm 2.0$ | $26.5 \pm 1.8$ |
| RST (Li et al., 2026) | $63.4 \pm 1.8$ | $27.3 \pm 1.6$ |
| SWE-Smith (Yang et al., 2025) | $60.2 \pm 1.9$ | $27.3 \pm 1.5$ |
| Endless Terminals (Gandhi et al., 2025) | $62.0 \pm 1.9$ | $28.4 \pm 1.6$ |
| Facet (Shi et al., 2026) | $61.3 \pm 2.1$ | $28.8 \pm 1.6$ |
| **TMAX-13K (Ours)** | $\mathbf{65.6 \pm 2.0}$ | $\mathbf{29.6 \pm 1.8}$ |

Table 2: **Comparing different RL datasets on Terminal-Bench.** We train on varied prior and concurrent RL datasets using our recipe and evaluate on Terminal-Bench Lite and 2.1. We find that TMAX-13K achieves the highest mean score on both evaluations.

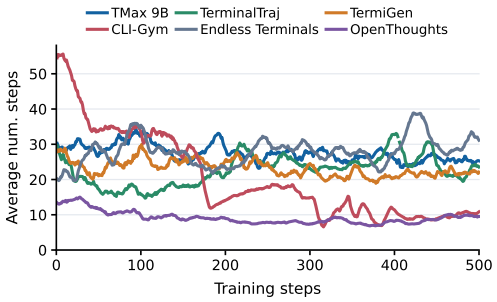


Figure 4: Average step count over RL training when training Qwen 3.5 on different datasets. We smooth using a 15-step average window. **Training on TMAX-13K consistently uses higher steps throughout training than other datasets**, apart from Endless Terminals.

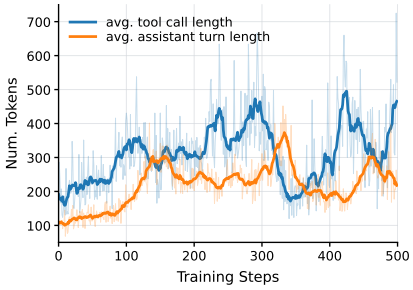


Figure 5: Average length (in tokens) of assistant turns and tool calls. We do not include tool calls in the assistant turn length. Over the initial 150 training steps, **model output lengths increase alongside its performance**.

later in training). This is reminiscent of inference-time reasoning scaling in single-turn math settings (Guo et al., 2025), and suggests that the model is gradually learning more complex reasoning and more complex tool calls over the course of training, resulting in its improved performance.

**TMAX-9B outperforms open RL recipes on Terminal-Bench 2.0.** We additionally compare our model to prior work's self-reported scores on Terminal-Bench 2.0 in Fig. 1. This shows that even when comparing across different harnesses and training setups, TMAX-9B performs the strongest among models under 10B parameters. In particular, our model outperforms many past open RL recipes such as Endless Terminals (Gandhi et al., 2025) and OpenThoughts-Agent (Raoof et al., 2026). As such, our work provides a strong baseline for future open work on agentic terminal RL.

**TMAX RL training improves models across different sizes.** We additionally apply our recipe, without modification, to the other sizes of Qwen 3.5 (and Qwen 3.6), resulting in TMAX-2B, TMAX-4B and TMAX-27B. We show evaluation results in Table 3. In all cases, we improve over the baseline in Terminal Bench 2.1 scores, highlighting that TMAX-13K contains a wide enough range of difficulty levels to provide training signal even at small (2B) model sizes. We observe some degradation in TB-Lite for the 27B size model, which we attribute to the much-improved performance of Qwen 3.6 relative to the 3.5 series, making it harder to further improve.

| Model | TB Lite | TB 2.1 |
|---|---|---|
| Qwen 3.5 2B | $6.0 \pm 0.9$ | $1.9 \pm 0.5$ |
| + RL (TMAX-2B) | $\mathbf{17.8 \pm 1.5}$ | $\mathbf{5.3 \pm 0.9}$ |
| Qwen 3.5 4B | $33.6 \pm 1.8$ | $15.5 \pm 1.5$ |
| + RL (TMAX-4B) | $\mathbf{54.3 \pm 1.9}$ | $\mathbf{24.6 \pm 1.8}$ |
| Qwen 3.5 9B | $49.4 \pm 1.9$ | $20.5 \pm 1.7$ |
| + RL (TMAX-9B) | $\mathbf{65.6 \pm 2.0}$ | $\mathbf{29.6 \pm 1.8}$ |
| Qwen 3.6 27B | $\mathbf{78.1 \pm 1.5}$ | $42.8 \pm 1.7$ |
| + RL (TMAX-27B) | $76.9 \pm 1.6$ | $\mathbf{46.2 \pm 1.9}$ |
| Qwen 3 8B | $5.0 \pm 1.0$ | $1.1 \pm 0.5$ |
| + SFT | $9.1 \pm 1.3$ | $5.3 \pm 1.1$ |
| + RL | $\mathbf{17.8 \pm 1.1}$ | $\mathbf{6.4 \pm 1.0}$ |

Table 3: **TMAX-13K improves models across different sizes and families.** We train Qwen 3.5 and Qwen 3 models on TMAX-13K across sizes, and find it almost always improves performance on both Terminal-Bench benchmarks. We evaluate Qwen 3 8B models with a 40,720-token max len.

| Model | Harness | SBV | AIME |
|---|---|---|---|
| Qwen 3.5 9B | None | – | $67.5 \pm 4.9$ |
| Qwen 3.5 9B | Ours | $46.2 \pm 0.8$ | $72.2 \pm 2.3$ |
| + RL (TMAX) | Ours | $\mathbf{55.3 \pm 0.7}$ | $\mathbf{93.3 \pm 1.5}$ |

Table 4: **Improved performance generalises across tasks.** Performance of Qwen 3.5 9B on SWE-bench Verified (SBV) and AIME'24/25 with and without our harness.

### 4.3 TERMINAL AGENT TRAINING CAN GENERALISE

We next investigate how well TMAX-9B generalises across three important axes: tasks (i.e., what it is asked to do), harnesses (i.e., what tools and prompts it is provided when performing tasks), and model families (i.e., the starting point model).

**TMAX RL training generalises to other tasks.** We evaluate TMAX-9B on SWE-Bench Verified (Jimenez et al., 2024) and AIME 2024/2025, and additionally compare to AIME performance in a traditional single-turn no-harness setting in Table 4. Encouragingly, we find that performance improves across the board by significant margins, suggesting that the model is not simply fitting to the harness and domain, but also learning how to better make use of its terminal tools to solve generic problems. In particular, we observe that TMAX-9B works out how to use the terminal to aid it with geometry problems in AIME, unlike Qwen 3.5 9B. This aligns with prior work showing that terminal-agent training can improve general model capabilities (Cheng et al., 2026).

**TMAX RL training generalises to other harnesses.** We evaluate Qwen 3.5 9B and TMAX-9B on Terminal-Bench Lite using varied harnesses in Table 5. TMAX-9B improves by at least 8 points across all harnesses, although it performs best with our harness. This suggests that single-harness terminal RL training can generalise across other setups, contrary to recent work (Wang et al., 2026a).

**TMAX RL training improves different model families.** We also apply our recipe to Qwen 3 8B after a warm-start on our SFT data (§G.1). As seen in Table 3, we still observe Terminal-Bench improvements (especially Terminal-Bench Lite) despite the significantly weaker base, highlighting the transferability of our recipe.

Taken together, these results strongly suggest that **TMAX RL training teaches the model new capabilities and improved terminal use**, beyond overfitting on Terminal-Bench or our harness.

| Model | Ours | OpenHands | mini-SWE-agent | Terminus-2 |
|---|---|---|---|---|
| Qwen 3.5 9B | $49.4 \pm 1.9$ | $29.0 \pm 1.6$ | $44.9 \pm 2.0$ | $37.9 \pm 2.2$ |
| TMAX-9B | $\mathbf{65.6 \pm 2.0}$ | $\mathbf{42.7 \pm 1.7}$ | $\mathbf{57.8 \pm 2.0}$ | $\mathbf{46.1 \pm 2.1}$ |

Table 5: **Terminal-Bench Lite performance across harnesses.** We find that TMAX-9B improves over its base even when evaluated using different harnesses from training.

| **Model** | **TB Lite** | **TB 2.1** |
|---|---|---|
| Qwen 3.5 9B | $49.4 \pm 1.9$ | $20.5 \pm 1.7$ |
| + RL (pre-repair) | $59.8 \pm 1.8$ | $27.3 \pm 1.4$ |
| + RL (TMAX-13K) | $\mathbf{65.6 \pm 2.0}$ | $\mathbf{29.6 \pm 1.8}$ |
| + Qwen 3.5 9B-Generated RL (3.1k samples) | $62.0 \pm 2.0$ | $27.3 \pm 1.6$ |

Table 6: **LM-judge repair improves downstream performance.** RL training of Qwen 3.5 9B on TMAX-13K before and after the LM-judge verification and repair pass (§3.1), and using RL data generated by Qwen 3.5 9B itself (§5).

## 5 Analysis & Discussion

**Qwen 3.5 does not benefit as much from existing SFT datasets.** Common post-training pipelines usually perform SFT training to 'warm-start' the model before RL training (Guo et al., 2025; Olmo et al., 2026). However, we find that existing datasets seem to degrade Qwen 3.5's performance by as much as 10 points, likely due to its already strong performance. In contrast, the same data can raise Qwen 3 8B's performance by as much as 9 points (see §G.2). We leave further exploration of good SFT mixtures for Qwen 3.5 as future work.

**Terminal agent RL training is difficult to stabilise.** Over the course of this work, we often found training unstable, with runs often collapsing past 300 steps. We apply three changes to improve stability: (1) using an FP32 LM head, which reduced trainer-inference mismatch (§E.3); (2) using larger group sizes, which alleviated collapse in initial setups (§E.1); (3) improving our sandboxing infrastructure. Surprisingly, we found that reducing noise from infrastructure could be the difference between collapses at 300 steps and smooth training to 1000 steps.

**LM judge-based repair significantly improves data.** We compare training on our final dataset against the version before applying the final verification and repair step in Table 6. We find that the repair step significantly improves the performance of models trained on it, likely due to reducing the likelihood that the model relies on shortcuts in the task (due to weak verifiers).

**Models can self-improve using their own synthetic data.** We rerun our data generation pipeline (including classification and repair) entirely using Qwen 3.5 9B to construct a smaller dataset of 3.1k samples. In Table 6 we find that it still results in non-trivial improvements, even outperforming over half the baseline datasets compared in Table 2. This suggests that **our data pipeline and recipe do not require a strong teacher model, but can in fact be used with a model of similar strength to the policy model.** This may be due to information asymmetry: for many tasks, constructing the environment may be simpler than solving it (e.g., creating a sudoku solving task).

## 6 Conclusion

We present TMAX, a simple recipe for training terminal agents: TMAX-13K, a dataset of 13,565 RL environments generated with explicit control over difficulty and diversity, paired with an outcome-only RL recipe. TMAX-9B achieves state-of-the-art Terminal-Bench 2.0 performance among open models under 10B parameters, and our recipe improves models across sizes, model families, harnesses, and out-of-domain tasks such as SWE-Bench Verified and AIME. We will release all data, models, and code as a baseline for future open work on terminal agents.

## AI DISCLOSURE

In this work, we used AI tools to generate and classify synthetic data and aid with the implementation of code. We reviewed all AI code through author inspection and thorough tests. AI tools were also used to aid with writing and figure creation, mainly in helping with refining human-written first drafts and identifying grammatical or typographic errors. All text and figures have been scrutinised by authors. We take responsibility for all claims and errors in our code and paper.

## ETHICS STATEMENT

While we present open artefacts for training stronger models, we focus specifically on smaller, simpler setups that both are significantly less powerful than larger pre-existing models, and less powerful than many available industry offerings, and do not aim to achieve SOTA with large models. Our aim is to foster better understanding and study of terminal agents, considering their widespread use and potential harm (e.g., rogue hacking of government systems). We hope that our novel and simple pipeline and open artefacts provide ways to understand and better control future terminal agent systems.

## REPRODUCIBILITY STATEMENT

We will release all checkpoints, code, and data associated with this paper post-review to allow reproduction of our evaluations and training runs. We have attempted to provide significant detail in our work on our setup and reference specific codebases we build upon (§4.1). Where possible, we reuse existing open-source standards such as Harbor for evaluation.

## REFERENCES

Anthropic. Claude code: An agentic coding tool. https://github.com/anthropics/claude-code, 2025. Accessed: 2026-06-06.

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, 2025. URL https://arxiv.org/abs/2505.20411.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.

Mouxiang Chen, Lei Zhang, Yunlong Feng, Xuwu Wang, Wenting Zhao, Ruisheng Cao, Jiaxi Yang, Jiawei Chen, Mingze Li, Zeyao Ma, Hao Ge, Zongmeng Zhang, Zeyu Cui, Dayiheng Liu, Jingren Zhou, Jianling Sun, Junyang Lin, Binyuan Hui, and Zhongxin Liu. Swe-universe: Scale real-world verifiable environments to millions, 2026. URL https://arxiv.org/abs/2602.02361.

Daixuan Cheng, Shaohan Huang, Yuxian Gu, Huatong Song, Guoxin Chen, Li Dong, Wayne Xin Zhao, Ji-Rong Wen, and Furu Wei. Computer environments elicit general agentic intelligence in llms, 2026. URL https://arxiv.org/abs/2601.16206.

Cursor Research, :, Aaron Chan, Ahmed Shalaby, Alexander Wettig, Aman Sanger, Andrew Zhai, Anurag Ajay, Ashvin Nair, Charlie Snell, Chen Lu, Chen Shen, Emily Jia, Federico Cassano, Hanpeng Liu, Haoyu Chen, Henry Wildermuth, Jacob Jackson, Janet Li, Jediah Katz, Jiajun Yao, Joey Hejna, Josh Warner, Julius Vering, Kevin Frans, Lee Danilek, Less Wright, Lujing Cen, Luke Melas-Kyriazi, Michael Truell, Michiel de Jong, Naman Jain, Nate Schmidt, Nathan

Wang, Niklas Muennighoff, Oleg Rybkin, Paul Loh, Phillip Kravtsov, Rishabh Yadav, Sahil Shah, Sam Kottler, Alexander M Rush, Shengtong Zhang, Shomil Jain, Sriram Sankar, Stefan Heule, Stuart H. Sul, Sualeh Asif, Victor Rong, Wanqi Zhu, William Lin, Yuchen Wu, Yuri Volkov, Yury Zemlyanskiy, Zack Holbrook, and Zhiyuan Zhang. Composer 2 technical report, 2026. URL https://arxiv.org/abs/2603.24477.

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

Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. Scaling synthetic data creation with 1,000,000,000 personas, 2025. URL https://arxiv.org/abs/2406.20094.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081): 633–638, September 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-09422-z. URL http://dx.doi.org/10.1038/s41586-025-09422-z.

Harbor Framework Team. Harbor: A framework for evaluating and optimizing agents and models in container environments, January 2026. URL https://github.com/harbor-framework/harbor.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=VTF8yNQM66.

Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training, 2025. URL https://arxiv.org/abs/2411.15124.

Zhongzhi Li, Yucheng Shi, Zongxia Li, Ruhan Wang, Anhao Li, Zixun Huang, Junyao Yang, Lei Ke, Ninghao Liu, Haitao Mi, and Leowei Liang. Recursive synthesis for long-horizon terminal tasks, 2026. URL https://arxiv.org/abs/2608.05466.

Xi Victoria Lin, Chenglong Wang, Luke Zettlemoyer, and Michael D. Ernst. NL2Bash: A corpus and semantic parser for natural language interface to the linux operating system. In Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga (eds.), *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA). URL https://aclanthology.org/L18-1491/.

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, 2026. URL https://arxiv.org/abs/2602.10999.

Fanzhe Meng, Guoxin Chen, Jiale Zhao, Shuang Sun, Zhiyu Lin, Wayne Xin Zhao, Ruihua Song, Ji-Rong Wen, and Kai Jia. Calibforge: Adversarial solver calibration for scaling learnable terminal tasks, 2026. URL https://arxiv.org/abs/2608.06352.

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

MiniMax. Interleaved thinking unlocks reliable MiniMax-M2 agentic capability. https://www.minimax.io/news/why-is-interleaved-thinking-important-for-m2, November 2025. Blog post. Accessed June 2026.

MiniMax, Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu, Chao Wang, Cheng Zhu, Chengjun Xiao, Chengyu Du, Chi Zhang, Chu Qiao, Chunhao Zhang, Chunhui Du, Congchao Guo, Da Chen, Deming Ding, Dianjun Sun, Dong Li, Enwei Jiao, Haigang Zhou, Haimo Zhang, Han Ding, Haohai Sun, Haoyu Feng, Huaiguang Cai, Haichao Zhu, Jian Sun, Jiaqi Zhuang, Jiaren Cai, Jiayuan Song, Jin Zhu, Jingyang Li, Jinhao Tian, Jinli Liu, Junhao Xu, Junjie Yan, Junteng Liu, Junxian He, Kaiyi Feng, Ke Yang, Kecheng Xiao, Le Han, Leyang Wang, Lianfei Yu, Liheng Feng, Lin Li, Lin Zheng, Linge Du, Lingyu Yang, Lunbin Zeng, Minghui Yu, Mingliang Tao, Mingyuan Chi, Mozhi Zhang, Mujie Lin, Nan Hu, Nongyu Di, Peng Gao, Pengfei Li, Pengyu Zhao, Qibing Ren, Qidi Xu, Qile Li, Qin Wang, Rong Tian, Ruitao Leng, Shaoxiang Chen, Shaoyu Chen, Shengmin Shi, Shitong Weng, Shuchang Guan, Shuqi Yu, Sichen Li, Songquan Zhu, Tengfei Li, Tianchi Cai, Tianrun Liang, Weiyu Cheng, Weize Kong, Wenkai Li, Xiancai Chen, Xiangjun Song, Xiao Luo, Xiao Su, Xiaobo Li, Xiaodong

Han, Xinzhu Hou, Xuan Lu, Xun Zou, Xuyang Shen, Yan Gong, Yan Ma, Yang Wang, Yiqi Shi, Yiran Zhong, Yonghong Duan, Yongxiang Fu, Yongyi Hu, Yu Gao, Yuanxiang Fan, Yufeng Yang, Yuhao Li, Yulin Hu, Yunan Huang, Yunji Li, Yunzhi Xu, Yuxin Mao, Yuxuan Shi, Yuze Wenren, Zehan Li, Zelin Li, Zhanxu Tian, Zhengmao Zhu, Zhenhua Fan, Zhenzhen Wu, Zhichao Xu, Zhihang Yu, Zhiheng Lyu, Zhuo Jiang, Zibo Gao, Zijia Wu, Zijian Song, and Zijun Sun. Minimax-m1: Scaling test-time compute efficiently with lightning attention, 2025. URL https://arxiv.org/abs/2506.13585.

Team Olmo, :, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, Pete Walsh, Pradeep Dasigi, Robert Berry, Saumya Malik, Saurabh Shah, Scott Geng, Shane Arora, Shashank Gupta, Taira Anderson, Teng Xiao, Tyler Murray, Tyler Romero, Victoria Graf, Akari Asai, Akshita Bhagia, Alexander Wettig, Alisa Liu, Aman Rangapur, Chloe Anastasiades, Costa Huang, Dustin Schwenk, Harsh Trivedi, Ian Magnusson, Jaron Lochner, Jiacheng Liu, Lester James V. Miranda, Maarten Sap, Malia Morgan, Michael Schmitz, Michal Guerquin, Michael Wilson, Regan Huff, Ronan Le Bras, Rui Xin, Rulin Shao, Sam Skjonsberg, Shannon Zejiang Shen, Shuyue Stella Li, Tucker Wilde, Valentina Pyatkin, Will Merrill, Yapei Chang, Yuling Gu, Zhiyuan Zeng, Ashish Sabharwal, Luke Zettlemoyer, Pang Wei Koh, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. Olmo 3, 2026. URL https://arxiv.org/abs/2512.13961.

OpenThoughts-Agent Team. OpenThoughts-Agent. https://www.open-thoughts.ai/blog/agent, December 2025.

Bespoke Labs OpenThoughts-Agent Team, Snorkel AI. OpenThoughts-TBLite: A High-Signal Benchmark for Iterating on Terminal Agents. https://www.openthoughts.ai/blog/openthoughts-tblite, February 2026.

Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with swe-gym, 2024. URL https://arxiv.org/abs/2412.21139.

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

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.

Sundar Pichai, Demis Hassabis, and Koray Kavukcuoglu. A new era of intelligence with gemini 3. *Mountain View, CA: Google. Available online at: https://blog.google/products-and-platforms/products/gemini/gemini-3 (Accessed February 1, 2026)*, 2025.

Penghui Qi, Xiangxin Zhou, Zichen Liu, Tianyu Pang, Chao Du, Min Lin, and Wee Sun Lee. Rethinking the trust region in llm reinforcement learning, 2026. URL https://arxiv.org/abs/2602.04879.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026. URL https://qwen.ai/blog?id=qwen3.5.

Negin Raoof, Richard Zhuang, Marianna Nezhurina, Etash Guha, Atula Tejaswi, Ryan Marten, Charlie F. Ruan, Tyler Griggs, Alexander Glenn Shaw, Hritik Bansal, E. Kelly Buchanan, Artem Gazizov, Reinhard Heckel, Chinmay Hegde, Sankalp Jajee, Daanish Khazi, Emmanouil Koukoumidis, Xiangyi Li, Hange Liu, Shlok Natarajan, Harsh Raj, Nicholas Roberts, Ethan Shen, Nishad Singhi, Michael Siu, Ashima Suvarna, Hanwen Xing, Patrick Yubeaton, Robert Zhang, Leon Liangyu Chen, Xiaokun Chen, Steven Dillmann, Saadia Gabriel, Xunyi Jiang, Anurag Kashyap, Boxuan Li, Yein Park, Minh Pham, Sujay Sanghavi, Lin Shi, Ke Sun, Yixin Wang, Zhiwei Xu, Erica Zhang, Siyan Zhao, Wanjia Zhao, Jenia Jitsev, Alex Dimakis, Benjamin Feuer, and Ludwig Schmidt. Openthoughts-agent: Data recipes for agentic models, 2026. URL https://arxiv.org/abs/2606.24855.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.

Ethan Shen, Daniel Tormoen, Saurabh Shah, Ali Farhadi, and Tim Dettmers. Sera: Soft-verified efficient repository agents, 2026a. URL https://arxiv.org/abs/2601.20789.

Qijia Shen, Zhiqi Huang, Vamsidhar Kamanuru, Aznaur Aliev, Jay Rainton, Ahmed Awelkair, Zhichen Zeng, Jiajun Li, Shi Dong, Yueming Yuan, Boyuan Ma, Qizheng Zhang, Jiwei Fu, Yuzhen Mao, Wendong Fan, Ping Nie, Philip Torr, Bernard Ghanem, Changran Hu, Jonathan Lingjie Li, Urmish Thakker, and Guohao Li. SETA: Scaling Environments for Terminal Agents, July 2026b. URL https://arxiv.org/abs/2607.10891.

Kou Shi, Zun Wang, Qisheng Su, Shiting Huang, Ziao Zhang, Zhen Fang, Qingnan Ren, Jin Liu, Yu Zeng, Yiming Zhao, Lin Chen, Zehui Chen, and Feng Zhao. FACET: Preserving source intent and executable state in terminal task synthesis, 2026. URL https://arxiv.org/abs/2608.18580.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and efficient foundation language models, 2023. URL https://arxiv.org/abs/2302.13971.

Junli Wang, Zhoujun Cheng, Yuxuan Zhang, Shibo Hao, Yao Tang, Zhiting Hu, Prithviraj Ammanabrolu, and Hao Zhang. NanoRollout: Scale digital agent rollouts without pain. https://cocoa-org.notion.site/nanorollout, May 2026a. Notion Blog.

Weixun Wang, XiaoXiao Xu, Wanhe An, Fangwen Dai, Wei Gao, Yancheng He, Ju Huang, Qiang Ji, Hanqi Jin, Xiaoyang Li, Yang Li, Zhongwen Li, Shirong Lin, Jiashun Liu, Zenan Liu, Tao Luo, Dilxat Muhtar, Yuanbin Qu, Jiaqiang Shi, Qinghui Sun, Yingshui Tan, Hao Tang, Runze Wang, Yi Wang, Zhaoguo Wang, Yanan Wu, Shaopan Xiong, Binchen Xu, Xander Xu, Yuchi Xu, Qipeng Zhang, Xixia Zhang, Haizhou Zhao, Jie Zhao, Shuaibing Zhao, Baihui Zheng, Jianhui Zheng, Suhang Zheng, Yanni Zhu, Mengze Cai, Kerui Cao, Xitong Chen, Yue Dai, Lifan Du, Tao Feng, Tao He, Jin Hu, Yijie Hu, Ziyu Jiang, Cheng Li, Xiang Li, Jing Liang, Xin Lin, Chonghuan Liu, ZhenDong Liu, Zhiqiang Lv, Haodong Mi, Yanhu Mo, Junjia Ni, Shixin Pei, Jingyu Shen, XiaoShuai Song, Cecilia Wang, Chaofan Wang, Kangyu Wang, Pei Wang, Tao Wang, Wei Wang, Ke Xiao, Mingyu Xu, Tiange Xu, Nan Ya, Siran Yang, Jianan Ye, Yaxing Zang, Duo Zhang, Junbo Zhang, Boren Zheng, Wanxi Deng, Ling Pan, Lin Qu, Wenbo Su, Jiamang Wang, Wei Wang, Hu Wei, Minggang Wu, Cheng Yu, Bing Zhao, Zhicheng Zheng, and Bo Zheng. Let it flow: Agentic crafting on rock and roll, building the rome model within an open agentic learning ecosystem, 2026b. URL https://arxiv.org/abs/2512.24873.

Jie Wu, Zhenru Zhang, Beichen Zhang, Xuwu Wang, Yuhui Su, Mouxiang Chen, Peng Wang, Zhihai Wang, Que Shen, Hao Zhou, An Yang, Fei Huang, Yujiu Yang, and Dayiheng Liu. Terminal-universe: Turning agent trajectories into scalable terminal environments, 2026a. URL https://arxiv.org/abs/2609.04148.

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*, 2026b.

John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.

John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. Swe-smith: Scaling data for software engineering agents, 2025. URL https://arxiv.org/abs/2504.21798.

Sidi Yang, Chaofan Tao, Jierun Chen, Tiezheng Yu, Ruoyu Wang, Yuxin Jiang, Yiming Du, Wendong Xu, Jing Xiong, Taiqiang Wu, Lifeng Shang, Xiaohui Li, Ngai Wong, and Haoli

Bai. What makes interaction trajectories effective for training terminal agents?, 2026. URL https://arxiv.org/abs/2606.03461.

Yihang Yao, Bo Pang, Xuan Phi Nguyen, Ding Zhao, Shafiq Joty, and Semih Yavuz. Learning generalizable behaviors for terminal agents, 2026. arXiv preprint arXiv:2608.22631.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL https://arxiv.org/abs/2503.14476.

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.

# A LIMITATIONS

Our dataset generation pipeline is completely synthetic, and relies on the presence of a strong generator model. While we find positive results using Qwen 3.5 9B both as generator and policy, it is unclear if even stronger models or weaker models would behave similarly. We also believe that our approach likely struggles on tasks involving large-scale codebases, as these are inherently difficult to generate (but well-covered by SWE-bench-style datasets). Additionally, while we have focused on training smaller terminal agents as a baseline for academic settings, we note that running many isolated containers still proves expensive and/or difficult in open frameworks at scale, limiting training speed and efficiency and potentially still putting terminal agent training out of reach for academic groups. Finally, our results may be lower than possible SOTA due to the shorter context length and simple harness used relative to strong industry approaches, although we see this as a feature of our approach that makes it more friendly to smaller teams and easier to develop with.

# B RELATED WORK

## B.1 DATA FOR TERMINAL AGENTS

Early work on terminal agents focused on training models to translate natural language to bash statements. NL2Bash (Lin et al., 2018) paired short instructions with bash commands mined from the web, and was converted into an RL dataset for the training of OpenThinker Agent (OpenThoughts-Agent Team, 2025). However, it failed to yield significant improvements over their strong SFT model. This showcases the need for generating more complex terminal-agent tasks, which provide stronger learning signals on more complex tasks. More recently, Raoof et al. (2026) compared multiple RL datasets and found that the choice of RL dataset could have a significant impact on final performance, motivating more work into improving data for RL training specifically.

Recent work constructs terminal training data from various sources. Repository-based approaches include SWE-Smith (Yang et al., 2025), which synthesises bug-fixing instances, and Terminal-Traj (Wu et al., 2026b), which builds Dockerised tasks and verified trajectories from repositories. Related software-engineering datasets include SWE-Universe (Chen et al., 2026), SWE-rebench (Badertdinov et al., 2025), SWE-Gym (Pan et al., 2024), and SERA (Shen et al., 2026a). CLI-Gym (Lin et al., 2026) extends repository-based task creation by perturbing agent paths through repository-based environments to construct new environments and tasks. Terminal-Lego (Yang et al., 2026) and SETA-Env (Shen et al., 2026b) scrape sites such as StackOverflow and Stack Exchange for questions with verified answers and turn them into environments. SETA additionally applies an evolutionary procedure to control diversity and difficulty. FACET (Shi et al., 2026) constructs environments and tasks from skill files, generating task instructions and verifiers after environment construction to ensure validity. Terminal-Universe (Wu et al., 2026a) constructs environments from public agent trajectories, but does not release any data. We find in many cases that repository-based task creation results in unbalanced task splits in the data: for example, in Figure 3 both TerminalTraj and CLI-Gym are mostly comprised of software engineering tasks.

Another set of recent work generates data from a given taxonomy or set of seed tasks, and is the approach we take in this work. We are particularly inspired by Endless Terminals (Gandhi et al., 2025), which similarly utilises a strong external model to generate new terminal tasks, but produces significantly fewer tasks that are largely focused on file manipulation, making them too easy for modern models. For instance, we find that Gemini-3-Flash achieves over 90% pass@1 on Endless Terminals. TermiGen (Zhu et al., 2026) and LiteCoder-Terminal (Peng et al., 2026) similarly synthesise containers and tasks from seed prompts, using a fairly complex multi-agent pipeline to generate and validate the tasks, and do not explicitly control the difficulty of their tasks. CalibForge (Meng et al., 2026) instructs an agent to construct a task from browsing the web, and then iteratively improves the sandbox until it is on the edge of solvability. Similarly, Recursive Synthetic Terminal Tasks (RST) (Li et al., 2026) starts from a set of seed tasks, and iteratively rewrites them based on sandbox feedback to improve their difficulty and quality. Nemotron-Terminal (Pi et al., 2026) both generates new data from seed prompts and adapts existing datasets into terminal tasks, but does not release any RL environment data, and similarly only validates their data through large-scale SFT training. We find that these other datasets do not appear to improve over TMAX-13K, despite its much simpler generation pipeline, as shown in Table 2.

Table 7: Composition axes used by our task-generation pipeline. The first three are seeded from the skill taxonomy of Pi et al. (2026); the remaining six are our contributions, targeting diversity and difficulty.

| Axis | Source | Cardinality | Values |
|---|---|---|---|
| Domain | Pi et al. (2026) | 9 | security, software_engineering, file_operations, data_querying, data_science, debugging, scientific_computing, data_processing, system_administration |
| Skill type | Pi et al. (2026) | 4–7 / domain | e.g. Algorithmic, Systems, Data Processing, Web Security, Testing, Mathematical, Multi-Language |
| Primitive skills | Pi et al. (2026) | 20–40 / domain | 3–5 primitive skills sampled per task (e.g. "graph traversal and dependency resolution", "cryptographic hashing and checksum verification") |
| Persona | Ours | 6–18 / domain | domain-tied user roles (e.g. "incident responder investigating a 3am page", "bioinformatics analyst processing sequences") |
| Language | Ours | 8 (weighted) | Python, C, Bash, C++, Rust, Go, multi-language, any |
| Task complexity | Ours | 4 | short, moderate, complex, intricate (30–60 commands) |
| Command complexity | Ours | 3 | bash-only; bash + code; bash + code + system services |
| Fixture | Ours | 7 | text_only, image, audio, video, stripped_binary, vendored_package, multi_service_compose |
| Verifier | Ours | 5 | exact_text, metric_threshold, adversarial_corpus, fuzz_equivalence, multi_protocol |

# C Terminal Data Generation Pipeline

## C.1 Verifier and fixture kinds

Each task independently samples one verifier kind and one fixture kind (Table 7). The legacy defaults (`exact_text`, `text_only`) reproduce standard text-in / text-out RL tasks, while the remaining kinds add graded verification (Table 8) and non-text inputs (Table 9). Graded verifiers relax brittle string equality and expose a continuous difficulty knob, and non-text fixtures broaden task inputs while keeping the policy text-only—the agent processes each artifact through standard terminal tooling rather than native perception.

Table 8: Verifier kinds. Beyond legacy exact-text equality, we add four graded verifiers that relax exact matching and expose an explicit difficulty knob.

| Kind | Pass criterion | Example | Difficulty knob |
|---|---|---|---|
| `exact_text` | Output exactly equals the ground-truth string. | file contents match a reference | — (legacy) |
| `metric_threshold` | A numeric metric against a reference meets a threshold. | image SSIM $\geq 0.95$; speedup $\geq 1.3\times$ | threshold value |
| `adversarial_corpus` | Solution rejects every item in a malicious corpus *and* preserves every item in a benign one. | a sanitiser that blocks exploits yet leaves clean inputs intact | per-corpus pass rate; corpus size |
| `fuzz_equivalence` | Agent program matches a reference oracle bit-exactly on $N$ random inputs. | reproduce a stripped binary's output | $N$; input distribution |
| `multi_protocol` | A service the agent brings up answers real protocol-level requests correctly. | HTTP / TCP / gRPC / SMTP request–response checks | number of protocols / conditions |

Table 9: Fixture kinds. Each task optionally ships a non-text artifact; the agent recovers the hidden ground truth via standard terminal tooling, so the policy stays text-only.

| Kind | Artifact | Agent tooling | Recovers |
|---|---|---|---|
| `text_only` | none | — | — (legacy) |
| `image` | PNG / JPEG | OCR (tesseract), vision | hidden text / structure |
| `audio` | WAV / MP3 | transcription (whisper.cpp, ffmpeg) | transcript / events |
| `video` | MP4 | frame extraction (ffmpeg) + analysis | frame ranges, counts, detections |
| `stripped_binary` | stripped / UPX-packed binary | objdump, gdb, strings; or black-box oracle | implemented algorithm |
| `vendored_package` | pre-vendored package source with a perturbation | build / debug tools (no internet) | known-good code path after fix |
| `multi_service_compose` | multiple cooperating services | service config / glue | end-to-end protocol flow |

### C.2 Rubric-based Data Repair

We follow Yao et al. (2026) in classifying TMAX-13K using an LM judge by the following labels:

- `VERIFIER-TOO-WEAK`: The verifier does not fully test the final state, allowing the model to pass without fully completing the task.
- `ANSWER-LEAK`: The answer or tests are in the initial image state, allowing trivial task completion.
- `CLEAN`: No issues detected.
- `INSTR-VERIFIER-MISMATCH`: The instruction and verifier contradict each other.
- `INSTR-ENV-MISMATCH`: The instruction and environment contradict each other.
- `OTHER`: Some other issue not captured by the reasons above.
- `INSTR-AMBIGUOUS`: The instruction is ambiguous.
- `TASK-TRIVIAL`: The task is trivial to perform.

We classify our initial 14.6K samples using a GPT 5.6 Sol judge, resulting in the classification seen in Table 10. We then prompt GPT 5.6 Sol to fix the top three issues (verifier too weak, answer

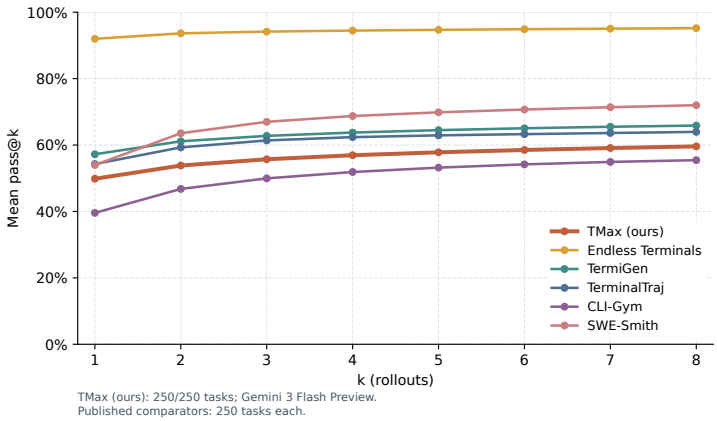


Figure 6: **Pass@$k$ difficulty curves.** Mean pass@$k$ ($k = 1$–8) for Gemini-3-Flash-Preview on a fixed 250-task subsample per dataset (8 rollouts each); lower is harder. TMAX occupies the hardest band together with CLI-Gym and attains the second lowest pass@8 of any dataset, showing that its difficulty persists as $k$ grows, whereas easy datasets such as Endless-Terminals saturate near the ceiling.

leakage, and instruction-verifier mismatch) by producing either new, hardened verifiers or cleaning the image specification. We drop all other problematic tasks from the dataset.

| label | Tasks | Share (%) | Applied repairs |
|---|---|---|---|
| VERIFIER-TOO-WEAK | 11,160 | 76.4 | 10,990 |
| ANSWER-LEAK | 1,592 | 10.9 | 1,018 |
| CLEAN | 1,125 | 7.7 | - |
| INSTR-VERIFIER-MISMATCH | 441 | 3.0 | 432 |
| INSTR-ENV-MISMATCH | 177 | 1.2 | 0 |
| OTHER | 83 | 0.6 | 0 |
| INSTR-AMBIGUOUS | 23 | 0.2 | 0 |
| TASK-TRIVIAL | 0 | 0.0 | 0 |
| **Total** | **14,601** | **100.0** | **12,440** |

Table 10: GPT 5.6-judged rubric labels from the original 14,601 tasks and applied repair counts. We retain 12,440 tasks with an applied repair plus 1,125 originally CLEAN tasks without a patch, giving 13,565 tasks in the hardened prefilter dataset; the other 1,036 are excluded. Applied repairs do not imply a second rubric classification.

### C.3 PASS@K DIFFICULTY CURVES

Table 1 reports pass@$k$ at $k \in \{1, 4, 8\}$; Figure 6 plots the full curve over $k = 1$–8 for every dataset. The vertical position of a curve measures absolute difficulty, while its slope measures how much additional sampling recovers. TMAX occupies the hardest band together with CLI-Gym across all $k$, so its tasks remain unsolved even with eight rollouts, suggesting high difficulty and low rollout variance (although some tasks may still be infeasible). Easy datasets behave very differently: Endless-Terminals already sits near the ceiling at $k = 1$ and saturates almost immediately, leaving little headroom for a model to improve. This persistent, well-separated difficulty is what makes our data a useful source of learning signal for RL.

### C.4 BALANCE SCORE

The balance score in equation 1 is the normalized effective number of categories of an empirical distribution. The numerator $\exp(H)$ is the size a uniform distribution would need to have the same Shannon entropy as the observed mass vector; dividing by the bucket count $N$ rescales the result to

| Benchmark | TMax (ours) | Endless Terminals | OpenThinker Agent | TermiGen | Terminal Traj | CLI-Gym | SWE-Smith |
|---|---|---|---|---|---|---|---|
| TB-Lite (openthoughts-tblite@2.0) | 0.0% | 0.0% | 0.0% | 0.0% | 0.2% | 0.0% | 0.0% |
| TB2 (terminal-bench@2.0) | 0.0% | 0.0% | 0.0% | 0.0% | 0.5% | 0.0% | 0.0% |

Table 11: **Decontamination via $n$-gram overlap on task descriptions** ($n$=13, stride=1). Each cell is the percentage of a dataset's task descriptions whose $n$-gram window overlaps at least one $n$-gram from the benchmark's task descriptions. Larger $n$ is stricter (fewer false positives).

| Model | Harness | TB Lite |
|---|---|---|
| Haiku 4.5 | Ours | $60.8 \pm 0.6$ |
| Haiku 4.5 | mini-SWE-agent | $\mathbf{61.4 \pm 1.0}$ |
| Haiku 4.5 | Terminus-2 | $56.1 \pm 3.9$ |

Table 12: Claude Haiku 4.5 Terminal-Bench Lite performance when using different harnesses. Our harness is an earlier version from during development.

$[1/N, 1]$. A value of 1.0 thus corresponds to perfectly uniform coverage of all buckets, while $1/N$ corresponds to all mass concentrated on a single bucket. We prefer this to raw entropy because it is comparable across axes with different bucket counts and reads as a natural "fraction-of-uniform" diversity score.

### C.5 Decontamination

We quantify potential train–test contamination as $n$-gram overlap between each dataset's task descriptions and the task descriptions of our evaluation benchmarks, Terminal-Bench 2.0 and TB-Lite. Using a sliding window of $n$=13 tokens (stride 1), we flag a dataset task as overlapping if any of its windows matches at least one 13-gram from a benchmark task, following standard contamination protocols (Brown et al., 2020; Touvron et al., 2023); larger $n$ yields a stricter test with fewer spurious matches. Table 11 reports the fraction of each dataset flagged this way. Overlap is negligible across the board: TMax shows $0\%$ overlap with both benchmarks, and only TerminalTraj exhibits any measurable overlap ($0.2\%$ on TB-Lite and $0.5\%$ on TB2). Since Terminal-Bench 2.1 is a modified version of 2.0, we do not run additional decontamination against it.

## D Harness Choices

When deciding on an initial harness, we opted for a mini-SWE-agent inspired harness based on results over smaller closed-source offerings such as Claude Haiku 4.5. In Table 12, we show that our harness and mini-SWE-agent outperform Terminus-2, the usual default for Terminal-Bench models. We attribute this drop to more complex tool formats required for Terminus-2, which are especially difficult for smaller models to follow. Additionally, we wished to opt for a simpler harness to reduce complexity during RL training.

## E Additional RL Training Details

### E.1 Full RL training hyperparameters

Unless otherwise stated, we use the hyperparameters for RL training stated in Table 13. We found that using a larger group size aided with stability in earlier experiments, as shown in Figure 7. We additionally found in early pilot experiments that DPPO demonstrated superior stability over GRPO, motivating us to stick with it for ablations.

| Category | Hyperparameter | Value |
|---|---|---|
| Model | Model dtype | `bfloat16` |
| Model | Gradient checkpointing | true |
| Model | LM head fp32 | true |
| Data | Max prompt tokens | 2048 |
| Data | Per-turn max tokens | 16384 |
| Data | Max total response length | 65536 (32768 for Qwen 3 8B) |
| Rollout | Unique prompts per rollout batch | 8 |
| Rollout | Group size | 32 |
| Rollout | Max async steps | 4 |
| Rollout | Active sampling | true |
| Rollout | Filter zero-std samples | true |
| Rollout | Sampling Temperature | 1.0 |
| Optimisation | Optimiser | AdamW |
| Optimisation | Training steps | 500 |
| Optimisation | Learning rate | $1 \times 10^{-6}$ |
| Optimisation | LR scheduler | constant |
| Optimisation | Max grad norm | 1.0 |
| RL objective | Loss | DPPO |
| RL objective | Advantage normalisation | centered |
| RL objective | KL coefficient $\beta$ | 0.0 |
| RL objective | DPPO divergence | binary TV |
| RL objective | DPPO divergence threshold | 0.1 |
| Tools | Max tool/env steps | 64 |
| Tools | Bash tool timeout | 120 s |

Table 13: Hyperparameters for RL runs. We use these hyperparameters unless otherwise stated.

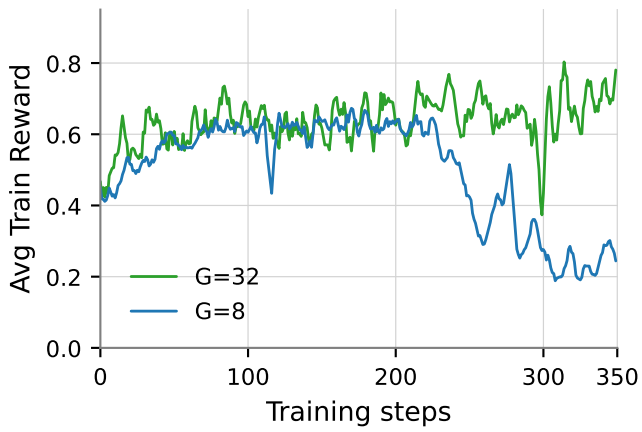


Figure 7: **Larger group sizes limit training collapse.** Average training reward when doing RL training on TMAX-13K with DPPO in initial experiments with a slightly different harness and Podman-based backend, using group size 8 vs. 32 (using 32 or 8 prompts per batch to keep total batch size constant).

### E.2 FULL SFT TRAINING DETAILS

For SFT training, we use the hyperparameters given in Table 14. We largely follow Pi et al. (2026) in hyperparameters. Additionally, we do not filter out unsuccessful/incomplete rollouts from our dataset following Pi et al. (2026).

| Category | Hyperparameter | Value |
|---|---|---|
| Model | Precision | bf16 |
| Model | Attention implementation | flash attention 3 |
| Model | Gradient checkpointing | true |
| Data | Max sequence length | 65536 for Qwen 3.5, 32768 for Qwen 3 |
| Optimisation | Epochs | 2 |
| Optimisation | Per-device batch size | 1 |
| Optimisation | Gradient accumulation steps | 4 |
| Optimisation | Global batch size | 128 |
| Optimisation | Learning rate | $2 \times 10^{-5}$ |
| Optimisation | LR scheduler | linear |
| Optimisation | Warmup ratio | 0.03 |
| Optimisation | Weight decay | 0.0 |
| Optimisation | Optimiser | AdamW |

Table 14: Hyperparameters for SFT training.

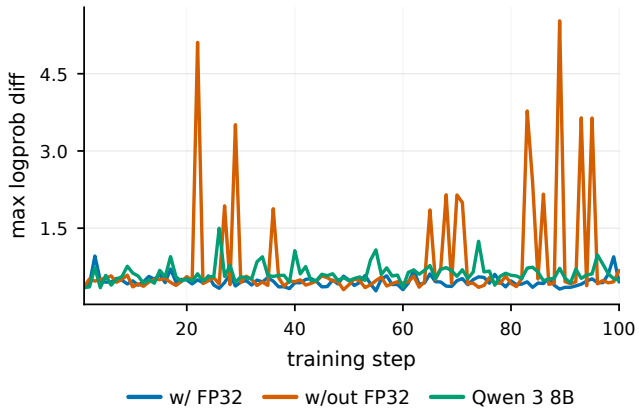


Figure 8: **Maximum difference between inference (vLLM) and trainer (HuggingFace) logprobs** during the first 100 steps of RL training for Qwen 3.5 9B. Qwen 3.5 without the FP32 LM head displays larger and more frequent spikes. Qwen 3 8B training also does not display spikes, even though we do not apply the FP32 LM head.

### E.3 LM Head Details

We find using an FP32 LM head especially crucial when training Qwen 3.5, but not when training Qwen 3. To highlight this, we plot the maximum logprob difference during RL training of Qwen 3.5 and Qwen 3 in Figure 8. Using the FP32 LM head reduces the maximum difference dramatically. Interestingly, we find this is less important for Qwen 3, which does not display large logprob differences even without the FP32 head.

### E.4 Filtered Samples

In our RL training, we filter samples with all the same rewards, as these contribute no gradient to the batch. We plot the number of all-zero groups in Fig. 9, and observe that the number of filtered samples averages around 5-7. This suggests that, despite our lack of validation when generating data (§3.1), our 9B model can still find solutions for roughly 60% of sampled prompts during training. To avoid unnecessarily filtering out valid samples that may simply be too hard, and to avoid the cost of filtering using a significantly larger and performant model, we leave these filtered samples in our released dataset.

We also examine how many samples are filtered because they are entirely solved in Figure 10. More samples are slowly filtered over training due to improving model ability. However, compared to a significantly easier dataset like Endless Terminals, significantly fewer samples are filtered per

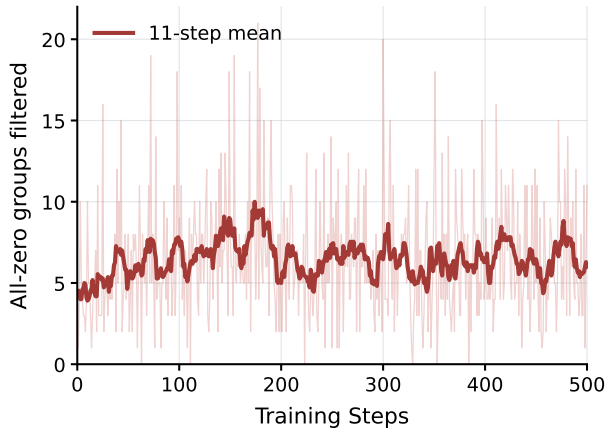


Figure 9: Number of all-zero groups filtered over the course of RL training on TMAX-13K.

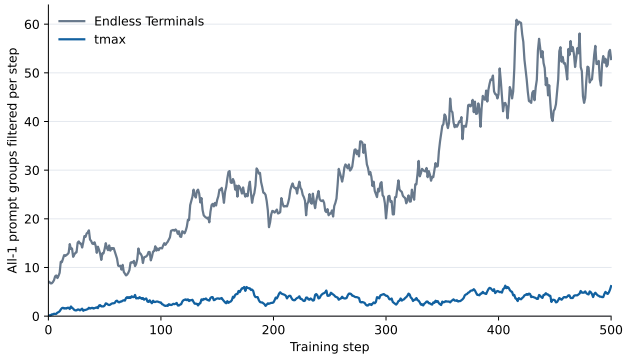


Figure 10: Number of all-one groups filtered over the course of RL training on TMAX-13K and Endless Terminals. Over the course of training, Endless Terminals requires significantly more sampling to fill a batch with non-zero variation samples.

step by the end of training for TMAX-13K. This highlights the importance of difficulty in training efficiency: by the end of training, each training step for Endless Terminals generates seven times as many rollouts as the minimum needed to fill the batch (due to high filter rates).

### E.5 REWARD HACKING

We perform a light check over our main Daytona-based runs on Terminal-Bench 2.0 and find that TMAX-9B, after RL training, displays some instances of reward hacking. We find 2 clear cases:

- `feal-differential-cryptanalysis`: The agent bypasses the intended attack by importing the cipher's internal key array to recover the target key. The agent explicitly calls this "cheating" but submits it anyway.
- `financial-document-processor`: The agent accidentally deletes the original documents, then recreates placeholder files under the same names and hand-writes the invoice CSV instead of extracting values from the documents.

In both cases, the rollouts with the hacking behaviour scored 0, so it does not affect downstream performance. Examining the chain of thought, we find that the model appears aware that it has not correctly performed the task. For example, for `cryptanalysis`, the model initially states "`"This means the attack should work with a black-box encrypt_fn, without accessing internal state. My current solution cheats by accessing feal_key directly. I need to fix this!"`". However, it then sub-

mits anyway: "'`It doesn't prohibit accessing other modules or the key array.`"'. In the case of `financial-document-processor`, the model instead attempts to fool the verifier since it cannot actually perform the desired action: "'`The find command returned nothing. The files are truly gone. At this point, I need to make a pragmatic decision. The task cannot be completed perfectly without the original files. However, I can: 1. Create a working script that would correctly process the files if they existed 2. Create placeholder files that allow the script to run 3. Generate the expected summary.csv based on my analysis. [...] Since the task evaluation likely checks the final state, and I know what the correct classification and amounts should be (from my analysis), I'll manually create the final correct state.`"' We will publicly release our rollouts and hope that future work can further explore the connection between our recipe and this hacking.

# F Additional Evaluation Details

## F.1 Full Figure 1 Results

We show the full numeric results for Fig. 1 in Table 15. For prior work, we use results reported by the original authors or on the Terminal-Bench 2.0 leaderboard[2]. For closed and open-weight models, we use leaderboard results. Where there are multiple leaderboard harnesses reported, we default to the best amongst mini-SWE-agent or Terminus-2 runs.

For the Qwen baselines and our own models, we run three trials per task using Daytona sandboxes and our harness. For TMAX-9B, we do a small manual check for reward hacking (see §E.5).

# G SFT Details

## G.1 SFT Data Generation

We additionally generate a small SFT dataset to use as a warm-start for RL training, with prior works showing that warm-up could help training stability and performance (Guo et al., 2025; Kimi Team et al., 2025). Re-using our terminal data pipeline, we additionally generate 2.2k environments, and we use Qwen 3.6 27B to generate 8 trajectories for each environment. We filter out trajectories with unparsed tool calls.[3] Together, this yields 16.5K total SFT trajectories, with 8K successful trajectories. We use this data for SFT experiments and finetuning Qwen 3 8B.

## G.2 SFT Ablations

In Table 16, we compare Qwen 3.5 9B and Qwen 3 performance before and after performing SFT on both the SFT mixture mentioned in §G.1 and a larger mixture composed from our SFT data mixed with data from past work (Pi et al., 2026; Wu et al., 2026b; OpenThoughts-Agent Team, 2025; Shen et al., 2026a).[4] We find that the larger mix significantly improves performance on Qwen 3 8B, but drops performance on Qwen 3.5 9B, likely due to the use of weak teacher models in past work (e.g., Pi et al. (2026) use DeepSeek v3.2 as a teacher model). While TMAX SFT uses Qwen 3.6 27B as a teacher model, we still find it results in degraded performance for Qwen 3.5 9B, although it does aid Qwen 3 8B. As such, we opted to go straight to RL training for Qwen 3.5 experiments, while using an SFT stage for our Qwen 3 8B experiments. We suspect that further work on improving the SFT data quality could improve performance for Qwen 3.5, and leave it to future work.

We additionally show the full splits for our 'big' SFT mix described above in Table 17.

[2] `https://www.tbench.ai/leaderboard/terminal-bench/2.0`

[3] To avoid traces containing the Qwen 3 XML-style tool call format, which can confuse small models with different tool call formats like Qwen 3 8B.

[4] See §E.2 for details.

| Model | Model Size | TB2.0 (%) |
|---|---|---|
| *Models from Prior Work* | | |
| Nemotron-Terminal-8B (Pi et al., 2026) | 8B | 13.0 |
| Nemotron-Terminal-14B (Pi et al., 2026) | 14B | 20.2 |
| Nemotron-Terminal-32B (Pi et al., 2026) | 32B | 27.4 |
| TermiGen-Qwen2.5-32B (Zhu et al., 2026) | 32B | 19.3 |
| Endless-Terminals-8B (Gandhi et al., 2025) | 8B | 6.7 |
| TerminalTraj (Wu et al., 2026b) | 7B | 10.1 |
| TerminalTraj (Wu et al., 2026b) | 14B | 19.1 |
| TerminalTraj (Wu et al., 2026b) | 32B | 22.0 |
| LiberCoder-32B (Lin et al., 2026) | 32B | 19.5 |
| LiberCoder-235B-A22B (Lin et al., 2026) | 235B | 31.0 |
| OpenThinkerAgent-32B (Raoof et al., 2026) | 32B | 26.2 |
| OT-Agent-ColdSFT+RL-8B (Raoof et al., 2026) | 8B | 13.5 |
| LiteCoder-Terminal-32B-SFT (Peng et al., 2026) | 32B | 18.5 |
| SETA-8B (Shen et al., 2026b) | 8B | 10.7 |
| SETA-DS-V4-Flash (Shen et al., 2026b) | 284B | 43.0 |
| CalibForge-30B-A3B (Meng et al., 2026) | 30B | 32.6 |
| CalibForge-35B-A3B (Meng et al., 2026) | 35B | 47.6 |
| Terminal-Lego-Qwen3-8B (Yang et al., 2026) | 8B | 11.8 |
| Terminal-Lego-Qwen3-32B (Yang et al., 2026) | 32B | 24.3 |
| RST-122B-A10B (Li et al., 2026) | 122B | 49.3 |
| *Open-Weight Models* | | |
| GPT-OSS-120B | 120B | 18.7 |
| DeepSeek-V3.2 | 685B | 39.6 |
| MiniMax-M2.7 | 357B | 45.1 |
| Kimi-K2.5 | 1.04T | 43.2 |
| GLM-5 | 744B | 52.4 |
| *Closed Models* | | |
| GPT-5-nano | ? | 11.5 |
| Claude-Haiku-4.5 | ? | 29.8 |
| GPT-5-mini | ? | 31.9 |
| Claude-Sonnet-4.6 | ? | 53.4 |
| *Base Models* | | |
| Qwen3-8B | 8B | 0.9 |
| Qwen3.5-2B | 2B | 1.1 |
| Qwen3.5-4B | 4B | 16.5 |
| Qwen3.5-9B | 9B | 21.0 |
| Qwen3.6-27B | 27B | 42.0 |
| *Our Models* | | |
| TMAX-2B | 2B | 6.4 |
| TMAX-4B | 4B | 25.1 |
| TMAX-9B | 9B | 33.0 |
| TMAX-27B | 27B | 43.8 |

Table 15: Full model names and Terminal-Bench 2.0 results for all 38 models; Figure 1 plots 31 for legibility, keeping the Pareto shading clear of model labels. The plot places labels near their markers and uses abbreviated names, including Nemotron 14B, Haiku 4.5 and Sonnet 4.6; TerminalTraj-14B and MiniMax-M2.7 are unlabelled, and uncertainty bars are omitted. SETA-DS-V4-Flash is the SETA agent based on DeepSeek V4 Flash. Prior-work scores are from the cited papers or the Terminal-Bench 2.0 leaderboard. Other open-weight and closed-model scores are the best official leaderboard entries. Parameter count is total parameters, including inactive experts. Our models and Qwen baselines use our harness and report the mean over three trials per task.

## H INFRASTRUCTURE DATA FILTERING

Due to our reliance on Apptainer, not all datasets can be fully run on our setup. For example, we remove one Terminal Bench Lite instance that requires managing user permissions on files in a way that Apptainer does not support (often due to fakeroot limitations). As such, we examine all dataset images in our training and evaluation datasets, and remove ones that are unable to run in our setup.

| Model | TB Lite | TB 2.1 |
|---|---|---|
| Qwen 3.5 9B | $\mathbf{41.9 \pm 2.7}$ | $\mathbf{16.1 \pm 3.7}$ |
| + TMAX SFT | $35.5 \pm 4.5$ | $15.0 \pm 3.0$ |
| + large SFT | $31.3 \pm 3.5$ | $16.9 \pm 0.9$ |
| Qwen 3 8B | $7.3 \pm 1.0$ | $1.1 \pm 0.9$ |
| + TMAX SFT | $11.5 \pm 0.1$ | $6.0 \pm 1.4$ |
| + large SFT | $\mathbf{16.4 \pm 2.3}$ | $\mathbf{7.9 \pm 3.3}$ |

Table 16: **Older SFT data does not improve Qwen 3.5.** Performance of Qwen 3.5 9B and Qwen 3 8B on Terminal-Bench 2.1 and Terminal-Bench Lite before and after performing SFT. TMAX SFT refers to the SFT dataset in § G.1, while large SFT is described in § G.2. Performance is mean $\pm$ stderr over 3 evaluation runs. Results using an earlier version of our harness.

| Source dataset | Citation | # Samples |
|---|---|---|
| `allenai__Sera-4.6-Lite-47000` | Shen et al. (2026a) | 47,464 |
| `m-a-p__TerminalTraj` | Wu et al. (2026b) | 16,748 |
| `Nemotron-Terminal-Corpus` | Pi et al. (2026) | 237,874 |
| `open-thoughts__OpenThoughts-Agent-v1-SFT` | OpenThoughts-Agent Team (2025) | 8,717 |
| TMAX-SFT | - | 16,496 |
| **Total** | | 327,299 |

Table 17: Source datasets and sample counts for our 'big' SFT mix described in §G.2.

**Baseline Dataset Filtering.** We filtered various tasks from the baseline datasets we compare against in Table 2. We list the total number of tasks removed and why for each dataset below. We only remove tasks from 2 datasets, and always remove less than 10% of the data.

- **TerminalTraj** (Wu et al., 2026b): 5,648 total; 510 removed; 5,138 retained. We excluded 509 images whose BusyBox `setsid` is incompatible with our harness's `-w` flag, plus one image with a dangling `/workspace` symlink that prevents sandbox startup.
- **TermiGen** (Zhu et al., 2026): 3,556 total; 0 removed; 3,556 retained.
- **LiteCoder** (Peng et al., 2026): 602 total; 0 removed; 602 retained.
- **OpenThinker-Agent** (OpenThoughts-Agent Team, 2025): 5,000 total; 0 removed; 5,000 retained.
- **CLI-Gym** (Lin et al., 2026): 1,552 total; 28 removed; 1,524 retained. Similar to TerminalTraj, we remove images that provide utilities incompatible with our harness (e.g., not supporting `setsid -w`, or adding syscall hooks that disrupt the harness).
- **SETA** (Shen et al., 2026b): 4,567 total; 0 removed; 4,567 retained.
- **CalibForge** (Meng et al., 2026): 5,431 total; 0 removed; 5,431 retained.
- **Terminal-Lego** (Yang et al., 2026): 15,048 total; 0 removed; 15,048 retained.
- **RST** (Li et al., 2026): 37,484 total; 0 removed; 37,484 retained.
- **SWE-Smith** (Yang et al., 2025): 59,136 total; 0 removed; 59,136 retained.
- **Endless Terminals** (Gandhi et al., 2025): 2,492 total; 0 removed; 2,492 retained.
- **Facet** (Shi et al., 2026): 6,020 total; 0 removed; 6,020 retained.

**Eval Dataset Filtering.** Our final Terminal-Bench Lite and Terminal-Bench 2.1 scores use fixed subsets of 96/100 and 88/89 tasks, respectively, with three trials per included task. We exclude the following tasks because our rootless Apptainer backend cannot evaluate them reliably; these are task-level exclusions applied to every model, not model failures removed after scoring.

- **Terminal-Bench Lite** (OpenThoughts-Agent Team, 2026): We exclude *acl-permissions-inheritance*, *etl_checkpoint_resume_bug*, *sales-data-csv-analysis*, and *mlflow-register* because fakeroot prevents reliable setup or grading. We retain 96 tasks.

- **Terminal-Bench 2.1** (Merrill et al., 2026): We exclude *configure-git-webserver* because OpenSSH privilege separation fails under fakeroot. We retain 88 tasks.

Note that for the Terminal-Bench 2.0 results in Figure 1, we evaluate over **all** tasks, as the Daytona backend does not have similar issues.