SCALING AGENTIC DATA FOR LONG-HORIZON TERMINAL INTELLIGENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Terminal agents are moving beyond local operations toward completing real-world professional tasks. Short trajectories can demonstrate commands and tool calls, but they do not fully show how agents accumulate evidence through exploration, use it to inform judgments, continually adjust strategies based on feedback, and turn intermediate results into final deliverables. High-quality long-horizon trajectories capture these *agentic behavioral patterns* and their connections within a single problem-solving process, providing richer experience for sustained autonomous work. However, environment synthesis may complete or simplify the key work underlying these patterns in advance and thus remove the need for exploration, implementation, and verification that the agent was meant to perform. To address this, we introduce **TerminalHorizon**, a fully automated data synthesis engine for long-horizon terminal agents. Starting from real-world professional work, it uses independent execution to identify key work omitted during construction and uses the resulting evidence to guide environment reconstruction. To extend these patterns across stages, it progressively expands tasks under the same goal, introducing new requirements that build on prior results and call for further exploration and improvement. Using this engine, we construct 1.5K task environments and perform supervised fine-tuning of Qwen3.5-27B on the resulting dataset to obtain **TerminalHorizon-27B**. We achieve substantial performance gains on all 12 benchmarks spanning coding-agent tasks, agentic tasks in real-world production settings, and general reasoning. Multiple RSI case studies, each with a 12-hour research budget, further demonstrate the model's ability to continue experimenting beyond an initial solution, improve its methods based on feedback, and retain effective improvements in its final deliverables.

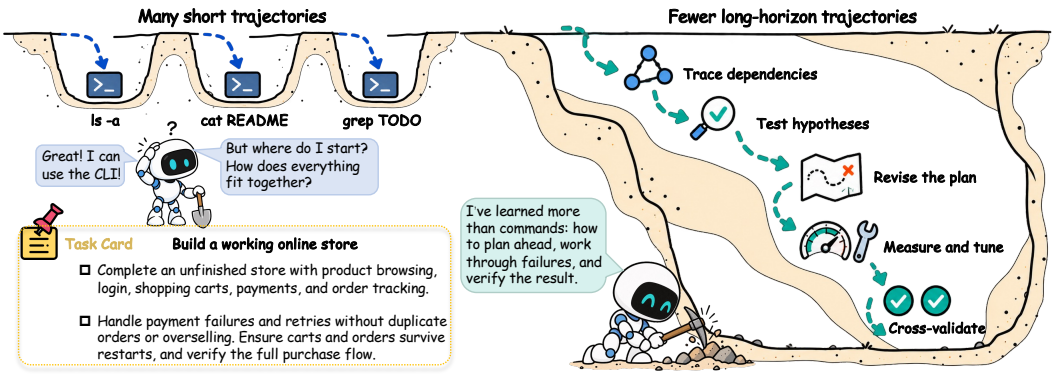

Figure 1: Motivation for TerminalHorizon. Short trajectories can teach agents to perform terminal operations, but not to autonomously complete complex, long-horizon tasks. Our long-horizon trajectories connect agentic behavioral patterns, such as exploration, strategy revision, and verification, into a complete autonomous problem-solving process.

1 INTRODUCTION

LLM agents are moving beyond answering questions in dialogue toward autonomous work in digital environments (Xu et al., 2026; Koh et al., 2026; Dong et al., 2026b; Zheng et al., 2026). The terminal

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.

provides a general interface for this work by bringing code, data, tools, and execution feedback into a shared workspace (Merrill et al., 2026), enabling agents to tackle complex tasks in software development, system administration, and scientific computing (Anthropic, 2025; OpenAI, 2025). As their scope expands from individual operations to long-horizon tasks (Dong et al., 2026b; Ma et al., 2026; Qwen Team, 2026; Ye et al., 2025), agents need to coordinate interdependent stages and adapt to execution feedback to complete these tasks reliably (Deng et al., 2025; Li et al., 2026b; Desai et al., 2026; Yang et al., 2026a). Developing these capabilities requires rich interaction experience during training, but high-quality terminal training data remains scarce.

Recent work on terminal environment synthesis constructs executable and verifiable tasks at scale through skill composition, task generation grounded in technical evidence, and recursive expansion, allowing agents to gain interaction experience through execution and feedback (Wu et al., 2026a; Cheng et al., 2026; Li et al., 2026a; Wu et al., 2026b; Fan et al., 2026). These methods expand task coverage, but their trajectories remain relatively short on average as shown in Figure 4, offering limited experience in sustained problem solving. Short trajectories can capture useful agentic behavioral patterns over a few steps, but the same patterns can extend across multiple stages or even an entire trajectory in long-horizon tasks. For example, agents may repeatedly revise their plans as evidence accumulates, build on intermediate results through a sequence of dependent steps, and revisit earlier work when later verification reveals problems that require further diagnosis and correction. This motivates our focus on high-quality long-horizon trajectories that capture these longer agentic behavioral patterns and show how they work together to complete a task (Section 3).

However, synthesizing these long-horizon trajectories presents two challenges. (i) Complex tasks do not necessarily lead to the intended agentic behavioral patterns. A task may require agents to investigate a problem, develop a solution, and verify it, yet supplied materials may already solve parts of the problem, allowing agents to skip exploration or implementation. These gaps can prevent the intended patterns from fully developing in the collected trajectories. (ii) Even when these patterns appear, long-horizon training data should also capture how they unfold over longer problem-solving processes and connect across stages. More difficult tasks or more interaction turns do not necessarily extend these patterns or strengthen their connections. The challenge is to identify and implement new requirements that go beyond what the original solution can achieve, while keeping the same goal. The added work needs to build on earlier judgments, actions, and results, so that these patterns continue to develop as the task progresses.

To this end, we introduce **TerminalHorizon**, a fully automated data synthesis engine with two stages: Challenge-Guided Long-Horizon Synthesis and Progressive Horizon Expansion. The first stage converts real-world professional problems into executable tasks and uses an agent uninvolved in their construction to solve them. Analysis of the resulting trajectories checks whether the intended agentic behavioral patterns occur and identifies where task construction has simplified or already completed the analysis, implementation, or verification expected of the agent. These findings guide environment reconstruction so that the agent must perform this work during execution. The second stage expands the original task under the same goal, introducing requirements that call for further analysis, implementation, and verification. It makes this added work depend on earlier results and execution feedback, allowing the intended patterns to extend across stages. Together, these stages support the generation of long-horizon trajectories in which the intended agentic behavioral patterns emerge and extend across stages toward task completion.

Using TerminalHorizon, we construct 1,535 task environments and collect 2,852 training trajectories to form **TerminalHorizon-3K**. The trajectories average 52.8 interaction turns, substantially longer than those from the methods compared in Figure 4. Supervised Fine-tuning Qwen3.5-27B on these 2,852 trajectories yields **TerminalHorizon-27B**, which improves over the base model by 14.2% and 13.85% on Terminal-Bench 2.0 and 2.1, respectively. These gains extend to real-world productivity tasks, with improvements of 11.7% in Pass Rate and 23.3% in Partial score on AutomationBench, and 11.3% on VitaBench. Furthermore, case studies on multiple long-horizon recursive self-improvement (RSI) tasks, each with a 12-hour research budget, show how the model continues experimenting beyond an initial solution, refines its methods based on feedback, and incorporates these improvements into final deliverables. Our main contributions are as follows:

- We propose a perspective on high-quality long-horizon terminal data centered on agentic patterns and their connections across stages, emphasizing the substantive work that environments must preserve to elicit them.

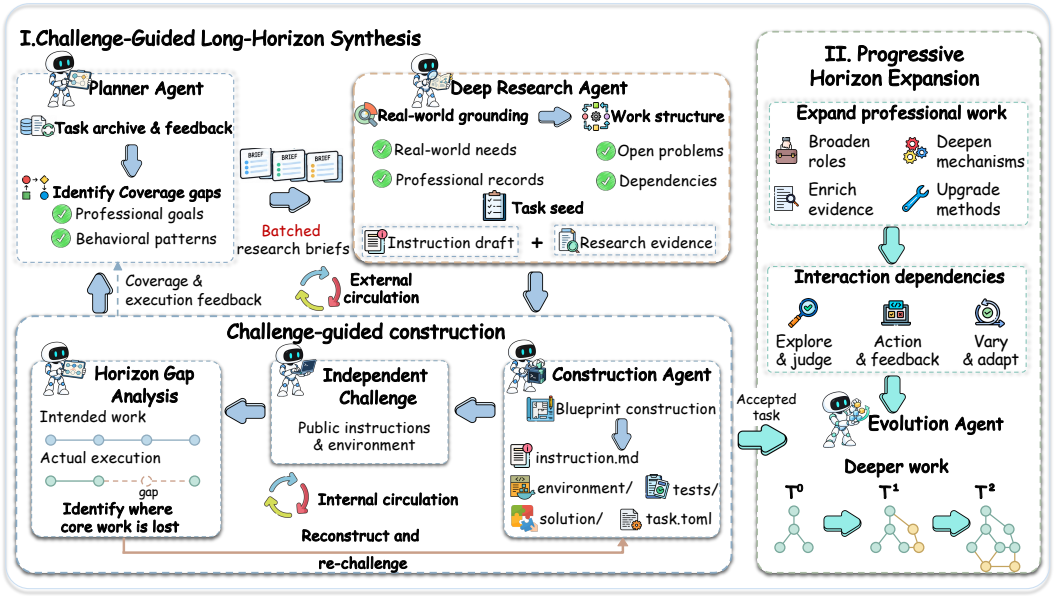

Figure 2: Overview of TerminalHorizon. Challenge-guided synthesis refines task environments through independent execution to preserve core problem-solving work. Progressive horizon expansion deepens tasks along the same goal to support richer agentic behavioral patterns.

- We introduce **TerminalHorizon**, a fully automated data synthesis engine that combines challenge-guided construction with progressive expansion to build **TerminalHorizon-3K**, comprising 1,535 task environments and 2,852 training trajectories.
- Supervised Fine-tuning on only about 3K trajectories improves performance across 12 benchmarks spanning coding, real-world productivity, and general reasoning, while multiple 12-hour RSI case studies demonstrate continued experimentation, improvement, and final delivery.

2 METHOD

TerminalHorizon builds executable environments from real-world professional problems to elicit agentic behavioral patterns (Section 3). In the first stage (Figure 2), independent execution reveals work simplified or completed in construction and guides reconstruction. The second stage extends accepted tasks under the same goal, making added work depend on prior results and feedback.

2.1 TASK FORMULATION AND DESIGN OBJECTIVE

Task formulation. Both stages construct and refine executable terminal task packages, each comprising five components (Merrill et al., 2026):

$$T = (I, E, S, V, M), \qquad \mathcal{P}(T) = (I, E). \tag{1}$$

The public instruction I specifies the goal and constraints, while the executable environment E provides the initial state, tools, and execution mechanisms. The reference solution S, private verifier V, and execution configuration M support task construction and evaluation. During execution, the agent has access only to $\mathcal{P}(T)$ and gathers information through interactions with the environment.

Long-horizon design objective. The work we aim to preserve connects exploration, diagnosis, construction, verification, and strategic coordination, rather than only demonstrating individual terminal operations. We describe these processes using 27 agentic behavioral patterns across five families (Section 3). For example, relationship tracing follows connections between objects to inform decisions. The taxonomy is shown in Figure 3d, with full definitions in Appendix E. Environment design must therefore retain evidence to explore, analysis and implementation left to complete, and feedback that helps agents revise earlier judgments.

2.2 CHALLENGE-GUIDED LONG-HORIZON SYNTHESIS

An inner loop uses independent execution to guide task reconstruction, while an outer loop uses the resulting tasks and review evidence to plan subsequent research batches.

Autonomous Batch Planning. Each batch begins by identifying professional work that would broaden coverage of agentic behavioral patterns. The Planner Agent uses the goals, solution approaches, and feedback in existing tasks, together with prior research and review records, to select the next research directions. For batch b, let $\mathcal{D}_b$ be the accepted task set, $\mathcal{F}_b$ the accumulated records, and n_b the number of available research slots. The Planner produces research briefs:

$$\mathcal{Q}_b = \mathrm{Plan}(\mathcal{D}_b, \mathcal{F}_b; n_b), \qquad \mathcal{Q}_b = \{q_{b,i}\}_{i=1}^{n_b}. \tag{2}$$

Each brief $q_{b,i}$ specifies a scope of work, requirements not yet covered by existing tasks, and real-world sources to investigate, providing the starting point for research.

Evidence-Grounded Deep Research. A research brief sets the direction, but task construction requires a concrete account of the work involved. Following $q_{b,i}$, the Deep Research Agent investigates documentation, code repositories, issue reports, and professional materials to establish real-world needs, technical constraints, and evidence needed to assess completion. The research identifies what the agent needs to discover, which problems require further analysis, implementation, or debugging, and how one result supports subsequent work. These findings form a task seed $z_{b,i}$, containing a draft public instruction and supporting evidence, which is passed to environment construction.

Joint Environment Construction. The task seed specifies the intended work, which construction must support without completing it in advance. From $z_{b,i}$, the Construction Agent synthesizes a concrete instance with a controlled initial state and verification conditions, specifying the available materials and tools, the work left to the agent, and the completion criteria. It then jointly constructs the instruction, environment, reference solution, verifier, and execution configuration around these requirements, producing a candidate package T. Evaluation must follow the public task requirements:

$$\mathcal{C}_{\mathrm{sem}}(V) \subseteq \mathrm{Sem}(I, \mathcal{O}_{\mathrm{pub}}(E))\,. \tag{3}$$

The left-hand side denotes the completion requirements checked by the verifier, and the right-hand side those supported by I and discoverable evidence $\mathcal{O}_{\mathrm{pub}}(E)$. Test data may remain private, but scoring cannot introduce requirements the agent cannot discover. The reference solution gives one feasible approach, while the verifier accepts other implementations meeting the public requirements.

Independent Challenge. An executable package does not establish that the intended work remains necessary during task solving. We therefore ask an Independent Challenger, uninvolved in construction, to attempt the candidate T_r at round r. Its policy π_B accesses only $\mathcal{P}(T_r)$, without research records, the blueprint, the reference solution, or private tests.

At each step, policy $\pi = \pi_B$ selects an action given the public instruction and history $h_t = (o_0, a_0, \ldots, o_t)$, and the environment dynamics P_E update the state and return an observation:

$$a_t \sim \pi(\cdot \mid I, h_t), \qquad (s_{t+1}, o_{t+1}) \sim P_E(\cdot \mid s_t, a_t). \tag{4}$$

The resulting trajectory $\tau_r = (o_0, a_0, \ldots, a_{H_r-1}, o_{H_r})$ records H_r interactions, and the verifier evaluates the final state and artifacts to compute $\rho_r = V_r(s_{H_r})$. The complete package, execution trace, and scoring evidence then support a diagnosis:

$$\tau_r \sim \mathrm{Rollout}(\pi_B; \mathcal{P}(T_r)), \qquad d_r = \mathcal{A}_{\mathrm{syn}}(T_r, \tau_r, \rho_r). \tag{5}$$

Here, $\mathcal{A}_{\mathrm{syn}}$ is the synthesis review, and d_r records runtime defects, information leakage, scoring inconsistencies, and whether the intended work is preserved, providing evidence for reconstruction.

Horizon Gap Analysis and Reconstruction. The trajectory reveals what the agent actually did, but reconstruction also requires identifying which gaps arise from task construction. Horizon Gap Analysis compares the work specified in the seed with the trajectory and traces omitted work to initial materials, object relationships, or tool behavior. For example, if a task asks the agent to implement a core function that the environment already provides, the remaining work may reduce to changing a single constant. Reconstruction restores the need to implement and debug the function while retaining the public specifications needed to do so. The review distinguishes efficient solutions that perform

the required work from work already completed during construction, targeting revisions at the latter. Using the seed $z = z_{b,i}$, trajectory, and diagnosis, the Construction Agent revises the task:

$$T_{r+1} = \mathrm{Refine}(T_r; z, \tau_r, d_r). \tag{6}$$

Refine preserves the goal while revising the initial state, object relationships, implementation requirements, or tool feedback and updating the instruction, reference solution, and verifier accordingly. The revised task returns to independent execution to check whether the restored requirements enter the actual solving process, closing the task-level challenge–reconstruction loop.

To inform subsequent planning, tasks that pass review and runtime checks form $\mathcal{A}_b$, while research, execution, review, and reconstruction records provide new feedback $\Delta\mathcal{F}_b$:

$$\mathcal{D}_{b+1} = \mathcal{D}_b \cup \mathcal{A}_b, \qquad \mathcal{F}_{b+1} = \mathcal{F}_b \cup \Delta\mathcal{F}_b. \tag{7}$$

The Planner uses the updated collection and records in the next batch to identify repeated work and decide which behavioral patterns to target next. Gaps found during execution thus inform both the current task's reconstruction and future research briefs, closing the batch-level outer loop.

2.3 Progressive Horizon Expansion

A solvable task is not the endpoint of our environment construction, but a starting point for exploring the professional problem more deeply. We extend accepted tasks with new requirements under the same goal, linking the added work to earlier judgments, actions, and results so that agentic behavioral patterns can continue across stages.

Progressive Content Expansion. The Evolution Agent examines an accepted parent task $T^{(g)}$ for valuable extensions under the same goal and forms a proposal $\delta^{(g)}$. It then checks whether the parent's solution can directly satisfy the new requirements, identifying the further analysis, implementation, or verification needed. We consider four content-expansion directions, including **broader roles**, using existing artifacts or states to support related responsibilities, and **deeper mechanisms**, making previously simplified relationships, dependencies, or uncertainties part of the work. **Richer evidence** adds records, observations, and professional materials that inform judgments and cross-source analysis, while **method upgrades** require substantial extensions or redesign of the original approach.

Implementing Interaction Dependencies. The Evolution Agent selects from three interaction designs to connect the added requirements to evidence, artifacts, and states encountered or produced during the task, so that later work depends on earlier progress. To support **exploration and judgment**, it makes later decisions depend on evidence gathered through public queries, probes, and observations. In **action and feedback**, observable tool or system effects link earlier operations to subsequent diagnosis, correction, and verification. With **instance variation and adaptation**, objects or states vary across runs under a fixed public specification, requiring the agent to inspect the current instance, adapt its actions, and verify the results. These designs are selected to serve the extension, not to impose a fixed sequence, and are implemented with the content proposal as a new candidate task:

$$\widetilde{T}^{(g+1)} = \mathrm{Expand}\Big(T^{(g)}; \delta^{(g)}\Big). \tag{8}$$

Expand produces a separate candidate $\widetilde{T}^{(g+1)}$ whose public specification combines the original and added requirements. The environment, reference solution, and verifier are updated accordingly, and new scoring conditions continue to follow the public completion contract in equation 3.

An accepted candidate becomes $T^{(g+1)}$, the parent task for the next expansion.

3 Analysis of Synthesized Data

The following analysis illustrates the training experience provided by TerminalHorizon: our model can learn not only how to execute terminal commands but, more importantly, how to organize sustained work toward a goal, observe feedback, revise strategies, and deliver final results.

Trajectory Collection. We use TerminalHorizon to automatically generate 1,535 task environment packages and collect two trajectories per task using DeepSeek-V4-Pro-0813 as the teacher model

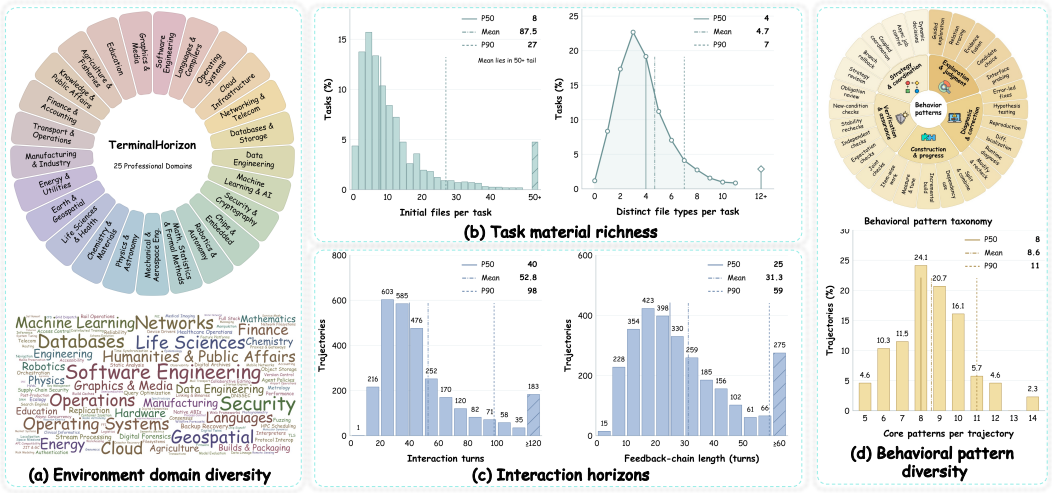

Figure 3: Comprehensive analysis of task environments and trajectories in TerminalHorizon.

and Terminus2 as the execution harness. We then apply the three filtering criteria described in Appendix F, retaining 2,852 trajectories to form TerminalHorizon-3K. We conduct rigorous benchmark decontamination before training (Appendix G).

Broad coverage of real-world domains. As shown in Figure 3a, TerminalHorizon-3K brings real-world work in areas such as chips, physics, and astronomy into a unified terminal interaction format, covering 25 professional domains. This diversity is enabled by the Planner Agent's scheduling and coordination throughout the automated construction process. See Appendix D for details.

Sustained interaction. For a fair comparison, we randomly sample 300 tasks from each of several open-source terminal-task synthesis methods and run experiments under the same teacher protocol and settings. Our trajectories have a mean interaction length of 52.8 turns and a 90th percentile of 98 turns, substantially exceeding those of other methods (Figure 4). We also measure the longest continuous action chain linked by terminal feedback to examine how long feedback remains part of the interaction. Figure 3c shows a median chain length of 25 turns. These results show that longer execution includes sustained sequences of actions and environment feedback, demonstrating how to connect observations with subsequent actions across multiple turns.

Dataset	Mean	P50	P90
Endless Terminals	3.87	4	5
Terminal-Lego	4.50	4	7
FACET	9.44	8	13
TMax	9.66	6	14
RST	10.76	9	17
TermiGen	10.79	7	20
CalibForge	12.22	9	22
TerminalTraj	12.63	6	26
CLI-Gym	15.60	12	30
TerminalHorizon	**52.82**	**40**	**98**



Figure 4: Interaction turns per trajectory: mean, P50, and P90.

Agentic behavioral patterns beyond individual commands. Beyond interaction length, we examine how these actions organize problem solving. As shown in Figure 3d, we group 27 behavioral patterns into five categories: information exploration and judgment, diagnosis and error correction, construction and work progress, verification and assurance, and strategy and execution coordination. These patterns describe how agents organize multi-step actions into a problem-solving process. For example, subproblem decomposition and integration breaks a task into subproblems and combines their results to achieve a shared goal. Detailed pattern definitions are provided in Appendix E. Our trajectories contain an average of 8.6 core patterns, providing diverse behavioral experience spanning exploration, construction, error correction, and verification.

4 Experiments

4.1 Experimental Setup

Evaluation Benchmarks. We evaluate the model on 12 benchmarks spanning terminal task solving, real-world productivity, and general reasoning. **(1) Coding:** Terminal-Bench 1.0, 2.0, and 2.1 (Merrill et al., 2026) assess terminal task-solving performance. **(2) Agentic:** AutomationBench (Shepard & Salimans, 2026), the Airline and Retail domains of τ^2-bench (Barrès et al., 2025), and the English Cross-domain main set of VitaBench (He et al., 2025) test whether experience learned

Table 1: Main results on Terminal-Bench 1.0 and 2.0. **Bold** marks the best score in each benchmark. - denotes unavailable scores, and † denotes results from official model cards.

Model	Base Model	Traj.	Scaffold	TB 1.0	TB 2.0
Terminal task synthesis methods					
CLI-Universe-8B (Hua et al., 2026)	Qwen3-8B	6k	Terminus2	19.1	10.9
TerminalTraj-14B (Wu et al., 2026b)	Qwen2.5-Coder-14B	50.7k	Terminus2	28.9	19.1
TermiGen-32B (Zhu et al., 2026)	Qwen2.5-Coder-32B	3.3k	BashAgent	–	19.3
LiberCoder-32B (Lin et al., 2026)	Qwen3-32B	48.3k	OpenHands	38.9	19.5
SkillSynth-14B (Fan et al., 2026)	Qwen3-14B	10.7k	Terminus2	22.9	19.9
Nemotron-Terminal-14B (Pi et al., 2026)	Qwen3-14B	490.5k	Terminus2	–	20.2
TerminalTraj-32B (Wu et al., 2026b)	Qwen2.5-Coder-32B	50.7k	Terminus2	35.3	22.0
CLI-Universe-14B (Hua et al., 2026)	Qwen3-14B	6k	Terminus2	27.5	23.0
Terminal-Lego-32B (Yang et al., 2026b)	Qwen3-32B	15.4k	Terminus2	–	24.3
OpenThinker-Agent-32B (Raoof et al., 2026)	Qwen3-32B	100k	Terminus2	–	26.2
Nemotron-Terminal-32B (Pi et al., 2026)	Qwen3-32B	490.5k	Terminus2	–	27.4
SkillSynth-32B (Fan et al., 2026)	Qwen3-32B	10.7k	Terminus2	33.8	29.6
LiberCoder-235B (Lin et al., 2026)	Qwen3-235B-A22B	48.3k	OpenHands	46.1	31.0
Terminal-World-32B (Cheng et al., 2026)	Qwen3-32B	5.7k	Terminus2	–	31.5
Meta-Task-32B (Pan et al., 2026)	Qwen3-32B	3.2k	Terminus2	–	31.8
CLI-Universe-32B (Hua et al., 2026)	Qwen3-32B	6k	Terminus2	38.8	33.4
TMax-27B (Ivison et al., 2026)	Qwen3.6-27B	16.5k	Vanillux2Agent	–	42.7
CalibForge-35B-A3B (Meng et al., 2026)	Qwen3.5-35B-A3B	5.4k	CalibForge-Eval	–	47.6
RST-27B (Li et al., 2026a)	Qwen3.5-27B	327.2k	Terminus2	–	49.4
Terminal-Universe-27B (Wu et al., 2026a)	Qwen3.5-27B	32.0k	Terminus2	–	52.8
Ours					
Qwen3.5-27B (base)	–	–	Terminus2	46.3	41.6†
TerminalHorizon-27B	Qwen3.5-27B	3k	Terminus2	**55.0**	**55.8**

from long-horizon trajectories transfers to real-world productivity tasks. **(3) General reasoning:** MATH-500 (Lightman et al., 2023), GSM8K (Cobbe et al., 2021), AIME 2024 (Zhang & Math-AI, 2024), AIME 2025 (Zhang & Math-AI, 2025), and KOR-Bench (Cipher) (Ma et al., 2025) measure how long-horizon interaction training affects general reasoning ability. Detailed evaluation settings are provided in Appendix A.2.

Implementation Details. We apply SFT to Qwen3.5-27B on TerminalHorizon-3K to obtain TerminalHorizon-27B. The maximum sequence length is 262,144 tokens. We use fused AdamW with a learning rate of 1×10^{-5}, weight decay of 0.1, and $(\beta_1, \beta_2) = (0.9, 0.95)$. The learning rate follows a cosine schedule after a 5% warmup, and the gradient clipping threshold is 1.0. Training runs for 3 epochs in BF16 with DeepSpeed ZeRO-3. Unless otherwise noted, we use 144 H20-141GB GPUs across 18 nodes. The full training configuration is provided in Appendix A.1.

4.2 Main Results

Tables 1 and 2 show that TerminalHorizon-27B outperforms its Qwen3.5-27B base model on all 12 benchmarks across three categories. We discuss these gains from three perspectives.

Stronger performance with fewer data. TerminalHorizon-27B achieves 55.0%, 55.8%, and 60.67% on Terminal-Bench 1.0, 2.0, and 2.1, respectively, improving over the base model by 8.7, 14.2, and 13.85%. As shown in Table 1, it also outperforms Terminal-Universe-27B by 3.0% using the same base model and harness but fewer than one-tenth as many training trajectories. This result demonstrates the data efficiency of long-horizon experience and highlights the high information density of our trajectories, which are constructed around sustained problem solving.

Transfer to real-world productivity tasks. Training on our high-quality long-horizon trajectories also improves performance on real-world enterprise automation, customer service, and service tasks spanning multiple scenarios, with gains across all four agentic benchmarks.

On AutomationBench, Pass Rate and Partial scores improve by 11.7% and 23.3%, respectively. VitaBench improves by 11.3%, while τ^2-bench Airline and Retail improve by 5.5% and 2.6%, respectively. These results show that our training benefits extend beyond terminal operations, enabling the model to reliably handle complex environmental feedback throughout real-world long-horizon tasks and apply this capability consistently across tasks.

General reasoning is preserved alongside agentic gains. Performance on the five general reasoning benchmarks is preserved and shows further gains. MATH-500 reaches 98.9%, while AIME 2024 and AIME 2025 reach 96.7% and 93.3%, improving over the base model by 3.4% and 6.6%, respectively. GSM8K and KOR-Bench also improve by 1.0% and 1.1%. Thus, long-horizon training improves agentic task performance while preserving and strengthening general reasoning.

Table 2: Performance across coding, agentic, and general reasoning benchmarks.

Benchmark	Qwen3.5-27B	TerminalHorizon-27B
Coding		
Terminal-Bench 2.1	46.82	**60.67** (+13.85)
Agentic		
AutomationBench (Pass Rate)	5.5	**17.2** (+11.7)
AutomationBench (Partial)	26.4	**49.7** (+23.3)
τ^2-bench Airline	79.0	**84.5** (+5.5)
τ^2-bench Retail	84.9	**87.5** (+2.6)
VitaBench Cross-domain	20.5	**31.8** (+11.3)
General reasoning		
MATH-500	96.9	**98.9** (+2.0)
GSM8K	96.3	**97.3** (+1.0)
AIME 2024	93.3	**96.7** (+3.4)
AIME 2025	86.7	**93.3** (+6.6)
KOR-Bench (Cipher)	90.1	**91.2** (+1.1)

4.3 Analysis and Ablation Studies

Effect of Teacher Model. We randomly sample 100 of the 1,535 tasks to compare interaction characteristics between two teacher models, GLM-5.2 and DeepSeek-V4-Pro-0813. As shown in Table 3, both teachers generate long trajectories on average, indicating that our tasks can provide long-horizon problem-solving experience across teachers. DeepSeek-V4-Pro produces longer feedback chains on average and has a lower rate of exact action repetition. These more sustained interactions are also associated with a higher task completion rate, supporting its selection as our teacher for collecting long-horizon trajectories.

Table 3: Interaction statistics and task success with two teacher models.

Metric	GLM-5.2	DeepSeek-V4-Pro
Mean turns	46.1	53.5
Mean feedback chain	19.2	30.3
Exact repeats (%)	15.8	13.2
Pass rate (%)	**44.0**	**50.0**

Performance across Model Sizes. To examine whether the training benefits of TerminalHorizon extend across model sizes, we also evaluate its effectiveness on Qwen3.5-35B-A3B and Qwen3.5-122B-A10B. As shown in Figure 5, all three configurations achieve double-digit percentage-point gains over their base models, improving by 13.85%, 13.11%, and 14.98%. These results show that the agentic behavioral patterns captured in our long-horizon trajectories support improvements across models, so their training value is not limited to a single base model.

Training Data Scaling. As shown in Figure 6, scaling from 1000 to approximately 3k training trajectories raises performance on Terminal-Bench 2.1 to 60.67%. This demonstrates that data generated by TerminalHorizon can effectively scale agent capabilities. We also observe no sign of performance saturation, suggesting that further expanding high-quality long-horizon training trajectories may yield additional gains.

4.4 Long-Horizon Task Applications

Autonomous long-horizon tasks require agents to continually improve their methods based on feedback and carry effective progress through to final delivery. To evaluate this system-level capability, we conduct experiments on several RSI-Exam cases under a 12-hour autonomous research budget. For the multi-view point cloud fusion task, Figure 7 shows both models' optimization processes and final results. We find that the base model consistently stops early. For example, it voluntarily stops after just 12.1 minutes based on a positive assessment of its current results, without pursuing further research or optimization. In contrast, our model continues experimenting beyond the initial solution,

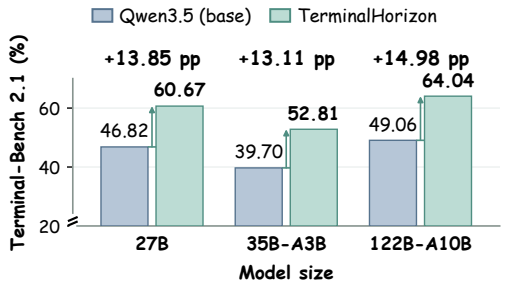

Figure 5: TB 2.1 performance across model sizes.

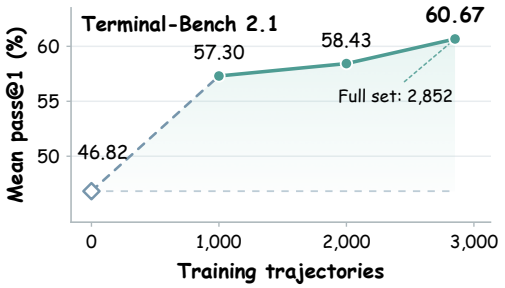

Figure 6: Training-data scaling on TB 2.1.

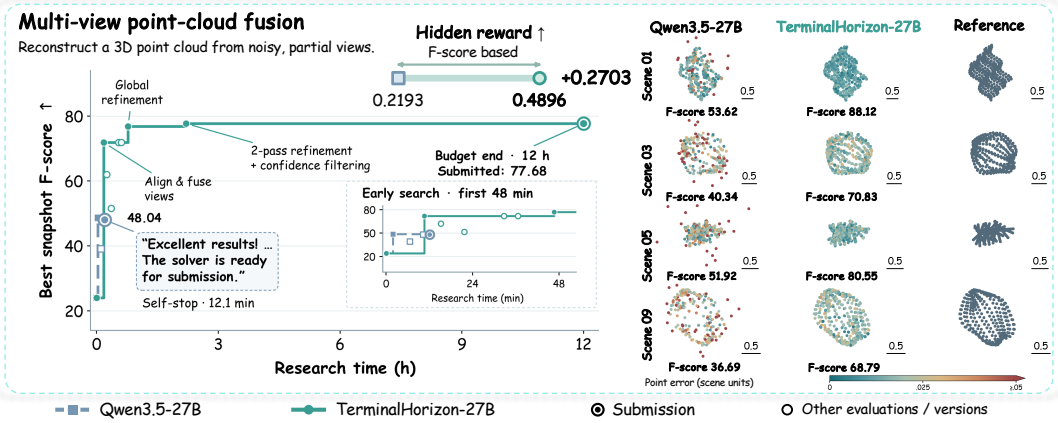

Figure 7: Long-horizon optimization on RSI-Exam. More cases in Appendix C.

revises its methods based on feedback, and retains effective improvements in its final submission. Its hidden evaluation score exceeds the base model's by 27.03%. This demonstrates the practical value of experience with long-horizon problem-solving patterns in helping agents keep seeking improvements and turn verified progress into final results. Additional cases are provided in Appendix C.

5 Related Work

Long-Horizon Agents. Long-horizon agents must retain information, track progress, and coordinate actions across stages (Dong et al., 2026b; Ren et al., 2026; Sun et al., 2026). Harnesses support execution through context, memory, and state management, helping maintain continuity during extended interactions (Ye et al., 2025; Ma et al., 2026; Choi et al., 2026). We complement this support by studying how training develops sustained problem-solving ability, particularly the agentic behavioral patterns that connect exploration, implementation, and verification around a shared goal.

Terminal Environment Synthesis. Environment synthesis scales interaction experience through programmatic construction and evolution (Song et al., 2026; Wang et al., 2026; Dong et al., 2026a). Terminal-task synthesis uses skill composition, technical evidence, and real-world needs (Wu et al., 2026a; Cheng et al., 2026; Wu et al., 2026b; Fan et al., 2026; Shi et al., 2026; Hua et al., 2026; Pi et al., 2026), while recursive reference-solution expansion increases complexity (Li et al., 2026a). Beyond coverage and complexity, we emphasize these patterns' less-studied cross-stage connections: findings inform decisions, feedback guides strategy revision, and intermediate results support delivery.

6 Conclusion

We introduced TerminalHorizon, a fully automated engine for synthesizing long-horizon trajectories that connect agentic behavioral patterns. Through challenge-guided reconstruction and progressive expansion, it links exploration, strategy revision, and verification around a shared goal. Training on these trajectories improves performance on all 12 benchmarks, while RSI-Exam cases show sustained experimentation, feedback-driven improvement, and progress retained through final delivery. These findings highlight the training value of long-horizon behavioral experience in advancing agents from local terminal operations toward sustained autonomous work.

AI USE STATEMENT

In this work, we used generative AI for automated task environment construction and solution trajectory generation as part of our research methodology, and for language polishing and the generation of selected illustration assets during manuscript preparation. The procedures and validation mechanisms for task environment construction and trajectory generation are described in the Method section and the appendix. We did not use generative AI to formulate research questions, develop the core research ideas, or design the research methodology. We manually reviewed all AI-assisted content. For language polishing, we checked that the revised wording preserved the intended meaning. For generated illustration assets, we checked that they accurately conveyed the relevant concepts without introducing misleading information. We take full responsibility for the final content of this work, including text, claims, and illustration assets produced with the aid of generative AI.

ETHICS STATEMENT

This work studies long-horizon task learning for terminal agents through synthetic task environments, generated solution trajectories, and benchmark evaluations, examining how sustained interaction experience affects agent capabilities. Within the scope of this study, we have identified no significant ethical concerns requiring specific disclosure. The procedures for task construction, data filtering, and experimental evaluation are detailed in the main text and the appendix.

REPRODUCIBILITY STATEMENT

To support reproducibility, we describe the task formulation and two-stage synthesis procedure in Section 2, and trajectory collection and experimental settings in Section 3 and Section 4. Appendix A.1 provides the training configuration, while Appendix A.2 specifies the benchmark protocols, including agent frameworks, decoding settings, interaction limits, and scoring procedures. Trajectory filtering and benchmark decontamination are described in Appendices F and G, with the corresponding prompt templates in Appendix J. For the long-horizon RSI case studies, Appendix C details the research setup, evaluation procedure, and evidence underlying the case analyses.

REFERENCES

Anthropic. Claude Code. GitHub repository, 2025. URL https://github.com/anthropics/claude-code. Accessed: 2026-09-03.

Victor Barrès, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2-bench: Evaluating conversational agents in a dual-control environment. *CoRR*, abs/2506.07982, 2025. doi: 10.48550/ARXIV.2506.07982. URL https://doi.org/10.48550/arXiv.2506.07982.

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. *CoRR*, abs/2605.20876, 2026. doi: 10.48550/ARXIV.2605.20876. URL https://doi.org/10.48550/arXiv.2605.20876.

Jae-Woo Choi, Hyungmin Kim, Hyobin Ong, Youngwoo Yoon, Minsu Jang, Dohyung Kim, and Jaehong Kim. Reactree: Hierarchical LLM agent trees with control flow for long-horizon task planning. In Viviana Mascardi, John Thangarajah, Chris Amato, and Louise A. Dennis (eds.), *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2026, Paphos, Cyprus, May 25-29, 2026*, pp. 319–328, 2026. doi: 10.65109/UCGT7089. URL https://dl.acm.org/doi/10.65109/UCGT7089.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021. URL https://arxiv.org/abs/2110.14168.

Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, Chetan Rane, Karmini Sampath, Maya Krishnan, Srivatsa Kundurthy, Sean Hendryx, Zifan Wang, Vijay Bharadwaj, Jeff Holm, Raja Aluri, Chen Bo Calvin Zhang,

Noah Jacobson, Bing Liu, and Brad Kenstler. Swe-bench pro: Can AI agents solve long-horizon software engineering tasks? *CoRR*, abs/2509.16941, 2025. doi: 10.48550/ARXIV.2509.16941. URL https://doi.org/10.48550/arXiv.2509.16941.

Rishi Desai, Jesse Hu, Joan Cabezas, Neel Harsola, Pratyush Shukla, Roey Ben Chaim, Adnan El Assadi, Omkaar Mukund Kamath, Fenil Faldu, Prannay Hebbar, Jiankai Sun, Yiyuan Li, Pramod Srinivasan, Ishan Gupta, Christopher Settles, Daniel Wang, Derek Chen, Pranav Raja, Albert Liu, Marek Suppa, Nevasini Sasikumar, Luyang Kong, Erik Quintanilla, Xiangyi Li, Ivan Bercovich, and Steven Dillmann. Swe-marathon: Can agents autonomously complete ultra-long-horizon software work? *CoRR*, abs/2606.07682, 2026. doi: 10.48550/ARXIV.2606.07682. URL https://doi.org/10.48550/arXiv.2606.07682.

Guanting Dong, Junting Lu, Junjie Huang, Wanjun Zhong, Longxiang Liu, Shijue Huang, Zhenyu Li, Yang Zhao, Xiaoshuai Song, Xiaoxi Li, Jiajie Jin, Yutao Zhu, Hanbin Wang, Fangyu Lei, Qinyu Luo, Mingyang Chen, Zehui Chen, Jiazhan Feng, Ji-Rong Wen, and Zhicheng Dou. Agent-world: Scaling real-world environment synthesis for evolving general agent intelligence. *CoRR*, abs/2604.18292, 2026a. doi: 10.48550/ARXIV.2604.18292. URL https://doi.org/10.48550/arXiv.2604.18292.

Guanting Dong, Xiaoshuai Song, Yuyang Hu, Jiajie Jin, Chenghao Zhang, Yifei Chen, Xiaoxi Li, Huaying Yuan, Xinyu Yang, Tongyu Wen, Jiejun Tan, Hongjin Qian, Shijue Huang, Junting Lu, Zhenyu Li, Wanjun Zhong, Yutao Zhu, Tat-Seng Chua, Zhicheng Dou, and Ji-Rong Wen. Towards long-horizon agents: A survey. *Preprints*, July 2026b. doi: 10.20944/preprints202607.1328.v1. URL https://doi.org/10.20944/preprints202607.1328.v1.

Zhiyuan Fan, Tinghao Yu, Yuanjun Cai, Jiangtao Guan, Yun Yang, Dingxin Hu, Jiang Zhou, Xing Wu, Zhuo Han, Feng Zhang, and Lilin Wang. Toward scalable terminal task synthesis via skill graphs. *CoRR*, abs/2604.25727, 2026. doi: 10.48550/ARXIV.2604.25727. URL https://doi.org/10.48550/arXiv.2604.25727.

Wei He, Yueqing Sun, Hongyan Hao, Xueyuan Hao, Zhikang Xia, Qi Gu, Chengcheng Han, Dengchang Zhao, Hui Su, Kefeng Zhang, Man Gao, Xi Su, Xiaodong Cai, Xunliang Cai, Yu Yang, and Yunke Zhao. Vitabench: Benchmarking LLM agents with versatile interactive tasks in real-world applications. *CoRR*, abs/2509.26490, 2025. doi: 10.48550/ARXIV.2509.26490. URL https://doi.org/10.48550/arXiv.2509.26490.

Zhanbo Hua, Yifan Yao, Weihao Xie, Yongchi Zhao, Minghao Liu, Ruizhi Qiu, Zhewei Huang, Zun Wang, Yiyan Ji, Yunhai Ye, Letian Zhu, Xinping Lei, Han Li, Zhiyuan Ma, Zili Wang, Zhaoxiang Zhang, and Jiaheng Liu. Cli-universe: Towards verifiable task synthesis engine for terminal agents. *CoRR*, abs/2606.22883, 2026. doi: 10.48550/ARXIV.2606.22883. URL https://doi.org/10.48550/arXiv.2606.22883.

Hamish Ivison, Junjie Oscar Yin, Rulin Shao, Teng Xiao, Nathan Lambert, and Hannaneh Hajishirzi. Tmax: A simple recipe for terminal agents. *CoRR*, abs/2606.23321, 2026. doi: 10.48550/ARXIV.2606.23321. URL https://doi.org/10.48550/arXiv.2606.23321.

Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Multi-agent computer use. *CoRR*, abs/2606.01533, 2026. doi: 10.48550/ARXIV.2606.01533. URL https://doi.org/10.48550/arXiv.2606.01533.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In Jason Flinn, Margo I. Seltzer, Peter Druschel, Antoine Kaufmann, and Jonathan Mace (eds.), *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pp. 611–626. ACM, 2023. doi: 10.1145/3600006.3613165. URL https://doi.org/10.1145/3600006.3613165.

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. *CoRR*, abs/2608.05466, 2026a. doi: 10.48550/ARXIV.2608.05466. URL https://doi.org/10.48550/arXiv.2608.05466.

Zongxia Li, Zhongzhi Li, Yucheng Shi, Ruhan Wang, Junyao Yang, Zhichao Liu, Xiyang Wu, Anhao Li, Yue Yu, Ninghao Liu, Lichao Sun, Haotao Mi, and Leowei Liang. Long-horizon-terminal-bench: Testing the limits of agents on long-horizon terminal tasks with dense reward-based grading. *CoRR*, abs/2607.08964, 2026b. doi: 10.48550/ARXIV.2607.08964. URL https://doi.org/10.48550/arXiv.2607.08964.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *CoRR*, abs/2305.20050, 2023. doi: 10.48550/ARXIV.2305.20050. URL https://doi.org/10.48550/arXiv.2305.20050.

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. *CoRR*, abs/2602.10999, 2026. doi: 10.48550/ARXIV.2602.10999. URL https://doi.org/10.48550/arXiv.2602.10999.

Kaijing Ma, Xinrun Du, Yunran Wang, Haoran Zhang, Zhoufutu Wen, Xingwei Qu, Jian Yang, Jiaheng Liu, Minghao Liu, Xiang Yue, Wenhao Huang, and Ge Zhang. Kor-bench: Benchmarking language models on knowledge-orthogonal reasoning tasks. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/c6f5851a0d9cb435ed8b50e87bd6a257-Paper-Conference.pdf.

Ziyu Ma, Hailang Huang, Shun Zou, Yong Wang, Shidong Yang, Yiming Hu, Fei Wei, and Xiangxiang Chu. Longhorizon-harness: Advancing long-horizon agents for real-world tasks. *CoRR*, abs/2608.01964, 2026. doi: 10.48550/ARXIV.2608.01964. URL https://doi.org/10.48550/arXiv.2608.01964.

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. *CoRR*, abs/2608.06352, 2026. doi: 10.48550/ARXIV.2608.06352. URL https://doi.org/10.48550/arXiv.2608.06352.

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

OpenAI. Codex cli, 2025. URL https://github.com/openai/codex. Accessed: 2026-09-03.

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

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

Qwen Team. Qwen3.8-max: A new bar for coding and cowork, August 2026. URL https://qwen.ai/blog?id=qwen3.8.

Negin Raoof, Richard Zhuang, Marianna Nezhurina, Etash Kumar Guha, Atula Tejaswi, Ryan Marten, Charlie F. Ruan, Tyler Griggs, Alexander Glenn Shaw, Hritik Bansal, Estefany 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. *CoRR*, abs/2606.24855, 2026. doi: 10.48550/ARXIV.2606.24855. URL https://doi.org/10.48550/arXiv.2606.24855.

Qingnan Ren, Shun Zou, Shiting Huang, Ziao Zhang, Kou Shi, Zhen Fang, Yiming Zhao, Yu Zeng, Qisheng Su, Lin Chen, Yong Wang, Zehui Chen, Xiangxiang Chu, and Feng Zhao. Saasbench: Exploring the boundaries of coding agents in long-horizon enterprise saas engineering. *CoRR*, abs/2605.17526, 2026. doi: 10.48550/ARXIV.2605.17526. URL https://doi.org/10.48550/arXiv.2605.17526.

Daniel Shepard and Robin Salimans. Automationbench. *CoRR*, abs/2604.18934, 2026. doi: 10.48550/ARXIV.2604.18934. URL https://doi.org/10.48550/arXiv.2604.18934.

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. *CoRR*, abs/2608.18580, 2026. doi: 10.48550/ARXIV.2608.18580. URL https://doi.org/10.48550/arXiv.2608.18580.

Xiaoshuai Song, Haofei Chang, Guanting Dong, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. Envscaler: Scaling tool-interactive environments for LLM agent via programmatic synthesis. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics, ACL 2026, San Diego, California, United States, July 2-7, 2026*, pp. 8326–8357. Association for Computational Linguistics, 2026. doi: 10.18653/V1/2026.FINDINGS-ACL.407. URL https://doi.org/10.18653/v1/2026.findings-acl.407.

Yiyou Sun, Xinyang Han, Weichen Zhang, Yuanbo Pang, Tianyu Wang, Yuhan Cao, Yixiao Huang, Chris Duroiu, Haoyun Zhang, Jeffrey Lin, Weishu Zhang, Tyler Zeng, Ying Yan, Bo Liu, Hanson Wen, Mingyang Xu, Xiaoyuan Liu, Zimeng Chen, Weiyan Shi, Amanda Dsouza, Vincent Sunn Chen, Patrick Bryant, Carl Boettiger, Yamini Rangan, Bradley Rothenberg, Kyle Steinfeld, Arvind Rao, Tapio Schneider, Georgios Yannakakis, Laure Zanna, Kaan Ozbay, Ida Sim, Tarek Zohdi, George Em Karniadakis, Jack Gallant, Teresa Head-Gordon, Yushan Li, Wenxi Deng, Tao Sun, Huiqi Wang, Zhun Wang, Justin Xu, Chris Yuhao Liu, Yafei Cheng, Rongwang Hu, Aras Bacho, Shengcao Cao, Zengyi Qin, Yixiong Chen, Hengduan Fan, Hao Liu, Lin Zeng, Shashank Muralidhar Bharadwaj, Litian Gong, Yingxuan Yang, Maojia Song, Ruheng Wang, Zongzheng Zhang, Honglin Bao, Shuo Lu, Jianhong Tu, Zhonghua Wang, Zheng Zhang, Zijiao Chen, Yanqiong Jiang, Zhendong Li, Bohan Lyu, Chang Ma, Peiran Xu, Benran Zhang, Shangding Gu, Haoyue Hua, Haoyang Li, Wanzhe Liao, Chengzhi Liu, Junbo Peng, Haoran Sun, Zechen Xu, Bo Chen, Jiayi Cheng, Yi Jiang, Keying Kuang, Yuan Li, Youbang Pan, Ziyan Rao, Alexander Schubert, Yifan Shen, Vincent Siu, Xiatao Sun, Kangqi Zhang, Xiaopan Zhang, Yuchen Zhu, Ishaan Singh Chandok, Lei Ding, Jingxuan Fan, Andrew Glover, Jiaming Hu, Yiran Hu, Wenbo Huang, and et al. Agents' last exam. *CoRR*, abs/2606.05405, 2026. doi: 10.48550/ARXIV.2606.05405. URL https://doi.org/10.48550/arXiv.2606.05405.

Zhaoyang Wang, Canwen Xu, Boyi Liu, Yite Wang, Siwei Han, Zhewei Yao, Huaxiu Yao, and Yuxiong He. Agent world model: Infinity synthetic environments for agentic reinforcement learning. *CoRR*, abs/2602.10090, 2026. doi: 10.48550/ARXIV.2602.10090. URL https://doi.org/10.48550/arXiv.2602.10090.

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. *arXiv preprint arXiv:2609.04148*, 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, and Chenghua Lin. Large-scale terminal agentic trajectory generation from dockerized environments. *CoRR*, abs/2602.01244, 2026b. doi: 10.48550/ARXIV.2602.01244. URL https://doi.org/10.48550/arXiv.2602.01244.

Zhangchen Xu, Junda Chen, Yue Huang, Dongfu Jiang, Jiefeng Chen, Hang Hua, Zijian Wu, Zheyuan Liu, Zexue He, Lichi Li, Shizhe Diao, Jiaxin Pei, Jinsung Yoon, Hao Zhang, Mengdi Wang, Radha Poovendran, Misha Sra, Alex Pentland, and Zichen Chen. Autolab: Can frontier models solve long-horizon auto research and engineering tasks? *CoRR*, abs/2606.05080, 2026. doi: 10.48550/ARXIV.2606.05080. URL https://doi.org/10.48550/arXiv.2606.05080.

John Yang, Kilian Lieret, Jeffrey Ma, Parth Thakkar, Dmitrii Pedchenko, Sten Sootla, Emily McMilin, Pengcheng Yin, Rui Hou, Gabriel Synnaeve, Diyi Yang, and Ofir Press. Programbench: Can language models rebuild programs from scratch? *CoRR*, abs/2605.03546, 2026a. doi: 10.48550/ARXIV.2605.03546. URL https://doi.org/10.48550/arXiv.2605.03546.

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? *CoRR*, abs/2606.03461, 2026b. doi: 10.48550/ARXIV.2606.03461. URL https://doi.org/10.48550/arXiv.2606.03461.

Rui Ye, Zhongwang Zhang, Kuan Li, Huifeng Yin, Zhengwei Tao, Yida Zhao, Liangcai Su, Liwen Zhang, Zile Qiao, Xinyu Wang, Pengjun Xie, Fei Huang, Siheng Chen, Jingren Zhou, and Yong Jiang. Agentfold: Long-horizon web agents with proactive context management. *CoRR*, abs/2510.24699, 2025. doi: 10.48550/ARXIV.2510.24699. URL https://doi.org/10.48550/arXiv.2510.24699.

Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2024, 2024. URL https://huggingface.co/datasets/math-ai/aime24.

Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2025, 2025. URL https://huggingface.co/datasets/math-ai/aime25.

Jingsheng Zheng, Xinyuan Fang, Jintian Zhang, Zhengke Gui, Huajun Chen, and Ningyu Zhang. Onedayagent: Towards a long-horizon harness for autonomous agents. *CoRR*, abs/2608.05013, 2026. doi: 10.48550/ARXIV.2608.05013. URL https://doi.org/10.48550/arXiv.2608.05013.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyan Luo. Llamafactory: Unified efficient fine-tuning of 100+ language models. In Yixin Cao, Yang Feng, and Deyi Xiong (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pp. 400–410. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-DEMOS.38. URL https://doi.org/10.18653/v1/2024.acl-demos.38.

Kaijie Zhu, Yuzhou Nie, Yijiang Li, Yiming Huang, Jialian Wu, Jiang Liu, Ximeng Sun, Zhenfei Yin, Lun Wang, Zicheng Liu, Emad Barsoum, William Yang Wang, and Wenbo Guo. Termigen: High-fidelity environment and robust trajectory synthesis for terminal agents. *CoRR*, abs/2602.07274, 2026. doi: 10.48550/ARXIV.2602.07274. URL https://doi.org/10.48550/arXiv.2602.07274.

A Implementation Details

A.1 Training Details

We use LLaMA-Factory (Zheng et al., 2024) to perform full-parameter supervised fine-tuning (SFT) on multi-turn trajectories from TerminalHorizon-3K for Qwen3.5-27B, Qwen3.5-35B-A3B, and Qwen3.5-122B-A10B[1]. All three models follow the same training recipe in Table 4, and we report the final checkpoint after three epochs.

Table 4: Shared SFT hyperparameters for all three model sizes.

Hyperparameter	Value
Optimizer	AdamW (β_1=0.9, β_2=0.95)
Peak learning rate	1.0×10^{-5}
LR scheduler	Cosine
Warmup ratio	0.05
Weight decay	0.1
Epochs	3
Maximum sequence length	262,144
Sequence packing	Disabled
Per-device train batch size	1
Precision	BF16
Distributed optimization	DeepSpeed ZeRO-3
Attention backend	FlashAttention-2
Random seed	42

A.2 Evaluation Protocols

General settings. All evaluated models are served through vLLM (Kwon et al., 2023) on eight H20 GPUs, with a tensor parallelism degree of 8 and a context window of 262,144 tokens. For each benchmark, the base and TerminalHorizon models use the same task set, agent framework, decoding settings, and scoring procedure. Terminal-Bench, AutomationBench, and general reasoning evaluations use a temperature of 1.0, top-$p = 0.95$, top-$k = 20$, and a maximum of 65,536 output tokens per response. Decoding settings for VitaBench and τ^2-bench are specified below.

Terminal-Bench. We evaluate on Terminal-Bench 1.0, 2.0, and 2.1 (Merrill et al., 2026) using Harbor[2] and Terminus-2 with a JSON parser. Terminal-Bench 1.0 contains 80 tasks and uses Daytona sandboxes, each allocated 2 CPUs, 4 GiB of memory, and 8 GiB of storage. Terminal-Bench 2.0 and 2.1 each contain 89 tasks and use E2B sandboxes. All three versions allow the agent three hours per task, while verifiers retain the time limits specified by each task. Terminus-2 triggers proactive summarization when the estimated remaining context capacity falls below 131,072 tokens. We run each task independently three times and report mean pass@1, with up to two retries for infrastructure errors.

AutomationBench. We evaluate AutomationBench v1.0.6 (Shepard & Salimans, 2026) with 600 rollouts over the full public set, using the official `api` tool set and a limit of 50 steps per trajectory. Each model is evaluated in three independent runs, and we report the mean official Partial score (the fraction of assertions passed) and Pass Rate (all assertions passed). Interrupted samples are retried for up to three rounds; samples that remain incomplete receive a score of zero.

VitaBench Cross-domain. We use the main English Cross-domain set of VitaBench 1.0 (He et al., 2025), which contains 100 tasks spanning food delivery, in-store services, and travel. Each task is

[1] https://huggingface.co/collections/Qwen/qwen35

[2] https://github.com/harbor-framework/harbor

run four times with a limit of 300 interaction steps, an error limit of 10, and a random seed of 300. The agent uses a temperature of 0 and a maximum of 65,536 output tokens per response. The user simulator and trajectory judge both use `deepseek-v4-pro-0813`, with a temperature of 0 and an output limit of 32,768 tokens. We use the official rubric-based sliding-window trajectory evaluator and report Avg@4, the mean success rate over all 400 runs.

τ^2-bench. We evaluate the Airline (50 tasks) and Retail (114 tasks) domains on the base split of τ^2-bench v1.0.1 (Barrès et al., 2025). We use the official `llm_agent` and user simulator, with four independent runs per task, a limit of 200 interaction steps, an error limit of 10, and a random seed of 300. The agent, user simulator, and natural-language assertion judge all use a temperature of 0; the latter two use `deepseek-v4-pro-0813`. We do not specify an additional per-response output limit for this evaluation. For each domain, we report the mean task success rate across all repeated runs. Model or API errors are retried up to three times, and these technical retries do not count as additional trials.

General reasoning. We evaluate on the 500 problems in MATH-500 (Lightman et al., 2023), the 1,319 test problems in GSM8K (Cobbe et al., 2021), the 30 problems each in AIME 2024 (Zhang & Math-AI, 2024) and AIME 2025 (Zhang & Math-AI, 2025), and the 250 problems in KOR-Bench Cipher (Ma et al., 2025). The first four benchmarks use the native evaluators in EvalScope 1.11.1[3], while KOR-Bench uses its official zero-shot prompts and answer normalization procedure. For AIME 2024 and AIME 2025, we use the ModelScope datasets `evalscope/aime24`[4] and `evalscope/aime25`[5], respectively, through EvalScope 1.11.1, evaluating all 30 problems in the test split of the default subset for each year. We generate one answer per problem and score it according to each benchmark's answer-matching rules, reporting mean accuracy over three independent runs.

B Additional Experimental Analysis

B.1 Performance across Task Categories

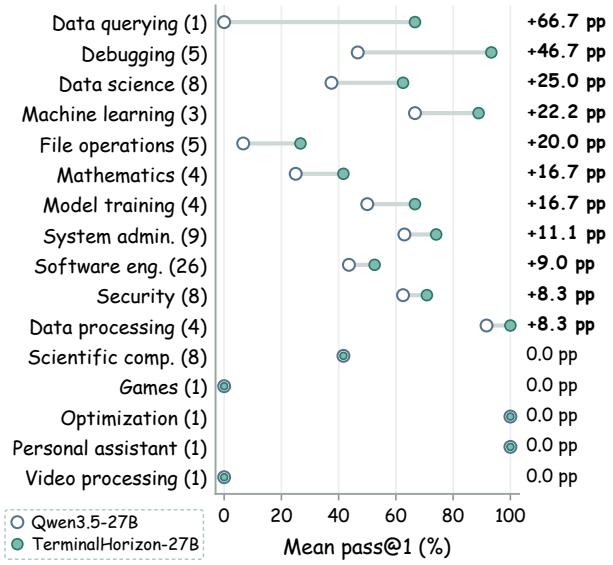

Figure 8: Category-wise mean pass@1 on Terminal-Bench 2.1.

[3] https://github.com/modelscope/evalscope
[4] https://modelscope.cn/datasets/evalscope/aime24
[5] https://modelscope.cn/datasets/evalscope/aime25

We compare category-wise mean pass@1 on Terminal-Bench 2.1 to examine how the training benefits are distributed across professional tasks. As shown in Figure 8, TerminalHorizon-27B matches or outperforms Qwen3.5-27B in all 16 categories, with improvements in 11. Gains are particularly large in debugging, data science, and machine learning, reaching 46.7, 25.0, and 22.2 percentage points, respectively, and also extend to system administration, security, and data processing. This distribution shows that the improvements are not concentrated in a single task type but extend across a broad range of professional problem-solving tasks.

B.2 Data Efficiency

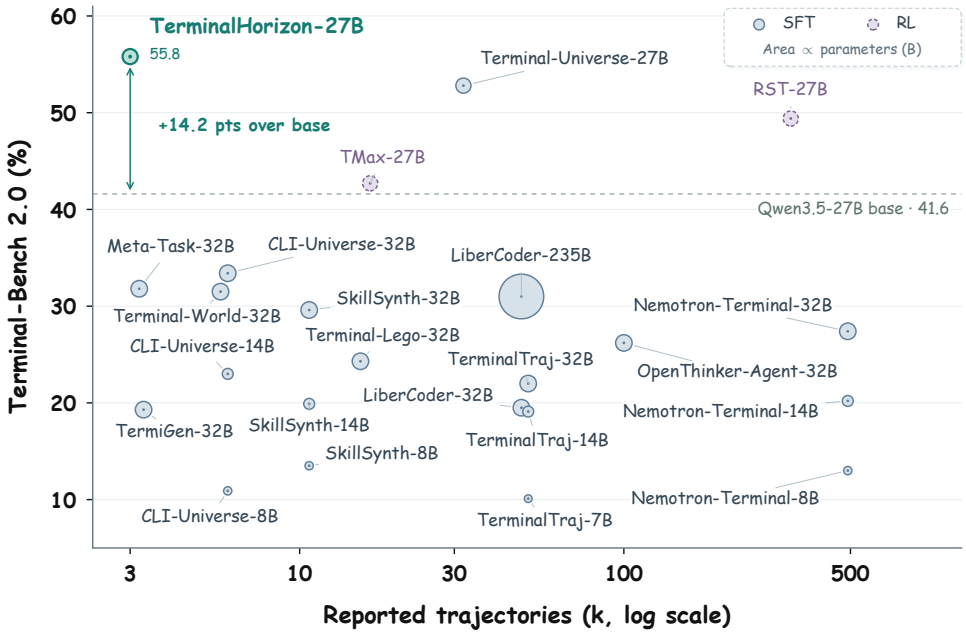

Figure 9: Data efficiency on Terminal-Bench 2.0. TerminalHorizon-27B achieves 55.8% with approximately 3k training trajectories.

Figure 9 compares Terminal-Bench 2.0 performance against the reported number of training trajectories. With 2,852 trajectories (approximately 3k), TerminalHorizon-27B achieves 55.8%, the highest score among the models shown. In particular, it outperforms Terminal-Universe-27B by 3.0 percentage points while using fewer than one-tenth as many trajectories (approximately 3k versus 32k). Both models use the Qwen3.5-27B base model and Terminus-2 harness, providing a more closely matched comparison within a plot that includes different model sizes and training methods. These results support the data efficiency of long-horizon experience organized around sustained problem solving, which enables strong terminal-task performance with a limited number of training trajectories.

C Additional Long-Horizon Task Applications

We examine four RSI-Exam[6] case studies to understand how terminal-task training transfers to sustained autonomous research. The cases span combinatorial optimization, physical inverse problems, game strategy, and geometric reconstruction. Their shared challenge is not merely to produce an executable solution, but to identify opportunities for further improvement, design informative experiments, and preserve effective changes in the final submission. These cases complement the aggregate benchmark results by showing how agents organize their work around an ongoing objective.

[6] https://github.com/aiming-lab/RSI-Exam

Table 5: RSI-Exam results (12-hour budget; one run per model per case). Normalized rewards $\times 100$. Time excludes setup and final verification.

Case	Hidden set score ↑			Research time (h)	
	Base	Ours	Gain	Base	Ours
Robot routing and macros	14.63	39.72	+25.09	11.42	12.00
Tidal friction inversion	27.41	48.15	+20.74	1.12	12.00
2048 strategy	37.54	45.25	+7.71	1.28	12.00
Multi-view point-cloud fusion	21.93	48.96	+27.03	0.20	12.00

For each case, we compare Qwen3.5-27B and TerminalHorizon-27B using the same task, research instructions, decoding settings, and evaluation procedure. Both models operate in an E2B environment with 8 GiB of memory through a Terminus-2-based research harness, with a research budget of 12 hours. Agents can inspect the supplied materials, revise their implementations, and run development experiments, but do not receive hidden evaluation results during research. The task's native verifier scores the final submission rather than a historical checkpoint selected afterward by the evaluator. Agents may also finish voluntarily before the deadline.

Table 5 summarizes the results for all cases. Each result corresponds to one research run and is based on the final submission's performance on the hidden test set. We convert each task's native normalized reward to a 0–100 scale, representing an optimization score rather than a task success rate. We use the original action–observation traces, experiment records, saved implementations, and final evaluation outputs to analyze how models sustain research and preserve effective improvements.

C.1 Warehouse Robot Routing and Macro Optimization

Task and source of difficulty. A robot must collect colored balls in a warehouse with obstacles and deliver each ball to its matching basket. The controller supports movement, turning, swapping, and the recording and replay of macros. The goal is to complete all deliveries with as few button presses as possible, subject to a limit on the number of primitive operations after macro expansion. This couples route planning with program compression: a shorter geometric route need not yield a shorter control program, and macros must remain valid as the robot's position, orientation, and carried object change. The public evaluator provides 100 instances, while the final verifier uses 200 unseen instances and penalizes incomplete or invalid solutions. Figure 10 connects method revisions with public-set progress and final hidden-set results.

From local improvements to a joint objective. TerminalHorizon-27B first builds an orientation-aware route planner, then adds single-macro compression and sequential macro replacement. Later experiments expand the route search with alternative initializations, route-segment reversals, and pairwise swaps. On the full public test set, mean cost falls from 651.42 in the initial measured version to 281.92 after pairwise swaps are introduced. The agent does not accept a candidate merely because it is more complex: it reverts a C++ local-search variant that takes longer and yields a slightly worse cost of 281.96. In contrast, it retains an exact dynamic-programming route optimizer for smaller instances as an additional candidate. The final method selects routes by their compressed button count, aligning search with the actual objective rather than uncompressed path length. The submitted implementation achieves a public mean cost of 281.11.

Preserving progress through final delivery. The base model also conducts substantial search, reaching a best recorded cost of 286.05 on the full public set. However, later revisions remove macro compression, and its actual submission has a cost of 466.35. It voluntarily stops after 11.42 hours, whereas TerminalHorizon-27B retains its stronger compressed solution through the research deadline. Their final hidden-set normalized scores are 14.63 and 39.72, respectively. Both agents discover competitive intermediate solutions, but only TerminalHorizon-27B preserves that progress in the evaluated implementation. Maintaining a reliable best-so-far solution and checking the final submission are part of long-horizon problem solving, not administrative steps after the substantive work is complete.

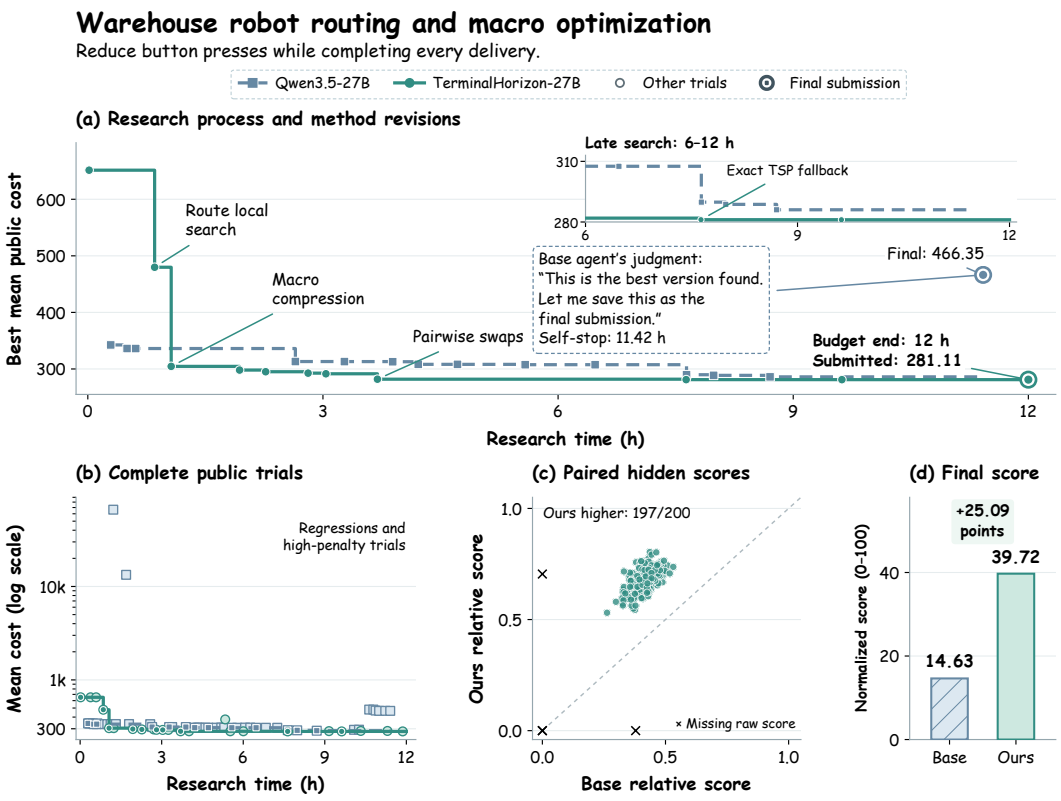

Figure 10: Robot routing and macro optimization. (a) Best mean button cost on 100 public instances. Rings mark final submissions. (b) All complete public trials, including regressions. (c) Paired relative scores (reference length / solver cost) on all 200 hidden instances. Crosses mark missing raw scores, mapped to zero by the grader. (d) Final normalized scores.

C.2 Tidal Seabed Friction Field Inversion

Task and source of difficulty. The agent must estimate a spatially varying seabed friction field from bathymetry, boundary forcing, and sparse observations of complex tidal elevation. Some gauges contain large, unmarked observation errors. Because the forward simulator smooths the effect of friction changes on tidal elevation, the inverse problem is ill-posed: fitting the observed gauges does not guarantee accurate predictions elsewhere. The submitted solver must return a bounded 24×24 log-friction field. Quality is measured by tidal-elevation root-mean-square error (RMSE) at separate gauges, rather than by direct agreement between the estimated and true friction fields. Development uses six public basins, and final evaluation uses eight hidden cases. Figure 11 shows the research process, spatial error comparisons, and results on all eight hidden basins.

Revising the model of the inverse problem. The base model recognizes the need for spatial regularization and outlier handling, adopting a 4×4 bilinear parameterization, coordinate descent, and iterative residual-based observation rejection. After achieving a public mean RMSE of 0.023616, it voluntarily stops after 1.12 hours. Before stopping, it judges the task complete because the solver has improved significantly over the initial implementation and is ready for evaluation. TerminalHorizon-27B instead continues to compare spatial bases, robust losses, and regularization methods rather than treating the initial parameterization as fixed. Its final method combines blockwise spatial coefficients with depth and log-depth features, penalizes differences between adjacent coefficients, and fits the complex-valued observations with a Huber loss. It then removes the two gauges with the largest residuals and refits. These choices address different sources of error, including limited spatial representation, unstable variation in the friction field, and corrupted observations.

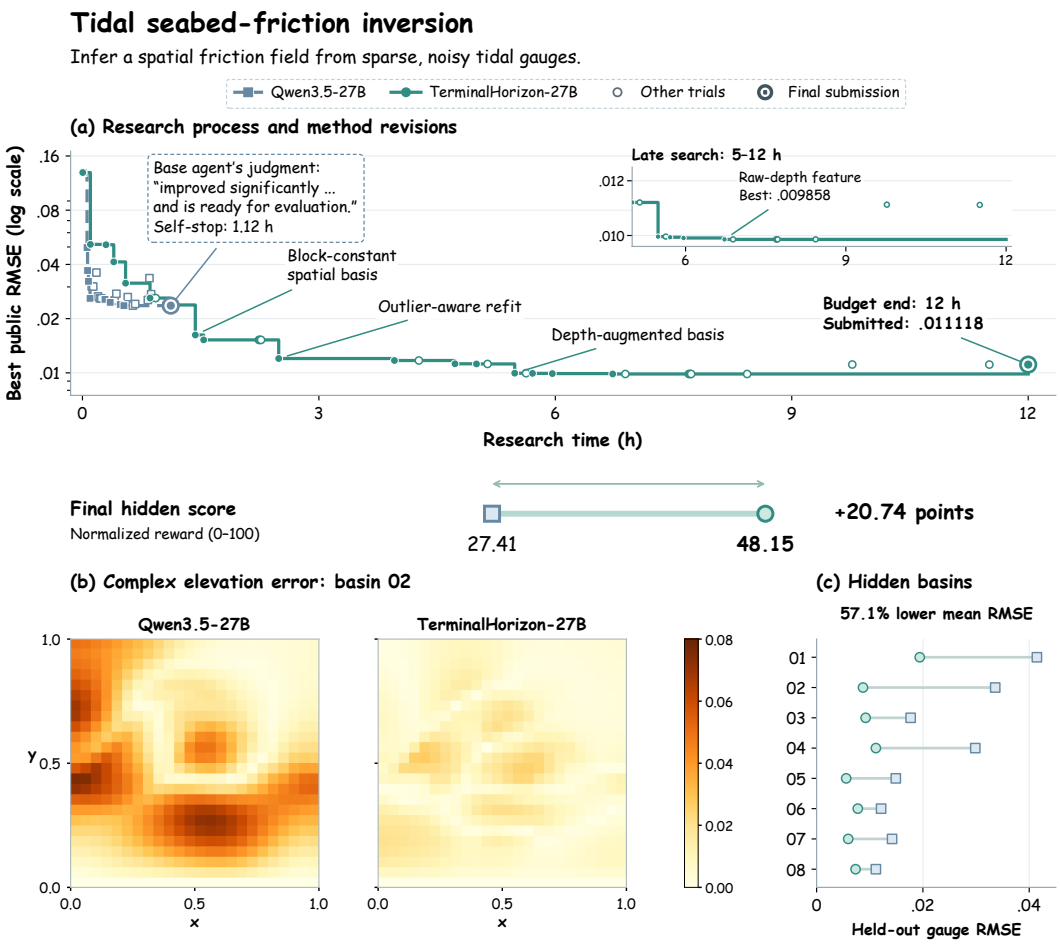

Figure 11: Tidal friction inversion. (a) Best mean public RMSE (six basins). (b) Unsmoothed 24×24 complex-elevation errors with a shared scale. Basin 02 is selected post hoc for the largest relative gauge-RMSE reduction. (c) Gauge RMSE on all eight hidden basins.

Generalization beyond the fitted observations. The trajectory includes comparisons on reserved public cases and additional synthetic basins, allowing the agent to test whether a configuration's improvements extend beyond the cases used for tuning. The final submission achieves a public mean RMSE of 0.011118, compared with 0.023616 for the base model. On hidden cases, mean RMSE decreases from 0.021873 to 0.009386, a reduction of approximately 57.1%, while the normalized score rises from 27.41 to 48.15. This case shows how sustained research goes beyond additional optimization iterations by using experimental feedback to revise the friction-field representation and outlier handling, then checking whether those changes generalize to cases not used for tuning.

C.3 2048 Strategy Optimization

Task and source of difficulty. The agent must improve a deterministic policy for the stochastic game 2048. Given a board, the policy must select a legal move using only the Python standard library while meeting a per-move latency limit and a shared CPU time budget. Development uses eight public game seeds, and the final policy is evaluated on 16 additional, non-overlapping seeds. These requirements create competing objectives: deeper search may improve decisions but increases computational cost, while heuristics that work well on a few seeds may not generalize to other random tile sequences. Figure 12 shows key search revisions and compares the final policies on each hidden game.

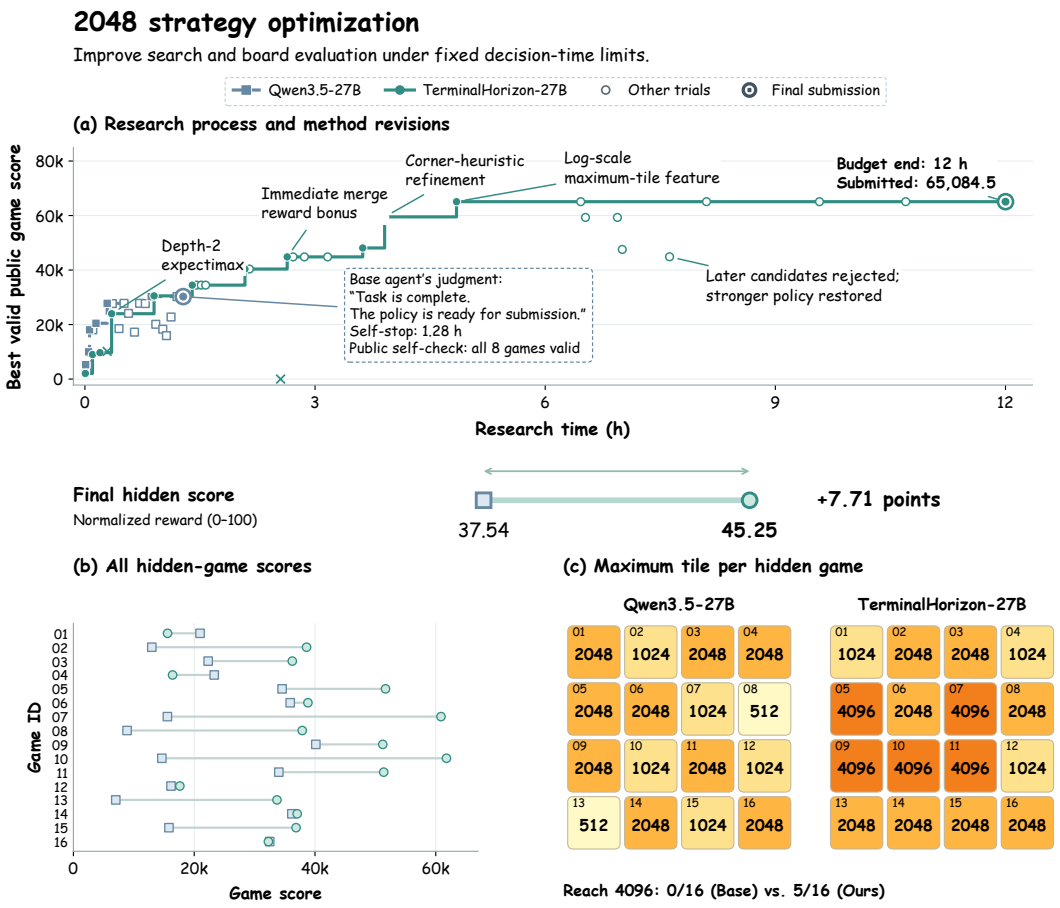

Figure 12: 2048 strategy optimization. (a) Best valid mean score on eight public seeds. Crosses mark invalid trials and rings mark submissions. (b–c) Scores and maximum tiles for all 16 hidden games in seed order. Each grid cell represents one game.

Search, evaluation, and selective revision. Both models improve on the weak initial implementation by introducing expectimax search and revising the board evaluation function. The base model explores lookahead search that accounts for random tile generation, adaptive search depth, and corner preferences, then restores its best public-set policy and voluntarily stops after 1.28 hours. It cites the improvement over the initial policy and the valid completion of all eight public games as grounds for submission, concluding that the task is complete and the policy is ready to submit. TerminalHorizon-27B continues to introduce an immediate merge reward and selective memoization, while adjusting the relative influence of empty-cell count, tile values, smoothness, and monotonicity. The development record also shows recovery from failed implementations: a version with unbounded cache growth fails to complete the full public evaluation, prompting a more selective caching design. Later, changing the maximum-tile feature to a logarithmic form raises the mean score on the full public set from 59,509.0 to 65,084.5.

Improvement includes rejecting attractive candidates. Several later configurations perform well on smaller development subsets but regress on the full public test set. Rather than replacing its best-so-far solution with the latest candidate, TerminalHorizon-27B reverts these changes and submits the stronger evaluated policy at the deadline. The final public mean scores are 30,257.0 for the base model and 65,084.5 for TerminalHorizon-27B. On hidden games, mean score increases from 23,183.75 to 38,626.25, while the normalized score rises from 37.54 to 45.25. Both final policies achieve a valid-run fraction of 1.0 on the hidden set, so the difference mainly concerns policy quality rather than avoidance of execution failures. This trajectory illustrates how sustained experimentation

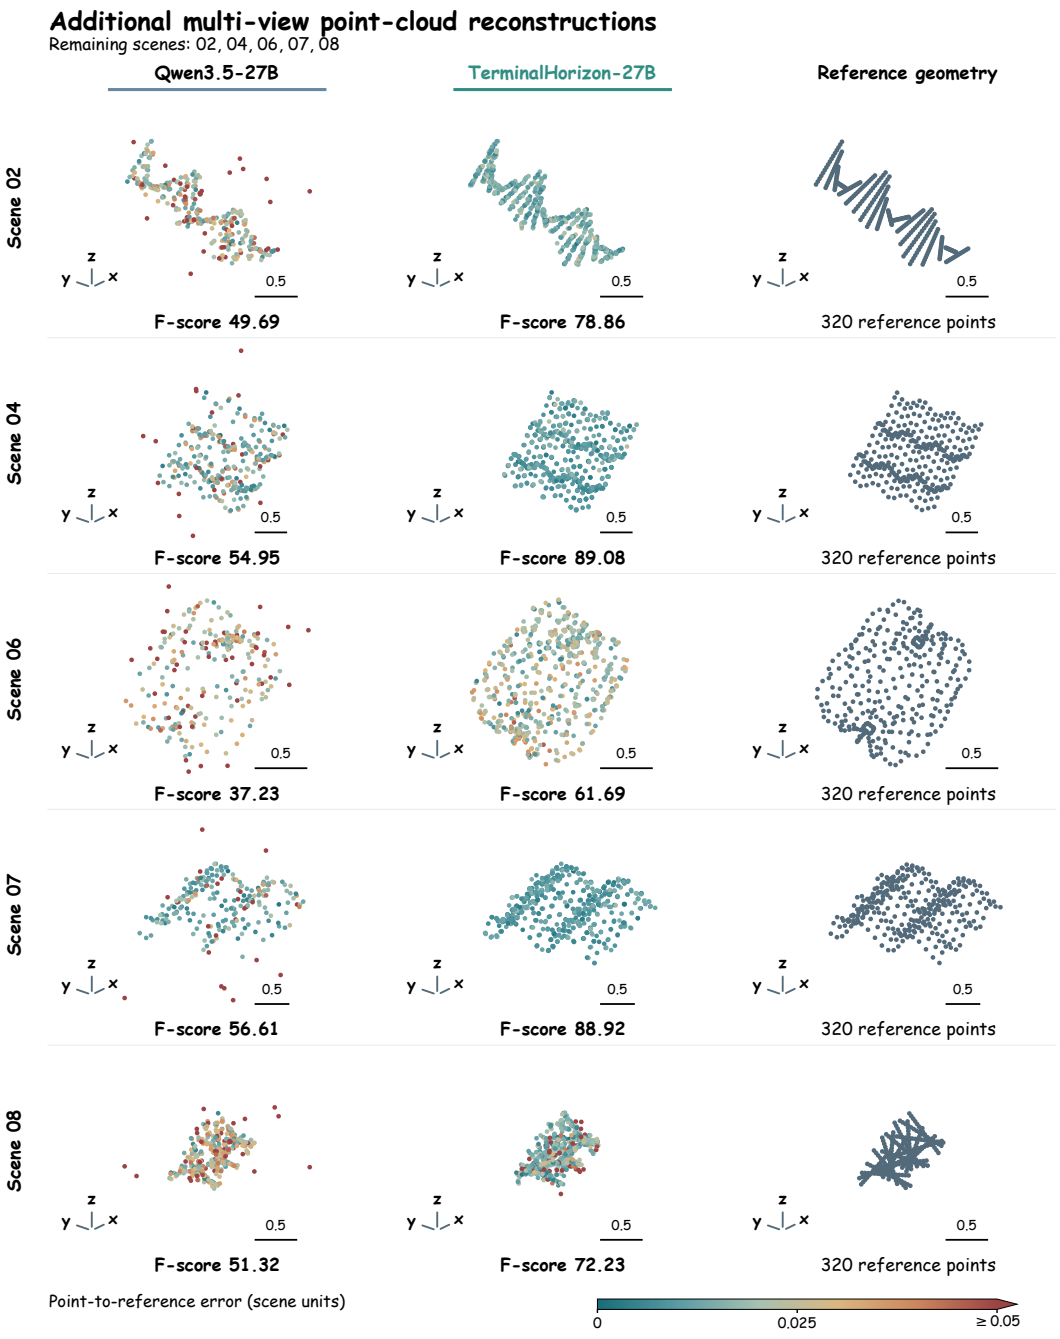

Figure 13: Additional point-cloud reconstructions for five scenes beyond Figure 7. Each row uses the same camera, scale, and point size, without post-hoc alignment, smoothing, or point removal. Colors indicate point-to-reference distance, saturated at 0.05 scene units. Scale bars represent 0.5 scene units, and F-scores are multiplied by 100.

combines algorithmic improvement with careful selection: failed attempts provide evidence for retaining, revising, or discarding a candidate.

C.4 Multi-View Point-Cloud Fusion

Task and source of difficulty. The agent must reconstruct a complete scene by fusing noisy point clouds from partially overlapping views. The first view defines the reference coordinate frame, while the remaining views contain unknown small rigid transformations; each observed point also has a confidence value. The agent must align the views, suppress unreliable points, and preserve surface coverage within a fixed output-size limit. The public self-check only verifies compliance with the output interface and provides neither ground-truth geometry nor reconstruction quality scores. Final evaluation uses nine unseen scenes and averages symmetric precision–recall F-scores over three distance tolerances, penalizing both inaccurate points and missing surface regions. Figure 13 shows the five remaining scenes (02, 04, 06, 07, and 08), complementing scenes 01, 03, 05, and 09 in Figure 7. These visualizations compare final submissions rather than intermediate research versions.

Questioning what a proxy actually measures. After trying iterative closest point alignment, the base model settles on a conservative approach: it translates the other views toward the reference view's centroid and retains only points close to that view. It obtains a high F-score on a self-designed proxy metric and voluntarily stops after 12.1 minutes. Specifically, it describes the results as excellent based on a mean proxy F-score of 0.8994 and a proxy F-score of 1.0 at the 0.05 distance threshold for every case, concluding that the solver is ready to submit. However, agreement with the reference observation is not equivalent to agreement with the complete scene. Selecting points close to a single view can raise this proxy score while omitting surfaces visible only in other views. TerminalHorizon-27B encounters a similar issue in early experiments. A variant that selects points by nearest distance to the reference view achieves a proxy score of about 0.989, but the agent inspects which views contribute the selected points and questions whether the metric rewards reconstruction or merely copying the reference view. It then introduces a leave-one-view-out diagnostic, identifies poor coverage of the excluded view, and revises its fusion strategy accordingly.

Aligning the method with geometric coverage. TerminalHorizon-27B's final solver filters nearest-point correspondences by a distance threshold, estimates rotation and translation through rigid alignment, and jointly refines each non-reference view against the other aligned views. It performs two rounds of global refinement and selects high-confidence points from the merged cloud, rather than restricting added points to the area near the first view. In the official hidden evaluation, the mean F-score across distance tolerances, expressed on a 0–100 scale, increases from 48.04 to 77.68; the separately calibrated normalized score rises from 21.93 to 48.96. The key behavioral difference is the agent's willingness to revise not only the solver but also the experiments used to assess progress. Interface validity and high proxy scores provide only intermediate evidence, while cross-view diagnostics help align the method with the actual reconstruction objective.

Together, these four cases illustrate complementary aspects of sustained autonomous research: jointly optimizing coupled objectives, revising scientific models, testing and reverting search heuristics, and identifying misleading proxy metrics. They share a close connection between observations, subsequent decisions, and the implementation ultimately submitted for evaluation. Longer interaction is valuable when it supports these connections, rather than serving as an objective in itself.

D Professional Domain Coverage

Our 1,535 tasks cover 25 primary professional domains. As illustrated in Figure 14, this coverage extends beyond software and infrastructure to education, agriculture, healthcare, finance, and scientific research.

The terminal provides a shared interaction interface rather than defining the scope of the work. Tasks such as educational assessment, payroll reconciliation, and radio-interferometric calibration bring different materials, constraints, and verifiable outputs into executable environments. Domain coverage describes what agents work on, while the behavioral patterns in Appendix E describe how they organize and advance that work.

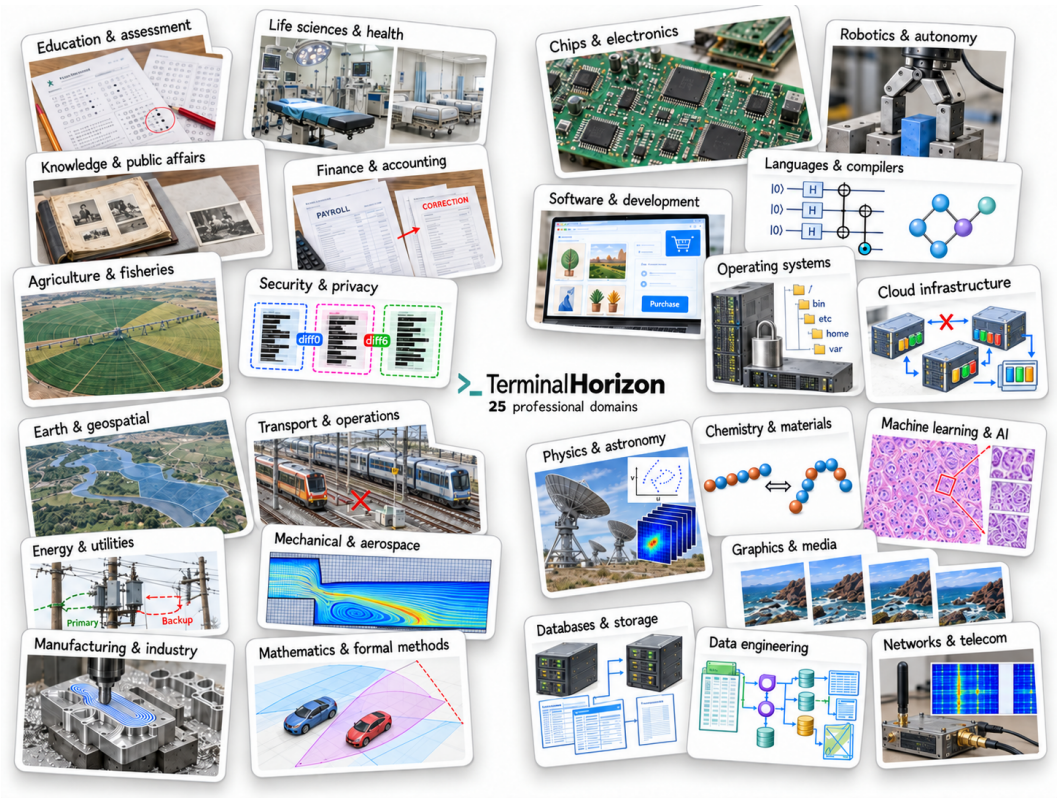

Figure 14: Overview of TerminalHorizon's 25 professional domains.

E Agentic Behavioral Patterns

E.1 Definitions and Taxonomy

We define an *agentic behavioral pattern* as the structure connecting actions, observations, and subsequent decisions around a specific goal during problem solving. Unlike classifications by task domain or tool, behavioral patterns describe how agents gather information, address problems, construct solutions, and verify their correctness, rather than characterizing individual commands or isolated operations.

We group 27 behavioral patterns into five categories: information exploration and judgment, diagnosis and error correction, construction and work progress, verification and assurance, and strategy and execution coordination. Tables 6 and 7 define each pattern. These categories describe different aspects of problem-solving behavior rather than mutually exclusive task types or fixed execution stages: a trajectory may contain multiple patterns, and tasks from different domains may exhibit the same patterns.

E.2 Organization of Behavioral Patterns within Trajectories

To examine how agentic behavioral patterns connect within a trajectory, we select three annotated cases with traceable links between observations, decisions, and subsequent work. Figure 15 presents a TLS gateway repair, a synthetic neutron-diffraction study, and an event-driven staffing task. The figure links selected episodes to observed patterns from the five categories above. These are qualitative examples of behavioral organization, rather than estimates of pattern frequency or duration.

TLS gateway repair: diagnosis returns as implementation progresses. The agent first distinguishes a working direct handshake from a failing retry path and corrects the route handling. Backend logs then expose a missing client identity, which requires a separate repair. Even after that repair, the

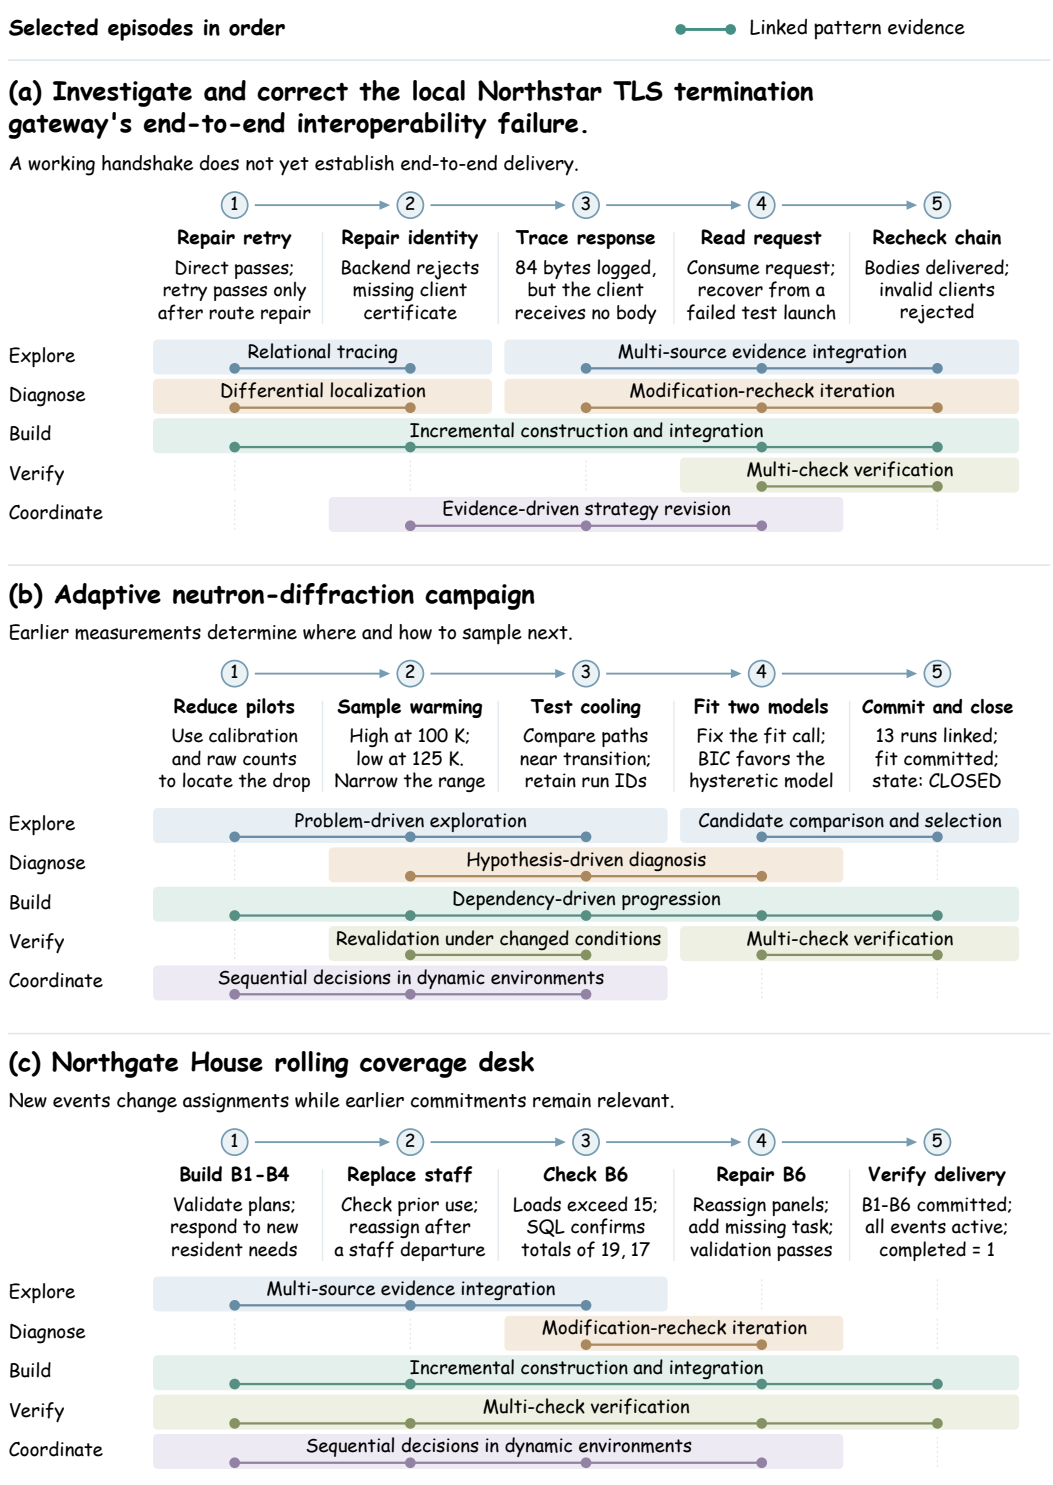

Figure 15: Behavioral-pattern organization in three trajectories. Columns show selected episodes in order. Dots mark supporting episodes, linked by pattern within each family. Horizontal spacing indicates order, not duration.

Table 6: Agentic behavioral patterns (1/2): exploration and diagnosis.

Pattern	Definition
Information exploration and judgment	
Problem-driven exploration	Using findings from earlier exploration to redirect subsequent searches toward different objects or scopes.
Relational tracing	Following call, data-flow, identity, or provenance links to understand how objects relate or where to intervene.
Multi-source evidence integration	Establishing correspondences across sources and combining their information to support a judgment.
Candidate comparison and selection	Comparing alternatives against explicit criteria and using the comparison to select a candidate or defer selection.
Tool and interface probing	Experimenting with unfamiliar tool or interface behavior and applying the resulting understanding in subsequent actions.
Diagnosis and correction	
Error-directed correction	Correcting an issue directly identified by an error or validation failure.
Hypothesis-driven diagnosis	Testing a specific explanation or proposition and using the outcome to strengthen, weaken, or revise it.
Reproduction-based diagnosis	Establishing and testing inputs or steps that reproduce an anomaly, then using them to investigate it further.
Differential and partition-based localization	Comparing or partitioning versions, inputs, or components to narrow the suspected source of a problem.
Runtime-guided localization	Using traces, profiles, breakpoints, or sampled internal behavior to narrow the source of a problem.
Modification–recheck iteration	Modifying a visible problem and rechecking that same problem on the modified artifact, regardless of the outcome.

downstream client receives no response body despite 84 bytes being recorded in the gateway log, a discrepancy that redirects the investigation toward request consumption and connection closure. After modifying the request-reading path and recovering from a failed test launch, the agent reruns the end-to-end checks, including invalid-client rejection. Relational tracing, diagnosis, and modification–recheck iteration thus support incremental construction rather than occurring as isolated stages. The final evidence is a passing local regression suite, not an exhaustive check of every TLS requirement.

Neutron diffraction: observations guide sampling and model selection. The agent combines instrument calibration with raw pilot measurements to locate a drop in the signal. New warming measurements narrow the temperature range, while cooling measurements provide a comparison near the transition. These observations support a test of thermal-path dependence and a BIC comparison between continuous and hysteretic models. The selected fit is committed with all 13 run identifiers before the campaign is closed. Problem-driven exploration, hypothesis-driven diagnosis, and dependency-driven progression are connected through the measured data. The native interface records the fit and closure, but the trajectory does not expose the independent verifier's refit results.

Care staffing: changing conditions connect planning, correction, and delivery. In the synthetic staffing environment, six consecutive windows require updates for resident transfers, changing care needs, admissions, and staff availability. Each committed window contributes to the cumulative schedule and constrains later assignments. For the final window, validation reports two panel loads of 19 and 17 against a capacity of 15, along with a missing licensed task. A separate SQL calculation confirms the loads before the agent redistributes residents, adds the task, and revalidates the plan. The final status records all six windows as committed. This case connects dynamic decision-making and incremental construction with a local correction loop and final delivery checks. The SQL calculation and validator use different checking paths but share the same database inputs.

Table 7: Agentic behavioral patterns (2/2): construction, verification, and coordination.

Pattern	Definition
Construction and work progression	
Subproblem decomposition and composition	Addressing distinct subproblems separately and combining their results to advance the main objective.
Dependency-driven progression	Establishing prerequisites that subsequent work actually uses.
Incremental construction and integration	Extending an existing partial capability or result and integrating additions into a more complete whole.
Measurement–comparison–tuning	Comparing measurements before and after an intervention and using them to retain, adjust, or replace it toward an improvement goal.
Item-wise processing and exception handling	Applying a common operation to a collection with identifiable item-level results, optionally routing exceptional items to separate handling.
Verification and result establishment	
Multi-check verification	Combining distinct, meaningful checks to support or qualify an outcome.
Reference-based verification	Comparing an actual result against an explicit expected relation to determine agreement or discrepancy.
Independent cross-verification	Checking the same claim through substantively different derivations, measurements, or implementations with distinct sources of error.
Stability rechecking	Repeating and comparing runs under a fixed setup to assess variability, intermittency, or repeatability.
Revalidation under changed conditions	Rechecking an existing claim after changing key inputs, environments, scales, operating conditions, or restart states.
Requirement and delivery reconciliation	Revisiting task requirements, matching them to actual results or gaps, and using that assessment to deliver or continue work.
Strategy and execution coordination	
Evidence-driven strategy revision	Changing an established solution method or consequential arrangement in response to new information.
Exploratory branching and backtracking	Trying a branch, then using feedback to undo its changes or return to an earlier decision state.
Coupled-object coordination	Coordinating operations across objects linked by interfaces, identities, timing, resources, or state, and checking their joint effects.
Asynchronous job management	Using identifiable background jobs' states to guide subsequent management, collect results, or establish unfinished status.
Sequential decision-making in dynamic environments	Selecting subsequent actions in response to evolving task-world conditions, including consequences of earlier actions.

Across these cases, a single episode can support several patterns, and one pattern can connect work across multiple episodes. The resulting experience lies in how evidence is used, partial results are extended, and checks determine the next action, rather than in a fixed sequence of behavioral stages.

F Trajectory Filtering Criteria

Final task success alone does not determine a trajectory's value for supervised learning. Successful trajectories may rely on answer leakage or circumvent task constraints, whereas unsuccessful ones may still demonstrate useful diagnosis, implementation, and error recovery. We therefore adopt process-oriented, trajectory-level filtering based on trajectory validity, environment-grounded progress, and multi-turn training value to assess whether a teacher trajectory provides a suitable interaction demonstration for the student model. We use `gpt-5.6-luna` with reasoning effort set to `high` as the reviewer. It reads the complete conversation without receiving the final reward

Table 8: Benchmark text used for decontamination. Paths are relative to each benchmark distribution.

Benchmark	Source file(s)	Compared content
Terminal-Bench 1.0 / 2.0 / 2.1	instruction.md	Complete task instruction
AutomationBench	automationbench/domains/ <domain>/tasks.py	User messages in the generated task prompt
VitaBench (EN cross-domain)	cross_domain/tasks_en.json	instructions, user_scenario, and nonempty message_history
τ^2-bench Airline	airline/tasks.json	Complete user_scenario
τ^2-bench Retail	retail/tasks.json	Complete user_scenario
MATH-500	test.jsonl	problem
GSM8K	test-00000-of-00001.parquet	question
AIME 2024	test-00000-of-00001.parquet	problem
AIME 2025	test.jsonl	problem
KOR-Bench (Cipher)	rule.jsonl; sample.jsonl	Rule title, rule content, and question

or a success/failure label as additional input, and must ground its judgments in the actions and observations recorded in the trajectory.

Trajectory validity. A valid trajectory should maintain coherent context and traceable action–observation correspondences, allowing decisions to be connected to their execution outcomes. We exclude trajectories whose core problem-solving process cannot be reliably interpreted because of severe truncation, persistent protocol corruption, or an incorrect execution environment. We also exclude trajectories that rely on leakage of private answers, unauthorized manipulation of internal state, or circumvention of prescribed interfaces. Local interaction errors do not, by themselves, invalidate a trajectory: if the model recognizes a parsing or command error, restores valid interaction, and continues making progress, the sequence can still provide a useful demonstration of error correction. This criterion distinguishes recoverable mistakes from substantive defects that compromise trajectory integrity or violate task constraints.

Environment-grounded progress. High-quality trajectories should demonstrate that actions are informed by evidence from the environment. We examine whether the model uses file contents, logs, command outputs, error messages, and test results to guide its decisions and, when necessary, revise its hypotheses or plans. Meaningful progress need not immediately produce a correct result: ruling out an infeasible approach, identifying a root cause, or uncovering a new constraint can also provide valuable learning signals. Conversely, we exclude trajectories dominated by persistent disregard for feedback, ungrounded trial and error, repeated ineffective actions, or unsupported adherence to incorrect conclusions. The assessment concerns the reliability of the overall action sequence, rather than requiring every attempt to be correct.

Multi-turn training value. We further assess whether a trajectory demonstrates substantively connected, multi-turn terminal work, including evidence gathering, implementation, verification, hypothesis revision, and error recovery. Its value lies in how interactions inform subsequent decisions, not in trajectory length, command count, or output volume. Waiting, repeated polling, and mechanical loops receive no additional credit merely for increasing the number of turns. Infrastructure operations used solely to connect to or manage the execution environment likewise do not count toward problem-solving depth. Rather than imposing a fixed turn-count threshold, we assess whether the trajectory provides a reliable, nonredundant interaction process worth learning from.

Overall decision. We retain trajectories that satisfy these criteria even if they do not complete the task or end at an unresolved obstacle. In particular, feedback-driven diagnosis, revision, and recovery following a failure should not be discarded simply because they contain mistakes. Local errors are

likewise insufficient grounds for rejection when they do not compromise the main action sequence. We instead reject a trajectory when its core solution relies on prohibited shortcuts, incorrect assumptions persist throughout the main problem-solving process, or interactions remain unproductive and fail to provide reliable learning signals. These process-oriented criteria preserve exploration and error correction in realistic problem solving while excluding systematic errors and ineffective interactions from supervised demonstrations.

G Detailed Benchmark Decontamination

Direct reuse or close paraphrasing of evaluation problems in training tasks can make it difficult to distinguish capability gains from memorization. To assess this risk, we compare all 1,535 training tasks with the evaluation benchmarks at the task level, combining exact-span matching, near-duplicate retrieval, and semantic assessment to identify potential task overlap.

G.1 Comparison Corpora

The comparison corpora cover 3,251 evaluation instances across 12 benchmark groups: Terminal-Bench 1.0, 2.0, and 2.1; AutomationBench; the English Cross-domain tasks in VitaBench; τ^2-bench Airline and Retail; MATH-500; GSM8K; AIME 2024; AIME 2025; and KOR-Bench Cipher. Each training task is represented by its full `instruction.md`, while each evaluation instance is represented by its problem statement or user request. Table 8 lists the sources and fields used.

The comparison texts retain the task objectives and necessary context. For KOR-Bench Cipher, we combine the cipher rules with the question. For VitaBench and τ^2-bench, we use the full user-request scenario, preserving its goals, known information, and constraints, including requirements that may be disclosed gradually during interaction. No text is summarized or truncated, and reference answers, private scoring criteria, and expected execution outcomes are excluded from the comparison.

G.2 Text Overlap and Near-Duplicate Retrieval

We apply Unicode NFKC normalization and lowercasing to the task descriptions, then tokenize them into contiguous sequences of Unicode letters, digits, and underscores. Exact matching covers both the full normalized text and contiguous 13-token and 14-token windows with a stride of 1. Let $G_n(x)$ denote the set of contiguous n-grams in task description x. Exact-span overlap between a training description x and an evaluation description y is defined as

$$E_n(x, y) = \mathbb{I}[G_n(x) \cap G_n(y) \neq \varnothing], \qquad n \in \{13, 14\}. \tag{9}$$

To retrieve near-duplicate tasks that do not share long verbatim spans, we also compute the Jaccard similarity between their 5-gram sets,

$$J_5(x, y) = \frac{|G_5(x) \cap G_5(y)|}{|G_5(x) \cup G_5(y)|}, \tag{10}$$

and the cosine similarity $S_{\mathrm{TF}}(x, y)$ between their unigram/bigram TF–IDF representations. The former captures local phrase reuse, while the latter captures lexical similarity under changes in word order. TF–IDF is fitted on the combined corpus of training tasks and the benchmark being compared, using logarithmic term frequency and smoothed inverse document frequency. For mathematical and symbolic expressions, we additionally retrieve contiguous 8-gram overlaps using a tokenizer that preserves LaTeX commands, numbers, and operators. Because these short spans may represent common formulas or interfaces, they supplement candidate retrieval rather than directly establish contamination. We also report base-2 unigram Jensen–Shannon divergence (JSD) to characterize differences in word distributions between the training corpus and each benchmark; this metric is not used to determine contamination.

G.3 Semantic Assessment of Candidate Pairs

Lexical similarity does not necessarily indicate task reuse, and a paraphrase of the same problem may not share long text spans. We therefore combine multiple retrieval methods to select candidate pairs

Table 9: Decontamination against 1,535 training tasks. N: benchmark instances, E_n/M_8: exact/math-aware overlap pairs, Contam./Reviewed: contaminated/reviewed pairs.

Benchmark	N	**Exact** E_{13}/E_{14}	M_8	**Max** J_5	**Max** S_{TF}	**JSD**	**Contam./ Reviewed**
Terminal-Bench 1.0	80	0/0	8	0.00175	0.07467	0.4334	0/19
Terminal-Bench 2.0	89	0/0	83	0.00289	0.08344	0.3858	0/19
Terminal-Bench 2.1	89	0/0	83	0.00289	0.08345	0.3848	0/19
AutomationBench	600	0/0	0	0.00196	0.07050	0.4447	0/14
VitaBench (EN cross-domain)	100	0/0	0	0.00096	0.02761	0.5907	0/12
τ^2-bench Airline	50	0/0	0	0.00000	0.03764	0.6179	0/10
τ^2-bench Retail	114	0/0	0	0.00177	0.03673	0.6068	0/13
MATH-500	500	0/0	15	0.00602	0.07990	0.5914	0/18
GSM8K	1,319	0/0	0	0.00252	0.07342	0.5486	0/15
AIME 2024	30	0/0	0	0.00179	0.06560	0.6280	0/13
AIME 2025	30	0/0	6	0.00381	0.06169	0.6547	0/18
KOR-Bench (Cipher)	250	0/0	210	0.00360	0.05116	0.5878	0/20
Total	3,251	0/0	405	–	–	–	0/190

and assess whether they are substantive near-duplicates based on the complete task descriptions. For each benchmark group, the candidate set includes all full-text and exact 13/14-gram matches, the top 10 pairs by TF–IDF similarity, the top 5 by 5-gram Jaccard similarity, and the top 5 by the number of shared math-aware 8-grams. Each retrieval method contributes only pairs with a nonzero score, and the combined set is deduplicated before semantic review.

Semantic assessment is performed by `gpt-5.6-luna` with reasoning effort set to `high`. The reviewer reads both task descriptions in full and compares their core objectives, initial conditions, key objects or data, distinctive combinations of constraints, required transformations, and completion criteria. A pair is classified as contaminated when one task is a direct copy or minor paraphrase of the other, or when renamed entities leave the distinctive instance structure and solution requirements unchanged. Shared tools, domain knowledge, terminology, or common work patterns alone do not constitute contamination. Each pair is labeled as non-contaminated, contaminated, or uncertain, with supporting comparisons drawn from the task descriptions.

G.4 Decontamination Results

Table 9 summarizes text similarity and semantic assessment results for each benchmark. None of the 12 benchmark groups contains an exact match to normalized full text or a contiguous 13/14-token window. The maximum 5-gram Jaccard similarity is 0.006024, and the maximum TF–IDF cosine similarity is 0.083445. Math-aware 8-gram retrieval identifies 405 pairs with short-span overlap, which are used for candidate selection alongside lexical nearest neighbors. After combining and deduplicating the retrieval results, 190 candidate pairs undergo semantic review; none is classified as contaminated or uncertain. These comparisons of task descriptions and semantic assessments find no evidence of evaluation-task reuse in the training tasks, so all 1,535 training tasks are retained.

H Case Studies of Challenge-Guided Reconstruction

We examine three cases of numerical calibration, distributed-service recovery, and household benefits administration in which independent challenges exposed work lost during environment construction. Each traces the initial execution, the resulting reconstruction, and the work observed in a subsequent independent challenge.

H.1 RADIO-INTERFEROMETRIC CALIBRATION

The agent must calibrate a synthetic Northstar-10 radio observation stored as a native MeasurementSet while preserving the raw measurements and metadata. Calibration is evaluated on holdout and target data excluded from fitting. The environment uses casacore, NumPy, and SciPy for data access and numerical computation.

Case 1 | Restoring numerical calibration work

01 Initial construction

The public `mscal` tool included `solve-bandpass` and `solve-gain`, which performed numerical fitting and wrote native calibration tables, so core implementation work was already packaged into the environment.

02 Independent challenge: successful, but compressed

After identifying the fields, the challenger invoked the bundled solvers, applied the calibration, and checked residuals. It adjusted the fitted field selection after a failed check, but did not need to implement the numerical calibration.

03 Reconstruction: leave the fit to the agent

Reconstruction removed the fitting routines but retained inspection, calibration application, and residual feedback. It introduced distinct time-varying responses for the two spectral windows and made the required solution axes and table conventions public.

04 Independent re-challenge: feedback changes the model

The fresh challenger inspected the data axes and public interfaces, implemented a fitting workflow, and revised its spectral-window-dependent parameters after observing high residuals. It wrote and applied native calibration tables, checked that the raw data remained unchanged, and evaluated holdout and target data without fitting them.

Final weighted RMS was 0.896 on the fitted calibrators, 1.784 on holdout, and 1.944 on target, all below the public threshold of 2.8. The independent run passed the verifier. These scores describe the reconstructed task, not a comparison between task versions.

What reconstruction restores. The initial challenger also passed, showing that task success alone does not reveal work completed in advance. Reconstruction moved numerical modeling and implementation from the supplied tool to the agent while retaining the interfaces and feedback needed to complete the task. The resulting trajectory connects model construction, residual-driven revision, and independent validation.

H.2 RABBITMQ RECOVERY ACROSS BROKER AND APPLICATION STATE

The agent must recover a three-node RabbitMQ service from a network partition while preserving confirmed messages, resolving failed or uncertain publications, and preventing duplicate business effects. The synthetic incident runs on real RabbitMQ brokers and AMQP clients, with firewall rules and persistent application state. Five reconstruction rounds and six independent challenges turned application recovery from a possible branch into work required by the running environment.

Case 2 | From broker recovery to application recovery

01 Initial challenge: recovery is largely prebuilt

The challenger used the supplied `partitionctl heal`, restarted one broker, and let the clients finish. Automatic retry and idempotent consumption handled much of the application recovery, leaving a short operational path.

02 Rounds 1–2: reconstruct the live recovery state

The redesigns removed the healing command and introduced three queues with different missing voting replicas. Missing members and network faults were selected from live broker state rather than a fixed repair map. Unresolved publications were kept in a durable outbox, and recovery had to preserve unrelated firewall policy.

03 Round 3: restore dependencies within application work

The workload changed from independent messages to 27 workflows of eight ordered operations. Publishers retained event identities for reconciliation and resumption; consumers tracked predecessor completion and prevented duplicate business effects. The intended recovery now depended on both broker health and the completion of earlier workflow steps.

04 Repeated challenges: the new work still does not occur

Intermediate challenges still finished without the intended reconciliation work. One run encountered no uncertain publications; in a later run, the client could reach its local queue leader despite the partition. The recovery branch existed but was not needed in either execution.

05 Rounds 4–5: connect the fault to its application consequence

The redesigns moved fault installation before client startup and connected the orders client to a runtime-selected follower whose leader was across the partition. The actual AMQP path now exposed the application failure that earlier versions could avoid.

06 Final independent challenge: recover the missing predecessor

After restoring connectivity and queue voters, the challenger found 71 confirmed orders rows and one `nacked` first-step event, with downstream consumer errors. It retrieved the stored message body and original event ID, then republished with AMQP publisher confirms. All 216 outbox rows became confirmed, every workflow reached step 8, and no duplicate business effects were observed. The independent run passed.

What reconstruction restores. Reconstruction linked network and replica repair to unfinished application work. Repeated challenges showed that a recovery branch adds no workload unless the running environment actually triggers it. The final run had to recover a specific missing event before dependent workflow steps could finish.

H.3 Family Enrollment Changes and Payroll Reconciliation

This case concerns a household coverage change following the loss of other insurance. The agent must assess dependent eligibility and keep benefits enrollment, payroll deductions, and insurer records consistent. The household, rules, and service are synthetic, modeling an employer's benefits workflow through stateful public interfaces. Four reconstruction rounds and five independent challenges developed a simple record-update task into evidence-based review and coordinated updates across systems.

Case 3 | From enrollment edits to evidence-based reconciliation

01 Initial challenge: a fixed update sequence

After inspecting the case, the challenger amended three plans, denied one dependent, and synchronized payroll and insurer records. Most extra interactions came from discovering command syntax, not from analyzing evidence or revising decisions.

02 Rounds 1–2: connect coverage dates to downstream records

Previously posted deductions now required retroactive payroll adjustments. A dependent's age limit split coverage into consecutive periods, changing pay-period calculations and insurer payloads. Legacy insurer plan codes also had to be updated before the new records could be accepted.

03 Round 3: revise an already accepted case

The starting state now contained committed enrollment, applied payroll, and an accepted insurer transaction, followed by a corrected dependent document. The agent had to reopen the case, create a replacement version, and update both downstream systems without erasing the earlier records.

04 Re-challenge: the evidence decision is still supplied

The corrected document was already marked verified, the affected person was obvious, and each downstream issue had one clear remedy. The challenger could still follow a fixed sequence once it had read the case. Versioned updates alone had not restored the intended evidence analysis.

05 Round 4: make review outcomes change the decision

Reconstruction introduced conflicting document claims and their source records. One dependent had both a mismatching document and a corroborating alternative; another had an authentic document whose loss date failed the event-date rule. Public review and validation returned feedback that determined which people could be included in the revised coverage.

06 Final independent challenge: revise and propagate the decision

The challenger rejected mismatching evidence, found the corroborating source, and revised its provisional enrollment after validation exposed the ineligible loss date. It approved D-4, denied D-3 and D-5, and preserved the age-bounded coverage for D-2. It then committed version 6, applied a $312 payroll adjustment, and obtained insurer acceptance while retaining version 5. The independent run passed.

What reconstruction restores. The key change is that document review now determines the next action rather than confirming a decision already supplied by the environment. An authentic document can still fail the eligibility rule, and an unsupported claim requires another source. Once the household decision changes, its effective dates and source version must remain consistent across enrollment, payroll, and insurer acceptance. The final trajectory therefore connects evidence comparison, feedback-driven revision, and verified updates across systems.

I Qualitative Analysis of Progressive Task Evolution

We examine three evolution lineages from two original tasks: two branches of photo-library reconciliation and one hotel event-planning lineage. For each generation, we examine what changes from its immediate parent, why the previous approach is no longer sufficient, and which additional operations and checks are needed. The photo branches show how one task can develop in different directions, while the hotel lineage shows how new constraints change earlier planning decisions.

I.1 Two Branches from the Same Photo-Library Task

The shared parent asks the agent to reconstruct a curated PhotoPrism library from media exports, catalogs, and sidecars. The work includes preserving distinct media files, organizing related representations, restoring curation and albums, and checking the native library after a complete rescan. Table 10 compares two three-generation branches that start from identical parent packages.

Branch A: reconciling changes in library state. This branch extends reconciliation from metadata edits to retirement and restoration, then to file grouping and primary-rendition selection. Each generation adds decisions that must precede ordinary updates. For example, later destination edits may prevent the agent from applying a source retirement directly. Likewise, splitting a file group in PhotoPrism can change file ownership and paths, so files must be re-identified before their curation and sidecars are assigned to the resulting groups.

Branch B: recovering identity across systems. This branch first requires the agent to infer file relationships from media evidence, then combines identity recovery with incremental synchronization into an independently edited destination. The final generation extends identity reasoning to both systems when transformed destination files no longer match the checkpoint bytes. The agent must

Table 10: Two photo-library evolution branches from the same parent task. Numbers indicate successive generations within each branch.

Version	Change from the parent	Additional work
Branch A: reconciling changes in library state		
A1	A one-way import becomes a second synchronization with independent edits on both sides. A shared photo's rating changes from 5 at the checkpoint to 4 at the source and 3 at the destination.	Compare checkpoint, source, and destination values. Preserve legitimate edits from either side; retain rating 3 and flag this concurrent conflict under the public policy.
A2	Reconciliation now includes retirement and restoration, rather than only edits to active photos. An explicit source retirement may conflict with later destination edits.	Check archive status alongside later edits. Keep photos active when retirement conflicts with new edits, apply native archive or restore operations, and check the state after rescanning.
A3	File grouping is no longer fixed. The child introduces a source-side split, an existing destination-side split to preserve, and a source-side change of the primary rendition.	Compare the three grouping states and primary choices. Split groups or switch primary files in PhotoPrism, re-identify files by digest, and update each group's curation and sidecars.
Branch B: recovering photo identity across systems		
B1	Classified file relations are replaced by evidence in media metadata. True companions may have different names, while similar names or timestamps may belong to distinct captures.	Extract capture identifiers, Live Photo identifiers, and document lineage. Resolve logical groups and reject misleading matches before choosing a target file layout.
B2	The destination now contains independently curated photos and target-only material. Resolving source identities no longer makes a source-authoritative overwrite safe.	Match photos across libraries, then compare checkpoint, source, and destination metadata. Apply source changes while preserving destination edits, independent collection changes, and target-only photos.
B3	Three destination representations no longer retain the exact baseline bytes. A transformed coastal PNG shares the source JPEG/RAW pair's identifier, while a similar target-only image has a different identifier.	Inspect identity evidence on both sides before merging curation. Link transformed representations, preserve the unrelated lookalike separately, and use native indexing results to check the grouping.

establish which files represent the same photo before deciding which edits to merge. Native indexing and rescanning then test whether these relationships and edits remain stable in the reconstructed library.

I.2 Hotel Event Scheduling

The original task asks the agent to coordinate lodging, activities, venues, catering, resources, and billing in a terminal-accessible hotel management system. All four versions retain the same ten activities, while three successive extensions introduce attendee flows, incident recovery, and resource logistics (Table 11).

A local disruption requires coordinated recovery. In the second generation, Harbor Hall's closure directly affects lunch and Workshop A. Moving lunch to the Grand Ballroom, however, conflicts with the plenary session's occupancy, including setup and teardown buffers. In the independent run, the agent moves the plenary to Seaview, lunch to the Ballroom, and Workshop A to Maple. Thus, two

Table 11: Three generations of hotel event-planning evolution, retaining the same ten activities.

Version	Change from the parent	Additional work
Hotel event scheduling		
1	Attendee groups and directed travel times are added to the existing venue, timing, and resource constraints. The end time and location of one activity affect arrival at the next.	Check travel time between each group's consecutive activities and adjust venues or times when needed. Verify that each group can complete its full schedule, not just attend individual activities.
2	Initial planning becomes recovery of a published schedule after Harbor Hall closes. Lodging and billing are completed baseline state to preserve, not work to repeat.	Trace the effects of venue changes and minimize disruption. Jointly revise affected assignments while preserving existing commitments, then commit a complete recovery plan atomically.
3	Resource quantities are refined into specific installed equipment, mobile units, and crew, with compatibility, external commitments, and strike, transfer, and installation times.	Check resource compatibility and calculate when equipment is ready after prior use, teardown, transfer, and installation. Revise venues, times, and resource bundles when these checks expose a conflict.

directly affected activities lead to three venue changes, which are committed in one atomic revision without changing any activity times.

Resource feedback changes the recovery plan. In the third generation, two complete candidates are rejected because shared equipment cannot serve the proposed activity sequence. A mobile display, for example, is released at 16:30 but needs another 15 minutes for teardown, 20 for transfer, and 20 for installation. Its earliest ready time of 17:25 cannot support Workshop A's original 17:15 start. The accepted plan moves the activity to Oak, combines its installed whiteboard with the mobile display, and shifts the start to 17:30 on the public 15-minute planning grid.

Across these cases, evolution changes both what the agent must determine and how those decisions depend on earlier work. New requirements connect identity recovery to metadata merging, venue changes to attendee flows, and resource availability to schedule revision, making dependency tracing, feedback-driven revision, and integrated verification part of completing the original task.

J Prompt Templates

J.1 Trajectory Filtering Prompt

The following prompt implements the three filtering criteria in Appendix F by reviewing complete trajectories and returning an evidence-backed `KEEP`/`DROP` decision for each input record.

Trajectory Filtering Prompt

You are an expert reviewer of terminal-agent trajectories for supervised fine-tuning. Your task is to determine whether each complete multi-turn trajectory provides a reliable and useful training example.

Read the entire conversation. Do not substitute excerpts, keyword counts, command counts, or programmatic summaries for the original trajectory. Base your judgment exclusively on the trajectory itself. Do not access task packages, private verifiers, external benchmarks, or the internet.

Understanding the trajectory

- The initial user message typically contains the task description and the agent interaction protocol.
- Subsequent user messages usually contain terminal observations produced by preceding commands, rather than new user requests.

- Assistant messages contain reasoning, plans, and terminal actions. One assistant turn may contain multiple commands.
- Reasoning, including `<think>`, analysis, and plan content, is not itself evidence of execution. Evaluate it together with the actual actions and subsequent observations.
- A Terminus-2 trajectory may end with the terminal observation from its final command batch. This is a normal ending and does not require an additional assistant confirmation.
- Final success or failure is not sufficient to determine training value. An incomplete trajectory or an honest stopping point at an unresolved obstacle may still provide a useful demonstration.

Evaluate only the following three dimensions.

1. Trajectory validity

Is the trajectory coherent and correctly ordered? Can actions be reliably associated with their terminal observations? Check for genuine context truncation, severe format corruption, execution in the wrong container, infrastructure failures, explicit refusal, or private-answer leakage that compromises the demonstration.

Do not reject an entire trajectory because of one or two isolated invalid JSON responses, malformed commands, or parser errors. If the agent recognizes the problem, restores valid interaction, and continues with a clear action–observation chain, this constitutes a learnable recovery process. Treat protocol errors as an overall validity defect when they recur without reliable recovery or make action–observation correspondence uninterpretable.

Protocol recovery does not excuse independent integrity violations. Reject trajectories whose core solution relies on private answers or hidden-test leakage, internal-state tampering, deleting runtime or database state to reset the task, or bypassing required task interfaces.

Distinguish failure to finish from lack of training value. An unfinished trajectory can be retained when its main content demonstrates reliable environment understanding, root-cause diagnosis, justified implementation, feedback-based elimination of alternatives, correction of assumptions, or identification of unresolved obstacles.

An isolated incorrect final judgment or unsupported completion claim does not automatically invalidate all preceding work. Determine whether the error compromises the core action chain. Retention is possible when the main process remains reliable and useful. Reject the trajectory when an incorrect strategy dominates the solution, central conclusions contradict observations, or the overall example primarily demonstrates an invalid way to complete the task.

2. Environment-grounded progress

Are important actions supported by actual terminal feedback? Does the agent use files, logs, command outputs, errors, and test results to advance the task and adjust its approach when necessary?

Identify persistent disregard of feedback, repeated ineffective commands, random attempts, unmotivated minor edits, and unsupported completion claims. Evaluate these behaviors in context: unsuccessful attempts can be valuable when their observations inform subsequent diagnosis, revision, or recovery.

3. Multi-turn training value

Does the trajectory demonstrate a multi-turn terminal problem-solving process worth learning? Do interactions contribute meaningful evidence, implementation, state changes, decisions, verification, or recovery?

Consider whether the apparent interaction could be trivially replaced by one obvious edit, a fixed script, or a mechanical loop. Do not count SSH, Docker, or Harbor infrastructure operations, waiting, repeated polling, or irrelevant directory browsing as evidence of task difficulty.

Do not judge training value solely from trajectory length, command count, the presence of mistakes, or final task success.

Decision

KEEP: The trajectory is valid, grounded in the environment, and provides sufficient multi-turn training value. Final failure, incomplete execution, or an honest stopping point does not preclude KEEP when the main process is reliable and instructive.

DROP: The trajectory is globally invalid; its core solution relies on leakage, state tampering, or interface bypass; its behavior substantially degenerates or lacks environmental grounding; or errors dominate the main action chain with little reliable multi-turn learning value.

Do not assign DROP solely because the task was not completed.

Disputed cases

If prior review claims are provided, treat them as hypotheses to verify, not as facts or votes. Return to the original conversation to establish whether the relevant commands were executed, whether observations support the claims, whether the issue was subsequently repaired, and whether it compromises the core solution. Explain how the trajectory evidence resolves the disagreement.

Output

Return exactly one review for each input record:

```
{
  "record_id": "<input record ID>",
  "verdict": "KEEP or DROP",
  "reason": "<concise justification>",
  "evidence": [
    "<specific evidence from the trajectory>",
    "<specific evidence from the trajectory>"
  ]
}
```

Copy `record_id` exactly. Provide two to four concrete pieces of evidence. Do not omit, duplicate, or invent records.

J.2 Benchmark Decontamination Prompt

The following prompt implements the candidate-pair assessment in Appendix G by comparing complete task descriptions and returning an evidence-backed `CLEAN`, `CONTAMINATED`, or `UNCERTAIN` verdict.

Benchmark Decontamination Prompt

Role and objective

You are an expert reviewer of contamination between training data and evaluation benchmarks. Compare each candidate Terminal/CLI task with its paired public evaluation problem to determine whether the candidate leaks, copies, or substantially reproduces that problem.

Read the complete descriptions of every pair in `assigned-pairs.jsonl`. Similarity scores are retrieval signals, not conclusions.

Assessment criteria

Determine whether exposure to the training candidate would make the evaluation problem substantially easier because the two are the same or nearly the same task. Compare their core objectives, initial problems, key objects or data, required transformations, distinctive constraints, key solution paths, and final completion criteria.

- **CONTAMINATED:** A direct copy or minor rewording, or a renamed task that retains sufficiently distinctive instance details, constraint combinations, and solution structure to remain substantially the same problem.
- **CLEAN:** The tasks share only a domain, tools, general terminology, or common professional work patterns; their specific instances, states, constraints, and required work are independent.
- **UNCERTAIN:** The evidence is genuinely close, but the public descriptions alone do not permit a reliable distinction. Do not classify ordinary topical similarity as UNCERTAIN.

Phrase overlap arising only from templates, common interfaces, mathematical notation, or standard industry expressions does not constitute contamination. Conversely, the absence of long n-gram overlap does not rule out a systematic rewrite of the same instance.

Output

Write exactly one record for each `pair_id` to `reviews.jsonl`:

```
{
  "pair_id": "pair-...",
  "verdict": "CLEAN",
  "reason": "<core justification>",
  "evidence": [
```

```
    "<specific comparative evidence>",
    "<specific comparative evidence>"
  ]
}
```

Copy `pair_id` exactly. The verdict must be `CLEAN`, `CONTAMINATED`, or `UNCERTAIN`. Write the complete output to a temporary file, then rename it atomically. Check that the JSONL is parseable and that no records are missing or duplicated.

Now read `assigned-pairs.jsonl` and complete the review.