CrossEnv: Synthesizing Verifiable Cross-Environment Tasks for Long-Horizon Agent Training

Zhengfeng Yu[1,2] **Tian Lan**[2] **Zhijing Wu**[1*] **Enbo Zhao**[2]
Xuannan Liu[3] **Feihu Jiang**[2] **Qianghuai Jia**[2] **Longyue Wang**[2]
Zhao Xu[2] **Weihua Luo**[2] **Xian-Ling Mao**[1]
[1] Beijing Institute of Technology
[2] Alibaba Group
[3] Beijing University of Posts and Telecommunications

arXiv:2611.22777v1 [cs.AI] 17 Nov 2026

Abstract

Executable synthetic environments provide scalable training tasks and verifiable feedback for long-horizon agents. Existing environment synthesis methods primarily generate tasks within individual environments. However, these tasks do not adequately capture real-world interactions across services, where actions in one service depend on information obtained from another. Synthesizing such tasks is challenging because data generated across environments must remain consistent to support executable and verifiable interactions. To address this challenge, we introduce **CrossEnv**, a framework that uses an entity–tool graph to identify connections between environments and aligns task-specific data to synthesize executable, verifiable cross-environment tasks. We further propose **CrossEnv-Curriculum**, which progresses from single-environment tool use to coordination across more environments and business domains, before introducing user interaction. Across four external benchmarks, combining cross-environment training data with this curriculum achieves competitive performance with fewer training examples. Further analyses support the benefits of cross-environment data and curriculum training, and provide evidence of transfer to new environment combinations, external tasks less similar to the training data, and longer workflows.

1 Introduction

Large language models (LLMs) are increasingly used as agents that invoke tools (Schick et al., 2023) and adapt their actions based on environmental feedback (Yao et al., 2023; Liu et al., 2025; Shinn et al., 2023). Training these agents requires diverse executable environments and verifiable tasks that provide reliable feedback (Xi et al., 2025). However, restricted access and high interaction costs limit training in real-world applications (Chen et al., 2025). Programmatic environment synthesis offers a scalable alternative by generating executable tools, stateful environments, and verifiable training tasks (Song et al., 2026; Tu et al., 2026).

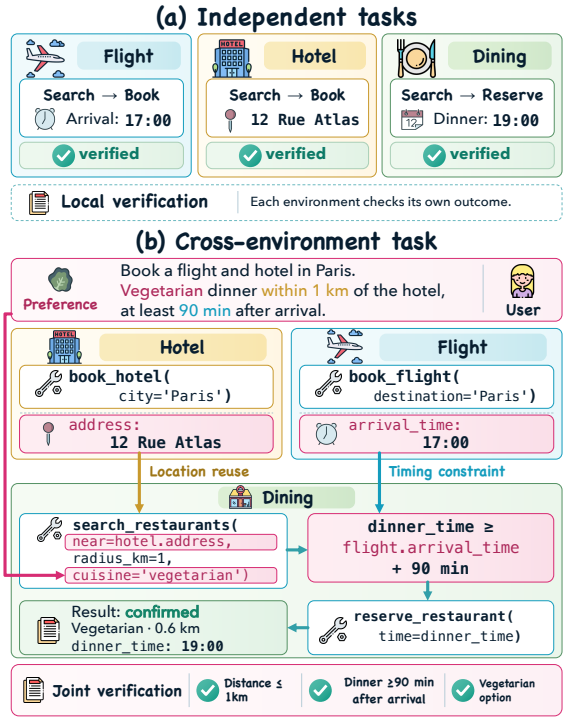

Figure 1: Independent tasks and a cross-environment task. (a) Flight, hotel, and dining tasks are completed and verified independently. (b) A travel-planning task uses flight arrival time, hotel location, and personal preferences to guide restaurant reservation, with joint verification of cross-environment requirements.

Real-world long-horizon workflows often span multiple services, where actions in one service depend on information obtained from an-

*Corresponding author.

other (Trivedi et al., 2024; Li et al., 2026; Xie et al., 2024). Figure 1(b) illustrates a travel-planning example: after booking flights and accommodation, an agent selects and reserves a restaurant based on the flight arrival time, hotel location, and personal preferences. Training agents for such workflows requires tasks that capture these information dependencies and the coordinated actions needed to achieve a shared goal.

However, existing programmatic environment synthesis methods primarily generate single-environment tasks, with limited support for combining independently constructed environments into joint tasks (Song et al., 2026; Wang et al., 2026). Cross-application benchmarks capture some cross-environment dependencies, but task and evaluator design is often manual or semi-automated (Trivedi et al., 2024; Li et al., 2026; Xu et al., 2025a; Zhou et al., 2024), limiting their scalability for training. This leaves a gap in automatically synthesizing cross-environment training tasks at scale.

A key difficulty is maintaining **data consistency across environments**. Data constructed independently for different environments may be valid in isolation but inconsistent with one another. For example, the same hotel may be recorded with different addresses in the booking and navigation environments. Synthesizing cross-environment tasks therefore requires aligning relevant data across environments to support executable, verifiable tasks.

To address this challenge, we introduce **CrossEnv**, a framework for automatically synthesizing executable, verifiable cross-environment tasks. CrossEnv uses an entity–tool graph to connect environments through shared entity types and guide the construction of task workflows. It then aligns task-specific data across environments to support execution and generates verification rules from task objectives and reference execution.

CrossEnv generates tasks with varying coordination complexity by controlling the number of participating environments and their business-domain scope. We introduce **CrossEnv-Curriculum** to establish basic tool-use skills before progressively training agents to coordinate operations across environments. Training begins with supervised fine-tuning on single-environment reference trajectories, followed by stagewise reinforcement learning that expands to multiple environments within a business domain, then across domains, and finally introduces user interaction.

Experiments on four external benchmarks show that combining cross-environment training data with this curriculum achieves competitive performance with fewer training examples than comparable environment-trained baselines. Further analyses support the benefits of cross-environment data and curriculum training, and provide evidence of transfer to new environment combinations, external environments, and longer workflows.

Method	General synthesis			Cross-environment		
	Exec.	Verif.	Ctrl.	Cross	Consist.	Ext.
Real-world environments						
ToolLLM (2023)	✓	✗	✓	✓	✗	✓
Toucan (2025c)	✓	✗	✓	✓	✗	✓
AgentTrek (2025b)	✓	✗	✗	✗	✗	✗
LLM-simulated environments						
Simia (2025)	✗	✗	△	✓	✗	✗
DreamGym (2025)	✗	✗	△	✗	✗	✗
EnvACE (2026c)	✗	✗	✗	✗	✗	✗
Programmatic environments						
EnvScaler (2026)	✓	✓	✗	✗	✗	✗
AWM (2026)	✓	✓	✗	△	✗	✗
ScaleEnv (2026)	✓	✓	✓	✗	✗	✗
Agent-World (2026)	✓	✓	✓	✗	✗	✗
Skill2Env (2026a)	✓	✓	✓	✗	✗	✗
CrossEnv	✓	✓	✓	✓	✓	✓

Table 1: Environment and task synthesis capabilities. Exec.: real-service or code execution; Verif.: programmatic outcome checks; Ctrl.: adjustable complexity; Cross: multi-environment tasks; Consist.: cross-environment data alignment; Ext.: new-environment integration into joint tasks. ✓: supported; △: limited; ✗: not explicitly provided.

Our main contributions are as follows:

1. We introduce **CrossEnv**, which combines entity–tool graph-guided workflows with cross-environment data alignment to synthesize executable, verifiable training tasks.

2. We propose **CrossEnv-Curriculum**, which organizes training from single-environment tool use to increasingly complex coordination across environments and business domains, and finally to tasks requiring user interaction.

3. We analyze stagewise capability development and retention, comparing curriculum versus mixed training and single- versus cross-environment tasks to assess the roles of training strategy and task composition.

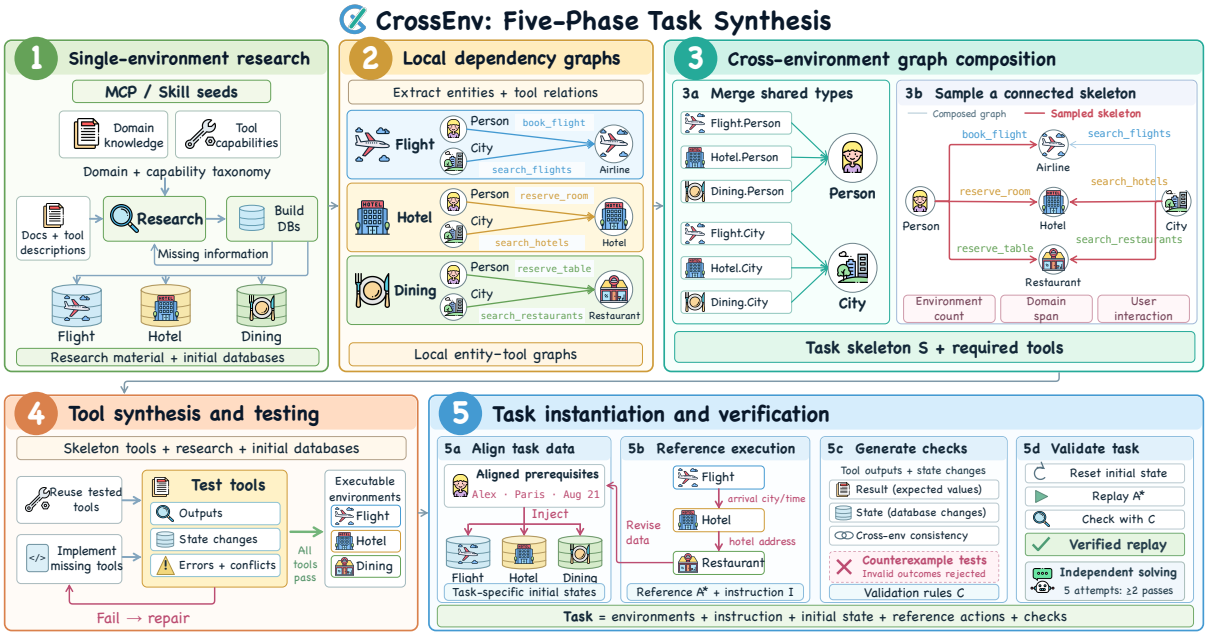

Figure 2: Overview of the CrossEnv task synthesis framework. Entity–tool graph composition forms task skeletons, while tool synthesis, data alignment, and reference execution produce executable tasks with verifiable outcomes. Tasks are validated through reference replay and independent agent attempts.

2 Related Work

2.1 Environment Synthesis

Table 1 compares approaches based on real-world access (Xu et al., 2025b), LLM simulation (Chen et al., 2025; Xu et al., 2026c), and programmatic construction. For general synthesis, real services and executable programs support action execution, while programmatic approaches provide task-outcome checks (Song et al., 2026; Yang et al., 2025c; Xu et al., 2026b). ScaleEnv (Tu et al., 2026), Agent-World (Dong et al., 2026), and Skill2Env (Xu et al., 2026a) explicitly control task complexity. Graph-guided synthesis supports long-horizon tool tasks in ToolVerse (Zhou et al., 2026), terminal skill composition in SkillSynth (Fan et al., 2026), and controllable subtask complexity in OmniBench (Bu et al., 2025). Across environments, ToolLLM (Qin et al., 2023) and Toucan (Xu et al., 2025c) support multi-service tasks and new-service integration; Simia (Li et al., 2025) simulates multi-environment workflows, while AWM (Wang et al., 2026) provides limited support within bundled service suites. ARE (Froger et al., 2025) addresses cross-app data coherence through dependency-aware generation. CrossEnv composes independently constructed environments through shared entity types and aligns task-specific records for jointly verifiable workflows.

2.2 Curriculum Learning

Curriculum learning organizes training examples by difficulty or learning progress to support progressive skill acquisition (Bengio et al., 2009; Kumar et al., 2010; Graves et al., 2017; Florensa et al., 2017, 2018). For LLM agents, AgentGym-RL (Xi et al., 2025) gradually increases interaction budgets, while WebRL (Qi et al., 2025) generates tasks from failures and filters them by difficulty. AgentGen (Hu et al., 2025) constructs planning tasks at varying difficulty levels through bidirectional evolution. More recently, GenEnv (Guo et al., 2025) and Tool-R0 (Acikgoz et al., 2026) use co-evolution to generate tasks adapted to the agent's current capabilities. CrossEnv-Curriculum organizes training around the coordination structure of cross-environment workflows, providing an explicit progression from foundational execution skills to joint task completion.

3 CrossEnv

CrossEnv automatically synthesizes executable, verifiable cross-environment tasks. Figure 2 presents an overview of its five-phase pipeline.

A cross-environment task involves tool use in at least two environments. Task completion requires using information from one environment to guide actions in another or meeting constraints on their

combined results. Each task is represented as

$$\tau = (\mathcal{E}_S, I, \mathbf{s}_{0,\tau}, A^*, C_\tau), \tag{1}$$

where $\mathcal{E}_S$ is the set of environments used in the task, I is the instruction, $\mathbf{s}_{0,\tau}$ is their joint initial state, A^* is a reference execution sequence, and C_τ contains executable verification rules. Task completion requires all rules to pass, evaluated from tool outputs and final environment states.

3.1 Single-Environment Research

Environment research. We collect environment seeds from MCP services and public skill documents, retaining 801 MCP seeds and 682 skill seeds after filtering. A research agent retrieves documentation and data, while a data construction agent builds initial databases and requests further retrieval when needed. These materials support entity extraction and tool implementation.

Environment taxonomy. We classify the seeds into 10 business domains and 84 capability categories, drawing on the Agent-World taxonomy (Dong et al., 2026) and seed functionality. This taxonomy supports within-domain and cross-domain task composition; the complete classification is provided in Appendix B.

3.2 Local Dependency Graphs

For each environment E_i, we extract entity types and attributes from its initial database and tool descriptions. We then construct a local entity–tool graph $G_i = (V_i, R_i)$, where nodes represent entity types and directed edges represent relations supported by tools. In Figure 2, for example, `reserve_room` supports a booking relation from `Person` to `Hotel`, while `search_hotels` supports a search relation from `City` to `Hotel`.

3.3 Cross-Environment Graph Composition

Graph composition. Different environments may support different operations involving the same entity types. We merge their local graphs through these shared types, retaining each relation's source environment and supporting tools. In Figure 2, shared `Person` and `City` nodes connect the flight, hotel, and dining graphs, bringing their operations into a cross-environment entity–tool graph.

Task skeleton sampling. From the composed graph, we sample connected subgraphs spanning at least two environments. We assess whether the selected operations involve related entities, complement one another, and form a coherent business workflow. For each selected combination, we define a task objective specifying how the operations work together, including which actions use results from other environments and what constraints their combined outcomes must satisfy. This produces a task skeleton S and its required tools. We control task composition through environment count, business-domain scope, and whether user interaction is required.

3.4 Tool Synthesis and Testing

Given a task skeleton, we reuse tested tool implementations and have a coding agent implement missing tools using their descriptions, research findings, and initial databases. We test required tools for correct outputs, intended state changes, and error handling, repairing and retesting failures. The tested implementations provide executable operations for the participating environments and support subsequent task instantiation.

3.5 Task Instantiation and Verification

Data alignment. Given a task skeleton and its objective, we select concrete entities and records, supplement missing data, and resolve conflicting information across environments. We plan initial data and reference actions jointly: required records must exist initially or be produced by earlier actions, while records the agent must create are not inserted in advance. This produces the aligned initial state $\mathbf{s}_{0,\tau}$.

Reference execution. We execute the planned actions from $\mathbf{s}_{0,\tau}$, recording outputs and state changes. Actions use results from other environments as required by the objective. Execution failures trigger revisions to the data and action plan. Successful execution provides the reference sequence A^*, from which we generate the task instruction I together with the task scenario.

Verification rule generation. We derive executable rules C_τ from the task objective and constraints, using reference execution results to specify expected outcomes. The rules check tool outputs, required state changes, and cross-environment consistency. Negative cases, such as omitted operations and incorrect values, test whether the rules reject incomplete or incorrect outcomes.

Task validation. We reset the environments and replay A^*, requiring reproducible outcomes that pass all verification rules. We then run five independent agent attempts from the same initial state, using the task instruction and available interaction interfaces.

We retain tasks with successful reference replay and at least two successful agent attempts.

4 CrossEnv-Curriculum

Curriculum organization. We organize CrossEnv tasks into five levels according to environment count, business-domain scope, and user interaction. **L1** contains single-environment tasks that establish basic tool-use skills. **L2** introduces coordination between two environments within one business domain, and **L3** expands to three or more environments within the same domain. **L4** extends coordination across multiple business domains while retaining three or more environments. **L5** retains L4's environment-count and domain-scope requirements and introduces user interaction. The curriculum progressively expands environment count and business-domain scope, then introduces user interaction under L4's structural requirements.

Training procedure. We first perform supervised fine-tuning on L1 reference trajectories to teach tool selection, argument specification, and the use of tool outputs. We then conduct reinforcement learning sequentially on L2–L5, initializing each stage from the preceding checkpoint. During RL, the agent acts on the task instruction and interaction history, receiving rewards from executable checks of tool outputs and final environment states.

5 Experiments

We present the experimental setup (§5.1) and overall performance (§5.2), then analyze curriculum learning (§5.3), generalization (§5.4), the value of cross-environment tasks (§5.5), and synthesis quality and construction cost (§5.6).

5.1 Experimental Setup

Benchmarks. Our main evaluation uses four benchmarks covering complementary capabilities: **interactive tool use**, assessed by τ^2-bench (Barres et al., 2025) through multi-step execution and user coordination; **long-horizon task completion**, assessed by DeepPlanning-Shopping (Zhang et al., 2026) through constrained planning and VitaBench (He et al., 2025) through coordinated service interactions; and **general assistant tasks**, assessed by Claw-Eval (Ye et al., 2026) through service orchestration and professional dialogue. Benchmark subsets and evaluation settings are detailed in Appendix E.

Level	Stage	Train	Val.	Avg. Turns
L1	SFT	800	40	4.73
L2	RL	579	40	4.84
L3	RL	620	50	8.70
L4	RL	600	50	12.66
L5	RL + User	355	40	13.76

Table 2: CrossEnv training and validation data by curriculum level. Turns count tool calls and, for L5, one required clarification question.

Models. We use Qwen3 backbones (Yang et al., 2025a) with 1.7B, 4B, and 8B parameters and compare against the corresponding base models to assess training gains. Environment-trained baselines include Toucan (Xu et al., 2025c), EnvACE (Xu et al., 2026c), EnvScaler (Song et al., 2026), Arctic-AWM (Wang et al., 2026), and Agent-World (Dong et al., 2026). We use in-house reproductions of EnvScaler-4B and Agent-World-4B/8B; the Agent-World implementation was not publicly available at the time of our experiments. We also include frontier closed-source models (OpenAI, n.d.; Anthropic, 2025; Google DeepMind, 2026; Qwen Team, n.d.) and open-weight general-purpose models (DeepSeek, 2026; Z.ai, n.d.; Kimi Team et al., 2026) as reference points. Model versions and sources are documented in Appendix D.

Training. We train on CrossEnv-generated data following the curriculum in Section 4. We first perform full-parameter SFT on 800 L1 reference trajectories, followed by stagewise trajectory-level outcome GRPO on 2,154 tasks across L2–L5, for a total of 2,954 training records. Table 2 reports training and validation counts and average reference lengths by level. All CrossEnv models in the main results table complete L1–L5 training. Capability tracking extends through L5; the Mixed comparison and longer-workflow tests use checkpoints through L4.

5.2 Overall Performance

Table 3 reports results on the four external benchmarks. We highlight three findings.

(1) Consistent gains over base models. Across all three model sizes, CrossEnv improves or matches the corresponding Qwen3 backbone on each benchmark. At 8B, scores increase from 1.7 to 10.8 on DeepPlanning, 7.0 to 14.0 on VitaBench, and 14.4 to 23.6 on Claw-Eval, showing gains across constrained planning, coordinated service interactions, and general assistant tasks.

Method	Train data	Source	τ^2-bench					DeepPlanning				VitaBench					Claw-Eval		
			Ret.	Air.	Tel.	Bnk.	**Avg.**	L1	L2	L3	**Avg.**	Del.	Ins.	OTA	Cr.	**Avg.**	Gen.	MT	**Avg.**
Frontier Closed-Source Models																			
GPT-5.6-Luna	—	☁	21.9	44.0	9.6	6.2	17.1	40.0	40.0	20.0	36.7	32.0	40.0	8.0	12.0	23.0	78.3	57.9	74.4
Claude-Haiku-4.5	—	☁	36.0	58.0	1.8	3.1	20.0	36.0	24.0	15.0	27.5	40.0	32.0	24.0	8.0	26.0	49.0	36.8	46.7
Gemini-3.8-Flash	—	☁	9.6	24.0	9.6	2.1	9.6	34.0	34.0	25.0	32.5	0.0	0.0	48.0	0.0	12.0	68.8	84.2	71.8
Qwen3.8-Max	—	☁	30.7	50.0	7.9	4.1	19.5	72.0	66.0	75.0	70.0	64.0	52.0	36.0	36.0	47.0	78.3	76.3	77.9
Open-Weight General-Purpose Models																			
DeepSeek-V4-Pro	—	☁	43.9	62.0	7.0	4.1	24.8	62.0	38.0	50.0	50.0	44.0	48.0	36.0	28.0	39.0	75.2	81.6	76.4
GLM-5.3	—	☁	16.7	36.0	4.4	3.1	12.0	62.0	42.0	30.0	48.3	64.0	64.0	60.0	40.0	57.0	77.1	71.1	75.9
Kimi-K3	—	☁	20.2	42.0	3.5	4.1	13.9	42.0	24.0	25.0	31.7	60.0	48.0	44.0	32.0	46.0	80.9	78.9	80.5
Qwen3-1.7B	—	↓	5.3	14.0	4.4	3.1	5.6	4.0	0.0	0.0	1.7	0.0	0.0	0.0	0.0	0.0	8.9	2.6	7.7
Qwen3-4B	—	↓	9.6	28.0	5.3	2.1	8.8	2.0	0.0	0.0	0.8	0.0	4.0	0.0	0.0	1.0	15.3	0.0	12.3
Qwen3-8B	—	↓	14.9	26.0	7.0	**5.2**	11.5	4.0	0.0	0.0	1.7	16.0	12.0	0.0	0.0	7.0	17.2	2.6	14.4
Environment-Trained Baselines																			
Toucan-7B	119,300	↓	9.6	8.0	0.0	3.1	4.8	6.0	2.0	0.0	3.3	0.0	0.0	0.0	0.0	0.0	26.1	2.6	21.5
EnvACE-1.7B	8,129	↓	4.4	18.0	0.9	2.1	4.5	2.0	0.0	0.0	0.8	0.0	0.0	0.0	0.0	0.0	6.4	0.0	5.1
EnvACE-8B	8,129	↓	12.3	**30.0**	3.5	3.1	9.6	4.0	2.0	0.0	2.5	8.0	0.0	0.0	0.0	2.0	16.6	**7.9**	14.9
EnvScaler-1.7B	11,572	↓	7.9	10.0	0.9	1.0	4.3	6.0	4.0	0.0	4.2	8.0	0.0	0.0	0.0	2.0	11.5	2.6	9.7
EnvScaler-4B	11,572	⚙	11.4	26.0	**10.5**	**5.2**	11.5	8.0	4.0	**10.0**	6.7	24.0	0.0	0.0	0.0	6.0	13.4	0.0	10.8
EnvScaler-8B	11,572	↓	**21.9**	22.0	7.0	4.1	**12.8**	12.0	6.0	5.0	8.3	24.0	4.0	4.0	4.0	9.0	24.8	**7.9**	21.5
Arctic-AWM-4B	3,315	↓	3.5	16.0	6.1	3.1	5.9	0.0	0.0	0.0	0.0	8.0	0.0	0.0	0.0	2.0	8.3	0.0	6.7
Arctic-AWM-8B	3,315	↓	12.3	26.0	**10.5**	4.1	11.5	2.0	2.0	0.0	1.7	8.0	12.0	0.0	0.0	5.0	15.9	**7.9**	14.4
Agent-World-4B	3,000	⚙	12.3	26.0	7.0	1.0	9.6	12.0	4.0	0.0	6.7	24.0	4.0	0.0	0.0	7.0	21.7	0.0	17.4
Agent-World-8B	3,000	⚙	14.0	18.0	7.0	2.1	9.3	**14.0**	6.0	5.0	9.2	**36.0**	8.0	**12.0**	0.0	**14.0**	24.2	5.3	20.5
Our Models																			
CrossEnv-1.7B	2,954	⚙	5.3	12.0	7.9	3.1	6.4	12.0	4.0	0.0	6.7	4.0	4.0	0.0	0.0	2.0	8.9	2.6	7.7
CrossEnv-4B	2,954	⚙	11.4	22.0	**10.5**	2.1	10.1	12.0	**10.0**	0.0	9.2	8.0	8.0	8.0	4.0	7.0	22.3	5.3	19.0
CrossEnv-8B	2,954	⚙	14.9	28.0	9.7	4.1	12.3	**14.0**	**10.0**	5.0	**10.8**	24.0	**16.0**	8.0	**8.0**	**14.0**	**27.4**	**7.9**	**23.6**

Table 3: Performance on four external benchmarks. Source: ☁ hosted API; ↓ public checkpoint deployed locally; ⚙ trained in this work. Avg. is weighted by subset size. Bold and underlined scores indicate the best and second-best among the displayed 1.7B–8B models.

(2) Competitive performance with fewer training records. Among the evaluated 1.7B–8B models, CrossEnv-8B achieves the highest aggregate scores on DeepPlanning and Claw-Eval and ties with Agent-World-8B on VitaBench. With 2,954 training records, approximately one quarter of EnvScaler's reported dataset, it outperforms EnvScaler-8B on three of the four benchmarks.

(3) Gains vary across capabilities and model scales. CrossEnv improves planning and service coordination across all three model sizes, with larger models also showing stronger gains on general assistant tasks. Improvements on interactive tool use are more modest. These results indicate benefits across multiple capabilities, with the extent of improvement varying by task type and model scale.

5.3 Curriculum Learning Analysis

Capability development and retention. We track Qwen3-8B from Base through L5 on fixed evaluation sets (Figure 3). L1 SFT improves basic tool use, while subsequent stages improve the model's ability to act across required environments and execute later workflow steps. These gains do not always translate immediately into correct outcomes: the model executes more required tools at L3 but passes fewer final-state checks. More final-state checks pass at L4 and L5, with user interaction also improving at L5. Although some execution scores decrease slightly at later stages, most earlier gains are retained. Overall, the model learns to coordinate more complex workflows while largely preserving previously acquired tool-use skills.

Comparison with mixed training. We compare curriculum and mixed training from the same L1 SFT checkpoint using the same L2–L4 task pools. At checkpoints with similar cumulative RL GPU-hours, curriculum training achieves higher success at every task level (Figure 4(a)). Starting with shorter two-environment tasks before introducing longer, more complex workflows may reduce the interaction cost per sample early in training. Curriculum training also accumulates more rollouts within

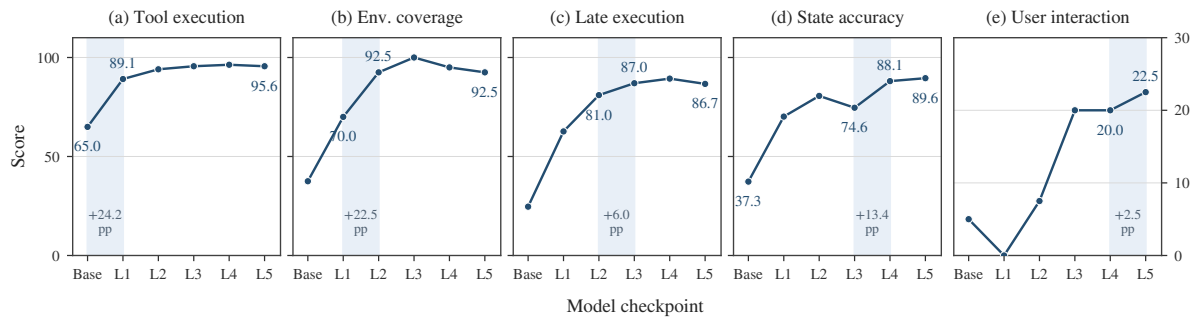

Figure 3: Qwen3-8B capability scores from Base to L5. Shading highlights the corresponding training stages. Panel (e) reports a composite interaction score, not a success rate (Appendix F.2).

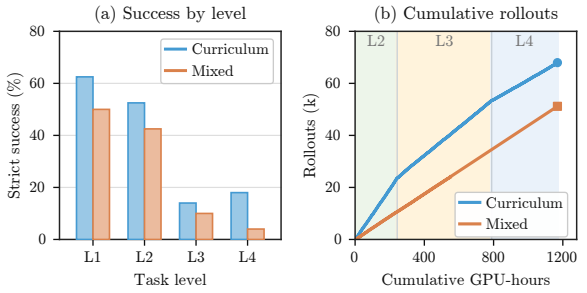

Figure 4: Curriculum vs. Mixed (8B): (a) strict success at similar GPU-hours; (b) cumulative rollouts.

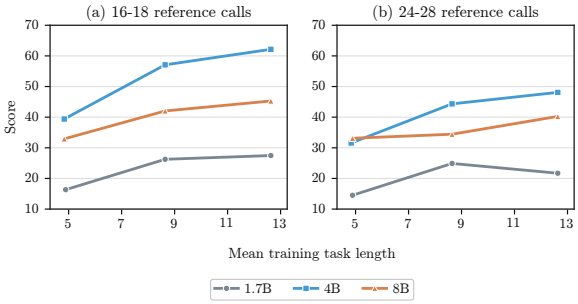

Figure 5: Generalization to longer workflows. Scores show the average proportion of task checks passed (%).

Model	Pair	Success	Obs.	State	Answer
1.7B	Seen	+6.7	+40.7	+37.5	+22.2
	Unseen	+5.3	+38.9	0.0	+11.1
4B	Seen	+30.0	+43.7	+48.5	+32.2
	Unseen	+24.0	+33.7	+26.9	+22.5
8B	Seen	+28.0	+44.3	+30.4	+26.4
	Unseen	+16.0	+34.3	+14.6	+13.0

Table 4: CrossEnv gains over Base on seen and unseen environment combinations (percentage points). Obs.: observation.

a comparable compute budget (Figure 4(b)). These results support the value of the progressive curriculum for both training effectiveness and efficient use of compute under the current configurations.

5.4 Generalization Analysis

We examine whether learned skills transfer to new environment combinations, longer workflows, and external tasks less similar to training data.

Generalization to unseen environment combinations. We evaluate combinations of familiar environments that were never used together in a training task. Table 4 compares gains over Base on seen and unseen combinations, with task success requiring all task requirements to be met. CrossEnv improves task success on both groups at all three model sizes, showing transfer to new combinations of familiar environments.

Generalization to longer workflows. We evaluate curriculum checkpoints on two sets of tasks with reference workflows of 16–18 and 24–28 tool calls, both exceeding the training maximum of 14 calls. Figure 5 shows the average proportion of task checks passed. From L1 to L4, both 4B and 8B improve on both sets, while 1.7B shows a late-stage decline on longer tasks. These results suggest that training on shorter workflows can improve performance on longer tasks, although gains vary by model size.

Transfer across task similarity. We examine whether CrossEnv learns transferable skills beyond its training tasks. Inspired by EnvScaler (Song et al., 2026), we use instruction-level semantic similarity as a proxy for task similarity. In Figure 6, both model sizes improve over Base on average in the lowest-similarity group of each benchmark, while higher similarity does not consistently yield larger gains. These results support generalization to tasks less similar to the training data.

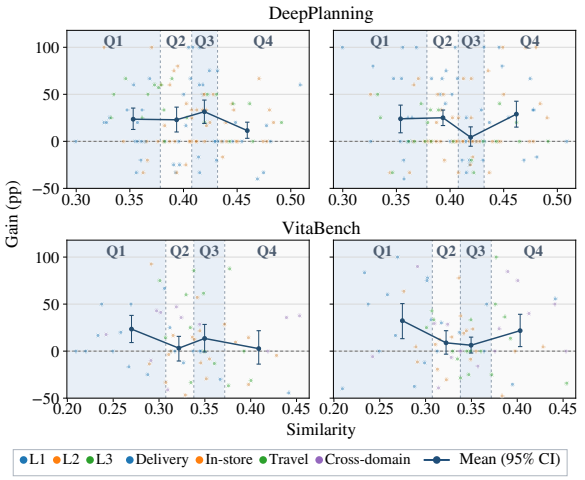

Figure 6: Task similarity and CrossEnv gains over Base at 4B (left) and 8B (right). Lines show mean gains in four similarity groups with 95% confidence intervals.

Metric	Base	SFT	SingleEnv	CrossEnv
Benchmark performance				
τ^2-bench	**11.5**	8.5	8.8	**11.5**
DeepPlanning	1.7	3.3	6.7	**7.5**
Claw-Eval	14.4	16.9	15.4	**20.0**
Execution capabilities				
Output reuse	6.7	8.3	1.7	**11.7**
Later-step success	23.3	18.3	28.3	**33.3**
Joint success	5.0	1.7	0.0	**8.3**

Table 5: Qwen3-8B comparison (%). Bold: row best.

5.5 The Value of Cross-Environment Tasks

Performance comparison. We compare single- and cross-environment RL using the same SFT cold-start initialization (Table 5). CrossEnv outperforms Base, SFT, and SingleEnv on DeepPlanning and Claw-Eval; on τ^2-bench, it exceeds SFT and SingleEnv and matches Base.

Execution analysis. To understand why cross-environment training yields larger gains than single-environment training, we systematically analyze evaluation trajectories across the four evaluation benchmarks. The differences center on three aspects: (1) **output reuse**, correctly using prior tool outputs in subsequent operations; (2) **later-step success**, continuing execution to complete required later operations; and (3) **joint success**, satisfying all service goals and task constraints together (Appendix H). CrossEnv outperforms SingleEnv on all three metrics (Table 5), suggesting that cross-environment training helps models coordinate multiple operations toward a shared goal.

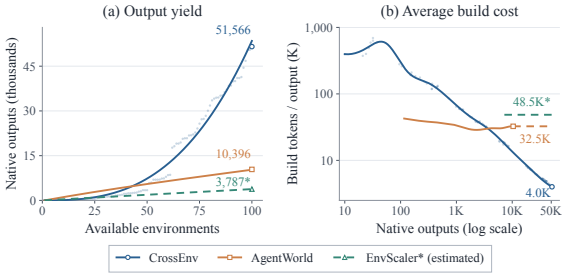

Figure 7: Output yield and average environment-build cost. *: estimates.

5.6 Synthesis Quality and Construction Cost

Execution and verification quality. We randomly sample 100 tasks, stratified by domain, with 50 each from MCP and Skill environments. All reference workflows run successfully and pass all checks. Further tests on subsets confirm that: (1) resetting the environments and rerunning the workflows yields the same results; (2) skipping a required operation prevents full credit; and (3) reordering queries without changing their results still earns full credit (Appendix I).

Scalability and construction cost. As the environment pool grows, CrossEnv combines new environments with existing ones, enabling more task combinations and accelerating growth in candidate task yield (Figure 7(a)). These combinations reuse environment data, tools, and entity relations without repeating environment construction, so the average construction cost per candidate generally decreases as output grows (Figure 7(b)). Although supporting long cross-environment workflows requires a higher initial investment, this reuse eventually reduces CrossEnv's average construction cost to approximately 4.0k tokens per candidate, below our reproduced AgentWorld result and the EnvScaler estimate based on reported figures.

6 Conclusion

We presented CrossEnv, a framework for synthesizing executable and verifiable cross-environment tasks with consistent data, and CrossEnv-Curriculum, a progressive training scheme for learning from these tasks. Across four external benchmarks, combining cross-environment data with curriculum training achieves competitive performance with fewer training examples. Our analyses suggest that cross-environment tasks improve workflow coordination, while progressive training achieves higher success and collects more rollouts

at comparable compute. Further evaluations provide evidence of transfer to new environment combinations, external environments, and workflows longer than those seen during training.

Limitations

Programmatically synthesized environments may simplify real-world service behavior, including changing interfaces, noisy responses, and complex business constraints. Although CrossEnv aligns task-specific data to support coherent workflows, transfer to real services with dynamically changing data and operational conditions requires further evaluation.

Our generalization evidence is bounded by the evaluated environments and workflow lengths. The combination experiments test new combinations of familiar environments, while the longer-workflow experiments extend to 24–28 reference tool calls and measure the proportion of task checks passed. These results do not establish reliable end-to-end completion in entirely new services or substantially longer workflows.

Some training comparisons do not fully isolate the factors contributing to performance gains. In particular, curriculum and mixed training use similar cumulative GPU-hours but differ in batch configuration, parallelism, and reference-model settings, limiting attribution of the gains solely to curriculum order. Likewise, comparisons of training-record counts across methods do not establish equivalent advantages in training compute.

Ethical Considerations

CrossEnv aims to reduce the effort required to construct cross-environment tasks for agent training. By synthesizing executable and verifiable workflows, it supports research on tool use and service coordination. This capability has potential dual-use implications: tasks designed for beneficial automation could also be adapted to train agents for harmful purposes. We encourage responsible use of the framework and synthesized tasks, and careful evaluation of trained agents before deployment in sensitive applications.

AI Use Statement

We used generative AI tools to assist with manuscript drafting and language editing. The authors reviewed the resulting content and take full responsibility for the final manuscript.

References

Emre Can Acikgoz, Cheng Qian, Jonas Hübotter, Heng Ji, Dilek Hakkani-Tür, and Gokhan Tur. 2026. Tool-R0: Self-Evolving LLM Agents for Tool-Learning from Zero Data. *Preprint*, arXiv:2602.21320.

Anthropic. 2025. Introducing Claude Haiku 4.5. Accessed September 21, 2026.

Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. 2025. τ^2-Bench: Evaluating Conversational Agents in a Dual-Control Environment. *Preprint*, arXiv:2506.07982.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pages 41–48. ACM.

Wendong Bu, Yang Wu, Qifan Yu, Minghe Gao, Bingchen Miao, Zhenkui Zhang, Kaihang Pan, Yunfei Li, Mengze Li, Wei Ji, Juncheng Li, Siliang Tang, and Yueting Zhuang. 2025. What Limits Virtual Agent Application? OmniBench: A Scalable Multi-Dimensional Benchmark for Essential Virtual Agent Capabilities. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 5725–5748. PMLR.

Zhaorun Chen, Zhuokai Zhao, Kai Zhang, Bo Liu, Qi Qi, Yifan Wu, Tarun Kalluri, Sara Cao, Yuanhao Xiong, Haibo Tong, Huaxiu Yao, Hengduo Li, Jiacheng Zhu, Xian Li, Dawn Song, Bo Li, Jason Weston, and Dat Huynh. 2025. Scaling Agent Learning via Experience Synthesis. *Preprint*, arXiv:2511.03773.

DeepSeek. 2026. DeepSeek-V4-Pro Update. Accessed September 21, 2026.

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. 2026. Agent-World: Scaling Real-World Environment Synthesis for Evolving General Agent Intelligence. *Preprint*, arXiv:2604.18292.

Zhiyuan Fan, Tinghao Yu, Yuanjun Cai, Jiangtao Guan, Yun Yang, Dingxin Hu, Jiang Zhou, Xing Wu, Zhuo Han, Feng Zhang, and Lilin Wang. 2026. Toward Scalable Terminal Task Synthesis via Skill Graphs. *Preprint*, arXiv:2604.25727.

Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. 2018. Automatic Goal Generation for Reinforcement Learning Agents. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1515–1528. PMLR.

Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. 2017. Reverse Curriculum Generation for Reinforcement Learning. In *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 482–495. PMLR.

Romain Froger, Pierre Andrews, Matteo Bettini, Amar Budhiraja, Ricardo Silveira Cabral, Virginie Do, Emilien Garreau, Jean-Baptiste Gaya, Hugo Laurençon, Maxime Lecanu, Kunal Malkan, Dheeraj Mekala, Pierre Ménard, Gerard Moreno-Torres Bertran, Ulyana Piterbarg, Mikhail Plekhanov, Mathieu Rita, Andrey Rusakov, Vladislav Vorotilov, and 5 others. 2025. ARE: Scaling Up Agent Environments and Evaluations. *Preprint*, arXiv:2509.17158.

Google DeepMind. 2026. Gemini 3.8 Flash. Accessed September 21, 2026.

Alex Graves, Marc G. Bellemare, Jacob Menick, Rémi Munos, and Koray Kavukcuoglu. 2017. Automated Curriculum Learning for Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1311–1320. PMLR.

Jiacheng Guo, Ling Yang, Peter Chen, Qixin Xiao, Yinjie Wang, Xinzhe Juan, Jiahao Qiu, Ke Shen, and Mengdi Wang. 2025. GenEnv: Difficulty-Aligned Co-Evolution Between LLM Agents and Environment Simulators. *Preprint*, arXiv:2512.19682.

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. 2025. VitaBench: Benchmarking LLM Agents with Versatile Interactive Tasks in Real-world Applications. *Preprint*, arXiv:2509.26490.

Mengkang Hu, Pu Zhao, Can Xu, Qingfeng Sun, Jianguang Lou, Qingwei Lin, Ping Luo, and Saravan Rajmohan. 2025. AgentGen: Enhancing Planning Abilities for Large Language Model based Agent via Environment and Task Generation. *Preprint*, arXiv:2408.00764.

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, M. C., Jianfeng Cai, Xinyuan Cai, Peizhou Cao, Yuxuan Cao, Ziwei Chai, Y. Charles, H. S. Che, Guanduo Chen, Guangyu Chen, Guanzheng Chen, Huarong Chen, Jia Chen, Jianlong Chen, Jun Chen, and 383 others. 2026. Kimi K3: Open Frontier Intelligence. *Preprint*, arXiv:2607.24653.

M. Pawan Kumar, Benjamin Packer, and Daphne Koller. 2010. Self-Paced Learning for Latent Variable Models. In *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.

Junlong Li, Wenshuo Zhao, Jian Zhao, Weihao Zeng, Haoze Wu, Xiaochen Wang, Rui Ge, Yuxuan Cao, Yuzhen Huang, Wei Liu, Junteng Liu, Zhaochen Su, Yiyang Guo, Fan Zhou, Lueyang Zhang, Juan Michelini, Xingyao Wang, Xiang Yue, Shuyan Zhou, and 2 others. 2026. The Tool Decathlon: Benchmarking Language Agents for Diverse, Realistic, and Long-Horizon Task Execution. *Preprint*, arXiv:2510.25726.

Yuetai Li, Huseyin A Inan, Xiang Yue, Wei-Ning Chen, Lukas Wutschitz, Janardhan Kulkarni, Radha Poovendran, Robert Sim, and Saravan Rajmohan. 2025. Simulating Environments with Reasoning Models for Agent Training. *Preprint*, arXiv:2511.01824.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, and 3 others. 2025. AgentBench: Evaluating LLMs as Agents. *Preprint*, arXiv:2308.03688.

OpenAI. n.d. GPT-5.6 Luna. Accessed September 21, 2026.

Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue Yang, Jiadai Sun, Shuntian Yao, Tianjie Zhang, Wei Xu, Jie Tang, and Yuxiao Dong. 2025. WebRL: Training LLM Web Agents via Self-Evolving Online Curriculum Reinforcement Learning. *Preprint*, arXiv:2411.02337.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. *Preprint*, arXiv:2307.16789.

Qwen Team. n.d. Qwen3.8-2.4T-A95B. Accessed September 21, 2026.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551. Curran Associates, Inc.

Quan Shi, Alexandra Zytek, Pedram Razavi, Karthik Narasimhan, and Victor Barres. 2026. τ-Knowledge: Evaluating Conversational Agents over Unstructured Knowledge. *Preprint*, arXiv:2603.04370.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652. Curran Associates, Inc.

Xiaoshuai Song, Haofei Chang, Guanting Dong, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. 2026. EnvScaler: Scaling Tool-Interactive Environments for LLM Agent via Programmatic Synthesis. *Preprint*, arXiv:2601.05808.

Harsh Trivedi, Tushar Khot, Mareike Hartmann, Ruskin Manku, Vinty Dong, Edward Li, Shashank Gupta, Ashish Sabharwal, and Niranjan Balasubramanian. 2024. AppWorld: A Controllable World of Apps and People for Benchmarking Interactive Coding Agents. *Preprint*, arXiv:2407.18901.

Dunwei Tu, Hongyan Hao, Hansi Yang, Yihao Chen, Yi-Kai Zhang, Zhikang Xia, Yu Yang, Yueqing Sun, Xingchen Liu, Furao Shen, Qi Gu, Hui Su, and Xunliang Cai. 2026. ScaleEnv: Scaling Environment Synthesis from Scratch for Generalist Interactive Tool-Use Agent Training. *Preprint*, arXiv:2602.06820.

Zhaoyang Wang, Canwen Xu, Boyi Liu, Yite Wang, Siwei Han, Zhewei Yao, Huaxiu Yao, and Yuxiong He. 2026. Agent World Model: Infinity Synthetic Environments for Agentic Reinforcement Learning. *Preprint*, arXiv:2602.10090.

Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2024. Agent Workflow Memory. *Preprint*, arXiv:2409.07429.

Zijian Wu, Xiangyan Liu, Xinyuan Zhang, Lingjun Chen, Fanqing Meng, Lingxiao Du, Yiran Zhao, Fanshi Zhang, Yaoqi Ye, Jiawei Wang, Zirui Wang, Jinjie Ni, Yufan Yang, Arvin Xu, and Michael Qizhe Shieh. 2025. MCPMark: A Benchmark for Stress-Testing Realistic and Comprehensive MCP Use. *Preprint*, arXiv:2509.24002.

Zhiheng Xi, Jixuan Huang, Chenyang Liao, Baodai Huang, Honglin Guo, Jiaqi Liu, Rui Zheng, Junjie Ye, Jiazheng Zhang, Wenxiang Chen, Wei He, Yiwen Ding, Guanyu Li, Zehui Chen, Zhengyin Du, Xuesong Yao, Yufei Xu, Jiecao Chen, Tao Gui, and 4 others. 2025. AgentGym-RL: Training LLM Agents for Long-Horizon Decision Making through Multi-Turn Reinforcement Learning. *Preprint*, arXiv:2509.08755.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments. In *Advances in Neural Information Processing Systems*, volume 37, pages 52040–52094. Curran Associates, Inc.

Binfeng Xu, Jian Hu, Hao Zhang, Shaokun Zhang, Yifan Zhang, Michael Demoret, Jan Kautz, and Yi Dong. 2026a. Reinforcing Agents with Collective Skills. alphaXiv preprint.

Minrui Xu, Zilin Wang, Mengyi DENG, Zhiwei Li, Zhicheng Yang, Xiao Zhu, Yinhong Liu, Boyu Zhu, Baiyu Huang, Chao Chen, Heyuan Deng, Fei Mi, Lifeng Shang, Xingshan Zeng, and Zhijiang Guo. 2026b. EnvFactory: Scaling Tool-Use Agents via Executable Environments Synthesis and Robust RL. *Preprint*, arXiv:2605.18703.

Tianqi Xu, Linyao Chen, Dai-Jie Wu, Yanjun Chen, Zecheng Zhang, Xiang Yao, Zhiqiang Xie, Yongchao Chen, Shilong Liu, Bochen Qian, Anjie Yang, Zhaoxuan Jin, Jianbo Deng, Philip Torr, Bernard Ghanem, and Guohao Li. 2025a. CRAB: Cross-environment agent benchmark for multimodal language model agents. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 21607–21647, Vienna, Austria. Association for Computational Linguistics.

Yiheng Xu, Dunjie Lu, Zhennan Shen, Junli Wang, Zekun Wang, Yuchen Mao, Caiming Xiong, and Tao Yu. 2025b. AgentTrek: Agent Trajectory Synthesis via Guiding Replay with Web Tutorials. *Preprint*, arXiv:2412.09605.

Zhangchen Xu, Adriana Meza Soria, Shawn Tan, Anurag Roy, Ashish Sunil Agrawal, Radha Poovendran, and Rameswar Panda. 2025c. TOUCAN: Synthesizing 1.5M Tool-Agentic Data from Real-World MCP Environments. *Preprint*, arXiv:2510.01179.

Zishan Xu, Zhiyuan Yao, Yuxin Chen, Yifu Guo, Zhengxi Lu, Yuquan Lu, Jinyang Huang, Yan Xu, Yasheng Wang, Weinan Zhang, Xingshan Zeng, and Weiwen Liu. 2026c. EnvACE: Internalizing Environment Dynamics via World Rehearsal for Agentic Reinforcement Learning. *Preprint*, arXiv:2608.06197.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. Qwen3 Technical Report. *Preprint*, arXiv:2505.09388.

Chen Yang, Ran Le, Yun Xing, Zhenwei An, Zongchao Chen, Wayne Xin Zhao, Yang Song, and Tao Zhang. 2025b. ToolMind Technical Report: A Large-Scale, Reasoning-Enhanced Tool-Use Dataset. *Preprint*, arXiv:2511.15718.

John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. 2025c. SWE-smith: Scaling Data for Software Engineering Agents. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference. Curran Associates, Inc.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. τ-bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains. *Preprint*, arXiv:2406.12045.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. *Preprint*, arXiv:2210.03629.

Bowen Ye, Rang Li, Qibin Yang, Yuanxin Liu, Linli Yao, Hanglong Lv, Zhihui Xie, Chenxin An, Lei Li, Lingpeng Kong, Qi Liu, Zhifang Sui, and Tong Yang. 2026. Claw-eval: Towards trustworthy evaluation of autonomous agents. *Preprint*, arXiv:2604.06132.

Z.ai. n.d. GLM-5.3. Accessed September 21, 2026.

Yinger Zhang, Shutong Jiang, Renhao Li, Jianhong Tu, Yang Su, Lianghao Deng, Xudong Guo, Chenxu Lv, and Junyang Lin. 2026. DeepPlanning: Benchmarking Long-Horizon Agentic Planning with Verifiable Constraints. *Preprint*, arXiv:2601.18137.

Shuaiyu Zhou, Fengpeng Yue, Zengjie Hu, Yuanzhe Shen, Chenyang Zhang, feng hong, Cao Liu, and Ke Zeng. 2026. ToolVerse: Unlocking Massive Environments and Long-Horizon Tasks for Agentic Reinforcement Learning. *Preprint*, arXiv:2607.15660.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. WebArena: A Realistic Web Environment for Building Autonomous Agents. In *International Conference on Learning Representations*, volume 2024, pages 15585–15606.

A Detailed CrossEnv Method

Building on the task definition in Equation 1, this appendix supplements Section 3 with formal interaction and graph definitions, seed sources and taxonomy statistics, tool testing procedures, and detailed data alignment and verification. It also specifies the reference replay and independent-agent filtering criteria used to retain tasks.

A.1 Agentic Interaction across Environments

We consider an agent that completes tasks by invoking tools across multiple environments and interacting with a user when required. Each environment $E_i = (\mathcal{S}_i, \mathcal{T}_i, \mathcal{O}_i, F_i)$ consists of a state space $\mathcal{S}_i$, a tool set $\mathcal{T}_i$, an observation space $\mathcal{O}_i$, and an execution function F_i. The state represents the current application data. Let $\mathcal{E} = \{E_1, \ldots, E_n\}$ denote the available environments. For a task τ, only its participating environments $\mathcal{E}_S \subseteq \mathcal{E}$ enter the joint state $\mathbf{s}_t = (s_t^{(i)})_{E_i \in \mathcal{E}_S}$, with $\mathbf{s}_0 = \mathbf{s}_{0,\tau}$. We suppress the task index on subsequent states for readability. At each step, the agent selects an action using the task instruction I and interaction history h_t. For a tool call, $a_t = (i_t, f_t, x_t)$, where $E_{i_t} \in \mathcal{E}_S$, $f_t \in \mathcal{T}_{i_t}$, and x_t denotes the tool arguments. The selected environment executes the call and returns an observation:

$$\begin{aligned}(s_{t+1}^{(i_t)}, o_t) &= F_{i_t}(s_t^{(i_t)}, (f_t, x_t)),\\ s_{t+1}^{(j)} &= s_t^{(j)} \quad (E_j \in \mathcal{E}_S,\ j \neq i_t).\end{aligned} \tag{2}$$

For tasks requiring user interaction, a_t may instead be a user-directed message. The task's user interaction interface supplies a reply observation o_t and determines any associated application-state changes. User exchanges are recorded in h_t alongside tool calls and their outputs. Equation 2 describes tool execution only.

The agent selects actions from the task instruction and interaction feedback, using earlier tool outputs to coordinate subsequent operations (Yao et al., 2023; Wang et al., 2024). Let $A = (a_0, \ldots, a_{T-1})$ denote the action sequence, including tool calls and user-directed interactions where applicable, and let $T = |A|$. The resulting observations $O_A = (o_0, \ldots, o_{T-1})$ include tool outputs and any user replies. The same convention applies to the reference sequence A^*. Task completion is evaluated using these observations and the final joint state $\mathbf{s}_T$ of the participating environments (Trivedi et al., 2024; Yao et al., 2024; Wu et al., 2025). Each executable rule $c \in C_\tau$ returns a binary result; task completion requires

$$c(O_A, \mathbf{s}_T) = 1, \qquad \forall c \in C_\tau. \tag{3}$$

A.2 Phase 1: Single-Environment Research

Environment theme collection. Diverse environment themes provide concrete application scenarios and operational context for environment synthesis. We collect environment seeds from MCP services and publicly available agent skill documents. MCP services provide functionality descriptions and tool interfaces, while skill documents describe task-specific procedures and workflows. After filtering the seeds and organizing their capabilities, we obtain **801 MCP seeds and 682 skill seeds**, covering **14,657 tool capability descriptions**. These serve as the starting points for environment research and construction.[1]

Environment taxonomy construction. To support coherent cross-environment task composition, we organize environment seeds using a two-level taxonomy of **10 business domains and 84 capability categories**, drawing on the Agent-World taxonomy (Dong et al., 2026) and the business characteristics and capabilities of our seeds. Each seed is assigned one primary domain and capability category, while a capability may be associated with multiple domains.

Environment research and data construction. Given an environment seed, a research agent retrieves relevant documentation and available data using the environment description and tool functionality descriptions. A data construction agent then organizes the environment data from these findings. When information is missing, it sends a request to the research agent for further retrieval. Once sufficient information is available, this process produces the initial database, which represents the initial state $s_0^{(i)}$ of environment E_i.

A.3 Phase 2: Local Dependency Graphs

Entity extraction. Using the initial database and tool descriptions, we abstract the application objects queried or manipulated by tools into entity types and organize their descriptive information as attributes. All environments share an entity type

[1] MCP seeds are collected from Smithery. Skill sources include Smithery Skills, SkillMD, and the GitHub repositories `sickn33/agentic-awesome-skills` and `github/awesome-copilot`.

set $\mathcal{C}$, where each type is defined as

$$c = (n_c, \mathcal{A}_c) \in \mathcal{C}, \qquad \mathcal{A}_c = \bigcup_{i:\, c \in V_i} \mathcal{A}_c^{(i)}, \quad (4)$$

where n_c is the entity type name, $\mathcal{A}_c^{(i)}$ is its attribute set in environment E_i, and $\mathcal{A}_c$ is the union of these attributes across environments. The entity types extracted from E_i form the node set $V_i \subseteq \mathcal{C}$ of its local dependency graph. This representation connects corresponding entity types across environments without selecting or equating concrete database records; task-specific records are aligned in Phase 5.

Relation extraction. We then represent tool functionality as relations between entity types and associate each relation with the tools that implement it. In Figure 2, for example, `reserve_room` supports a booking relation from `Person` to `Hotel`, while `search_hotels` supports a search relation from `City` to `Hotel`. The figure labels these relations with their supporting tool names. This yields a local dependency graph

$$G_i = (V_i, R_i), \qquad r = (u, \rho, v, \mathcal{T}_r) \in R_i, \quad (5)$$

where V_i is the entity type set and R_i is the set of directed relations. Each relation r contains a source entity type $u \in V_i$, a relation label ρ describing its semantics, a target entity type $v \in V_i$, and a set of supporting tools $\mathcal{T}_r \subseteq \mathcal{T}_i$.

A.4 Phase 3: Cross-Environment Graph Composition

Connecting environments through shared entity types. Tools in different environments may provide different operations on the same entity type. We merge the local graphs through shared entity types while retaining the environment associated with each relation and its tools. The resulting cross-environment dependency graph is

$$\begin{aligned} G^{\times} &= (V^{\times}, R^{\times}), \\ V^{\times} &= \bigcup_i V_i, \\ R^{\times} &= \bigcup_i \{(i, r) \mid r \in R_i\}. \end{aligned} \quad (6)$$

Each shared entity type corresponds to a single node in the composed graph. Its attributes aggregate information from different environments, while its incident edges retain the operations supported by each environment. To incorporate a new environment, we extract its entities and relations and connect its local graph to the existing graph through shared entity types. We can then sample task combinations that include the new environment.

Task skeleton sampling. We select connected relations and their associated tools from the cross-environment graph to form task skeletons. Candidate combinations are assessed along three dimensions: **entity relationships**, which indicate whether operations in different environments can involve common objects; **complementary operations**, which indicate whether create, read, update, and delete (CRUD) operations can form a coherent workflow; and **business relevance**, which indicates whether the capabilities naturally belong in the same task. Filtering candidates along these dimensions yields a task skeleton

$$\begin{aligned} S &= (V_S, R_S), \\ V_S \subseteq V^{\times}, &\qquad R_S \subseteq R^{\times}, \end{aligned} \quad (7)$$

where (V_S, R_S) is a connected subgraph specifying the entity types, relations, and supporting tools. Its participating environments are $\mathcal{E}_S = \{E_i \mid (i, r) \in R_S\}$, and $\mathcal{T}_S$ denotes their selected supporting tools. The index i in each relation identifies its source environment. We then assign a task objective to the selected combination, specifying what the operations should accomplish together. This determines which operations must use earlier results and which cross-environment constraints the final outcomes must satisfy. The subgraph specifies the skeleton's structure, while the objective and constraints supply its task semantics. If the workflow requires additional operations, we record the corresponding tool requirements.

Structural difficulty control. We control task skeleton complexity along two dimensions: the number of environments and their semantic scope. Restricting the number of environments controls the scale of coordination required by a task. Restricting their business domains controls its semantic scope, from coordination within a single domain to coordination across domains.

User interaction. We additionally specify whether a task requires interaction with the user. In CrossEnv-Curriculum, L5 introduces user interaction while retaining L4's environment-count and business-domain-scope requirements, as described in Section 4.

A.5 Phase 4: Tool Synthesis and Testing

Tool synthesis. Given a task skeleton, we group the required tools by environment, reuse existing implementations that have passed testing, and use a coding agent to implement missing tools. The coding agent uses tool descriptions, research findings, and the initial database to generate executable code that reads or modifies data and returns results. This turns the tool capabilities associated with graph relations into executable operations. For environment E_i, the tested implementations define the execution function F_i introduced in Section A.1.

Tool testing. We invoke the generated tools and check whether their outputs and data changes match the expected behavior. For read operations, we check that the returned information is correct and the data remains unchanged. For create, update, and delete operations, we check that the intended changes occur and confirm them through subsequent queries. We also test invalid inputs and conflicting operations to assess error handling. Failures are returned to the coding agent for correction, followed by another round of testing. The executable environments are considered complete only when all tools required by the skeleton pass these tests.

A.6 Phase 5: Task Instantiation and Verification

Task instantiation. Given a task skeleton S, its objective, and its constraints, we select concrete entities and data for the task. Because the initial states $s_0^{(i)}$ are constructed independently, records for the same object may be missing or contain conflicting information across relevant environments. We supplement required records and resolve conflicts so that the selected objects and cross-environment constraints are consistent. Initial data and reference actions are planned jointly: records needed by read, update, and delete operations must either exist initially or be created by earlier actions; records that the task requires the agent to create are not inserted in advance. This produces a task-specific initial state $s_{0,\tau}^{(i)}$ for each environment and the joint initial state

$$\mathbf{s}_{0,\tau} = \left(s_{0,\tau}^{(i)}\right)_{E_i \in \mathcal{E}_S}. \tag{8}$$

Reference execution and task generation. Starting from $\mathbf{s}_{0,\tau}$, we execute the planned actions, including user interaction where required, and record their observations and data changes. Later actions can use results produced by earlier operations. If execution diverges from expectations, we revise the task data and action plan based on the feedback and execute the plan again. Once execution succeeds, we retain the reference action sequence A^* and generate the task instruction I from the scenario and execution results.

Verification rule generation. We derive executable verification rules C_τ from the task objective and constraints, using the tool outputs and data changes observed during reference execution to specify expected outcomes. These rules check whether the required records have been created, updated, or deleted; whether returned results satisfy the specified conditions; and whether related information is consistent across environments. We then construct negative cases, such as omitted operations and incorrect values, to check whether the rules reject incomplete or incorrect task outcomes.

Task verification. For an action sequence A, let $T(A) = |A|$ and define its binary verification outcome as

$$\mathrm{Pass}_\tau(A) = \prod_{c \in C_\tau} c(O_A, \mathbf{s}_{T(A)}), \tag{9}$$

where execution starts from $\mathbf{s}_{0,\tau}$ and $\mathbf{s}_{T(A)}$ contains the final states of $\mathcal{E}_S$. This is one exactly when every verification rule passes. We use two levels of verification: reference action replay and independent agent attempts. First, we reset all participating environments to $\mathbf{s}_{0,\tau}$, replay A^*, and require reproducible results with $\mathrm{Pass}_\tau(A^*) = 1$. We then run five independent agent attempts using the task instruction, available tools, and user interaction where required, with every attempt starting from the same initial state. Let $A^{(k)}$ denote the action sequence from attempt k. Following successful reference replay, we retain the task if $\sum_{k=1}^{5} \mathbf{1}[\mathrm{Pass}_\tau(A^{(k)})] \geq 2$. These two checks assess the reproducibility of the reference solution and whether an agent can find a valid solution from the task instruction and interfaces, respectively.

The retained task follows Equation 1, with aligned initial states $\mathbf{s}_{0,\tau}$, a successful reference sequence A^*, and executable verification rules C_τ.

B Environment Taxonomy

Table 6 lists 10 business domains (C_1) and 84 capability categories (C_2). Categories may span domains; each seed has one primary domain and category, assigned by its main purpose.

Business domain (C_1)	Capability categories (C_2)
Travel	Flight booking; Accommodation booking; Food and dining; Transit and navigation; Earth science and environmental monitoring; Maps and location services; Weather and forecasting; Itinerary planning; Currency exchange; Messaging and notification; Task and project management; Local search; Calendar and scheduling.
Work & Office	Decision support; Messaging and notification; Task and project management; Calendar and scheduling; Email management; Document processing; Presentation and design; Knowledge management; Web search; Customer support.
Life Services	Food and dining; Transit and navigation; Earth science and environmental monitoring; Maps and location services; Weather and forecasting; Messaging and notification; Calendar and scheduling; Shopping; Home services; Local news; Medical appointment booking; Billing and payments; Arts, culture, and museums; Events and entertainment; Gaming and interactive experiences; Real estate services.
Finance & Investment	Currency exchange; Billing and payments; Crypto market data; Market intelligence; Stock and equity investment; Blockchain infrastructure and wallets; DeFi and DEX protocols; Accounting; Risk management and compliance.
Health & Wellness	Medical appointment booking; Biomedical research; Medical information lookup; Fitness tracking; Nutrition and diet; Mental wellness.
Academic Research	Earth science and environmental monitoring; Knowledge management; Biomedical research; Literature search; Citation analysis; Data collection and processing; Arithmetic and general calculation; Scientific and symbolic computing; Statistical analysis; Academic writing and collaboration; Data visualization; Experiment design.
Data Analytics	Earth science and environmental monitoring; Trend analysis and analytics; NoSQL and document stores; Relational databases; Arithmetic and general calculation; Scientific and symbolic computing; Statistical analysis; Data visualization; Agent orchestration and multi-agent workflows; Model hubs and LLM APIs; Machine learning; Data pipelines; Knowledge graphs and vector retrieval; Web scraping; Report generation.
Content Creation	Trend analysis and analytics; Video and audio production; Image generation and editing; Social network integration; Video and content platforms; Writing and copywriting; Graphic design; Content scheduling; Agent orchestration and multi-agent workflows; Model hubs and LLM APIs.
Software Engineering	Code development; Package and dependency management; Cloud services and deployment; CI/CD and DevOps; Browser control and interaction; Software testing and quality assurance; NoSQL and document stores; Relational databases; API integration and testing; Shell and infrastructure; Frameworks and libraries; Technical documentation.
Cybersecurity	Compliance and audit; Log analysis; Security assessment; Threat intelligence; Network monitoring and response; Identity and access management; Security training.

Table 6: Complete environment taxonomy. The table covers all 10 business domains and all 84 unique capability categories. Categories shared across domains are repeated with identical names; the number of domain–capability associations therefore exceeds 84.

C Curriculum and Training Details

This appendix specifies task partitions, checkpoint initialization, and training objectives for Section 4. Appendix F documents curriculum evaluation, and Appendix G documents generalization tests.

C.1 Curriculum Organization

Task levels. Tasks are grouped into L1–L5 according to environment count, business-domain scope, and user interaction. L1 contains single-environment tasks, which use the representation in Equation 1 with $|\mathcal{E}_S| = 1$. Within a single business domain, L2 introduces tasks involving two environments, and L3 extends coordination to three or more environments. L4 requires coordination across multiple business domains and three or more environments. L5 retains L4's environment-count and domain-scope requirements and introduces user interaction. Let $\mathcal{D}$ denote the synthesized training tasks. The task set at each level is

$$\mathcal{D}_k = \{\tau \in \mathcal{D} \mid \text{level}(\tau) = L_k\}, \quad k = 1, \ldots, 5. \tag{10}$$

C.2 Single-Environment Supervised Fine-Tuning

We perform SFT on reference execution trajectories from $\mathcal{D}_1$ to teach tool selection, argument specification, and the use of tool outputs to guide subsequent actions (Qin et al., 2023; Yang et al., 2025b; Xu et al., 2025b). For task τ, let $A^* = (a_0^*, \ldots, a_{T_\tau - 1}^*)$

denote the reference action sequence, with length $T_\tau = |A^*|$. Here, I denotes the task instruction and h_t contains the preceding interactions. The policy π_θ predicts actions conditioned on I and h_t. The training objective is

$$\mathcal{L}_{\mathrm{SFT}}(\theta) = -\mathbb{E}_{\tau\sim\mathcal{D}_1} \sum_{t=0}^{T_\tau-1} \log \pi_\theta(a_t^* \mid I, h_t). \quad (11)$$

The resulting checkpoint θ_1 initializes subsequent cross-environment RL.

C.3 Stagewise Reinforcement Learning on Cross-Environment Tasks

Stagewise training. Starting from the L1 SFT checkpoint, we train sequentially on L2–L5. For $k = 2, \ldots, 5$, stage k uses $\mathcal{D}_k$ and is initialized with $\theta_k^{\mathrm{init}} = \theta_{k-1}$. We denote the checkpoint obtained after this stage by θ_k.

Execution feedback. For a sampled task τ, the agent starts from the joint initial state $\mathbf{s}_{0,\tau}$ and selects tool calls based on the instruction I and interaction history h_t. At L5, execution also includes interaction with a user. Let A denote the agent's action sequence, including tool calls and user-directed interactions where applicable. The verification rules C_τ evaluate its outputs and final states, producing the execution reward $r_\tau(A)$. The training objective at level k is

$$J_k(\theta) = \mathbb{E}_{\substack{\tau\sim\mathcal{D}_k \\ A\sim\pi_\theta(\cdot|\tau)}} [r_\tau(A)], \quad k = 2, \ldots, 5. \quad (12)$$

Here, $\pi_\theta(\cdot \mid \tau)$ denotes the rollout distribution induced by the policy interacting with task τ; the policy observes only the instruction and interaction history, not the underlying environment states or verification rules.

C.4 Training Configuration

The mainline uses Qwen3 backbones at 1.7B, 4B, and 8B. Table 2 gives the stagewise data counts. The settings below document L1 SFT and L2–L5 RL: single-environment supervision is followed by two-environment tasks within one domain, three-environment tasks within one domain, three-environment tasks across domains, and tasks requiring user interaction. At each transition, we load the selected policy weights from the preceding stage and reinitialize the optimizer and task sampler. The paper's L4 corresponds to L5 in the historical experiment directories; this naming is distinct from the paper's subsequent user-interaction stage.

Parameter	L1 SFT	L2–L5 RL
Learning rate	5×10^{-7} (peak)	5×10^{-7} (constant)
LR schedule	≈10% warmup, cosine decay	Constant
Global batch	8 assistant responses	16 task slots
Samples per task	Reference supervision	8 trajectories
PPO epochs	—	1
KL coefficient	—	0.001 in the loss
Temperature	—	Train: 1.0 Validation: 0

Table 7: Training configuration for L1 SFT and L2–L5 RL. Each RL batch contains 128 sampled trajectories.

L5 user interaction. L5 retains the L4 optimization settings and reward formulation, including the interaction-turn budgets. Qwen3.8-Flash generates user replies during interaction. The agent initially receives only the public task instruction, with user-specific facts withheld. It can request this information by calling `ask_user` with a `question` argument. The reply is appended to the context as a tool-role message, after which the agent continues execution. Each task allows at most three questions. These calls consume assistant turns, context, and wall-clock time, but do not count toward the environment's business-action budget.

Supervised loss. SFT supervises only the current assistant response. Preceding dialogue and tool feedback are masked from the loss, so the global batch size counts assistant-response samples rather than complete trajectories.

GRPO advantages. RL uses trajectory-level outcome GRPO. For the eight trajectories sampled for a task, let R_i be trajectory i's terminal reward, and let μ_R and σ_R be the within-group reward mean and standard deviation. The normalized advantage is

$$\hat{A}_i = \frac{R_i - \mu_R}{\sigma_R + 10^{-6}}. \quad (13)$$

This advantage is assigned to all assistant-generated tokens in that trajectory. The KL term is added to the optimization loss with coefficient 0.001.

C.5 Terminal Reward

For L2–L5, the terminal reward combines progress on task checks, constraint violations, invalid actions, and completion bonuses:

$$R = 0.85(P_T - P_0) - 0.20(V_T - V_0) - 0.01n_{\text{invalid}} + 0.10I_F + 0.05I_F I_D e. \quad (14)$$

Here, P is the weighted fraction of task checks passed, and V is the weighted fraction of invariant constraints violated; subscripts 0 and T denote the initial and terminal values. The count n_{invalid} records invalid actions. The indicator I_F requires all execution checks to pass and a valid final answer, while I_D indicates satisfaction of the tool-discovery protocol. The efficiency term is

$$e = \text{clip}\left(\frac{T_{\max} - T}{T_{\max} - T_{\text{ref}}}, 0, 1\right), \quad (15)$$

where T is the number of interaction turns. The budget parameters $(T_{\max}, T_{\text{ref}})$ are $(32, 18)$ for L2 and $(40, 30)$ for L3–L5. Thus, the reward gives partial credit for progress and is not a binary task-success signal.

D Evaluation Model Versions

Table 8 lists all models in Table 3, their sources, and available evaluation identifiers. The GPT-5.6-Luna and Gemini-3.8-Flash API IDs are recorded in the evaluation configurations and request code. IDs marked Claw-Eval are recorded for that benchmark; their use on other benchmarks remains to be verified. API IDs record requested model names; backend snapshots remain unverified. Hugging Face repositories identify public releases, but their evaluation revisions remain to be verified against the run configurations. Agent-World entries include the frozen checkpoint identifiers recorded for our in-house reproductions. For EnvScaler-4B and CrossEnv, we report the Qwen3 backbone; CrossEnv entries also specify the final curriculum stage. The Qwen3.8 link describes a related open-weight release and does not establish equivalence to the Max API.

E Evaluation Protocol Notes

E.1 Benchmark Subsets and Metrics

τ^2-bench. We evaluate Retail (114 tasks), Airline (50), Telecom (114), and the Banking Knowledge extension (97) (Shi et al., 2026), for a total of 375 tasks. These domains assess conversational task completion under domain-specific rules; Telecom additionally requires coordination with a user who can act in the environment. We report pass@1.

DeepPlanning-Shopping. We evaluate the L1, L2, and L3 difficulty subsets, containing 50, 50, and 20 tasks, respectively. These 120 tasks assess information gathering and planning under local and global constraints. We report the average case score. These benchmark difficulty labels are separate from the CrossEnv curriculum levels.

VitaBench. We evaluate 25 tasks from each of Delivery, In-store, OTA, and the cross-scenario subset, for a total of 100 tasks. They assess service execution, information integration, and adaptation to user requests, with the cross-scenario subset requiring coordination across services. We report pass@1.

Claw-Eval. We evaluate the adapted text subset of 157 General and 38 Multi-Turn tasks, excluding tasks involving multimodal content. General covers service orchestration, while Multi-Turn covers professional dialogue. We report pass@1; the task-score threshold and protocol adaptations are specified below.

The four benchmarks comprise 790 tasks per model: 375 from τ^2-bench, 120 from DeepPlanning, 100 from VitaBench, and 195 from Claw-Eval. Within each benchmark, Avg. is weighted by subset size. Capability groups describe the main evaluation emphasis rather than mutually exclusive abilities: VitaBench and Claw-Eval Multi-Turn also involve user interaction.

E.2 Evaluation Settings

Checkpoints and decoding. We evaluate CrossEnv models after L1–L5 training. Capability tracking extends through L5. The Mixed comparison and longer-workflow tests use checkpoints through L4. Agent-World-4B and Agent-World-8B use our reproduced Qwen3 RL step50 checkpoints. The recorded October baseline campaigns use temperature 0, thinking enabled, and no replay of historical reasoning; Toucan uses its native Qwen2.5 non-thinking template. GPT-5.6-Luna and Claude-Haiku-4.5 use temperature 1 on τ^2-bench, DeepPlanning, and VitaBench. Claw-Eval settings are specified separately below.

τ^2-bench. The four API reference models Qwen3.8-Max, DeepSeek-V4-Pro, GLM-5.3, and Kimi-K3 use a 128K context budget. Qwen3 Base, EnvScaler, Arctic-AWM, and Toucan use 64K. Local runs use the Hermes tool-call adapter with ex-

Model	Version / source
Frontier closed-source models	
GPT-5.6-Luna	OpenAI model documentation API ID: `gpt-5.6-luna`
Claude-Haiku-4.5	Anthropic model announcement Claw-Eval API ID: `bedrock/claude-haiku-4_5`
Gemini-3.8-Flash	Google DeepMind model card API ID: `gemini-3.8-flash`
Qwen3.8-Max	Claw-Eval API ID: `qwen3.8-max-0902` Qwen/Qwen3.8-2.4T-A95B (related model card)
Open-weight general-purpose models	
DeepSeek-V4-Pro	deepseek-ai/DeepSeek-V4-Pro-0813 Recorded ID: `deepseek-v4-pro-0813`
GLM-5.3	zai-org/GLM-5.3 Claw-Eval API ID: `glm-5.3`
Kimi-K3	moonshotai/Kimi-K3 Claw-Eval API ID: `kimi-k3`
Qwen3-1.7B	Qwen/Qwen3-1.7B
Qwen3-4B	Qwen/Qwen3-4B
Qwen3-8B	Qwen/Qwen3-8B
Environment scaling methods	
EnvScaler-1.7B	XXHStudyHard/EnvScaler-Qwen3-1.7B
EnvScaler-4B	Our reproduction; Qwen3-4B backbone
EnvScaler-8B	XXHStudyHard/EnvScaler-Qwen3-8B
Arctic-AWM-4B	Snowflake/Arctic-AWM-4B
Arctic-AWM-8B	Snowflake/Arctic-AWM-8B
Toucan-7B	Agent-Ark/Toucan-Qwen2.5-7B-Instruct-v0.1
EnvACE-1.7B	Team-ACE/EnvACE-Qwen3-1.7B
EnvACE-8B	Team-ACE/EnvACE-Qwen3-8B
Agent-World-4B	Our reproduction; Qwen3-4B RL step50 `AgentWorld-Qwen3-4B-dlcxnsnwrsoefb16-step50`
Agent-World-8B	Our reproduction; Qwen3-8B RL step50 `AgentWorld-Qwen3-8B-dlc13s4sujlwtirs-step50`
CrossEnv	
CrossEnv-1.7B	Trained in this work (Qwen3-1.7B); final L5 checkpoint
CrossEnv-4B	Trained in this work (Qwen3-4B); final L5 checkpoint
CrossEnv-8B	Trained in this work (Qwen3-8B); final L5 checkpoint

Table 8: Model versions and sources.

plicit model-format handling. Fixed-budget failures receive zero credit and remain in the task denominator.

VitaBench. October runs use a 4,096-token judge output budget. The EnvACE-1.7B and Arctic-AWM-4B runs use a 16,384-token judge budget and a YaRN original-position baseline of 40,960. Tool outputs and user interactions are scored using the benchmark's task rubrics.

Claw-Eval scoring. We perform one trial per task on the 157 General and 38 Multi-Turn tasks specified above. A native task score of at least 0.75 counts as success. All 195 tasks remain in the denominator, including model errors and fixed-budget failures. The native score is

$$\begin{aligned} q &= \text{safety}\,(0.8\,\text{completion} \\ &\qquad + 0.2\,\text{robustness}), \\ \text{task_score} &= \text{round}(q, 4). \end{aligned} \tag{16}$$

Gen. and MT report subset success rates, and the table's average is

$$\text{Avg.} = \frac{157 s_{\text{Gen.}} + 38 s_{\text{MT}}}{195}. \tag{17}$$

The mean native task score is a separate metric from this success rate.

Claw-Eval local models. Qwen3 and EnvScaler use temperature 0, thinking enabled, no replay of historical reasoning, a 131,072-token context, and an 8,192-token Agent output budget. Toucan-7B, EnvACE-1.7B/8B, and Arctic-AWM-4B/8B use

the same context and Agent output budgets; Toucan uses its native non-thinking template. Native per-task turn and time limits are retained. The judge and user simulator use Qwen3.8-Max with output budgets of 8,192 and 65,536 tokens, respectively. Toucan-7B, EnvACE-8B, and Arctic-AWM-8B use eight task workers; Arctic-AWM-4B uses two TP1 replicas with four workers each.

Claw-Eval API models. The recorded identifiers for Qwen3.8-Max, DeepSeek-V4-Pro, GLM-5.3, and Kimi-K3 are `qwen3.8-max-0902`, `deepseek-v4-pro-0813`, `glm-5.3`, and `kimi-k3`. These models use temperature 0, thinking enabled, and no replay of historical reasoning. Agent output limits follow upstream defaults, with a 131,072-token input-plus-output budget per response; judge and user output budgets are 8,192 and 65,536 tokens. Full evaluations use eight workers per model. GPT-5.6-Luna uses temperature 1 and Responses medium reasoning. Gemini-3.8-Flash uses temperature 0 with native signature handling. Claude-Haiku-4.5 uses the `bedrock/claude-haiku-4_5` Chat-compatible API at temperature 0 with thinking requested. These three models have no explicit Agent output cap and use the same 131,072-token input-plus-output budget, with at most eight task workers per model.

Claw-Eval tool interfaces. Tasks run in separate gVisor environments. IQS replaces Google via Novada search. OCR and long-document text are retained in files and exposed through previews and segmented reads; the Read preview limit is 16,000 characters. Adapter configurations include model-specific tool-call parsing, web-fetch transport handling, PDF rendering, and non-UTF-8 output handling. The revised web-fetch configuration uses a cancellable total deadline of at most 30 seconds and a 35-second outer RPC wait.

F Curriculum Learning Analysis

We define the metrics used in Section 5.3. Each metric uses a fixed evaluation set across checkpoints. For a set of evaluated requirements $\mathcal{U}$, let $b_m(u) \in \{0, 1\}$ indicate whether model m satisfies requirement u, and define

$$r_m(\mathcal{U}) = \frac{100}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} b_m(u). \tag{18}$$

F.1 Stagewise Capability Metrics

Tool execution coverage. This measures whether the agent executes the tools required by reference workflows, capturing basic tool-use coverage. For the evaluated reference tool requirements $\mathcal{U}_{\text{tool}}$,

$$S_{\text{tool}}(m) = r_m(\mathcal{U}_{\text{tool}}). \tag{19}$$

Coverage does not by itself establish correct arguments or task success.

Complete environment coverage. This measures whether the agent acts in every required environment, detecting tasks where it omits an entire service. For task τ, let $\mathcal{E}_\tau$ and $\widehat{\mathcal{E}}_\tau(m)$ be the required and actually used environment sets. On evaluation tasks $\mathcal{D}_{\text{env}}$,

$$S_{\text{env}}(m) = \frac{100}{|\mathcal{D}_{\text{env}}|} \sum_{\tau \in \mathcal{D}_{\text{env}}} \mathbf{1}[\mathcal{E}_\tau \subseteq \widehat{\mathcal{E}}_\tau(m)]. \tag{20}$$

Late-workflow tool coverage. This measures execution of tool requirements in the later portion of reference workflows, assessing whether the agent continues beyond early steps. For the evaluated late-workflow requirements $\mathcal{U}_{\text{late}}$,

$$S_{\text{late}}(m) = r_m(\mathcal{U}_{\text{late}}). \tag{21}$$

Target-state check pass rate. This measures whether execution produces the required business states, complementing action coverage with outcome correctness. For target-state checks $\mathcal{U}_{\text{state}}$ evaluable across the compared models,

$$S_{\text{state}}(m) = r_m(\mathcal{U}_{\text{state}}). \tag{22}$$

Unavailable checks are excluded consistently across models, rather than counted as failures.

F.2 User Interaction Metric

This metric assesses three aspects of user interaction: obtaining required information through a relevant question, using the reply in appropriate tool arguments with normal tool execution, and passing native strict task checks. Their combination captures intermediate progress as well as task completion. Let n_q, n_u, and n_s count tasks satisfying these criteria on the same $N = 40$ held-out interaction tasks:

$$S_{\text{int}} = 100 \frac{n_q + n_u + n_s}{N}. \tag{23}$$

The score ranges from 0 to 300. The criteria may hold on different tasks, so it is neither a success percentage nor the rate of completing all three steps together.

F.3 Curriculum and Mixed RL Metrics

Both schedules initialize from the same Qwen3-8B L1 SFT checkpoint. The compared curriculum and Mixed checkpoints use 243.04 and 238.85 cumulative RL GPU-hours, respectively. Mixed training samples from the L2, L3, and L4 pools of 579, 620, and 600 tasks in a 1:1:1 ratio. Each batch contains eight tasks from each level, for 24 tasks in total, with eight rollouts per task (192 rollouts per batch). Tasks are shuffled independently within each level, and the combined batch is shuffled again. The runs differ in parallelism, batch configuration, and reference-model settings; this comparison therefore assesses the two training schedules at similar compute budgets under their respective configurations, rather than isolating curriculum order.

Strict success measures complete task attainment; cumulative rollouts measure sampling volume at a given compute budget. Let $y_m(\tau)$ indicate that all task requirements pass and n_j be the rollout count for completed update j. We report

$$\begin{aligned} \mathrm{Success}(m) &= \frac{100}{|\mathcal{D}|}\sum_{\tau\in\mathcal{D}} y_m(\tau), \\ \mathrm{Rollouts}(h) &= \sum_{j:\, h_j \le h} n_j, \end{aligned} \tag{24}$$

where h_j is cumulative RL GPU-hours at update j. The comparison uses the same 180 validation tasks at similar GPU-hours. Rollout volume does not measure trajectory validity or task success; the comparison reflects the current training configurations.

G Generalization Analysis

We define the tests and metrics used in Section 5.4. Combination tests evaluate transfer to new environment pairs, and long-horizon tests evaluate transfer to longer workflows; task-similarity analysis compares CrossEnv and Base at 4B and 8B.

G.1 Combination Generalization

This test asks whether an agent can coordinate familiar environments in new combinations. Seen pairs co-occur in training; unseen pairs never co-occur, including within larger training combinations. Each group contains 150 tasks from 30 environment pairs. Strict success uses all tasks, retaining model failures in the denominator. Let $y_m(\tau)$ indicate that model m passes all task requirements. For group g,

$$S_g(m) = \frac{100}{|\mathcal{D}_g|}\sum_{\tau\in\mathcal{D}_g} y_m(\tau). \tag{25}$$

Table 4 reports $S_g(\mathrm{CrossEnv}) - S_g(\mathrm{Base})$ in percentage points.

Observation, State, and Answer scores localize gains to business observations, final states, and final answers. For category c, let $\mathcal{D}_{g,c}$ contain applicable tasks with native scores available for both models, and $\mathcal{C}_{\tau,c}$ contain the task's category-specific hard checks. With binary check outcomes $b_m(\tau,k)$,

$$S_{g,c}(m) = \frac{100}{|\mathcal{D}_{g,c}|}\sum_{\tau\in\mathcal{D}_{g,c}}\prod_{k\in\mathcal{C}_{\tau,c}} b_m(\tau,k). \tag{26}$$

Category gains are also CrossEnv minus Base. They require all checks in a category to pass and describe the jointly scored subset, not the full test population.

G.2 Long-Horizon Generalization

This test measures transfer to workflows longer than those used in training. Two sets of 40 tasks have reference lengths of 16–18 and 24–28 tool calls, exceeding the training maximum of 14. To capture partial progress, we average each task's fraction of passed observation and target-state checks:

$$\mathrm{Score}(m) = \frac{100}{|\mathcal{D}|}\sum_{\tau\in\mathcal{D}}\frac{\sum_{k\in\mathcal{C}_\tau} b_m(\tau,k)}{|\mathcal{C}_\tau|}, \tag{27}$$

where $\mathcal{C}_\tau$ contains those checks. Tasks receive equal weight; this is not strict task success. Reference length describes the supplied solution, not a minimum for all valid solutions.

G.3 External-Task Similarity

Following the similarity-analysis motivation of EnvScaler (Song et al., 2026), we test whether gains are concentrated on instructions close to training data. We encode instructions with `sentence-transformers/all-MiniLM-L6-v2`. For test instruction x, training-instruction set $\mathcal{X}_{\mathrm{train}}$, and embedding function f, similarity is

$$s(x) = \max_{z\in\mathcal{X}_{\mathrm{train}}}\frac{f(x)^\top f(z)}{\|f(x)\|\,\|f(z)\|}. \tag{28}$$

For normalized native task scores $q_m(x)$, per-task and similarity-group gains are

$$\begin{aligned} \Delta(x) &= 100\big(q_{\mathrm{CrossEnv}}(x) - q_{\mathrm{Base}}(x)\big), \\ \overline{\Delta}_b &= \frac{1}{|\mathcal{B}_b|}\sum_{x\in\mathcal{B}_b}\Delta(x), \end{aligned} \tag{29}$$

where $\mathcal{B}_b$ is one of four similarity-quantile groups. Figure 6 shows group means with 95% confidence intervals. The VitaBench-8B panel uses 71 paired tasks, native rubric pass fractions, and paired-task bootstrap intervals; this score differs from strict task success. Results describe the plotted subsets, and instruction similarity alone does not establish unseen environments or domains.

H Cross-Environment Training Analysis

This appendix defines the execution diagnostics in Table 5. SingleEnv and CrossEnv RL start from a shared SFT checkpoint and differ in whether their training tasks involve single or multiple environments. Base execution scores are historical references.

H.1 Execution Diagnostics and Measurement Rules

The three diagnostics describe data dependencies, continuation of execution, and final outcomes. They are related views: one downstream failure can affect multiple metrics, so they do not identify three independent mechanisms.

Use of prior tool outputs. Annotate required cross-service dependencies from task instructions and verification rules, then check whether actual objects, identifiers, amounts, or states returned by earlier services are correctly bound to later operation parameters. Both models must use the same applicable tasks and dependency edges. Failing to obtain required information or omitting its downstream operation counts as unsatisfied. Any conditional rate given successful upstream retrieval must identify that restricted denominator.

Later-step completion. Pre-annotate the business operations required after earlier steps, including retries, notifications, and actions after clearing a blockage. An operation counts as complete only if it executes successfully and produces the required state change. A tool call, mention of its name, or verbal completion claim alone is insufficient; omitted necessary operations remain unsatisfied.

All task constraints met. On tasks with joint goals, require all necessary service outcomes and constraints to hold simultaneously, including correct orders and notifications, specified cleanup objects, and restrictions on extra writes. Score each task by the conjunction of its requirements rather than the mean fraction of atomic checks passed.

Applicable tasks, dependencies, actions, and constraints use the same annotation rules and denominators across models. Infrastructure anomalies and missing scores are recorded separately; execution failures and omitted required operations remain in the denominator.

I Synthesis and Verification Quality

This appendix supplements Section 5.6 with the current task-quality evaluation and construction-cost accounting.

I.1 Execution and Verification Checks

We sample 100 retained L5 tasks, with 50 each from MCP and Skill environments and stratification by domain. All reference workflows execute successfully and receive full scores. On the tested subsets, replay from a reset environment reproduces these outcomes, omission of a required operation prevents a full score, and semantically equivalent query reorderings preserve full scores. These checks concern reference execution, required operations, and equivalent query orders.

I.2 Construction-Cost Accounting

Figure 7 compares candidate output yield as the environment pool grows and average environment-construction cost per output. CrossEnv composes new environments with existing ones and reuses their data, tools, and entity relations. The reported CrossEnv pool of 100 environments yields 51,566 candidate tasks, with a recorded environment-construction cost of 206,496,276 tokens. The average cost is therefore

$$\frac{206{,}496{,}276}{51{,}566} \approx 4{,}004.5 \tag{30}$$

tokens per candidate, consistent with the approximately 4.0k value in Section 5.6.

CrossEnv and AgentWorld outputs are screened candidates; EnvScaler outputs are released scenarios. AgentWorld values come from our reproduction, whereas EnvScaler values are estimates derived from its reported figures. Costs cover environment construction rather than end-to-end task generation, and exclude subsequent task instantiation and acceptance. The denominator is candidate output, not the final set of training examples. Dashed extensions in the figure indicate projections under unchanged mean yield and construction cost per additional environment.

J Environment and Task Synthesis Comparison Criteria

Table 1 compares explicitly reported synthesis capabilities rather than empirical guarantees or only the default training recipe. A supported optional branch counts as supported; an unsupported mark means that the described pipeline does not explicitly provide the capability.

Executable (Exec.). Actions run through real services or executable code. Execution during data collection counts, without implying that an online RL training environment is released.

Verifiable (Verif.). Executable checks evaluate task outcomes, rather than only syntax or tool functionality. Pure-code verification branches count even when another branch uses an LLM judge.

Controllable (Ctrl.). Explicit settings adjust task complexity through tool count, graph length, branching, constraints, environment count, or discrete complexity levels. Limited support denotes qualitative workflow expansion or adaptive difficulty targeting; a fixed request for diverse or challenging tasks does not constitute an adjustable control.

Cross-environment (Cross). A synthesized task coordinates multiple environments, applications, or services. AWM's limited support covers bundled service suites, without demonstrated general composition across independently constructed environments.

Consistent (Consist.). The pipeline explicitly aligns entities or data across environments, or enforces cross-environment constraints. Local database consistency alone does not establish this capability.

Extensible (Ext.). New environments or services can be integrated with reusable existing ones to generate joint tasks. Interface-based integration counts; integration need not require zero adaptation.