

# PAYING LESS GENERALIZATION TAX: A CROSS-DOMAIN GENERALIZATION STUDY OF RL TRAINING FOR LLM AGENTS

**Zhihan Liu** \*†  
Northwestern University  
zhihanliu2027@u.northwestern.edu

**Lin Guan**‡ & **Yixin Nie**  
Meta Superintelligence Labs  
linguan@meta.com

**Kai Zhang**  
The Ohio State University

**Zhuoqun Hao**  
University of Pennsylvania

**Lin Chen**  
Meta Superintelligence Labs

**Asli Celikyilmaz**  
FAIR at Meta  
linguan@meta.com

**Zhaoran Wang**  
Northwestern University  
zhaoranwang@gmail.com

**Na Zhang**  
Meta Superintelligence Labs  
nazhang@meta.com

## ABSTRACT

Generalist LLM agents are often post-trained on a narrow set of environments but deployed across far broader, unseen domains. In this work, we investigate the challenge of agentic post-training when the eventual test domains are unknown. Specifically, we analyze which properties of reinforcement learning (RL) environments and modeling choices have the greatest influence on out-of-domain performance. First, we identify two environment axes that strongly correlate with cross-domain generalization: (i) state information richness, i.e., the amount of information for the agent to process from the state, and (ii) planning complexity, estimated via goal reachability and trajectory length under a base policy. Notably, domain realism and text-level similarity are not the primary factors; for instance, the simple grid-world domain Sokoban leads to even stronger generalization in SciWorld than the more realistic ALFWorld. Motivated by these findings, we further show that increasing state information richness alone can already effectively improve cross-domain robustness. We propose a randomization technique, which is low-overhead and broadly applicable: add small amounts of distractive goal-irrelevant features to the state to make it richer without altering the task. Beyond environment-side properties, we also examine several modeling choices: (a) SFT warmup helps prevent catastrophic forgetting during RL but undermines generalization to domains that are not included in the mid-training datamix; and (b) turning on step-by-step thinking during RL, while not always improving in-domain performance, plays a crucial role in preserving generalization.

## 1 INTRODUCTION

Post-training for *generalist* large language model (LLM) agents is typically carried out on a narrow slice of tasks, while deployment spans far broader and shifting environments. Even with well-maintained data flywheels, the full cycle of sampling, post-processing, and running large-scale train-

---

\*Work done during the internship at Meta Superintelligence Labs.

†Equal contribution.

‡Equal contribution.

ing remains compute and time intensive. As a result, post-training can only incorporate a limited set of domains and tasks (Luo et al., 2024). Empirically, this skew shows up when agents that look strong on popular leaderboards only deliver subpar user experience: recent studies find sizable gaps between benchmarked strengths and performance in real-world “office” tasks (Xu et al., 2024; Miller & Tang, 2025; Jimenez et al., 2023; Jain et al., 2024). Therefore, when selecting or constructing post-training datasets, a key question is what types of task the agent should be trained on to ensure that the model not only achieves strong in-domain performance but also maintains generalization in unseen and broader scenarios.

On the other hand, online RL has become a principled tool for post-training LLMs, but it introduces its own constraints. Typical online RL training demands interactive environments that have reward signals and support the sampling of multiple rollouts. These simulation environments are costly to build and even infeasible for proprietary domains (Cao et al., 2024; Wang et al., 2024a; Samsami et al., 2024). The state of public agentic environments reflects this difficulty: despite notable progress, most benchmarks remain narrow (e.g., specific websites, platforms, or workflows) and expose only slices of the broader task distribution (Deng et al., 2023; Zhou et al., 2024; Drouin et al., 2024; Boisvert et al., 2024; Le Sellier De Chezelles et al., 2024; Wang et al., 2024b). Furthermore, even for production models deployed in real traffic, the resulting flywheel data is predominantly single-rollout and off-policy with respect to the updated policy in RL. Therefore, we aim to understand what characteristics make a training environment effective in preserving broader out-of-domain performance, so that we can prioritize constructing such environments and tasks for LLM agent post-training (Guan et al., 2023; Hu et al., 2025).

Overall, given the limited coverage in post-training data and the practical difficulty of constructing simulators at scale for online RL, we offer a data-centric perspective to study:

*How can we strategically sample tasks or selectively build domain simulators for RL training that maximally preserves or transfers agentic capabilities of LLMs to unseen domains?*

Before moving forward, we clarify that “generalization” has been assessed at different levels in prior work (Fu et al., 2025; Huan et al., 2025; Zhang et al., 2025b). The first is in-domain: models are tested on held-out cases where the task space is similar but not identical to that seen in training (e.g., shopping on Amazon vs. eBay). The second is cross-domain generalization: transferring to tasks with substantially different tools, states, or goals (e.g., from shopping to mathematical problem solving). Much of the reported evidence for “RL generalization” concerns the first type, while our interest here lies in the second. In our preliminary experiments on multiple agentic domains, as RL training proceeds, performance on held-out (unseen) environments drops while in-domain scores improve. Therefore, this work is more prevention-focused: identify training- and generation-time choices that maximally preserve performance in environments not observed during post-training.

By performing RL training on four distinct environments, namely Webshop (Yao et al., 2022a), Sokoban (Junghanns & Schaeffer, 2001), ALFWorld (Shridhar et al., 2021), and SciWorld (Wang et al., 2022), we find that two axes exhibit high correlation with cross-domain performance: (a) **State information richness**, i.e., the volume of information that the agent must perceive and process from the state. Higher richness increases the *perception load*, forcing the policy to actively extract task-relevant signals from dense or noisy inputs. (b) **Planning complexity**, i.e., the difficulty of composing long and contingent action sequences, which imposes a higher *reasoning load*. We estimate this via goal reachability under a base policy. Greater complexity exercises decomposition, progression, and potentially correction over a long horizon.

Given that state information richness and planning complexity are key factors for generalization, one natural follow-up question is whether we can make training environments more effective by manipulating either of these properties. In this work, we begin with state information richness as it is not intrinsic to a domain. Specifically, we apply state randomization by injecting small amounts of goal-irrelevant or domain-irrelevant content into the agent’s observations during training. This increases the information richness of states without altering other components in task space, and consistently improves cross-domain generalization. These results not only validate our hypothesis but also provide a practical method for augmenting any existing training environments.

Alongside the main experiments, we also examine two modeling choices relevant to generalization. First, we find that mid-training has a strong and lasting influence on generalization. Specifically, it deeply infuses knowledge into the model, such that performance on covered domains remains stable

even after subsequent RL training on other domains. However, the performance drops significantly on domains not included in mid-training, revealing a clear trade-off: while mid-training consolidates knowledge and improves robustness to downstream training, it risks erasing prior knowledge in uncovered domains. Moreover, we find that enabling step-by-step thinking during both training and evaluation does not improve in-domain scores, but plays a key role in preserving generalization. Taken together, all these findings offer concrete guidance for post-training when the target environment is unknown: select or construct training environments that have greater state information richness and planning complexity, enable explicit reasoning, and apply lightweight state randomization to better preserve generalization.

## 2 RELATED WORK

### 2.1 CROSS-DOMAIN ANALYSIS FOR LLM TRAINING

As LLMs are applied to a wide range of tasks, recent studies have examined how well RL fine-tuning generalizes to different domains. Cheng et al. (2025) empirically measure the transferability of skills acquired in a training domain to other unseen domains. Besides, Huan et al. (2025) compare the cross-domain generalization between SFT and RL on single-turn tasks and find that RL preserves more generalization than SFT. These findings highlight the importance of understanding how an RL-tuned LLM performs on tasks beyond its training domains. Our work aligns with these perspectives, but with a particular emphasis on multi-turn agentic tasks, which are not well covered in previous studies. Moreover, previous analyses largely focus on quantifying the correlation between training domains and performance on held-out domains. Our work moves beyond this by dissecting these correlations and explaining them through relevant environmental factors and modeling choices, which further informs the derivation of practical solutions for better generalization.

### 2.2 STATE AUGMENTATION FOR GENERALIZATION AND ROBUSTNESS

Domain and state augmentation have long been explored as strategies to improve generalization, particularly in sim-to-real settings where agents trained in simulation must transfer to real-world deployments. These methods typically inject variations into the training environment by modifying background textures, object properties, lighting, or physics parameters (Tobin et al., 2017; Sadeghi & Levine, 2017; Muratore et al., 2018; Loquercio et al., 2019). More recent work has applied similar ideas in visual RL by augmenting observations or introducing distractors to improve robustness across domains or simulator variants (Pinto et al., 2017; Packer et al., 2018; Hansen et al., 2021; Guan et al., 2021; Yue et al., 2019; Chen et al., 2021; Wang et al., 2025a). It is important to note that our contribution lies beyond the technique of state randomization itself; we utilize it essentially as a tool to validate our systematic findings. Nevertheless, adapting it to LLM agent training is also meaningful: our work is the first to demonstrate semantic textual noise and controlled distraction degree are helpful for preventing OOD performance degradation in agentic RL.

Due to the page limit, we postpone the discussion of related work about multi-turn RL training for LLM agents in Appendix B.

## 3 PRELIMINARY

### 3.1 RL TRAINING FOR LLM AGENTS

**Environment and Agent.** We consider a reinforcement learning framework in which a language model functions as an agent interacting with an environment to solve multi-step tasks described by a textual prompt  $x \sim p(X)$ . At each discrete timestep  $t \in \{1, \dots, T\}$ , the agent perceives a state  $s_t \in \mathcal{S}$  and produces a textual action  $a_t \in \mathcal{V}^n$ , where  $\mathcal{V}$  represents the token set and  $n$  is the maximum number of tokens generated. The environment then emits a scalar feedback  $r_t \in \mathbb{R}$  and transitions to a subsequent state  $s_{t+1}$ . An entire episode can be represented by the trajectory  $\tau = \{(s_1, a_1, r_1), (s_2, a_2, r_2), \dots, (s_T, a_T, r_T)\}$ . The agent’s behavior is dictated by a conditional policy  $\pi_\theta(a_t | s_t, x)$  parameterized by  $\theta$ , which defines the distribution over textual outputs given the current state and task prompt. The goal of the LLM agent is to find a policy  $\pi$  that maximizes the corresponding expected cumulative reward  $\mathbb{E}[R(\tau) | \tau \sim \pi]$ , where  $R(\tau) = \sum_{t=1}^T r_t$  denotes

the sum of rewards along the trajectory, and the expectation is taken with respect to the stochastic policy  $\pi$ . In many practical agentic cases, the environment provides only binary success indicators at the end of an episode. As a result, the cumulative reward  $R(\tau)$  is also sparse, which enables direct adoption of single-turn optimization approaches like GRPO (Shao et al., 2024).

### 3.2 CROSS-DOMAIN GENERALIZATION OF RL TRAINING

In this subsection, we provide a definition of cross-domain generalization and clarify its connection to other types of generalization. Let  $e_i$  denote a text-based agentic environment, and let  $\mathcal{T}_i$  be the set of tasks associated with  $e_i$ . The most basic form of generalization evaluation tests a model’s performance on unseen tasks *within the same environment*. Formally, for each environment  $e_i$ , we partition its task set into two non-overlapping subsets,  $\mathcal{T}_i = \mathcal{T}_i^{\text{train}} \cup \mathcal{T}_i^{\text{val}}$ ,  $\mathcal{T}_i^{\text{train}} \cap \mathcal{T}_i^{\text{val}} = \emptyset$ . Let  $\mathcal{I}_{\text{env}}$  be the index set of all environments. We then define the training and evaluation task collections  $\mathcal{E}_{\text{ID}}^{\text{train}} = \bigcup_{i \in \mathcal{I}_{\text{env}}} \mathcal{T}_i^{\text{train}}$ ,  $\mathcal{E}_{\text{ID}}^{\text{eval}} = \bigcup_{i \in \mathcal{I}_{\text{env}}} \mathcal{T}_i^{\text{val}}$ . We refer to this setup as *in-domain (ID) validation*.

Beyond this, one can evaluate a generalist agent on *unseen environments* that may differ substantially in states, actions, and dynamics. Let  $\mathcal{I}_{\text{env}}^{\text{train}}$  and  $\mathcal{I}_{\text{env}}^{\text{eval}}$  denote the index sets of environments used for training and evaluation, respectively, with  $\mathcal{I}_{\text{env}}^{\text{train}} \cap \mathcal{I}_{\text{env}}^{\text{eval}} = \emptyset$ . We define the corresponding sets of environments  $\mathcal{E}^{\text{train}} = \{e_i \mid i \in \mathcal{I}_{\text{env}}^{\text{train}}\}$ ,  $\mathcal{E}^{\text{eval}} = \{e_i \mid i \in \mathcal{I}_{\text{env}}^{\text{eval}}\}$ . We refer to this setup as the evaluation of *cross-domain generalization* or *out-of-domain (OOD) generalization*. Due to the limited availability of public agentic domains, for cross-domain evaluation, we use one environment for training at a time and use the remaining for evaluation.

## 4 ANALYTICAL FRAMEWORK

### 4.1 DOMAINS AND METRICS

We conduct our analysis on four representative agentic environments: (a) **WebShop**, a text-based e-commerce simulator requiring the agent to navigate pages, filter attributes, and select products matching complex user queries (Yao et al., 2022a). (b) **Sokoban**, a classic grid-based puzzle where the agent needs to push boxes to targets while avoiding irreversible lock-states (Junghanns & Schaeffer, 2001). Unlike the matrix representation used in prior work (Wang et al., 2025c), we use explicit coordinate representations of walls, players, objects, and the target as the state representation. We find this necessary because RL training with `Llama-8B-instruct` fails to yield meaningful performance gains when using the original matrix representation (Wang et al., 2025c). (c) **ALFWorld**, a widely adopted benchmark that requires the agent to generate symbolic actions (e.g., “open fridge”, “take apple”) to interact with a household environment (Shridhar et al., 2021). (d) **SciWorld**, an interactive simulator of a scientific laboratory (Wang et al., 2022). We include this as a significantly more challenging extension of ALFWorld. SciWorld features a larger action space and the requirements for agents to manipulate specialized tools and materials to execute complex, multi-stage experiments. Examples that compare between ALFWorld and SciWorld are provided in Appendix E.

These domains cover distinct types of common agentic tasks, ranging from web navigation and information retrieval to spatial reasoning and embodied interaction. While these environments vary across multiple dimensions—such as domain realism, action space size, and task horizon—our study identifies *state information richness* and *planning complexity* as the two primary factors influencing generalization. We will provide concrete definitions and measurements for these two axes in Section 5 to quantitatively analyze their correlation with out-of-domain performance.

**Evaluation Protocol.** We perform RL training on one source domain and evaluate the performance on the remaining three to assess OOD generalization. Training runs for 150 steps, and we save checkpoints every 15 steps. We report the mean success rate (pass@1 per task) averaged over the last four checkpoints (i.e., checkpoints from steps 105 to 150). All experiments are repeated three times with different seeds.

## 4.2 MODELS

We adopt `Llama-3.1-8B-Instruct` (Grattafiori et al., 2024) as the base model. To establish a viable starting point for analyzing out-of-domain performance, the initial policy must exhibit non-trivial performance on all evaluation domains such that we can observe how the success rate changes over the course of subsequent RL on other training domains. However, the base model has a near-zero success rate in two domains, i.e., WebShop and SciWorld.

For WebShop, conducting a 20-step RL training on it can increase the in-domain success rate to 34.4%. The resulting policy, Checkpoint V1 (`Ckpt V1`), serves as an initial policy for all downstream experiments to be presented. However, for SciWorld, both the base model and `Ckpt V1` remain at a near-zero success rate even after extensive RL training. As a result, we resort to warming up the base model using SFT with expert demonstrations. This setup is similar to mid-training (Wang et al., 2025b; Mo et al., 2025), which imparts sufficient prior knowledge for effective subsequent RL training. To ensure balanced initial performance on all domains, we also mixed self-generated trajectories from ALFWorld and WebShop into the SFT dataset. The details of the SFT dataset and training can be found in Appendix D. We refer to this mid-training checkpoint as Checkpoint V2 (`Ckpt V2`). In summary, `Ckpt V1` and `Ckpt V2` serve as the two initial policies for our downstream experiments, with `Ckpt V2` additionally enabling us to investigate the influence of SFT warm-up or mid-training.

## 5 CHARACTERIZING EFFECTIVE TRAINING DOMAINS FOR PRESERVING GENERALIZATION

Table 1: OOD success rates after RL training. Shaded cells correspond to ID evaluation.  $\Delta_{ID}$  is the change in ID success rate relative to the corresponding base model. Note that SciWorld evaluations for `Ckpt V1` are omitted as the success rates remain near zero as explained in Section 4.2.

| Initial Policy | Training Domain | Evaluation Domain  |                    |                    |                    |
|----------------|-----------------|--------------------|--------------------|--------------------|--------------------|
|                |                 | ALFWorld           | Webshop            | Sokoban            | SciWorld           |
| <b>Ckpt V1</b> | Base Model      | 25.8               | 34.4               | 12.5               | 3.1                |
|                | ALFWorld        | $\Delta_{ID}=47.2$ | 15.2               | 14.0               | -                  |
|                | Webshop         | 10.8               | $\Delta_{ID}=23.6$ | 12.8               | -                  |
|                | Sokoban         | 20.5               | 15.0               | $\Delta_{ID}=36.0$ | -                  |
| <b>Ckpt V2</b> | Base Model      | 21.9               | 38.3               | 14.1               | 18.0               |
|                | ALFWorld        | $\Delta_{ID}=60.6$ | 30.5               | 9.8                | 10.0               |
|                | Webshop         | 17.0               | $\Delta_{ID}=28.7$ | 9.0                | 13.8               |
|                | Sokoban         | 20.0               | 34.0               | $\Delta_{ID}=35.7$ | 13.0               |
|                | SciWorld        | 19.8               | 35.8               | 12.0               | $\Delta_{ID}=20.2$ |

This section presents the initial experiment results based on the setup outlined in Section 4. We aim to identify the properties of training environments that best align with the preservation of agent capabilities in unseen domains. We first report the performance of agents trained on each domain and evaluated on all others in Table 1. Consistent with prior studies (Luo et al., 2024), while RL training significantly improves in-domain performance, it leads to performance drops in OOD tasks.

For a concise and clear comparison of how different training domains affect robustness, we introduce a ranking-based metric. For each training domain, we calculate its performance rank on every OOD domain (e.g., 1 represents the best performance compared against checkpoints trained on other training domains) and compute the sum of these ranks as OOD Ranking Score.

The aggregated results are shown in Table 1, with detailed ranks provided in Table 6 in Appendix. Recall from Section 4.2 that `Ckpt V1` and `Ckpt V2` are applicable to different training domains due to SciWorld’s requirement for SFT warmup. Consequently, we report two sets of aggregated results: one based solely on `Ckpt V2`, and the other combining `Ckpt V1` and `Ckpt V2`. The former compares all training domains (including SciWorld), while the latter compares the other three domains across a broader set of experimental trials. With `Ckpt V2`, SciWorld yields the most robust training outcome (ranking score: 3), closely followed by Sokoban (ranking score: 5). WebShop ranks third (ranking score: 6), while ALFWorld performs the worst (ranking score: 8). A similar pattern is observed in the combined results (`Ckpt V1` and `Ckpt V2`), where Sokoban is more effective (ranking score: 7) than both WebShop and ALFWorld.

Table 2: Correlation between properties of training environments and OOD robustness. **Avg. Character Count** denotes the average character count of states, which is a proxy for state information richness. **Avg. Trajectory Length** denotes the average trajectory length, which indicates planning horizon and complexity. **OOD Ranking Score** ( $\downarrow$ ) represents cross-domain robustness, which is defined in Section 5. A lower OOD Ranking Score corresponds to a better overall rank, indicating stronger generalization.

| Training Domain | Avg. Character Count | Avg. Trajectory Length | OOD Ranking Score ( $\downarrow$ ) |         |
|-----------------|----------------------|------------------------|------------------------------------|---------|
|                 |                      |                        | Ckpt V1 and V2                     | Ckpt V2 |
| SciWorld        | Medium (2851)        | High (43.5)            | –                                  | 3 (1st) |
| Sokoban         | High (3114)          | High (44.0)            | 7 (1st)                            | 5 (2nd) |
| WebShop         | High (3063)          | Low (33.9)             | 10 (2nd)                           | 6 (3rd) |
| ALFWorld        | Low (1572)           | Medium (42.0)          | 10 (2nd)                           | 8 (4th) |

Based on these rankings, our goal is to identify factors that correlate with stronger OOD performance. Given the complexity and the exploratory nature of such an analysis, we do not aim to exhaustively identify all factors, but rather to reveal at least some key insights that can inform the selection or construction of training domains. Overall, we find two planning-related concepts below better explain the effectiveness of environments. For a quantitative analysis, we ground these two concepts into two quantifiable metrics, which are formally defined and measured as follows: (a) **State Information Richness**, i.e., the volume of information contained in the environment states. A higher state information richness means that the agent must process a larger amount of content from the state to extract useful information for the current decision-making. We note that the observations in all four studied domains are highly compact, which convey each piece of state information in a similarly concise manner (see Appendix E for examples). In this case, a feasible way of quantifying information richness would be to use the *average character count of environment states*. Practically, we employ Ckpt V2 to collect 128 trajectories per domain with a maximum step limit  $T_{\max} = 50$  and calculate the average length of the states within the sampled trajectories. (b) **Planning Complexity**, i.e., the difficulty of composing a multi-step plan to solve the task. A higher planning complexity indicates that the task requires a longer, more sustained planning chain and has a lower goal reachability. Although planning complexity is often measured by properties such as branching factors or search depth in the classical planning literature (de Roo, 2016), these metrics are less applicable to LLM agents that leverage strong priors rather than relying on conventional search space traversal. Instead, we use *average trajectory length* as an empirical proxy of search depth to measure the *reasoning effort* in order to reach the terminal states. Specifically, we also employ Ckpt V2 to collect 128 trajectories per domain with a maximum step limit  $T_{\max} = 50$  and compute this metric by averaging the trajectory lengths, where failed trajectories are assigned the maximum length  $T_{\max}$  to reflect the inability to solve the task within the budget. The average trajectory length in this case captures both task horizon and goal reachability.

Table 2 shows how the state information richness and the planning complexity are correlated to the robustness rankings. Specifically, Sokoban exhibits high state richness (3114 characters) and high planning complexity (44.0 steps), while SciWorld similarly combines substantial planning complexity (43.5 steps) with moderate-to-high information richness. This indicates that environments that demand both high perception loads and deep reasoning chains foster better generalization preservation. In contrast, the two less effective domains, namely WebShop and AlfWorld, exhibit lower complexity in at least one dimension. Specifically, although WebShop has a richer state space characterized by lengthy product descriptions and feature lists, it demands relatively low planning complexity. A WebShop agent generally only needs to iterate through the products to identify a match based on attributes. Similarly, while AlfWorld involves more complicated task logic, its state space remains relatively simple and compact.

From Table 1, we also observe that another potentially correlated factor could be the change in ID success rates (i.e.,  $\Delta_{ID}$ ) over the course of RL. The results seem to suggest a naive trade-off: intuitively, the more specialized knowledge absorbed through RL training, the lower the model’s OOD performance. This raises the question of whether knowledge absorption is a more dominant factor than the environment properties discussed above. To answer this, we conducted an additional experiment where we extended the training steps for Sokoban and SciWorld to further increase their  $\Delta_{ID}$ , comparing the final checkpoints to the ALFWorld and WebShop checkpoints at 150 steps. As shown in Table 3, even when the  $\Delta_{ID}$  of Sokoban is increased to match or exceed that of ALFWorld and WebShop, its OOD performance remains higher. These results confirm the impact of the

two environment properties and suggest that OOD performance is not simply determined by  $\Delta_{ID}$ . Additionally, we also identify several qualitative factors that have low correlations with OOD generalization. For instance, domain realism does not explain the performance differences. Sokoban, an abstract grid-world puzzle, enables strong generalization even to realistic domains like SciWorld. In contrast, ALFWorld, which simulates realistic household tasks, ranks lowest for generalization.

Table 3: OOD performance of Sokoban and SciWorld with extended training steps. The SciWorld checkpoints were trained with fewer extra steps due to earlier convergence. The full results, including OOD success rates and ranking scores, are provided in Table 7 and Table 8 in Appendix.

| Training Domain | Training Steps | Initial Policy     | $\Delta_{ID}$ | OOD Ranking Score ( $\downarrow$ ) |         |
|-----------------|----------------|--------------------|---------------|------------------------------------|---------|
|                 |                |                    |               | Ckpt V1 and V2                     | Ckpt V2 |
| Sokoban         | 150            | Ckpt V1<br>Ckpt V2 | 36.0<br>35.7  | 7 (1st)                            | 5 (2nd) |
|                 | 300            | Ckpt V1<br>Ckpt V2 | 67.9<br>68.4  | 8 (1st)                            | 6 (2nd) |
| SciWorld        | 150            | Ckpt V2            | 20.2          | -                                  | 3 (1st) |
|                 | 240            | Ckpt V2            | 41.8          | -                                  | 5 (1st) |

As an implication, one could strategically select or synthesize training environments with greater state information richness and planning complexity to better preserve out-of-domain generalization. Nonetheless, given the limited number of available domains, our current evidence should be interpreted as indicative rather than conclusive. We acknowledge that the initial characterization of effective domains is relatively ad hoc and may not causally isolate the critical factors. Therefore, in the next section, we extend the analysis through additional experiments designed to test some of these hypotheses and provide practical approaches for better generalization.

## 6 FROM CORRELATION TO PRACTICAL SOLUTION FOR PRESERVING GENERALIZATION

The analysis in Section 5 identifies two environmental properties, namely planning complexity and state information richness, which align with cross-domain generalization. A natural next step is to intervene along either of these axes to test whether modifying them can causally enhance generalization. Since planning complexity is largely intrinsic to the task logic and difficult to manipulate without altering the environment’s underlying dynamics, we focus on increasing *state information richness* through a technique we call *State Information Augmentation*. This approach involves enriching observations with goal-irrelevant information while keeping task goals, action spaces, and reward functions unchanged. Intuitively, this approach forces the agent to more actively extract task-relevant signals while ignoring distractive noise, which reduces the model’s reliance on domain-specific patterns.

Formally, let  $s$  denote the original state and  $g$  the goal. For each state  $s$ , we sample a short goal-irrelevant text fragment  $\delta$  from a distribution of content that does not interfere with the original optimal policy conditioned on  $g$ . We thereby construct an augmented state  $s' = \text{Augment}(s, \delta, \epsilon)$  by injecting a controlled volume  $\epsilon$  of goal-irrelevant information  $\delta$  into the textual state. Note that the augmentation is applied only to the agent’s input; components such as transitions and rewards still operate on the original state.

Practically, there can be several options for coming up with valid augmentations that do not interfere with the original optimal plan. With minimal domain knowledge, one could inject observations from a dramatically different domain into the training environment. For instance, we can inject Webshop observation into an Aleworld state by appending “there is a monitor showing a shopping webpage: <random webshop state>.” With slightly more understanding of a domain, one can come up with more contextual-relevant augmentation. For example, for the ALFWorld task “slice an apple”, valid augmentations could involve describing unrelated static objects (e.g., “the wall is painted blue”) or redundant features (e.g., “the apple is red”). Examples of augmentation for all training domains can be found in Appendix F. Nevertheless, we note that this study emphasizes verifying how increased state information richness can facilitate cross-domain generalization, instead of focusing extensively on the exact technique for constructing state augmentation.

Another practical consideration is setting the volume ( $\epsilon$ ) of distractive information, which can also be quantified using the character count of the injected text. Our experiments reveal insufficient aug-

mentation may not yield a noticeable change in OOD performance, while an excessive amount of noise can make the tasks overly difficult for the agent to learn, thereby negatively affecting ID performance. In practice, the appropriate  $\epsilon$  can be determined by interactively increasing its value until a significant drop in ID success rate is observed. Furthermore, we found it beneficial to stochastically apply the augmentation to only a portion of trajectories within a batch. This allows for better control over learning difficulty on top of adjusting the noise volume. In our experiments, we apply the augmentation to each trajectory with a probability of 50%, and the results suggest that this strategy helps maintain ID performance.

Table 4: Impact of state information augmentation on OOD success rates after RL training. Here, **+ Info. Aug.** denotes the model trained with state information augmentation. Shaded cells correspond to in-domain (ID) evaluation.  $\Delta_{ID}$  is the change in ID success rate relative to the corresponding base model. The **OOD Change** column reports the difference in OOD performance between the model using augmented states and the baseline using original states. A positive value indicates improvement.

| Initial Policy | Training Domain | Evaluation Domain  |                    |                    |          | OOD Change |
|----------------|-----------------|--------------------|--------------------|--------------------|----------|------------|
|                |                 | ALFWorld           | Webshop            | Sokoban            | SciWorld |            |
| Ckpt v1        | Base Model      | 25.8               | 34.4               | 12.5               | 3.1      | -          |
|                | ALFWorld        | $\Delta_{ID}=47.2$ | 15.2               | 14.0               | -        | -          |
|                | + Info. Aug.    | $\Delta_{ID}=47.7$ | 23.8               | 11.2               | -        | +32.6%     |
|                | Webshop         | 10.8               | $\Delta_{ID}=23.6$ | 12.8               | -        | -          |
|                | + Info. Aug.    | 17.8               | $\Delta_{ID}=27.6$ | 11.0               | -        | +35.5%     |
|                | Sokoban         | 20.5               | 15.0               | $\Delta_{ID}=36.0$ | -        | -          |
| + Info. Aug.   | 21.2            | 24.8               | $\Delta_{ID}=36.7$ | -                  | +42.5%   |            |
| Ckpt v2        | Base Model      | 21.9               | 38.3               | 14.1               | 18.0     | -          |
|                | ALFWorld        | $\Delta_{ID}=60.6$ | 30.5               | 9.8                | 10.0     | -          |
|                | + Info. Aug.    | $\Delta_{ID}=56.6$ | 30.3               | 11.0               | 12.5     | +7.0%      |
|                | Webshop         | 17.0               | $\Delta_{ID}=28.7$ | 9.0                | 13.8     | -          |
|                | + Info. Aug.    | 25.8               | $\Delta_{ID}=27.7$ | 11.8               | 15.5     | +33.4%     |
|                | Sokoban         | 20.0               | 34.0               | $\Delta_{ID}=35.7$ | 13.0     | -          |
| + Info. Aug.   | 20.8            | 37.0               | $\Delta_{ID}=37.4$ | 13.0               | +5.7%    |            |

Table 4 presents the empirical results of this intervention. To rigorously quantify the effect of state information augmentation, we introduce the **OOD Change** metric ( $\Delta_{OOD}$ ). Let  $\pi_{w/o\ Aug}^{RL}$  and  $\pi_{w/ Aug}^{RL}$  denote the policies trained with the original states and augmented states, respectively. Following the notations in Section 3.2, let  $\mathcal{E}^{eval}$  be the set of unseen evaluation environments for a given training session, and  $G(\pi, e)$  be the success rate of policy  $\pi$  on environment  $e$ .

We define the aggregate OOD change as:

$$\Delta_{OOD} = \frac{\sum_{e \in \mathcal{E}^{eval}} (G(\pi_{w/ Aug}^{RL}, e) - G(\pi_{w/o Aug}^{RL}, e))}{\sum_{e \in \mathcal{E}^{eval}} G(\pi_{w/o Aug}^{RL}, e)} \times 100\%.$$

A positive  $\Delta_{OOD}$  indicates that state information augmentation effectively improves the agent’s generalization. For most combinations of training domains and base policies, state information augmentation yields positive  $\Delta_{OOD}$ . In the experiments with Ckpt v1, where augmentation universally improves OOD performance across all training domains: ALFWorld (+7.0%), WebShop (+35.5%), and Sokoban (+42.5%). We observe a similar positive trend in the Ckpt v2 setting for ALFWorld (+7.0%), WebShop (+33.4%), and Sokoban (+5.7%). These results causally validate the hypothesis that state information richness is a determinant of agent robustness. Moreover, the technique of state information augmentation can also serve as a low-cost solution to preserve generalization.

## 7 OTHER INFLUENTIAL FACTORS ON CROSS-DOMAIN GENERALIZATION

During the experiments on the environment-side intervention, we also determined two modeling-side design choices that have a significant influence on OOD performance. This section proceeds by first discussing how SFT warmup impacts knowledge retention and forgetting during the subsequent RL phase, followed by an analysis of the necessity of explicit reasoning for both ID and OOD performance.

### 7.1 SFT WARMUP RESHAPES KNOWLEDGE RETENTION AND FORGETTING OF SUBSEQUENT RL

Recall from Section 4.2 that  $Ckpt_{V2}$  is obtained via SFT warmup to ensure a viable starting point for OOD analysis. This specific setup further provides an unique setup to study how SFT warmup reshapes retention and forgetting by comparing  $Ckpt_{V2}$  and  $Ckpt_{V1}$ , which does not undergo extra SFT training. Specifically, we compare the OOD performance of these two checkpoints after identical downstream RL training. Our results (Table 9) indicate that SFT warmup has its distinct advantages and disadvantages from the perspective of generalization.

Table 5: Impact of SFT warmup on OOD success rates (%) after RL training. We compare OOD performance of models initialized with  $Ckpt_{V1}$  against  $Ckpt_{V2}$  after identical downstream RL training. **Rel. Change Diff. (Ckpt V1 to V2)** shows the difference of percentage change in success rate relative to the initial policy of  $Ckpt_{V2}$  and of  $Ckpt_{V1}$ . A positive value indicates the checkpoint trained from  $Ckpt_{V2}$  being more robust than that trained from  $Ckpt_{V1}$ . The column “**In V2 SFT-warmup datamix?**” indicates whether the evaluation domain is covered in the SFT warmup datamix. Full results are provided in Tab. 9.

| RL Training Domain | Eval Domain | In V2 SFT-Warmup Datamix? | Rel. Change Diff. (Ckpt V1 to V2) |
|--------------------|-------------|---------------------------|-----------------------------------|
| <b>Sokoban</b>     | WebShop     | Yes                       | +45.2%                            |
|                    | ALFWorld    | Yes                       | +11.8%                            |
| <b>ALFWorld</b>    | WebShop     | Yes                       | +15.1%                            |
|                    | Sokoban     | No                        | -42.5%                            |
| <b>WebShop</b>     | ALFWorld    | Yes                       | +35.7%                            |
|                    | Sokoban     | No                        | -38.6%                            |

On the positive side, SFT warmup provides a strong lever for making domain knowledge more robust to downstream RL training. As shown in the “In V2 SFT-warmup datamix” column in Table 9, the SFT-warmup datamix of  $Ckpt_{V2}$  consists of data from three domains: ALFWorld, WebShop, and SciWorld. While SFT warmup does not change the fact that OOD performance declines throughout RL training, this decline is consistently and substantially smaller on evaluation domains covered by SFT warmup. For example, with Sokoban as the RL training domain and WebShop as the evaluation domain, the success rate of  $Ckpt_{V1}$  drops by  $-56.4\%$  whereas that of  $Ckpt_{V2}$  only drops by  $-11.2\%$ . Likewise, with WebShop as the RL training domain and ALFWorld as the evaluation domain, the success rate of  $Ckpt_{V1}$  drops by  $-58.1\%$  whereas that of  $Ckpt_{V2}$  only drops by  $-22.4\%$ . The same pattern holds in all experiments with evaluation domains covered by SFT warmup. Moreover, we note that the warmup data for ALFWorld and WebShop is directly sampled from the base model itself without extra filtering (e.g., with success indicator). This suggests that simply exposing the model to environment dynamics during SFT warmup can effectively consolidate domain knowledge.

Nevertheless, SFT warmup also sacrifices performance on domains outside its coverage, making knowledge from these domains more vulnerable to forgetting. This is evident when evaluating on Sokoban where  $Ckpt_{V2}$  consistently exhibits worse performance than  $Ckpt_{V1}$ . For instance, when training on WebShop,  $Ckpt_{V1}$ ’s OOD performance on Sokoban slightly increases by  $+2.4\%$  whereas that of  $Ckpt_{V2}$  decreases by a large margin ( $-36.2\%$ ).

These results suggest an inherent trade-off. SFT warmup or mid-training can effectively infuse and consolidate knowledge within a policy model, but it also carries a significant risk of erasing or diminishing prior capabilities in domains that are not covered. When the deployment environment is unknown, it is crucial to carefully control the extent of mid-training and to ensure that the datamix is as broad as possible to mitigate the risk of catastrophic forgetting.

Due to the page limit, we postpone the discussion of the necessity of explicit reasoning for both ID and OOD performance in Appendix A.

## 8 CONCLUSION

In this work, we take a data-centric perspective on RL post-training for *generalist* LLM agents, and probe how to strategically select or construct training domains so that agentic capabilities can be maximally preserved in unseen domains. By analyzing correlations between cross-domain perfor-

mance and properties of training environments, we find that *state information richness* and *planning complexity*, which respectively control perception load and reasoning load, align more with OOD performance.

Building upon these observations, we causally validate the importance of state information richness by directly intervening on this dimension. We introduce a state randomization method that injects small amounts of goal-irrelevant content into the agent’s observations, increasing information volume without changing the task space. Beyond environment-side designs, we also examine two modeling choices. Our results reveal that mid-training or SFT warmup can deeply infuse domain knowledge into the policy model, and the infused knowledge can retain even after subsequent RL that is conducted on other domains. However, mid-training can also lead to larger generalization degradation on uncovered domains compared to RL. We further find that, while enabling step-by-step thinking during RL and evaluation does not increase ID success rates, it plays a crucial role in transferring agentic capabilities to unseen environments.

## REFERENCES

- Léo Boisvert, Megh Thakkar, Maxime Gasse, Massimo Caccia, Thibault Le Sellier De Chezelles, Quentin Cappart, Nicolas Chapados, and Alexandre Drouin. Workarena++: Towards compositional planning and reasoning-based common knowledge work tasks. In *NeurIPS Datasets and Benchmarks Track*, 2024.
- Y. Cao et al. Survey on large language model-enhanced reinforcement learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. arXiv:2404.00282.
- Xiaoyu Chen, Jiachen Hu, Chi Jin, Lihong Li, and Liwei Wang. Understanding domain randomization for sim-to-real transfer. *arXiv preprint arXiv:2110.03239*, 2021.
- Zhoujun Cheng, Shibo Hao, Tianyang Liu, Fan Zhou, Yutao Xie, Feng Yao, Yuexin Bian, Yonghao Zhuang, Nilabjo Dey, Yuheng Zha, Yi Gu, Kun Zhou, Yuqi Wang, Yuan Li, Richard Fan, Jian-shu She, Chengqian Gao, Abulhair Saparov, Haonan Li, Taylor W. Killian, Mikhail Yurochkin, Zhengzhong Liu, Eric P. Xing, and Zhiting Hu. Revisiting reinforcement learning for LLM reasoning from a cross-domain perspective. *arXiv preprint arXiv:2506.14965*, 2025.
- Gert de Roo. Planning and complexity: An introduction. In *A planner’s encounter with complexity*, pp. 1–18. Routledge, 2016.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In *NeurIPS*, 2023.
- Alexandre Drouin, Léo Boisvert, Megh Thakkar, Maxime Gasse, Massimo Caccia, Thibault Le Sellier De Chezelles, Quentin Cappart, Nicolas Chapados, Alexandre Drouin, et al. Workarena: How capable are web agents at solving realistic enterprise tasks? *arXiv preprint arXiv:2403.07718*, 2024.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. Group-in-group policy optimization for llm agent training. *arXiv preprint arXiv:2505.10978*, 2025.
- Dayuan Fu, Keqing He, Yejie Wang, Wentao Hong, Zhuoma Gongque, Weihao Zeng, Wei Wang, Jingang Wang, Xunliang Cai, and Weiran Xu. Agentrefine: Enhancing agent generalization through refinement tuning. *arXiv preprint arXiv:2501.01702*, 2025.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Lin Guan, Mudit Verma, Suna Sihang Guo, Ruohan Zhang, and Subbarao Kambhampati. Widening the pipeline in human-guided reinforcement learning with explanation and context-aware data augmentation. *Advances in Neural Information Processing Systems*, 34:21885–21897, 2021.
- Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning. *Advances in Neural Information Processing Systems*, 36:79081–79094, 2023.

- Nicklas Hansen, Sébastien Racanière, Josip Djolonga, Esmail Danier, Matteo Hessel, Lasse Espeholt, and Mohammad Babaeizadeh. Soft data augmentation for visually grounded reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pp. 23393–23406, 2021.
- Mengkang Hu, Pu Zhao, Can Xu, Qingfeng Sun, Jian-Guang Lou, Qingwei Lin, Ping Luo, and Saravan Rajmohan. Agentgen: Enhancing planning abilities for large language model based agent via environment and task generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pp. 496–507, 2025.
- Maggie Huan, Yuetai Li, Tuney Zheng, Xiaoyu Xu, Seungone Kim, Minxin Du, Radha Poovendran, Graham Neubig, and Xiang Yue. Does math reasoning improve general llm capabilities? understanding transferability of llm reasoning. *arXiv preprint arXiv:2507.00432*, 2025.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023.
- Andreas Junghanns and Jonathan Schaeffer. Sokoban: Enhancing general single-agent search methods using domain knowledge. *Artificial Intelligence*, 129(1-2):219–251, 2001.
- Thibault Le Sellier De Chezelles, Maxime Gasse, Alexandre Drouin, Massimo Caccia, Léo Boisvert, Megh Thakkar, et al. The browsergym ecosystem for web agent research. *arXiv preprint arXiv:2412.05467*, 2024.
- Zhihan Liu, Hao Hu, Shenao Zhang, Hongyi Guo, Shuqi Ke, Boyi Liu, and Zhaoran Wang. Reason for future, act for now: A principled framework for autonomous llm agents with provable sample efficiency. *arXiv preprint arXiv:2309.17382*, 2023.
- Antonio Loquercio, Elia Kaufmann, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. Deep drone racing: From simulation to reality with domain randomization. *IEEE Transactions on Robotics*, 36(1):1–14, 2019.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Qingwei Lin, Jianguang Lou, Shifeng Chen, Yansong Tang, and Weizhu Chen. Arena learning: Build data flywheel for llms post-training via simulated chatbot arena. *arXiv preprint arXiv:2407.10627*, 2024.
- Justin K. Miller and Wenjia Tang. Evaluating llm metrics through real-world capabilities. *arXiv preprint arXiv:2505.08253*, 2025.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Kaixiang Mo, Yuxin Shi, Weiwei Weng, Zhiqiang Zhou, Shuman Liu, Haibo Zhang, and Anxiang Zeng. Mid-training of large language models: A survey. *arXiv preprint arXiv:2510.06826*, 2025.
- Fabio Muratore, Felix Treede, Michael Gienger, and Jan Peters. Domain randomization for simulation-based policy optimization with transferability assessment. In *Conference on Robot Learning*, pp. 700–713. PMLR, 2018.
- Charles Packer, Katelyn Gao, Vineet Kosaraju, Lucas Beyer, Vladlen Koltun, and Jason Yosinski. Assessing generalization in deep reinforcement learning. *arXiv preprint arXiv:1810.12282*, 2018.
- Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 2817–2826, 2017.
- Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. In *Robotics: Science and Systems (RSS)*, 2017.

- Mohammad Reza Samsami, Artem Zhohus, Janarthanan Rajendran, and Sarath Chandar. Mastering memory tasks with world models. In *International Conference on Learning Representations (ICLR)*, 2024.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv: 2409.19256*, 2024.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. Aleworld: Aligning text and embodied environments for interactive learning. In *ICLR*, 2021.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Haotian Sun, Yuchen Zhuang, Ling kai Kong, Bo Dai, and Chao Zhang. Adaplaner: Adaptive planning from feedback with language models. *Advances in neural information processing systems*, 36:58202–58245, 2023.
- Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 23–30, 2017.
- Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. ScienceWorld: Is your agent smarter than a 5th grader? In *EMNLP*, 2022.
- S. Wang, Z. Tang, et al. Reinforcement learning enhanced llms: A survey. *arXiv preprint arXiv:2412.10400*, 2024a.
- Zaitian Wang, Pengfei Wang, Kunpeng Liu, Pengyang Wang, Yanjie Fu, Chang-Tien Lu, Charu C Aggarwal, Jian Pei, and Yuanchun Zhou. A comprehensive survey on data augmentation. *IEEE Transactions on Knowledge and Data Engineering*, 2025a.
- Zengzhi Wang, Fan Zhou, Xuefeng Li, and Pengfei Liu. Octothinker: Mid-training incentivizes reinforcement learning scaling. *arXiv preprint arXiv:2506.20512*, 2025b.
- Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Xing Jin, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, et al. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*, 2025c.
- Zilong Wang, Yuedong Cui, Li Zhong, Zimin Zhang, Da Yin, Bill Yuchen Lin, and Jingbo Shang. Officebench: Benchmarking language agents across multiple applications for office automation. *arXiv preprint arXiv:2407.19056*, 2024b.
- 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, Zuxuan Wu, Qi Zhang, Xuanjing Huang, and Yu-Gang Jiang. AgentGym-RL: Training LLM agents for long-horizon decision making through multi-turn reinforcement learning. *arXiv preprint arXiv:2509.08755*, 2025.
- Frank F. Xu, Yufan Song, Boxuan Li, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, Zora Z. Wang, Xuhui Zhou, Zhitong Guo, Murong Cao, et al. Theagentcompany: Benchmarking llm agents on consequential real world tasks. *arXiv preprint arXiv:2412.14161*, 2024.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. In *NeurIPS*, 2022a.

- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*, 2022b.
- Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2100–2110, 2019.
- Hanchen Zhang, Xiao Liu, Bowen Lv, Xueqiao Sun, Bohao Jing, Iat Long Iong, Zhenyu Hou, Zehan Qi, Hanyu Lai, Yifan Xu, Rui Lu, Hongning Wang, Jie Tang, and Yuxiao Dong. AgentRL: Scaling agentic reinforcement learning with a multi-turn, multi-task framework. *arXiv preprint arXiv:2510.04206*, 2025a.
- Kai Zhang, Xiangchao Chen, Bo Liu, Tianci Xue, Zeyi Liao, Zhihan Liu, Xiyao Wang, Yuting Ning, Zhaorun Chen, Xiaohan Fu, et al. Agent learning via early experience. *arXiv preprint arXiv:2510.08558*, 2025b.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand, 2024. Association for Computational Linguistics. URL <http://arxiv.org/abs/2403.13372>.
- 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. Webarena: A realistic web environment for building autonomous agents. In *ICLR*, 2024.
- Yifei Zhou, Song Jiang, Yuandong Tian, Jason Weston, Sergey Levine, Sainbayar Sukhbaatar, and Xian Li. SWEET-RL: Training multi-turn LLM agents on collaborative reasoning tasks. *arXiv preprint arXiv:2503.15478*, 2025.

Table 6: Performance ranking of models trained on different source domains when evaluated on each unseen domain. For each evaluation column, the training domains are ranked from best (1) to worst (3 or 4) based on their success rate. Domains with success rate differences  $< 0.5$  are considered tied and assigned the same rank. In-domain entries are marked with  $-$ .

| Initial Policy | Training Domain | Evaluation Domain |         |         |          |
|----------------|-----------------|-------------------|---------|---------|----------|
|                |                 | ALFWorld          | Webshop | Sokoban | SciWorld |
| Ckpt V1        | ALFWorld        | -                 | 1       | 1       | -        |
|                | Webshop         | 2                 | -       | 2       | -        |
|                | Sokoban         | 1                 | 1       | -       | -        |
| Ckpt V2        | ALFWorld        | -                 | 3       | 2       | 3        |
|                | Webshop         | 2                 | -       | 3       | 1        |
|                | Sokoban         | 1                 | 2       | -       | 2        |
|                | SciWorld        | 1                 | 1       | 1       | -        |

Table 7: OOD success rates of Sokoban and SciWorld with extended training steps.

| Training Domain | Total Training Steps | Initial Policy | Evaluation Domain |         |                    |                    |
|-----------------|----------------------|----------------|-------------------|---------|--------------------|--------------------|
|                 |                      |                | ALFWorld          | Webshop | Sokoban            | SciWorld           |
| Sokoban         | 150                  | Ckpt V1        | 20.5              | 15.0    | $\Delta_{ID}=36.0$ | -                  |
|                 |                      | Ckpt V2        | 20.0              | 34.0    | $\Delta_{ID}=35.7$ | 13.0               |
|                 | 300                  | Ckpt V1        | 19.2              | 15.8    | $\Delta_{ID}=67.9$ | -                  |
|                 |                      | Ckpt V2        | 17.5              | 36.9    | $\Delta_{ID}=68.4$ | 9.7                |
| SciWorld        | 150                  | Ckpt V2        | 19.8              | 35.8    | 12.0               | $\Delta_{ID}=20.2$ |
|                 | 240                  | Ckpt V2        | 16.9              | 33.6    | 11.8               | $\Delta_{ID}=41.8$ |

Table 8: OOD performance of Sokoban and SciWorld with extended training steps. The SciWorld checkpoints were trained with fewer extra steps due to earlier convergence. Domains with success rate differences  $< 0.5$  are considered tied and assigned the same rank.

| Training Domain | Total Training Steps | Initial Policy | Evaluation Domain |         |         |          |
|-----------------|----------------------|----------------|-------------------|---------|---------|----------|
|                 |                      |                | ALFWorld          | Webshop | Sokoban | SciWorld |
| Sokoban         | 150                  | Ckpt V1        | 1                 | 1       | -       | -        |
|                 |                      | Ckpt V2        | 1                 | 2       | -       | 2        |
|                 | 300                  | Ckpt V1        | 1                 | 1       | -       | -        |
|                 |                      | Ckpt V2        | 2                 | 1       | -       | 3        |
| SciWorld        | 150                  | Ckpt V2        | 1                 | 1       | 1       | -        |
|                 | 240                  | Ckpt V2        | 2                 | 2       | 1       | -        |

## A ENABLING STEP-BY-STEP REASONING PRESERVES OOD PERFORMANCE

Another important modeling-side consideration is the use of explicit step-by-step reasoning during RL training. Although explicit reasoning has become the default strategy for many LLM-driven agentic systems, it is worth noting that “reactive” architectures, such as DQN (Mnih et al., 2015; Silver et al., 2016), can also achieve human-level performance on target domains without explicit reasoning. Here, we provide an additional perspective: explicit reasoning is essential for transferring or preserving agentic capabilities in novel domains that are not seen during training.

Table 10 in Appendix demonstrates that models trained with explicit reasoning consistently maintain higher OOD performance compared to those reactive baselines that output actions directly. In experiments with Ckpt V2, disabling reasoning causes the OOD performance to collapse by over 200% in the ALFWorld, Sokoban, and SciWorld training runs. For instance, when training on ALFWorld and evaluating on WebShop, the success rate of the policy trained without explicit reasoning drops to 1.0%, whereas the policy with explicit reasoning maintains a success rate of 30.5%.

Table 9: Impact of SFT warmup on OOD success rates (%) after RL training. We compare OOD performance of models initialized with Ckpt V1 against Ckpt V2 after identical downstream RL training. **Rel. Change** indicates the percentage change in success rate relative to the corresponding initial policy. **Rel. Change Diff. (Ckpt V1 to V2)** shows the difference of the Rel. Change of Ckpt V2 and the Rel. Change of Ckpt V1. The column “**In V2 SFT-warmup datamix?**” indicates whether the evaluation domain is covered in the SFT warmup datamix of Ckpt V2.

| RL Training Domain | Eval Domain | In V2 SFT-Warmup Datamix? | Ckpt V1      |             | Ckpt V2      |             | Rel. Change Diff. (Ckpt V1 to V2) |
|--------------------|-------------|---------------------------|--------------|-------------|--------------|-------------|-----------------------------------|
|                    |             |                           | success rate | Rel. Change | success rate | Rel. Change |                                   |
| Sokoban            | WebShop     | Yes                       | 15.0         | -56.4%      | 34.0         | -11.2%      | +45.2%                            |
|                    | ALFWorld    | Yes                       | 20.5         | -20.5%      | 20.0         | -8.7%       | +11.8%                            |
| ALFWorld           | WebShop     | Yes                       | 10.8         | -68.6%      | 17.8         | -53.5%      | +15.1%                            |
|                    | Sokoban     | No                        | 14.0         | +12.0%      | 9.8          | -30.5%      | -42.5%                            |
| WebShop            | ALFWorld    | Yes                       | 10.8         | -58.1%      | 17.0         | -22.4%      | +35.7%                            |
|                    | Sokoban     | No                        | 12.8         | +2.4%       | 9.0          | -36.2%      | -38.6%                            |

Table 10: Impact of Step-by-Step Reasoning on OOD success rates (%) after RL training. We compare the baseline RL runs (which use thinking) against runs where thinking is disabled (**w/o Thinking**). Shaded cells correspond to in-domain (ID) evaluation. To align with the discussion in Section 5, we report the change in ID success rate relative to the base model (i.e.,  $\Delta_{ID}$ ). The **OOD Change** column reports the relative performance drop of the non-thinking setup compared to the baseline. A large negative value indicates that disabling thinking severely hurts generalization.

| Initial Policy | Training Domain | Evaluation Domain  |                    |                    |                    | OOD Change |
|----------------|-----------------|--------------------|--------------------|--------------------|--------------------|------------|
|                |                 | ALFWorld           | Webshop            | Sokoban            | SciWorld           |            |
| Ckpt V1        | Base Model      | 25.8               | 34.4               | 12.5               | 3.1                | -          |
|                | ALFWorld        | $\Delta_{ID}=47.2$ | 15.2               | 14.0               | -                  | -          |
|                | w/o Thinking    | $\Delta_{ID}=51.2$ | 0.8                | 7.5                | -                  | -119.0%    |
|                | Webshop         | 10.8               | $\Delta_{ID}=23.6$ | 12.8               | -                  | -          |
|                | w/o Thinking    | 8.2                | $\Delta_{ID}=27.6$ | 17.8               | -                  | +16.9%     |
|                | Sokoban         | 20.5               | 15.0               | $\Delta_{ID}=36.0$ | -                  | -          |
| Ckpt V2        | w/o Thinking    | 16.0               | 17.2               | $\Delta_{ID}=23.0$ | -                  | -9.1%      |
|                | Base Model      | 21.9               | 38.3               | 14.1               | 18.0               | -          |
|                | ALFWorld        | $\Delta_{ID}=60.6$ | 30.5               | 9.8                | 10.0               | -          |
|                | w/o Thinking    | $\Delta_{ID}=51.3$ | 1.0                | 6.5                | -                  | -234.2%    |
|                | Webshop         | 17.0               | $\Delta_{ID}=28.7$ | 9.0                | 13.8               | -          |
|                | w/o Thinking    | 4.5                | $\Delta_{ID}=14.5$ | 12.5               | 11.0               | -82.5%     |
| Ckpt V2        | Sokoban         | 20.0               | 34.0               | $\Delta_{ID}=35.7$ | 13.0               | -          |
|                | w/o Thinking    | 9.2                | 25.8               | $\Delta_{ID}=60.7$ | 7.8                | -216.5%    |
|                | SciWorld        | 19.8               | 35.8               | 12.0               | $\Delta_{ID}=20.2$ | -          |
| w/o Thinking   | 6.2             | 32.5               | 14.0               | $\Delta_{ID}=17.5$ | -216.9%            |            |

Notably, enabling reasoning does not always benefit in-domain performance. In several cases, such as WebShop and ALFWorld, policies trained without thinking achieve comparable or even slightly higher in-domain success rates. For instance, the in-domain performance on ALFWorld using `ckpt v1` increases from 73.0% to 77.0% when reasoning is disabled. This suggests reactive policies can sufficiently learn the optimal policies by discovering domain-specific patterns tailored for the training environments. However, these patterns or heuristics can be brittle and not transferrable. By forcing the model to articulate its reasoning process step-by-step, the model is less prone to overfitting to shallow heuristics and thereby preserving more cross-domain generalization.

## B RELATED WORK IN MULTI-TURN RL TRAINING FOR LLM AGENTS

Recent works tackle the challenges of training LLMs as interactive agents via RL in multi-turn settings (Yao et al., 2022b; Sun et al., 2023; Liu et al., 2023; Wang et al., 2025c). For example, Wang et al. (2025c) and Feng et al. (2025) extend the popular RL training framework VeRL (Sheng et al., 2024) from the standard single-turn settings (as in coding and math tasks) to the multi-turn settings. Besides, Zhang et al. (2025a) introduce a scalable framework for agentic multi-turn RL that features an asynchronous generation-training pipeline and Xi et al. (2025) propose a unified platform for training LLM agents in diverse environments. Zhou et al. (2025) focus on the credit assignment problem in long dialogues. They develop a benchmark for collaborative multi-turn tasks such as team-based software development. These advances in LLM-agent frameworks, together with prior works such as AlphaGo (Silver et al., 2016), demonstrate that with well-designed training environments and optimization strategies, RL agents can acquire highly specialized knowledge. Our work further expands the perspective of training specialized agents by investigating what types of domains are most effective for training generalist RL agents.

## C LIMITATIONS AND FUTURE WORK

This work establishes a practical framework for analyzing and improving cross-domain generalization in agentic post-training. Looking ahead, several promising directions merit further exploration. First, expanding this study to a broader suite of domains and diverse model scales would help validate the universality of factors identified as critical for OOD performance. A larger-scale investigation may also reveal additional factors that influence generalization. Second, there is room to further refine the empirical measures used to characterize effective training environments. Developing more formal and causally grounded definitions would support more principled environment design. Methodology-wise, future work could investigate automated environment augmentation techniques, the relationship between reasoning structure and generalization, and adaptive SFT warmup schedules that dynamically balance knowledge consolidation and flexibility. Lastly, connecting these empirical findings to theoretical interpretations of generalization and policy optimization remains an important step toward establishing quantitative design principles for scalable and robust agent training.

## D EXPERIMENTAL DETAILS

### D.1 RL IN ALFWORLD

We use `verl-agent` (Feng et al., 2025) as the codebase and use 8 NVIDIA A100 (80G) GPUs for the training and evaluation of ALFWorld. We perform the evaluation using the out-of-domain split of the ALFWorld dataset and perform the training using the train split.

**Hyperparameters for ALFWorld.** The maximum prompt length is 2048 tokens, and the maximum response length is 512 tokens. Each episode allows up to 50 environment steps. The learning rate is set to  $1e-6$  for the actor. We adopt a rule-based reward, assigning a reward of 10 for success and 0 for failure. To handle invalid actions generated by the agent, we apply a reward penalty of -0.1. For GRPO, we sample 8 trajectories per task prompts (i.e., group size = 8) and use 16 task prompts per training step. The rollout temperature is set to 1.0, while the validation temperature is set to 0.4. The training mini-batch size consists of 256 (representing the state count), and the KL-divergence regularizer coefficient is set to 0.01.

## D.2 RL IN WEBSHOP

We use verl-agent (Feng et al., 2025) as the codebase and use 8 NVIDIA A100 (80G) GPUs for the training and evaluation of Webshop. We use the first 500 goals from the “small dataset” with 1000 products for evaluation and the remaining “small dataset” for training.

**Hyperparameters for Webshop.** The maximum prompt length is 4096 tokens, and the maximum response length is 1024 tokens. Each episode allows up to 15 environment steps. The learning rate is set to  $1e-6$  for the actor. We adopt a rule-based reward, assigning a reward of 10 for success and 0 for failure. To handle invalid actions generated by the agent, we apply a reward penalty of -0.1. For GRPO, we sample 8 trajectories per task prompts (i.e., group size = 8) and use 16 task prompts per training step. The rollout temperature is set to 1.0, while the validation temperature is set to 0.4. The training mini-batch size is 256 (representing the state count), and the KL-divergence regularizer coefficient is set to 0.01.

## D.3 RL IN SOKOBAN

We use verl-agent (Feng et al., 2025) as the codebase and use 8 NVIDIA A100 (80G) GPUs for the training and evaluation of Sokoban.

**Hyperparameters for Sokoban.** The maximum prompt length is 4096 tokens, and the maximum response length is 512 tokens. Each episode allows up to 15 environment steps. The learning rate is set to  $1e-6$  for the actor. We adopt a rule-based reward, assigning a reward of 10 for success and 0 for failure. To handle invalid actions generated by the agent, we apply a reward penalty of -0.1. For GRPO, we sample 8 trajectories per task prompts (i.e., group size = 8) and use 16 task prompts per training step. The rollout temperature is set to 1.0, while the validation temperature is set to 0.4. The training mini-batch size is 256 (representing the state count), and the KL-divergence regularizer coefficient is set to 0.01.

## D.4 RL IN SCIWORLD

Since the SciWorld environment is incompatible with the verl-agent, we modify verl-agent (Feng et al., 2025) to implement the interaction interface via a proxy-client framework to decouple agent interaction, LLM inference, and model training. We use 8 NVIDIA A100 (80G) GPUs for the training and evaluation of SciWorld.

**Hyperparameters for SciWorld** The maximum prompt length is 2048 tokens, and the maximum response length is 1024 tokens. Each episode allows up to 50 environment steps. The learning rate is set to  $1e-6$  for the actor. We adopt a rule-based reward, assigning a reward of 10 for success and 0 for failure. To handle invalid actions generated by the agent, we apply a reward penalty of -0.1. For GRPO, we sample 8 trajectories per task prompts (i.e., group size = 8) and use 16 task prompts per training step. The rollout temperature is set to 1.0, while the validation temperature is set to 0.4. The training mini-batch size is 256 (representing the state count), and the KL-divergence regularizer coefficient is set to 0.01.

## D.5 SFT WARMUP

We use LlamaFactory (Zheng et al., 2024) as the codebase for SFT warmup. The SFT dataset is constructed by mixing expert SciWorld data (Wang et al., 2022), ALFWorld data generated by Ckpt v1, and WebShop data generated by Ckpt v1 in a 19:3:1 ratio. This ratio is defined over the count of state-action pairs.

**Hyperparameters for SFT Warmup** The learning rate is  $5e-6$  with the cosine learning rate scheduler and warmup ratio 0.1. The batch size is 32. We train the model for 100 steps to best balance the performance on SciWorld and the performance on ALFWorld, Webshop, and Sokoban.

## E CONCRETE EXAMPLES

In this section, we provide example states and actions from ALFWorld, Webshop, Sokoban, and SciWorld for comparison.

## E.1 ALFWORLD

### Example State from ALFWorld

You are an expert agent operating in the ALFRED Embodied Environment. Your current observation is: == Welcome to TextWorld, ALFRED! ==

You are in the middle of a room. Looking quickly around you, you see a cabinet 16, a cabinet 15, a cabinet 14, a cabinet 13, a cabinet 12, a cabinet 11, a cabinet 10, a cabinet 9, a cabinet 8, a cabinet 7, a cabinet 6, a cabinet 5, a cabinet 4, a cabinet 3, a cabinet 2, a cabinet 1, a coffeemachine 1, a countertop 2, a countertop 1, a diningtable 1, a drawer 5, a drawer 4, a drawer 3, a drawer 2, a drawer 1, a fridge 1, a garbagecan 1, a microwave 1, a safe 1, a sinkbasin 1, a stoveburner 4, a stoveburner 3, a stoveburner 2, a stoveburner 1, and a toaster 1.

Your task is to: find two knife and put them in drawer. Your admissible actions of the current situation are: ['go to cabinet 1' 'go to cabinet 10' 'go to cabinet 11' 'go to cabinet 12' 'go to cabinet 13' 'go to cabinet 14' 'go to cabinet 15' 'go to cabinet 16' 'go to cabinet 2' 'go to cabinet 3' 'go to cabinet 4' 'go to cabinet 5' 'go to cabinet 6' 'go to cabinet 7' 'go to cabinet 8' 'go to cabinet 9' 'go to coffeemachine 1' 'go to countertop 1' 'go to countertop 2' 'go to diningtable 1' 'go to drawer 1' 'go to drawer 2' 'go to drawer 3' 'go to drawer 4' 'go to drawer 5' 'go to fridge 1' 'go to garbagecan 1' 'go to microwave 1' 'go to safe 1' 'go to sinkbasin 1' 'go to stoveburner 1' 'go to stoveburner 2' 'go to stoveburner 3' 'go to stoveburner 4' 'go to toaster 1' 'inventory' 'look'].

Now it's your turn to take an action. You should first reason step-by-step about the current situation. This reasoning process MUST be enclosed within <think> </think> tags. Once you've finished your reasoning, you should choose an admissible action for current step and present it within <action> </action> tags.

### Example Action (with Thinking) from ALFWorld

<think> The current situation is that I am in the middle of a room, and I need to find two knives and put them in a drawer. I have a list of admissible actions that I can take to achieve this goal. The first step is to determine which drawer I should go to. Since I need to put the knives in a drawer, I should go to the drawer that is closest to the knives. I have already observed the room, and I see that there are several drawers, but I don't know which ones contain knives. I need to go to the drawer that is most likely to contain knives.

The most likely drawer to contain knives is the one that is closest to the kitchen area, where the knives are typically kept. I see that there are several cabinets and appliances in the room, but I don't see any knives. I need to go to the drawer that is closest to the kitchen area. The closest drawer to the kitchen area is drawer 5. </think> <action>go to drawer 5</action>

## E.2 WEBSHOP

### Example State from Webshop

You are an expert autonomous agent operating in the WebShop e-commerce environment. Your task is to: Find me slim fit, loose fit men's tuxedo shirts with long sleeve, short sleeve, contrast color, classic fit for teen girls with color: b-bk1, and size: small, and price lower than 30.00 dollars. Prior to this step, you have already taken 14 step(s). Below are the most recent 2 observations and the corresponding actions you took: [Observation 13: ''Back to Search' [SEP] '< Prev' [SEP] 'size' [SEP] 'medium' [SEP] 'large' [SEP] 'x-large' [SEP] 'xx-large' [SEP] '3x-large' [SEP] 'color' [SEP] 'white' [SEP] 'yellow' [SEP] 'Mens Linen Shirt,Men's Striped Shirts Casual Short Sleeve Button Down Shirts Regular Fit Hawaiian Shirts Beach Tees Tops' [SEP] 'Price: \$3.78 to \$11.38' [SEP] 'Rating: N.A.' [SEP] 'Description' [SEP] 'Features' [SEP] 'Reviews' [SEP] 'Buy Now'', Action 13: 'click[back to search]' [Observation 14: ''Search'', Action 14: 'search[query: \"slim fit loose

```

fit men's tuxedo shirts long sleeve short sleeve contrast color classic
fit teen girls b-bk1 small lower than 30.00\"])'] You are now at step 15
and your current observation is: 'Back to Search' [SEP] 'Page 1 (Total
results: 50)' [SEP] 'Next >' [SEP] 'B09S3BN15C' [SEP] 'Mens Linen Shirt,
Men's Striped Shirts Casual Short Sleeve Button Down Shirts Regular Fit
Hawaiian Shirts Beach Tees Tops' [SEP] '$3.78 to $11.38' [SEP] '
B09NNMV9LN' [SEP] 'JSPOYOU Mens Short Sleeve Crewneck 3D Graphic Tunic
Shirts Big & Tall Tie Dye Summer Top Basic Designed Classic Cotton Shirt
' [SEP] '$3.99 to $10.99' [SEP] 'B09Q67H373' [SEP] 'JSPOYOU 2022 Newly T
-Shirt for Mens Funny 3D Graphics Pattern Crewneck Short Sleeve Tees Big
and Tall Summer Casual Comfy Top' [SEP] '$1.99 to $8.99' [SEP] '
B09QQP3356' [SEP] 'HAUKLIE Men's Sports Waffle Ribbed Polo Shirts Summer
Short Sleeve Cotton Muscle Quarter-Zip Henley T-Shirt Tunics Tops' [SEP]
'$10.99' [SEP] 'B07F2G93BJ' [SEP] 'Amazon Essentials Men's Slim-fit
Long-Sleeve Solid Pocket Oxford Shirt' [SEP] '$18.5' [SEP] 'B09QGK5XHZ'
[SEP] 'WENKOMG1 Men's Long Sleeve Undershirt with Mask Turtleneck Hooded
T-Shirt Solid Color Workout Tops Zipper Side Slit Shirts Slim Fit
Sweatshirt Spring/Summer Tee Shirts(Gray,)' [SEP] '$8.39' [SEP] '
B09HGFV91W' [SEP] 'ZHDD Tops for Mens, Men's Knight Punk Vintage Style
Long Sleeve Asymmetric Hem Cowl Neck Teen Boys Gothic Pullover' [SEP] '
$4.98 to $7.99' [SEP] 'B09R9YCM6R' [SEP] 'XXBR Summer T-shirts for Mens,
Soldier Short Sleeve 3D Street Vintage Printed Shirt Slim Fit Muscle
Casual Tee Tops' [SEP] '$8.98 to $11.99' [SEP] 'B09T756KQ5' [SEP] '
WENKOMG1 Men's Crewneck Goth Tee Shirts Skull Print Tops Spring/Summer
Long Sleeve Sports T-Shirt Baggy Y2K Soft Streetwear' [SEP] '$4.09 to $8
.09' [SEP] 'B09ND9DP7J' [SEP] 'InterestPrint Gold Horse Pattern Men's 2-
Piece Sleepwear Set, Long Sleeve Shirt with Pants Loungewear' [SEP] '$43
.59'. Your admissible actions of the current situation are: [ 'click[
back to search]', 'click[next >]', 'click[b09s3bn15c]', 'click[
b09nnmv9ln]', 'click[b09q67h373]', 'click[b09qqp3356]', 'click[
b09qqp3356]', 'click[b07f2g93bj]', 'click[b09qgk5xhz]', 'click[b09hgf91w
]', 'click[b09r9ycm6r]', 'click[b09t756kq5]', 'click[b09nd9dp7j]', ].

```

Note: Prioritize click a product in the current page over going to next page. Now it's your turn to take one action for the current step ( respond in English). You should first reason step-by-step about the current situation. This reasoning process MUST be enclosed within <think > </think> tags. Once you've finished your reasoning, you should choose an admissible action for current step and present it within <action> </action> tags.

### Example Action (with Thinking) from Webshop

```

<think> The current situation is that I have already searched for slim
fit loose fit men's tuxedo shirts long sleeve short sleeve contrast
color classic fit teen girls b-bk1 small lower than 30.00. The search
returned 50 results, but none of them match the specified criteria of
having a long sleeve, short sleeve, contrast color, classic fit for teen
girls with color: b-bk1, and size: small, and price lower than 30.00
dollars.

```

```

The current page displays a list of products that do not match the
specified criteria. The first product is a men's linen shirt, which does
not have a long sleeve, short sleeve, contrast color, classic fit for
teen girls, or a color of b-bk1. The other products also do not match
the specified criteria. I need to review the products on the current
page to see if any of them match the specified criteria. If not, I will
need to go to the next page to review the remaining products. However, I
need to prioritize reviewing the products on the current page first, as
it is more efficient than going to the next page.</think>
<action> click[b09s3bn15c] </action>

```

### E.3 SOKOBAN

#### Example State from Sokoban

```

You are an expert agent operating in the Sokoban environment.

# Objective

```

Your objective is to push all boxes onto their corresponding target locations. Once every box is correctly placed on a target, the puzzle is solved and you win.

#### # Environment Setup

- The environment is a 6x6 grid. - Each cell in the grid can contain one of the following objects: **player (you)**, **box**, **target**, **wall**, or be **empty**. - All objects occupy exactly one grid cell and are specified using their (row, column) coordinates, where (0, 0) is the top-left corner.

#### # Rules

1. **Movement and Constraints** - You can move up, down, left, or right by one cell at a time. - You **cannot** move through walls. - You **cannot** move onto or through boxes unless you are pushing them. - You **cannot** pull boxes, only push them.

2. **Pushing Boxes** - To push a box, it must be directly adjacent to your current position in the direction you are moving. - The cell immediately beyond the box (in the same direction) must be empty (i.e., not a wall, another box, or the boundary). - When you push a box, both you and the box move one cell forward.

**Example:** If you are at (2, 3), a box is at (2, 4), and (2, 5) is empty, then moving right pushes the box to (2, 5), and you move to (2, 4).

3. **Deadlocks and Planning** - Avoid pushing boxes into corners or against walls where they cannot be moved again.

#### # Observation Format

You will be given a list of all objects and their coordinates, including:  
: - Player position - Box positions - Target positions - Wall positions  
All unspecified cells are empty.

#### # Current Step

Prior to this step, you have already taken 27 step(s). Below are the most recent 1 observations and the corresponding actions you took: [Text

Observation 27: Wall at (0, 0) Wall at (0, 1) Wall at (0, 2) Wall at (0, 3) Wall at (0, 4) Wall at (0, 5) Wall at (1, 0) Wall at (1, 5) Wall at (2, 0) Goal at (2, 1) Wall at (2, 5) Wall at (3, 0) Box at (3, 3) Wall at (3, 5) Wall at (4, 0) Wall at (4, 1) Wall at (4, 2) Player at (4, 3) Wall at (4, 4) Wall at (4, 5) Wall at (5, 0) Wall at (5, 1) Wall at (5, 2) Wall at (5, 3) Wall at (5, 4) Wall at (5, 5) Action 27: 'Still'] You are now at step 28 and your current observation is: Wall at (0, 0) Wall at (0, 1) Wall at (0, 2) Wall at (0, 3) Wall at (0, 4) Wall at (0, 5) Wall at (1, 0) Wall at (1, 5) Wall at (2, 0) Goal at (2, 1) Wall at (2, 5) Wall at (3, 0) Box at (3, 3) Wall at (3, 5) Wall at (4, 0) Wall at (4, 1) Wall at (4, 2) Player at (4, 3) Wall at (4, 4) Wall at (4, 5) Wall at (5, 0) Wall at (5, 1) Wall at (5, 2) Wall at (5, 3) Wall at (5, 4) Wall at (5, 5) Your admissible actions are ["up", "down", "left", "right"].

Now it's your turn to make a move (choose ONE action only for the current step). You should first reason step-by-step about the current situation. This reasoning process MUST be enclosed within <think> </think> tags. Once you've finished your reasoning, you should choose an admissible action for current step and present it within <action> </action> tags.

### Example Action (with Thinking) from Sokoban

```
<think>The current situation is that I am at (4, 3) and I have a box at (3, 3) in front of me. The goal is to push the box to (2, 1) where the target is located. To achieve this, I need to move the box to the target location. I can see that the cell immediately beyond the box (in the same direction) is empty. This means that I can push the box to the target location. The current admissible actions are ["up", "down", "left", "right"]. Since the box is in front of me, I need to move to the right to push the box to the target location. </think>
<action>right</action>
```

## E.4 SCIWORLD

### Example State from SciWorld

You are an expert agent operating in the SciWorld Embodied Environment. ## Begin of an Example: Your task is to boil water. For compounds without a boiling point, combusting the substance is also acceptable. First, focus on the substance. Then, take actions that will cause it to change its state of matter. This outside location is called the outside. Here you see: the ground an axe the agent a substance called wood a fountain (containing nothing) a substance called air a fire pit (containing nothing) You also see: A door to the green house (that is open) A door to the foundry (that is open) A door to the kitchen (that is open)

<think> Going to the kitchen seems like a sensible choice because it is likely the place where I can find tools for boiling like a pot and a heat source. I am heading to the kitchen as it might contain equipment or a setup suitable for boiling water. </think><action>go kitchen</action> ## End of an Example

Now, your task is: Your task is to determine if unknown substance B is electrically conductive. The unknown substance B is located around the workshop. First, focus on the unknown substance B. If it is electrically conductive, place it in the blue box. If it is electrically nonconductive, place it in the orange box.. Prior to this step, you have already taken 7 step(s). Below are the most recent 2 observations and the corresponding actions you took: [Observation 6: 'You move the unknown substance B to the workshop.', Action 6: 'drop unknown substance in blue box'] [Observation 7: 'No known action matches that input.', Action 7: 'look around'] You are now at step 8 and your current observation is: This room is called the workshop. In it, you see: the agent a substance called air a blue box (containing nothing) a green light bulb, which is off a orange box (containing nothing) a table. On the table is: a battery, a green wire, a orange wire, a red light bulb, which is off, a red wire, a switch, which is off, a violet light bulb, which is off, a yellow light bulb, which is off. a ultra low temperature freezer. The ultra low temperature freezer door is closed. unknown substance B You also see: A door to the hallway (that is open) Your admissible actions of the current situation are: ['activate OBJ' 'close OBJ' 'connect OBJ to OBJ' 'deactivate OBJ' 'disconnect OBJ' 'dunk OBJ in OBJ' 'eat OBJ' 'flush OBJ' 'focus on OBJ' 'go OBJ' 'inventory' 'look around' 'look at OBJ' 'look in OBJ' 'mix OBJ' 'move OBJ to OBJ' 'open OBJ' 'pick up OBJ' 'pour OBJ in OBJ' 'put down OBJ' 'read OBJ' 'use OBJ on OBJ' 'wait'], where you should replace the 'OBJ' with the specific objects from the admissible objects list: ['agent' 'air' 'anode in battery' 'anode in component' 'anode in green light bulb' 'anode in red light bulb' 'anode in violet light bulb' 'anode in yellow light bulb' 'battery' 'battery cathode' 'blue box' 'cathode in component' 'cathode in green light bulb' 'cathode in red light bulb' 'cathode in violet light bulb' 'cathode in yellow light bulb' 'door' 'freezer' 'green light bulb' 'green wire' 'green wire terminal 1' 'green wire terminal 2' 'hallway' 'orange box' 'orange wire' 'orange wire terminal 1' 'orange wire terminal 2' 'red light bulb' 'red wire' 'red wire terminal 1' 'red wire terminal 2' 'switch' 'table' 'unknown substance' 'violet light bulb' 'workshop' 'yellow light bulb'].

Now it's your turn to take an action. Do remember that do NOT repeat the same action in your previous steps such as 'look around' and try to explore different actions to achieve the goal. You should first reason step-by-step about the current situation. This reasoning process MUST be enclosed within <think> </think> tags. Once you've finished your reasoning, you should choose an admissible action for current step and present it within <action> </action> tags.

### Example Action (with Thinking) from SciWorld

<think> The goal is to determine if unknown substance B is electrically conductive. The unknown substance B is located around the workshop. If it is electrically conductive, place it in the blue box. If it is electrically nonconductive, place it in the orange box. The current situation is that I have already moved the unknown substance B to the

```
workshop, and I am now in the workshop. The unknown substance B is
currently in the blue box, which means that I have already placed it
there. The current observation is that the unknown substance B is in the
blue box, and I need to determine if it is electrically conductive. To
do this, I need to check if the green light bulb is on when the switch
is activated. If it is, then the unknown substance B is electrically
conductive. If not, then it is electrically nonconductive. The current
action is to drop the unknown substance B in the blue box, which means
that I need to activate the switch to check if the green light bulb is
on. </think><action>activate switch</action>
```

## F INFORMATION AUGMENTATION DETAILS

In this section, we provide the implementation details of state augmentation. We also cover the selection of information volume for ALFWorld, Webshop, and Sokoban.

### F.1 ALFWORLD

For ALFWorld, we increase state information richness by introducing synthetic distractor objects. To introduce distractor objects that do not interfere with the original tasks, we employ a multi-step synthesis process. First, we sample an object type from a predefined list: {bowl, cup, pan, spoon, fork, knife, plate, mug, bottle, can, glass, basket, tray, vase, jar, container, jug, thermos, pitcher, dish, tumbler, sponge, towel, napkin, rack, coaster, utensil, lid}. Adhering to ALFWorld’s naming convention of (type, ID), if the sampled object type currently exists in the scene, we will assign a distinct ID to ensure uniqueness.

Finally, we sample a “trivial” object descriptor from the list: {cracked, dirty, slightly burnt, covered in dust, wrapped in foil, sticky, still warm, rusty, filled with water, emits a faint smell, has some liquid inside, tipped over, seems unused, greasy, wet on the surface, has a faint label, smudged, with faded color, shows fingerprints}. The objects and their descriptors are combined using short natural-language templates such as “You notice a ... that looks ...” or “Near the sink, there is a ... that is ...”. Such distractive sentences are inserted at random positions in the textual states (i.e., observations).

The augmentation strength is controlled by an information volume parameter  $\epsilon$ . We estimate the length of each sentence as 12 tokens and inject  $n_{\text{distractive}} = \lfloor \epsilon/12 \rfloor$  sentences per state. For example,  $\epsilon = 120$  yields about 10 extra sentences, while  $\epsilon = 360$  yields about 30.

### F.2 WEBSHOP

We consider injecting distractive information into two types of pages in WebShop, namely product-details pages and search-result pages.

On both page types, we inject irrelevant advertising information in the form of (clothing categories, promotional descriptors). For clothing categories, we sample from the following list: {rompers, jumpsuits, blouses, cozy cardigans, pajamas, loungewear sets, button-down tops, long sleeve tees, casual jackets, knitwear, soft sweaters, thermal layers, fleece pullovers, v-neck tunics, red outfits, plus-size picks, cropped hoodies, layered looks, textured knits, everyday basics}. For promotional descriptors, we sample from the following list: {30% off, new arrivals, trending now, selling fast, limited edition, back in stock, editor’s pick, members-only deal, must-have looks, staff favorite, top-rated, recommended for daily wear, cozy & cute, perfect for layering, style spotlight, hand-picked for you, fan favorites, most loved by shoppers, inspired by you, perfect for your capsule wardrobe}. Categories and descriptors are combined into short ad sentences via the template: [AD] <descriptor> --- shop <category> today!.

On product-detail pages, we introduce trivial features sampled from the following list: `{{fabric: machine wash cold}, {shipping note: ships within 5-7 days}, {sizing tip: runs half size large}, {model info: model is 5'8'' wearing size medium}, {brand description: known for cozy minimalism}, {care label: made in Indonesia}, {origin: eco-conscious packaging}, {batch code: Spring 2024 release}, {warehouse tag: popular for travel & errands}, {return policy: soft fleece interior}, {manufacturing month: delicate stitching on seams}, {recommendation: designed for colder climates}, {intended use: do not iron over print}, {layering advice: packaged folded in bag}}`.

On search-result pages, we also create trivial products by sampling non-targeted product types paired with non-targeted features. The product types include `{romper, jumpsuit, blouse, tunic, pajama top, sweater, tank top, hoodie, cardigan, bodysuit, thermal shirt, dress, coverall, camisole}`. Non-targeted features include `{sleeveless, misleadingly tagged, overpriced, heavily discounted, out of stock, not available in large size, no reviews, size runs small, fabric mismatch, zipper broken, category mislabel, style mismatch, faded color, missing product images, customer complaints}`. The information is presented with templates such as `Women <desc> <obj> --- likely unsuitable or Warning: <desc> <obj> (cannot be selected)`.

The information volume  $\epsilon$  determines how many of such units are inserted. On each result page, we allocate up to 10 additional item slots and fill  $k = \min(\lfloor (\epsilon/100) \alpha \cdot 10 \rfloor, 10)$  of them with a mix of advertisements and trivial products (e.g., with  $\epsilon = 100$  and  $\alpha = 0.5$ , we insert  $k = 5$  entries). On each detail page, we append  $k = \max(1, \lfloor \epsilon/25 \rfloor)$  sentences about product features, and  $k = \max(1, \lfloor \epsilon/30 \rfloor)$  sentences about advertisements.

### F.3 SOKOBAN

**Augmentation Implementation** For Sokoban, the augmentation operates at the level of the textual description of the grid. Given the true room dimensions  $H \times W$ , we sample coordinates that lie strictly outside the truly reachable locations and associate each with a random object drawn from the list: `{Broken crate, Sealed door, Pile of rubble, Faded sign, Leaking pipe, Loose tile, Stacked boxes}`. The locations are described using phrases from `{outside the room, beyond the wall, inaccessible, behind a collapsed wall, on a blocked corridor}`.

These elements are combined with the template `(r, c) shows a <obj> (<desc>; unreachable) ..`. The final sentences are appended as additional lines after the original textual state. Here, the information volume  $\epsilon$  directly controls the number of appended sentences: we estimate the length of each sentence as 10 tokens and generate  $k = \max(1, \lfloor \epsilon/10 \rfloor)$  distractor lines (e.g.,  $\epsilon = 50$  yields about 5 lines of unreachable locations, while  $\epsilon = 150$  yields about 15). Again, all distractors are explicitly placed outside the grid.

### F.4 INFORMATION AUGMENTATION VOLUME SELECTION

In Tab. 11, we report the information augmentation volume  $\epsilon$  used for the experiments in Section 6.

Table 11: Information augmentation volume  $\epsilon$  used for the experiments in Section 6.

| <b>Initial Policy</b> | <b>Training Domain</b> | <b>Information Augmentation Volume <math>\epsilon</math></b> | <b>Probability of Applying Augmentation</b> |
|-----------------------|------------------------|--|---|
| Ckpt V1               | ALFWorld               | 300  | 100%  |
| Ckpt V1               | Webshop                | 30   | 100%  |
| Ckpt V1               | Sokoban                | 80   | 100%  |
| Ckpt V2               | ALFWorld               | 200  | 50%   |
| Ckpt V2               | Webshop                | 40   | 50%   |
| Ckpt V2               | Sokoban                | 80   | 50%   |