

Enhancing Agent Learning through World Dynamics Modeling

Zhiyuan Sun^{1*} †, Haochen Shi^{1*}

Marc-Alexandre Côté², Glen Berseth^{1,3}, Xingdi Yuan^{2‡}, Bang Liu^{1,3‡}

¹ Université de Montréal & Mila, Montréal, Canada

² Microsoft Research, Montréal, Canada

³ Canada CIFAR AI Chair

Abstract

Large language models (LLMs) have been increasingly applied to tasks in language understanding and interactive decision-making, with their impressive performance largely attributed to the extensive domain knowledge embedded within them. However, the depth and breadth of this knowledge can vary across domains. Many existing approaches assume that LLMs possess a comprehensive understanding of their environment, often overlooking potential gaps in their grasp of actual world dynamics. To address this, we introduce **Discover, Verify, and Evolve (DiVE)**, a framework that **discovers** world dynamics from a small number of demonstrations, **verifies** the accuracy of these dynamics, and **evolves** new, advanced dynamics tailored to the current situation. Through extensive evaluations, we assess the impact of each component on performance and compare the dynamics generated by DiVE to human-annotated dynamics. Our results show that LLMs guided by DiVE make more informed decisions, achieving rewards comparable to human players in the Crafter environment and surpassing methods that require prior task-specific training in the MiniHack environment¹.

1 Introduction

By absorbing internet-scale knowledge autoregressively, large language models (LLMs) develop a broad understanding of the world (Achiam et al., 2023; Team et al., 2023; Brown et al., 2020; Touvron et al., 2023). This understanding enables them to perform well across a variety of tasks (Brown et al., 2020; Huang et al., 2022; Yao et al., 2022, 2024). However, achieving this level of comprehension requires a training dataset that is diverse,

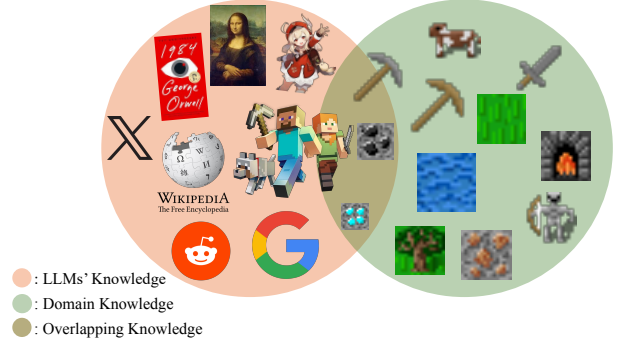


Figure 1: The knowledge gap between LLMs and downstream domains. Although LLMs have a broad understanding of the world, they may struggle to grasp the complex dynamics of specific downstream domains.

precise, and in-depth to cover essential domain information. Otherwise, a knowledge gap may emerge between LLMs and the target domain, as illustrated in Figure 1.

Since LLMs tend to follow the most common patterns in the dataset (Gunasekar et al., 2023) and are prone to hallucination, it is crucial that the necessary information not only appears frequently on the internet but is also reliable (Dziri et al., 2022). For tasks and environments that are newly emerging (Hafner, 2021; Samvelyan et al., 2021), these requirements are often difficult to meet due to the noisy nature of internet data, leading to potential knowledge gaps.

While LLMs may have a general understanding of a domain, optimal decision-making requires in-depth, state-specific knowledge. Providing LLMs with such tailored information enhances their grasp of both the environment and the current state. For example, mastering Go involves not just knowing the rules but also applying strategies for specific board states. This specialized knowledge varies across domains and states, making offline collection impractical.

To address these challenges, we propose **Discover, Verify, and Evolve (DiVE)**. Building on

* Equal contribution.

† zhiyuan.sun@umontreal.ca

‡ Equal advising.

¹ The code is available at <https://github.com/ZhiyuanS/DiVE>

the concept of the World Model (Ha and Schmidhuber, 2018), DiVE not only discovers and verifies primitive world dynamics from demonstrations but also evolves state-specific knowledge for the downstream domain. By providing LLMs with this comprehensive set of dynamics, DiVE bridges the knowledge gap between LLMs and the downstream domain, thereby enhancing their decision-making abilities.

DiVE consists of three distinct components: The **Discoverer**: this component iteratively uncovers the environment’s dynamics from demonstrations using a curriculum learning approach. The **Verifier**: this component eliminates unreliable dynamics caused by LLMs’ tendency to hallucinate. The **Evolver**: this component reasons through in-depth, state-specific strategies tailored to the current situation based on the learned dynamics.

In the Crafter environment (Hafner, 2021) and the MiniHack environment (Samvelyan et al., 2021), DiVE learns comprehensive and reliable dynamics from demonstrations, guiding the agent’s decision-making process by evolving in-depth strategies. This enables the agent to outperform all baselines, achieving rewards comparable to human players in the Crafter environment and surpassing the performance of methods requiring task-specific training in the MiniHack environment. To gain deeper insights into DiVE’s behavior, we provide both quantitative and qualitative analyses.

In summary, our primary contribution is a framework that learns world dynamics from demonstrations, guiding LLMs in the decision-making process by online evolving contextual strategies. This approach bridges potential knowledge gaps, resulting in a more optimal decision-making process for LLMs.

2 The Knowledge Gap

In this work, we consider a setting where a pre-trained LLM-based agent is employed to solve tasks in downstream domains. Conceptually, we define \mathcal{K}_{LLM} as the set of knowledge embedded in the LLMs through their training process, and $\mathcal{K}_{\text{target}}$ as the universal set of knowledge relevant to the downstream domains.

To ensure the effectiveness of the LLMs, we aim for $\mathcal{K}_{\text{relevant}}$, the subset of \mathcal{K}_{LLM} that is relevant to $\mathcal{K}_{\text{target}}$, to cover the broadest possible scope. Additionally, we seek for $\mathcal{K}_{\text{relevant}}$ to contain more reliable knowledge \mathcal{K}_+ than inaccurate knowledge

\mathcal{K}_- , where $\mathcal{K}_{\text{relevant}} = \mathcal{K}_+ \cup \mathcal{K}_-$. Thus, we define three desirable properties to guide our system design:

- **Recall** $R = \frac{|\mathcal{K}_+|}{|\mathcal{K}_{\text{target}}|}$ measures the extent to which the knowledge required to solve tasks in the target domain is covered by the LLMs. A low recall typically indicates a significant knowledge gap between the LLMs’ training data and the target domain.
- **Precision** $P = \frac{|\mathcal{K}_+|}{|\mathcal{K}_{\text{relevant}}|}$ measures the accuracy of the LLMs’ knowledge when applied to the target domain. An example of inaccurate dynamics is shown in Figure 2 (left), where the LLMs hallucinate that defeating a skeleton will drop items, which is not true in Crafter. This may be due to the LLMs memorizing Minecraft-specific data.
- **Depth** D measures the abstraction levels of knowledge representations. Knowledge can be represented at varying levels of abstraction, from basic game rules to higher-level strategies.

We recognize that precisely quantifying the knowledge overlap between LLMs and a downstream domain is challenging. In Section 3, we provide a mathematical analysis showing how DiVE improves across all three dimensions. In Section 5.4, we measure the knowledge overlap to demonstrate that DiVE effectively learns useful game dynamics and, to some extent, bridges the knowledge gap.

3 DiVE: Discover, Verify, and Evolve

In an ideal scenario, one could bridge the knowledge gap by fine-tuning the LLM to adapt to the target domain. However, this approach is often less practical due to its reliance on large amounts of annotated data and significant computational overhead (Hu et al., 2021; Zheng et al., 2024; Carta et al., 2023; Ouyang et al., 2022). Our framework, DiVE, is designed to address the knowledge gap while considering all three desirable properties—recall, precision, and depth—without requiring extensive data collection from the target domain. It is a prompt-based method that learns world dynamics \mathcal{W} directly from the environment.

As shown in Figure 2, DiVE is initially bootstrapped using a small set of human demonstration trajectories H , each consisting of observations o_t , actions a_t , and rewards r_t at each timestep t . We then transform each observation o_t into

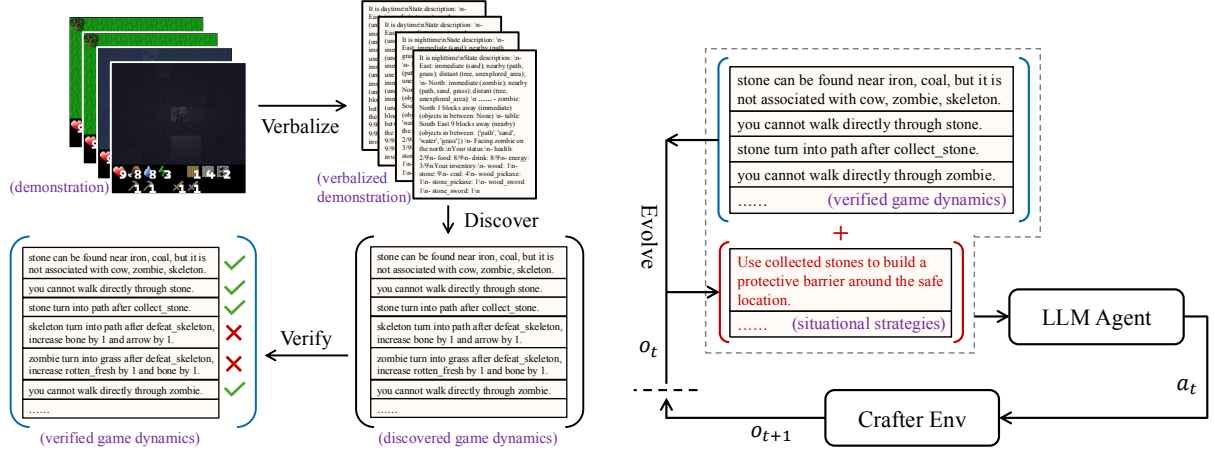


Figure 2: Overall pipeline of DiVE. **Left:** Learning basic game dynamics from offline demonstrations (Section 3.1). We want to highlight the incorrect game dynamics being identified by the Verifier (labeled by \times), they are evidence of the LLMs hallucinate false facts perhaps because of memorizing Minecraft data. **Right:** Learning situational strategies from online interactions (Section 3.2). For simplicity, we omit the verbalization process in the right figure.

the language space as \tilde{o}_t using a Verbalizer, resulting in trajectories represented by \tilde{o}_t , a_t , and r_t at each timestep t . Subsequently, the **Discoverer** extracts a set of world dynamic candidates, $\tilde{\mathcal{W}} = \{\tilde{\mathcal{W}}_+, \tilde{\mathcal{W}}_-\}$, from human demonstrations H , where $\tilde{\mathcal{W}}_+$ and $\tilde{\mathcal{W}}_-$ represent the correct and inaccurate world dynamic sets, respectively.

Empirically, we find that the inclusion of $\tilde{\mathcal{W}}_-$ in $\tilde{\mathcal{W}}$ is often unavoidable, either due to LLM’s difficulties in extracting meaningful knowledge from trajectory data or its tendency to hallucinate. To address this, we employ the **Verifier** to filter out potentially invalid and conflicting world dynamic candidates from $\tilde{\mathcal{W}}$, leaving only the valid dynamics \mathcal{W} . Lastly, we use the **Evolver**, designed to derive advanced dynamics \mathcal{I} tailored to the verbalized observation \tilde{o}_t based on the filtered world dynamics \mathcal{W} .

The final decision-making process on primitive actions $a_t \in \mathcal{A}$ is hierarchically decomposed as planning tasks on sub-goals SG , sub-tasks ST , and actions \mathcal{A} . The planning procedure is further guided by both \mathcal{W} and \mathcal{I} . In cases where $\mathcal{W} \neq \emptyset$, Recall, Precision and Depth are guaranteed to increase as formulated below:

$$\begin{aligned} \text{Recall} &: \frac{|\mathcal{K}_+|}{|\mathcal{K}_{\text{target}}|} \xrightarrow[H]{\text{Discoverer}} \frac{|\mathcal{K}_+| + |\mathcal{W}|}{|\mathcal{K}_{\text{target}}|} \\ \text{Precision} &: \frac{|\mathcal{K}_+|}{|\mathcal{K}_{\text{relevant}}|} \xrightarrow[\tilde{\mathcal{W}}, H]{\text{Verifier}} \frac{|\mathcal{K}_+| + |\mathcal{W}|}{|\mathcal{K}_{\text{relevant}}| + |\mathcal{W}|} \\ \text{Depth} &: \emptyset \xrightarrow[\text{Verifier}]{\text{Discoverer}} \mathcal{W} \xrightarrow[\mathcal{W}, H]{\text{Evolver}} \mathcal{I} \cup \mathcal{W} \end{aligned}$$

The DiVE framework can be divided into two stages: an offline dynamics learning stage and an

online strategy learning stage.

3.1 Offline Dynamics Learning

The offline dynamics learning procedure aims to bridge the knowledge gap between LLMs’ understanding and the basic dynamics of downstream domains by learning the world dynamics \mathcal{W} as a prior for the decision-making process. Instead of relying on human-authored game manuals or handbooks to extract world dynamics, as in (Wu et al., 2024b,a), which are not only difficult to obtain in many real-world scenarios but also often lack critical details (as demonstrated in Table 6), we propose learning world dynamics \mathcal{W} directly from experiences H , which are more accessible and provide richer information.

Hierarchical Curriculum Learning Given the varying complexities of learning the dynamics of different elements in downstream domains, we adopt a curriculum learning approach (Bengio et al., 2009). Our method follows a sequential learning strategy that progresses from simpler to more complex dynamics, thereby enabling more effective learning. Specifically, we propose a method for learning the dynamics of each element within the task decomposition hierarchy, denoted as $TD = \{\mathcal{A} \cup O, ST, SG\}$, where O represents the set of objects in downstream domains.

Our approach begins with elements of lower abstraction, such as actions $a \in \mathcal{A}$ and objects $o \in O$, and gradually progresses to higher-level elements, such as sub-tasks $st_i \in ST$. The sub-tasks ST are represented as nodes in the achieve-

ment graph $\mathcal{G} = (V, E)$ within the downstream domains, i.e., $ST = V$. Finally, we transition to the subgoals $sg_i \in SG$. The subgoal sequence $SG = [sg_1, sg_2, \dots]$ is an ordered list used to unlock achievements in the achievement graph \mathcal{G} , where $SG = \text{TopologicalSort}(\mathcal{G})$, and each sg_i corresponds to a vertex in V .

We leverage the Discoverer to extract this order from human demonstrations H . Achieving a subgoal sg_i may involve completing several sub-tasks multiple times. This approach ensures a logical progression through tasks, thereby enabling a deeper understanding and integration of downstream domain dynamics.

Discoverer The Discoverer is designed to identify dynamic candidates $\tilde{\mathcal{W}}$ related to elements within the task decomposition hierarchy TD . A single dynamics discovery step for an element $E \in TD$ involves the following three main steps:

1. **Construction of the Semantic Experience Bank:** For each element E , we construct a semantic experience bank B^E using demonstrations H . This bank stores experiences that are transformed from H into a suitable granularity for analyzing dynamics related to E . The transformation process involves chunking and summarizing the verbalized demonstrations to capture essential semantic details.
2. **Sampling of Relevant Experiences:** For each attribute of an instance $e \in E$, a subset of experiences B_e^E that are relevant to the instance e is sampled from B^E .
3. **Identification of Dynamic Candidates:** A dynamic candidate \tilde{w} is identified from the subset B_e^E by recognizing patterns that are consistent across all experiences within B_e^E .

The action-level semantic experience bank, denoted as B^A , stores transition tuples derived from verbal demonstrations and is represented as: $B^A = \{\{\tilde{o}_t, a_t, \tilde{o}_{t+1}\}_i\}_{i=1}^{|B^A|}$. Similarly, the object-level semantic experience bank, denoted as B^O , gathers individual observations contain an specific object and is represented as: $B^O = \{\{\tilde{o}_i\}_{i=1}^{|B^O|}\}$. The sub-task-level semantic experience bank, denoted as B^{ST} , aggregates trajectory segments representing the completion of sub-tasks and is formatted as: $B^{ST} = \{\{\tilde{o}_t, \dots, a_{t_{st}}, \tilde{o}_{t_{st}+1}\}_i\}_{i=1}^{|B^{ST}|}$, where t_{st} denotes the timestep at which a sub-task $st \in ST$ is completed. For sub-goals, the

sub-goal-level experience bank, denoted as B^{SG} , records sequences of sub-tasks that culminate in the completion of sub-goals, expressed as: $B^{SG} = \{\{\{st_t, \dots, st_{t_{sg}}\}_i\}_{i=1}^{|B^{SG}|}\}$, where t_{sg} is the timestep at which the sub-goal sg is achieved.

For action-level dynamics learning, the relevant experiences, denoted as B_a^A , are compiled by randomly sampling transition tuples from B^A where the action a has been successfully executed. A similar approach is employed for dynamics learning of other elements within the task decomposition hierarchy TD .

For action-level dynamics discovery, we identify the prerequisites and outcomes of each action (e.g., *The action MakeWoodPickaxe requires 1 wood*). For object-level dynamics, we focus on co-occurrence relationships between objects and their temporal patterns. The attribute set for a sub-task generally encompasses the necessary steps for completion, as well as its prerequisites, outcomes, and termination conditions. In contrast, the primary attribute of interest for a sub-goal is its correct position within the sub-goal sequence SG .

Verifier Dynamic discovery processes are vulnerable to noise from various sources, including confounders, hallucinations by the LLMs, and difficulties in the LLMs' ability to extract meaningful insights from trajectory data. To address these challenges, we introduce a dynamic verifier designed to filter out noisy dynamic candidates from $\tilde{\mathcal{W}}$. For each dynamic candidate $w \in \tilde{\mathcal{W}}$, the verifier begins by sampling a subset of relevant semantic experiences, denoted as B_e^E , from the corresponding semantic experience bank B_E . Here, w represents a dynamic candidate associated with a specific attribute of the instance $e \in E$, where $E \in \{\mathcal{A} \cup \mathcal{O}, ST, SG\}$ corresponds to an element of the task decomposition hierarchy TD . The verification of w is conducted as follows: w is deemed inaccurate and filtered out if it does not consistently hold across experiences within B_e^E or if it conflicts with any established dynamics. The dynamics that pass this verification process are classified as verified dynamics and are denoted as \mathcal{W} .

3.2 Online Strategy Learning

To effectively incorporate the learned world dynamics \mathcal{W} into the downstream domains, we deploy an LLM-based agent defined by $\pi : \mathcal{S} \times \mathcal{W} \rightarrow \mathcal{P}(\mathcal{A})$. Here, \mathcal{S} represents the state space, \mathcal{A} denotes the action space, and \mathcal{P} symbolizes the probability dis-

tribution over the action space. Instead of directly mapping the world dynamics \mathcal{W} and the current state observation o_t to the action a_t , we tackle the challenge of long-horizon planning by integrating an online strategy learning method. This approach decomposes the planning process into three distinct stages: sub-goal planning, sub-task planning, and action planning.

Sub-goal Planning Given that the sub-goal sequence $SG = [sg_1, sg_2, \dots]$ is derived from human demonstrations H and treated as a fixed sequence, we utilize a straightforward heuristic for sub-goal planning. When a sub-goal is completed, the current sub-goal is updated to the first uncompleted sub-goal in SG .

Sub-task Planning For a given current sub-goal sg_i , we have developed an LLM-based sub-task planner. This planner evaluates and ranks all sub-tasks $st \in ST$ based on the learned world dynamics \mathcal{W} , the verbalized current observation \tilde{o}_t , and the most recently executed sub-task st_{t-1} . The highest-ranked sub-task is then designated as the current sub-task st_t . To ensure accurate execution, the completion of a sub-task st is contingent upon satisfying its specific termination condition. This condition is verified by querying an LLM using the current verbalized observation, the observation at the time the sub-task began, and the termination conditions of the current sub-task.

Learning Strategies In addition to learning the fundamental rules of the downstream domains, we also focus on developing advanced game-playing strategies based on these dynamics. Unlike world dynamics learning, the strategy space is often too expansive for exhaustive exploration. To address this challenge, we propose evolving the dynamics into strategies, denoted as \mathcal{I} , using an online learning approach.

This method reduces the search space by conditioning not only on the dynamics \mathcal{W} , but also on the verbalized current observation \tilde{o}_t and the sub-task st_i . This targeted approach enables the generation of strategies that are more contextually grounded and responsive to current game scenarios compared to those developed through offline methods. To support this process, we have designed an LLM-based Evolver that generates strategy candidates $\tilde{\mathcal{I}}$ by applying deductive reasoning to the learned dynamics \mathcal{W} . Specifically, the Evolver derives strategy candidates using rules of inference, such as modus

ponens. These strategy candidates, denoted as $\tilde{\mathcal{I}}$, are then evaluated for validity and ranked based on their utility by an LLM-based critic. Finally, the valid and useful candidates are incorporated into the situational strategy set \mathcal{I} .

Action Planning The final action selection process is executed in two main steps:

1. **Invalid Action Masking:** This step involves filtering out actions that are infeasible under the current situation, based on the verified dynamics \mathcal{W} and current verbalized observation \tilde{o}_t .
2. **Action Selection:** From the set of valid actions, a specific primitive action a is chosen based on multiple factors: the current sub-task st_i , the verbalized current observation \tilde{o}_t , the world dynamics \mathcal{W} , a windowed history of previously planned actions and observations, and the derived strategies \mathcal{I} .

4 Experiment Setup

To demonstrate the effectiveness of DiVE in bridging the knowledge gap, we evaluate its performance within the Crafter and MiniHack environments. For a fair comparison, all LLM-based agents, including ours, utilize the GPT-4o model and the same environment seed in the Crafter setting. Further details on the setups can be found in Appendix A.1 and B.1.

4.1 Crafter

Crafter (Hafner, 2021) is an open-world survival game set on 64×64 grid-based maps, featuring a diverse array of materials such as `tree`, `stone`, and `coal`, as well as entities including `cow`, `zombie`, and `skeleton` semi-randomly spawn on the maps. The games include an achievement graph with 22 unique achievements across 7 levels. The agent perceives its surroundings through a local 7×9 observation window and maintains awareness of its status within the game environment.

The text description generated by the verbalizer includes: the nearest object of each type within the accumulated observations, the objects situated between these nearest objects, the objects in each direction, as well as the agent’s current inventory and status. An example of the verbalization process is provided in Appendix B.5.

The agent is evaluated using two primary metrics: reward and score. Agents receive a +1 reward for each new achievement unlocked (e.g., `make`

wood pickaxe, place furnace) and a ± 0.1 reward for every health point gained or lost. The score is calculated by aggregating the success rates s_i for each achievement and is formulated as:

$$S \doteq \exp \left(\frac{1}{N} \sum_{i=1}^N \ln(1 + s_i) \right) - 1.$$

We compare DiVE against the following baselines:

- **LLM-based approaches:** SPRING (Wu et al., 2024b), ELLM (Du et al., 2023), and Chain-of-Thought (CoT) (Wei et al., 2022).
- **Reinforcement Learning (RL) approaches:** DreamerV3 (Hafner et al., 2023), PPO (Schulman et al., 2017), and AD (Moon et al., 2024).
- **Human Players:** Expert performance on the Crafter environment.

4.2 MiniHack

MiniHack (Samvelyan et al., 2021) is a grid-based environment built on the video game NetHack (Kütler et al., 2020). Unlike the Crafter environment, it supports the creation of tasks that target specific agent capabilities. In our work, we focus on the Skill Acquisition Tasks subset, which evaluates the agent’s ability to leverage the rich diversity of NetHack’s objects, monsters, and dungeon features, as well as their interactions. These tasks introduce complex world dynamics, where actions are factorized autoregressively and require executing a sequence of follow-up actions for the initial action to produce the desired effect.

We use the Lava Crossing, Wand of Death, and Quest tasks as testbeds to evaluate DiVE. We provide detailed descriptions of these task in Appendix A.2. The agent is rewarded for unlocking achievements within each task. To facilitate interaction with LLMs, we use the NLE language wrapper (Goodger et al., 2023) to verbalize both observations and actions from the environment.

We compare DiVE against SSO (Nottingham et al., 2024) and Reflexion (Shinn et al., 2024), both of which require prior training on the tasks. Following the official guidelines, we run each method for 30 iterations, evaluating performance after every 10 iterations by attempting the task 10 times using a fixed set of skills or reflections. The final performance is determined based on the results from the last 10 evaluation attempts.

Method	Score	Reward
Human Experts	50.5 \pm 6.8%	14.3 \pm 2.3
DiVE	35.9 \pm 3.2%	14.5 \pm 2.4
SPRING*	8.2 \pm 2.3%	6.9 \pm 1.8
CoT	1.3 \pm 0.3%	2.5 \pm 0.5
AD	21.79 \pm 1.4%	12.6 \pm 0.3
ELLM	N/A	6.0 \pm 0.4
DreamerV3	14.5 \pm 1.6%	11.7 \pm 1.9
PPO	4.6 \pm 0.3%	4.2 \pm 1.2
Random	1.6 \pm 0.0%	2.1 \pm 1.3

Table 1: Performance comparison of DiVE against baseline models in the Crafter environment. Methods with * indicate that they were obtained using the official code implementation, executed with the same five random seeds and model configurations as used for DiVE.

5 Experimental Results

We evaluate the performance of DiVE in the Crafter and MiniHack environments. In Section 5.1, we present the overall results to demonstrate DiVE’s effectiveness in bridging the knowledge gap in these tasks. Section 5.2 provides a detailed analysis of the contributions of individual components through controlled experiments. Then, in Section 5.3, we evaluate the effectiveness of the dynamics learned by DiVE in the Crafter environment. Finally, in Section 5.4, we further analyze the learned dynamics in both the Crafter environment and a modified MiniHack setting with altered dynamics.

5.1 DiVE’s Performance

Table 1 and 2 demonstrate that DiVE surpasses all other baselines in the Crafter and MiniHack environments. In the Crafter environment, DiVE exceeds the previous state-of-the-art (SOTA) LLM-based method, SPRING, by a substantial margin, achieving a 337.8% relative improvement in score and a 110.1% enhancement in reward. Additionally, DiVE also surpasses the prior SOTA RL-based approach, DreamerV3, with a 21.4% absolute improvement in score and a 2.8 absolute increase in reward. Notably, DiVE achieves rewards comparable to human players using 10 demonstrations.

In the MiniHack environment, with only a single demonstration, DiVE matches the performance of SSO and Reflexion (both of which require 30 iterations of training) on the Lava Crossing task, and outperforms both baselines on the Wand of Death and Quest tasks. Specifically, DiVE achieves a 68% improvement over Reflexion on the Wand of Death task and a 30% and 36% improvement over SSO and Reflexion, respectively, on the Quest task.

Method	Lava Crossing	Wand of Death	Quest
Reflexion	0.8 ± 0.0	0.32 ± 0.15	0.87 ± 0.43
SSO*	0.8 ± 0.0	0.52 ± 0.31	0.91 ± 0.29
DiVE	0.8 ± 0.0	0.54 ± 0.34	1.18 ± 0.38

Table 2: Comparison of DiVE performance against baseline models in the MiniHack environment. Results marked with * were obtained using the official code implementation and evaluated over 10 runs.

Methods	Score	Reward
Component analysis		
DiVE	$35.9 \pm 3.2\%$	14.5 ± 2.4
w/o E	$21.1 \pm 9.7\%$	11.3 ± 4.3
w/o V	$9.8 \pm 1.0\%$	10.1 ± 0.7
w/o V&E	$11.5 \pm 4.9\%$	8.3 ± 3.8
w/o D&V&E	$0.9 \pm 0.1\%$	2.5 ± 1.3
CoT	$1.3 \pm 0.3\%$	2.5 ± 0.5
CoT + D&V	$3.6 \pm 0.9\%$	3.9 ± 2.3
Dynamics from distinct sources		
DiVE	$35.9 \pm 3.2\%$	14.5 ± 2.4
w/o D&V+S [†]	$15.7 \pm 5.3\%$	8.9 ± 5.1
w/o D&V&E+S [†]	$12.1 \pm 4.6\%$	8.7 ± 3.0
w/o D&V&E+H [†]	$34.2 \pm 2.8\%$	14.5 ± 0.9

Table 3: Impact of different components on performance: Crafter. D, V, and E represent Discover, Verifier, and Evolver, respectively; S[†] refers to dynamics derived from the game manual, H[†] refers to human-annotated dynamics.

5.2 Contribution of Individual Components

We conduct a series of ablation studies to clarify the contribution of each individual element to DiVE’s overall performance. We report the results on Crafter in the first section of Table 3, and the results on MiniHack in Table 4.

Crafter The significant performance gap between DiVE and its variant without the Evolver component empirically demonstrates the Evolver’s effectiveness in developing gameplay strategies based on world dynamics \mathcal{W} , thereby enhancing the agent’s overall proficiency in this environment. Similarly, the performance decline observed in the variant without the Verifier underscores the importance of formulating strategies \mathcal{I} based on accurate world dynamics \mathcal{W} . Moreover, the further performance drop in the version lacking both the Verifier and Evolver components highlights their complementary roles—the Verifier ensures precision in capturing dynamics, while the Evolver focuses on strategy development.

The performance of DiVE without the Discoverer, Verifier, and Evolver components reverts to the CoT baseline, indicating that simply decomposing the task according to the hierarchy \mathcal{H} without integrating domain knowledge $\mathcal{K}_{\text{target}}$ provides

Components	Lava Crossing	Wand of Death	Quest
DiVE	0.80 ± 0.00	0.54 ± 0.34	1.18 ± 0.38
w/o E	0.72 ± 0.24	0.52 ± 0.30	1.16 ± 0.48
w/o V	0.69 ± 0.25	0.52 ± 0.38	1.18 ± 0.38
w/o V&E	0.72 ± 0.24	0.40 ± 0.40	1.11 ± 0.48
w/o D&V&E	0.72 ± 0.17	0.20 ± 0.00	0.41 ± 0.29

Table 4: Impact of different components on performance: MiniHack. D, V, and E represent Discover, Verifier, and Evolver, respectively.

no performance benefit. The substantial gap between CoT + D&V and DiVE w/o E further demonstrates that an LLM-based agent struggles with long-horizon planning tasks in the absence of task decomposition, underscoring the importance of the decomposition hierarchy \mathcal{H} .

MiniHack Since MiniHack is built on the popular video game NetHack (Goodger et al., 2023), LLMs already possess a certain level of understanding of the environment. For example, DiVE without the Discover, Verifier, and Evolver components (CoT baseline) can successfully solve the Lava Crossing task. However, as the complexity of the tasks increases, using LLMs alone is insufficient for solving more challenging scenarios like the Wand of Death and Quest tasks, highlighting a potential knowledge gap for these more intricate problems.

The primary performance gains of DiVE in the MiniHack environment can be attributed to the Discover component, which accurately identifies the underlying environment dynamics required for each task. Due to the embedded knowledge of LLMs about the MiniHack environment, they can identify these dynamics effectively, as shown by the minimal performance drop when DiVE is used without the Verifier. However, in the modified MiniHack setting described in Section 5.4, we highlight the critical role of the Verifier in maintaining reliability. Additionally, since the tasks have a relatively short horizon, the dynamics uncovered by the Discover component are sufficient to complete them successfully.

5.3 Evaluation of Learned Dynamics

We investigate the performance of DiVE leveraging world dynamics derived from different sources in Crafter. As shown in the second section of Table 3, DiVE significantly outperforms variants that utilize the dynamics S^{\dagger} from the game manual (Wu et al., 2024b). This performance improvement indicates that the learned dynamics \mathcal{W} are more ad-




Figure 3: Recall of learned dynamics over discovery steps, presented with mean and standard deviation, in the Crafter environment.

vantageous than S^\dagger , likely because S^\dagger lacks certain beneficial details that are captured in \mathcal{W} . The performance gap between methods using S^\dagger with and without the Evolver further highlights the importance of strategy evolution, whose effectiveness is closely tied to the quality of the underlying world dynamics.

In addition to dynamics learned from human demonstrations and game manual, we have explored a third source: human-annotated dynamics. The results show that DiVE performs comparably to the variants using human-annotated dynamics, demonstrating the robustness and effectiveness of DiVE’s approach to dynamic learning.

As previously mentioned, it is difficult to quantify the desired properties because we cannot precisely measure the domain-relevant information $\mathcal{K}_{\text{relevant}}$ in LLMs or the exact amount of domain knowledge $\mathcal{K}_{\text{target}}$ required. However, by using human-annotated dynamics H^\dagger as a reference benchmark for $\mathcal{K}_{\text{target}}$, we can estimate the precision and recall of the learned dynamics \mathcal{W} , enabling us to effectively assess the progress of LLM-based dynamic learning. Specifically, we define recall as $R = \frac{|\mathcal{W} \cap H^\dagger|}{|H^\dagger|}$ and precision as $P = \frac{|\mathcal{W} \cap H^\dagger|}{|\mathcal{W}|}$.

As illustrated in Figure 3, both the discovered dynamics $\tilde{\mathcal{W}}$ and the verified dynamics \mathcal{W} exhibit an increase in recall as the number of discovery steps progresses, indicating that the richness of the learned dynamics improves over time. Moreover, the narrowing gap in recall between $\tilde{\mathcal{W}}$ and \mathcal{W} suggests that the Verifier effectively filters out ‘noisy’ dynamic candidates while preserving those that generalize across different trajectory segments.

To evaluate whether the Verifier preserves correct world dynamic candidates $\tilde{\mathcal{W}}+$ while filtering out unreliable ones $\tilde{\mathcal{W}}-$, we analyze the precision

Correctness	Outcome
Correct (\checkmark)	none
Confounder (\times)	1 health
In-domain Hallucination (\times)	1 wood
Out-domain Hallucination (\times)	1 bone

Table 5: The dynamics underlying the outcome of defeating zombie in the Crafter environment.

of both the discovered dynamics $\tilde{\mathcal{W}}$ and the verified dynamics \mathcal{W} . As shown in Figure 4, the precision of the verified dynamics consistently and significantly exceeds that of the discovered dynamics, demonstrating the Verifier’s effectiveness in identifying and eliminating inaccurate candidates. This confirms the Verifier’s role in enhancing the reliability of the dynamics used for decision-making.

5.4 Analysis of Learned Dynamics

The correctness of the learned and verified dynamics is classified as either correct or erroneous, with errors stemming from confounders, in-domain hallucinations, or out-of-domain hallucinations. As shown in Table 5, an example of a confounder-related mistake occurs in Crafter when a simultaneous increase in health points is incorrectly attributed to the act of defeating a zombie. In this scenario, the Discoverer misclassifies the health increase as a direct result of defeating the zombie. In the case of in-domain hallucinations, the Discoverer incorrectly associates an increase in wood with defeating the zombie, even though it is not possible for wood to increase during this event, despite its presence in the observation. Lastly, out-of-domain hallucinations involve the discovery of dynamics that reference nonexistent objects in the observation or elements not present in the Crafter environment.

Compared to the dynamics from the game manual, as shown in Table 6, we found that DiVE’s dynamics are not only more precise but also more detailed. For instance, while SPRING only identified that placing a stone requires stones, DiVE determined that it specifically requires exactly one stone and the precise facing condition needed for successful placement. Moreover, using this information, the Evolver can infer advanced dynamics for placing a stone, such as its potential to serve as a barrier between the agent and dangerous creatures.

In a customized MiniHack Lava Crossing setting, we modified the task such that it requires the agent

Sources	Dynamics
Manual	Place stone requires stones
Discoverer	Place stone requires 1 stone and faces paths, grass, sand, water, and lava
Evolver	Place stone to block zombies and skeletons, preventing them from reaching the player

Table 6: Comparing SPRING and DiVE on place stone’s dynamics in the Crafter environment.

to use the Wand of Death to freeze the lava for crossing, even though the wand’s original purpose is to zap monsters. Based on the demonstration, the Discoverer mistakenly identifies the precondition for using the wand as **The player must be adjacent to a wall or obstacle that can be destroyed or altered by the wand of death, with the outcome being The wand of death was used, altering the environment by removing walls and revealing a dark area to the east, southeast, and south.** This unreliable dynamic is successfully identified and filtered out by the Verifier, highlighting its crucial role in maintaining the reliability of learned dynamics.

6 Related Work

Language Models Language models (LLMs) are trained autoregressively in a left-to-right sequence, predicting each token based on its preceding context from an internet-scale corpus. Through this training, LLMs develop a comprehensive understanding of both language and the world it represents (Achiam et al., 2023; Touvron et al., 2023; Dubey et al., 2024; The, 2024), enabling them to perform competently across a wide range of tasks (Yao et al., 2024; Shinn et al., 2024).

Embodied Agent Building an embodied agent using LLMs is challenging because LLMs lack embodied experience (Valmeekam et al., 2022; Wang et al., 2024) in downstream environments (Weir et al., 2022; Shridhar et al., 2020; Côté et al., 2019). However, LLMs can still provide foundational world knowledge that serves as a prior for the agent (Shi et al., 2024; Colas et al., 2023; Zhong et al., 2024; Fu et al., 2024a). A notable distinction of our work is that we do not assume LLMs have the necessary knowledge to solve tasks in specific domains.




Figure 4: Precision of learned dynamics before and after verification in the Crafter environment.

Discover Dynamics LLMs can discover knowledge by inducing and deducing rules for reasoning tasks (Zhu et al., 2023) and by extracting underlying domain knowledge from prior trajectories and interactive experiences (Colas et al., 2023; Majumder et al., 2023; Fu et al., 2024b; Zhao et al., 2024). However, the knowledge obtained through these methods is often unstructured, not suited for addressing long-horizon planning problems, and lacks verification for reliability, as it overlooks LLMs’ tendency to hallucinate (Zhang et al., 2023b).

Evolve Dynamics LLMs can refine their decision-making process by reflecting on past trajectories (Shinn et al., 2024). Leveraging this ability, studies such as (Wang et al., 2023; Stengel-Eskin et al., 2024; Zhang et al., 2023a; Nottingham et al., 2024) focus on evolving new and advanced skills from pre-defined ones. However, these approaches often assume that LLMs already possess comprehensive domain knowledge and can derive new dynamics based solely on their understanding.

7 Conclusion

We introduce DiVE, a framework that bridges the knowledge gap between LLMs and downstream domains. The Discoverer extracts world dynamics, while the Verifier filters unreliable candidates. In an online setting, the Evolver reasons strategies through interaction. Our experiments demonstrate that DiVE effectively bridges the knowledge gap between LLMs and downstream domains.

Limitations

Evaluating DiVE in embodied environments may not fully capture the diversity and complexity of real-world dynamics that an agent might encounter in practice. If the environment’s dynamics change across episodes, the offline-learned dynamics from DiVE may struggle to adapt, introducing potential biases in the agent’s understanding and hindering the development of advanced strategies. Moreover, acquiring these dynamics from human demonstrations can be difficult or impractical in certain scenarios.

Ethical Concerns

We do not anticipate any immediate ethical or societal impact from our work. This study aims to bridge the knowledge gap between LLMs and the target domain. However, despite our efforts, DiVE may still exhibit hallucinations due to the inherent tendency of LLMs to hallucinate.

Acknowledgements

This work is supported by the Canada CIFAR AI Chair Program and the Canada NSERC Discovery Grant (RGPIN-2021-03115).

References

2024. The claude 3 model family: Opus, sonnet, haiku. <https://api.semanticscholar.org/CorpusID:268232499>.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. 2023. [Grounding large language models in interactive environments with online reinforcement learning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 3676–3713. PMLR.
- Cédric Colas, Laetitia Teodorescu, Pierre-Yves Oudeyer, Xingdi Yuan, and Marc-Alexandre Côté. 2023. [Augmenting autotelic agents with large language models](#). In *Proceedings of The 2nd Conference on Lifelong Learning Agents*, volume 232 of *Proceedings of Machine Learning Research*, pages 205–226. PMLR.
- Marc-Alexandre Côté, Akos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, et al. 2019. Textworld: A learning environment for text-based games. In *Computer Games: 7th Workshop, CGW 2018, Held in Conjunction with the 27th International Conference on Artificial Intelligence, IJCAI 2018, Stockholm, Sweden, July 13, 2018, Revised Selected Papers 7*, pages 41–75. Springer.
- Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, Abhishek Gupta, and Jacob Andreas. 2023. [Guiding pretraining in reinforcement learning with large language models](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 8657–8677. PMLR.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane, and Siva Reddy. 2022. [On the origin of hallucinations in conversational models: Is it the datasets or the models?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5271–5285, Seattle, United States. Association for Computational Linguistics.
- Haotian Fu, Pratyusha Sharma, Elias Stengel-Eskin, George Konidaris, Nicolas Le Roux, Marc-Alexandre Côté, and Xingdi Yuan. 2024a. Language-guided skill learning with temporal variational inference. *arXiv preprint arXiv:2402.16354*.
- Yao Fu, Dong-Ki Kim, Jaekyeom Kim, Sungryull Sohn, Lajanugen Logeswaran, Kyunghoon Bae, and Honglak Lee. 2024b. Autoguide: Automated generation and selection of state-aware guidelines for large language model agents. *arXiv preprint arXiv:2403.08978*.
- Nikolaj Goodger, Peter Vamplew, Cameron Foale, and Richard Dazeley. 2023. [A nethack learning environment language wrapper for autonomous agents](#). *Journal of Open Research Software*, 11.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.
- David Ha and Jürgen Schmidhuber. 2018. [Recurrent world models facilitate policy evolution](#). In *Advances in Neural Information Processing Systems 31*, pages 2451–2463. Curran Associates, Inc. <https://worldmodels.github.io>.
- Danijar Hafner. 2021. Benchmarking the spectrum of agent capabilities. In *International Conference on Learning Representations*.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2023. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR.
- Heinrich Küttler, Nantas Nardelli, Alexander H. Miller, Roberta Raileanu, Marco Selvatici, Edward Grefenstette, and Tim Rocktäschel. 2020. The NetHack Learning Environment. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

- Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Peter Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang, Chris Callison-Burch, and Peter Clark. 2023. Clin: A continually learning language agent for rapid task adaptation and generalization. *arXiv preprint arXiv:2310.10134*.
- Seungyong Moon, Junyoung Yeom, Bumsoo Park, and Hyun Oh Song. 2024. Discovering hierarchical achievements in reinforcement learning via contrastive learning. *Advances in Neural Information Processing Systems*, 36.
- Kolby Nottingham, Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Sameer Singh, Peter Clark, and Roy Fox. 2024. [Skill set optimization: Reinforcing language model behavior via transferable skills](#). *arXiv*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Parker-Holder, Minqi Jiang, Eric Hambro, Fabio Petroni, Heinrich Kuttler, Edward Grefenstette, and Tim Rocktäschel. 2021. [Minihack the planet: A sandbox for open-ended reinforcement learning research](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Haochen Shi, Zhiyuan Sun, Xingdi Yuan, Marc-Alexandre Côté, and Bang Liu. 2024. Opex: A component-wise analysis of llm-centric agents in embodied instruction following. *arXiv preprint arXiv:2403.03017*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2020. Alfworld: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768*.
- Elias Stengel-Eskin, Archiki Prasad, and Mohit Bansal. 2024. Regal: Refactoring programs to discover generalizable abstractions. *arXiv preprint arXiv:2401.16467*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2022. Large language models still can’t plan (a benchmark for llms on planning and reasoning about change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Ruoyao Wang, Graham Todd, Ziang Xiao, Xingdi Yuan, Marc-Alexandre Côté, Peter Clark, and Peter Jansen. 2024. Can language models serve as text-based world simulators? *arXiv preprint arXiv:2406.06485*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Nathaniel Weir, Xingdi Yuan, Marc-Alexandre Côté, Matthew Hausknecht, Romain Laroché, Ida Momennejad, Harm Van Seijen, and Benjamin Van Durme. 2022. One-shot learning from a demonstration with hierarchical latent language. *arXiv preprint arXiv:2203.04806*.
- Yue Wu, Yewen Fan, Paul Pu Liang, Amos Azaria, Yuanzhi Li, and Tom M Mitchell. 2024a. Read and reap the rewards: Learning to play atari with the help of instruction manuals. *Advances in Neural Information Processing Systems*, 36.
- Yue Wu, So Yeon Min, Shrimai Prabhumoye, Yonatan Bisk, Russ R Salakhutdinov, Amos Azaria, Tom M Mitchell, and Yuanzhi Li. 2024b. Spring: Studying papers and reasoning to play games. *Advances in Neural Information Processing Systems*, 36.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

- Jesse Zhang, Jiahui Zhang, Karl Pertsch, Ziyi Liu, Xiang Ren, Minsuk Chang, Shao-Hua Sun, and Joseph J Lim. 2023a. [Bootstrap your own skills: Learning to solve new tasks with large language model guidance](#). In *7th Annual Conference on Robot Learning*.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023b. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2024. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19632–19642.
- Jiawei Zheng, Hanghai Hong, Xiaoli Wang, Jingsong Su, Yonggui Liang, and Shikai Wu. 2024. Fine-tuning large language models for domain-specific machine translation. *arXiv preprint arXiv:2402.15061*.
- Victor Zhong, Dipendra Misra, Xingdi Yuan, and Marc-Alexandre Côté. 2024. Policy improvement using language feedback models. *arXiv preprint arXiv:2402.07876*.
- Zhaocheng Zhu, Yuan Xue, Xinyun Chen, Denny Zhou, Jian Tang, Dale Schuurmans, and Hanjun Dai. 2023. Large language models can learn rules. *arXiv preprint arXiv:2310.07064*.

Appendix

A	MiniHack	15
A.1	Configuration	15
A.2	Des-Files for Each MiniHack Task	15
A.3	Prompts for MiniHack Tasks	16
B	Crafter	19
B.1	Configuration	19
B.2	An Example of Learned Dynamics	19
B.3	An Example of Trajectory	19
B.4	Changes in Dynamics Precision Over Time	21
B.5	An Example of Verbalized Observation	21
B.6	Tech Tree in the Crafter Environment	23
B.7	Prompts for Crafter Environment	24

A MiniHack

This section offers further details on the MiniHack environment, along with a comprehensive description of the prompts utilized by DiVE within this environment.

A.1 Configuration

For this environment, we employ the gpt-4o-2024-08-06 model.

A.2 Des-Files for Each MiniHack Task

In this subsection, we present the Des-Files utilized for the Lava Crossing, Wand of Death, and Quest tasks.

A.2.1 Lava Crossing

```
MAZE: "mylevel", ' '
FLAGS:hardfloor
INIT_MAP: solidfill, ' '
GEOMETRY:center,center
MAP
-----
|.....L.....|
|.....L.....|
|.....L.....|
|.....L.....|
|.....L.....|
|.....L.....|
-----
ENDMAP
REGION:(0,0,12,6),lit,"ordinary"
$left_bank = selection:fillrect (1,1,5,5)
$right_bank = selection:fillrect (7,1,11,5)
OBJECT:('=', "levitation"),rndcoord($left_bank),blessed
STAIR:rndcoord($right_bank),down
BRANCH:(1,1,5,5),(0,0,0,0)
```

A.2.2 Wand of Death

```
MAZE: "mylevel", ' '
FLAGS: hardfloor
INIT_MAP: solidfill, ' '
GEOMETRY: center, center
MAP
|-----|
| .....|
| .....|
| .....|-----|
| .....|
| .....|
| .....|
|-----|
ENDMAP
REGION: (1,1,5,5), lit, "ordinary"
REGION: (6,1,26,1), lit, "ordinary"
REGION: (26,1,27,1), lit, "ordinary"
$safe_room = selection: fillrect(1,1,5,5)
OBJECT: ('/', "death"), (1,1), blessed
MONSTER: ('H', "minotaur"), (26,1)
STAIR: (27,1), down
BRANCH: (1,1,5,5), (1,1,2,2)
```

A.2.3 Quest

```
MAZE: "mylevel", ' '
FLAGS:hardfloor
INIT_MAP: solidfill,' '
GEOMETRY:center,center
MAP
-----|
```

```

|.....L.....|--          |-----|
|.....L.....|-----|.....|
|.....L.....|-----|.....|
|.....L.....|-----|.....|
|.....L.....|--          |-----|
|.....L.....|--          |-----|
|-----|
ENDMAP
REGION:(0,0,28,6),lit,"ordinary"
$left_bank = selection:fillrect (1,1,5,5)
$right_bank = selection:fillrect (7,1,11,5)
$goal_room = selection:fillrect (25,2,27,4)

OBJECT:('/',"cold"),(2,3),blessed
BRANCH:(1,1,3,3),(0,0,0,0)
MONSTER:random, (25,3)
STAIR:(26,3),down

```

A.3 Prompts for MiniHack Tasks

This section includes the prompts used to solve the MiniHack tasks.

A.3.1 Offline Dynamics Learning

The prompts for offline learning the dynamics from demonstrations.

Discover Action Dynamics

The player is playing a MiniHack game and would like you to analyze the game's dynamics. Specifically, they need help identifying the general preconditions and outcomes of actions by comparing consecutive observations.

- In the previous state, the player observed: {previous_state['state_description']} and took the action: {action}.
- In the current state, the player observed: {state['state_description']}.

Definitions:

- ****General precondition****: The condition that must be met for the action to succeed, regardless of the specific game state.
- ****General outcome****: The effect of the action, regardless of the specific game state.

Here is the action-to-text mapping: {state['action_prompt']}.

Discover Object Dynamics

The player is playing a MiniHack game and would like you to analyze the game's dynamics. Specifically, they need help identifying the general preconditions and outcomes of actions by comparing consecutive observations.

- In the previous state, the player observed: {previous_state['state_description']} and took the action: {action}.
- In the current state, the player observed: {state['state_description']}.

Definitions:

- ****General precondition****: The condition that must be met for the action to succeed, regardless of the specific game state.
- ****General outcome****: The effect of the action, regardless of the specific game state.

Here is the action-to-text mapping: {state['action_prompt']}.

Discover Subtask Dynamics

The player is playing a MiniHack game and would like your help in discovering the detailed steps to complete a subtask, and the preconditions and outcomes of the subtask. The player has provided the following information:

- In previous states, the player observed: {previous_state_description} and took the following actions: {previous_action}.

Definition:


```

- **Steps for completing the subtask**: The general and detailed sequence of actions
  , represented in text, that the player must follow to achieve the subtask. These
  steps should be general and not tied to the specific state of the environment.
- **Preconditions**: The general conditions that must be met for the subtask to be
  completed.
- **Outcomes**: The general effects of completing the subtask.

Here is the mapping of actions to their text representations: {state['action_prompt
']].

```

A.3.2 Online Dynamics Learning

The prompts for online evolving the dynamics from existing dynamics.

Discover Subtask Dynamics

```

You are playing the MiniHack game.
Your current subtask is: {subtask}.
Your current observation is: {observation}.
The primitive dynamics of the game are: {action_dynamics}.
The object dynamics of the game are: {object_dynamics}.

You are asked to evolve useful strategies that aid in completing the subtask based
on the provided information.

Instructions for evolving strategies:
- Do not introduce any new objects that are not part of the primitive dynamics.
- The evolved dynamics should not contradict the primitive dynamics.
- If a difficulty cannot be resolved by existing dynamics, evolve new and advanced
  dynamics by combining only existing dynamics using deductive reasoning.

Now, consider:
- All the potential difficulties that may arise, based on the primitive dynamics and
  current observation.
- All the potential dangers that may be encountered, based on the primitive dynamics
  and current observation.
- The strategies should be general that do not include specific actions, directions
  and dis

List all the potential difficulties, dangers and the evolved strategies to resolve
them.

```

A.3.3 Grounding

The prompts for grounding the agent.

Task Termination

```

You are playing the MiniHack game.
Your observation is: {description}.
Your current subtask is: {current_subtask}.
Your previous history is: {previous_history}.
The action and character mapping is: {action_prompt}.

You are asked to decide whether the current subtask should be terminated:
- If the termination conditions are met or the player is in danger or the player is
  in deadlock, the subtask should be terminated.

Output yes if the subtask has been completed, and no if it has not been completed.

```

Task Termination

```

You are playing the MiniHack game.
Your observation is: {description}.
Your current subtask is: {current_subtask}.
Your previous history is: {previous_history}.
The action and character mapping is: {action_prompt}.

You are asked to decide whether the current subtask should be terminated:

```

- If the termination conditions are met or the player is in danger or the player is in deadlock, the subtask should be terminated.

Output yes if the subtask has been completed, and no if it has not been completed.

Task Selection

You are playing the MiniHack game.
Your goal is: {task_description}.
Your observation is: {description}.
Your previously completed subtask is: {previous_subtask}.
All the subtasks are: {subtask_dynamics}.
All the environment dynamics are: {action_dynamics}.
All the object dynamics are: {object_dynamics}.

You are asked to select the next best subtask for completing the goal, based on all the provided information.

Action Selection

You are playing the MiniHack game.
Your current subtask is: {subtask}.
Your current observation is: {description}.
You also observe the feedback from the environment for the previous action: {observation}, which you need to consider for the next action.
The action and character mapping is: {action_prompt}.
Here are the strategies for the subtask: {evolved_strategy}.
Here are the primitive dynamics of the game: {action_dynamics} and {object_dynamics }.

Your past action and thought history: {history}

Now, you are asked to:

- List the potential difficulties and dangers that may arise based on the **current observation** and primitive dynamics for completing the subtask.
- Output the thoughts on the future actions how to resolve the difficulties and dangers for completing the subtask.
- Select the best action for the next step based on the thoughts.
- Do not copy the previous thoughts.

In the game, the distance is only measured as very far > far > near

B Crafter

This section provides a detailed overview of the Crafter environment and a comprehensive description of the prompts used by DiVE.

B.1 Configuration

For this setup, we utilize the gpt-4o-2024-05-13 model in the Crafter environment with seeds 1, 2, 3, 4, and 5.

B.2 An Example of Learned Dynamics

```
grass
grass can be found near ['tree', 'water', 'path'], but it is not associated with
    ['diamond', 'coal', 'iron']
You can walk directly through grass.
grass can only be used for: ['collect_sapling', 'eat_plant'] ,
coal
coal can be found near ['stone', 'iron', 'diamond'], but it is not associated
    with ['grass', 'cow', 'skeleton']
You cannot walk directly through coal.
coal turn into path after collect_coal
coal can only be used for: ['make_iron_pickaxe', 'make_iron_sword', '
    collect_coal']
coal can be collected by ['wood_pickaxe'] ,
```

B.3 An Example of Trajectory

```
subgoal
Collect 4 units of wood and maintain a
healthy level of health, food, drink,
and energy.

subtask
collect_wood": {
  "General Plan": {
    "step_1": "Locate a tree, ensuring 'tree' is within immediate distance.",
    "step_2": "Move closer to the tree if it's not within immediate distance.
      This may involve navigating around obstacles such as grass, stone, or
      paths.",
    "step_3": "Face the tree to meet the precondition of the action.",
    "step_4": "Execute the 'collect_wood' action to gather wood from the tree."
  },
  "termination_condition": "The inventory's wood amount increases by 1, or if
    health, food, drink, or energy levels drop critically low, or a threat (such
    as skeletons or zombies) is detected nearby."
}

state description
It is daytimeState description:
- East: immediate (grass); nearby (unexplored_area);
- North: immediate (grass); nearby (unexplored_area);
- North East: immediate (grass); nearby (unexplored_area);
- North West: immediate (grass); nearby (unexplored_area);
- South: immediate (grass); nearby (unexplored_area);
- South East: immediate (grass); nearby (cow, unexplored_area);
- South West: immediate (grass); nearby (unexplored_area);
- West: immediate (grass); nearby (unexplored_area);
Closest:
- grass: North 1 blocks away (immediate) (objects in between: None)
- cow: South East 5 blocks away (nearby) (objects in between: {'grass'})
- Facing grass on the south.
Your status:
- health: 9/9
- food: 9/9
- drink: 9/9
- energy: 9/9
You have nothing in your inventory.
```

evolve dynamics

Expand the exploration area to discover new surroundings that could potentially include a tree.

action

move_north

B.4 Changes in Dynamics Precision Over Time




Figure 5: Precision of verified dynamics over verified steps

As shown in Figure 5, the precision of the verified dynamics \mathcal{W} increases significantly during the initial verification step. However, as the process continues, further improvements in precision are minimal, indicating that the first step successfully filters out most of the unreliable dynamics.

B.5 An Example of Verbalized Observation



Figure 6: The observation within the Crafter environment

Verbalized Observation

It is daytime

State description:

- East: immediate (grass); nearby (unexplored_area);
- North: immediate (grass); nearby (unexplored_area);
- North East: immediate (grass); nearby (unexplored_area);
- North West: immediate (grass); nearby (tree, unexplored_area);
- South: immediate (grass); nearby (unexplored_area);
- South East: immediate (grass); nearby (unexplored_area, tree);
- South West: immediate (grass); nearby (tree, unexplored_area);
- West: immediate (grass); nearby (tree, unexplored_area);

Closest:

- grass: North 1 blocks away (immediate) (objects in between: None)
- tree: West 4 blocks away (nearby) (objects in between: {'grass'})
- Facing grass on the south.

Your status:

- health: 9/9
- food: 9/9
- drink: 9/9
- energy: 9/9

You have nothing in your inventory.

B.6 Tech Tree in the Crafter Environment

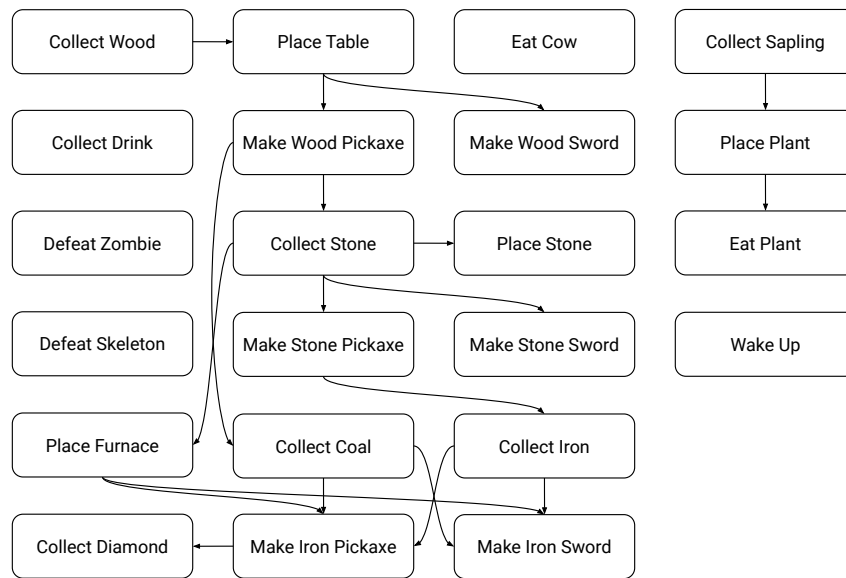


Figure 7: The tech tree within the Crafter environment

B.7 Prompts for Crafter Environment

This section includes the prompts used in the Crafter environment.

B.7.1 Online Dynamics Learning

```
Given the following details:
- Primitive dynamics: {env_dynamics}
- Current subtask: {transition['subtask']}
- Current observation: {transition['state_description']}

You are asked to identify the difficulties in completing the current subtask and
provide 3 primitive or evolved dynamics to solve each of these difficulties.

Instructions for identifying difficulties:
- List the objects required to complete the subtask, specify their locations, and
  explain where to find them if they are not in the current observation.
- Outline all possible obstacles that may be encountered along the way.

Instructions for evolving advanced dynamics:
- Do not introduce any new objects that are not part of the primitive dynamics.
- The evolved dynamics should not contradict the primitive dynamics.
- If a difficulty cannot be resolved by existing dynamics, evolve new and advanced
  dynamics by combining only existing dynamics using deductive reasoning.

Instructions for providing deductive reasoning steps:
- For each evolved dynamics, provide the used primitive dynamics.
- For the deductive reasoning steps, provide the steps to combine the primitive
  dynamics to evolve the advanced dynamics and the rule of inference used (Modus
  Ponens, Modus Tollens, ..... )

Last, for each situation and dynamics should be general and do not contain details
about specific locations about the objects.
Output in the following format:{output_format} and leaves 'None' for
deductive_reasoning_steps if the dynamics are primitive.
```

Given the following details:

- Primitive dynamics: {env_dynamics}
- Evolved dynamics: {transition['reformatted_dynamics']}
- Current subtask: {transition['subtask']}

You are asked to examine the validity of the evolved dynamics and provide feedback.

Instructions for examining the validity of the evolved dynamics:

- The evolved dynamics should only be a combination of existing dynamics using deductive reasoning and should not introduce new dynamics. Output 'True' if the evolved dynamics introduce new dynamics; otherwise, output 'False'.
- The evolved dynamics should not introduce any new objects that are not mentioned in the primitive dynamics. Output 'True' if the evolved dynamics introduce new objects; otherwise, output 'False'.
- Each deductive reasoning step should not contradict any of the primitive dynamics. Output 'True' if the evolved dynamics contradict the primitive dynamics; otherwise, output 'False'.

Instructions for examining the usefulness of the evolved dynamics:

- The difficulties should be directly related to the current subtask.
- The evolved dynamics should be useful in solving the difficulties identified in the subtask.
- Evaluating the usefulness of the evolved dynamics on a scale of 1 to 5, where 5 is the most useful and 1 is the least useful.

Last, output the validity of the evolved dynamics in the following format: {output_format}.

B.7.2 Grounding

The prompts for grounding the agent.

Subtask Termination

Given the following details:

- Subtask description: {self._transition['subtask']},
 - Current observation: {self._transition['state_description']},
 - Initial observation: {self._transition['initial_state_description']},
 - Previous executed actions: {self._transition['previous_actions'][-3:]}
- you are asked to decide whether the subtask should be terminated or not.

For deciding whether to terminate the subtask, consider:

- The previous action, provided it was executed successfully.
- The difference between the initial and current observations, including the inventory changes.

The subtask should be terminated, and the output should be 'True' only if any of its termination conditions are met.

Otherwise, if none of the termination conditions are met, the subtask should continue running, and the output should be 'False'.

Justify whether the termination conditions are met or not first, and then provide the termination decision.

Output in this format: {output_format}.

Action Selection

Given the following details:

- Current observation: {self._transition['state_description']}
- Current subtask's description: {self._transition['subtask']}
- Previous actions: {self._transition['previous_actions'][-3:]}
- Primitive dynamics: {self._transition['primitive_dynamics']}

You are asked to:

- identify the objects related to the current subtask and provide their locations and dynamics.
- select the top 3 actions that contributes to the subtask by either moving closer to the object or interacting with the object; and provide all the objects and dynamics directly related with each action.

- based on each action's related objects, provide the rationale and detailed consequences of executing each action on the objects.
- select the best action to execute next and provide the justification for your choice.

Lastly, select the action only from the available actions: {self._transition['available_actions']}; {feedback}.

Note: Avoid unnecessary crafting and placement if the items are within reachable distance.

Please format your response in the following format: {action_format}

B.7.3 Offline Dynamics Learning

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the materials used in the inventory for executing this action.

To discover the materials used in the inventory, consider the following aspects:

- List all the materials used not gained for each state transition and their corresponding quantities;
- List all the common materials used in all state transitions and their corresponding quantities.

If no materials are used in common, simply output as "None".

materials are: wood, stone, coal, iron, diamond, sapling.

Output as many pre-conditions as you can think of in the following format: {output_format}

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the objects within immediate distance required for executing this action.

To discover the objects within immediate distance, consider the following aspects:

- List all the objects within immediate distance for each state transition before executing the action, which you can only choose from the object list.
- Identify and list the objects that are present in all state transitions within the immediate distance before executing the action, which you can only choose from the object list.
- The common object ****must**** be present in all state transitions' immediate objects.

If no objects within immediate distance are in common, simply output as "None".

Object list: [coal, cow, diamond, furnace, iron, lava, skeleton, stone, table, tree, water, zombie, plant]

Output as many pre-conditions as you can think of in the following format: {output_format}

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the facing object required for executing this action.

To discover the objects within immediate distance, consider the following aspects:

- List the facing object for each state transition before this action executed, not the direction.
- List the union of all the facing object across all the state transitions before this action executed.
- The union of all the facing object must be present in at least one state transitions' facing object.

Object list: [coal, cow, diamond, furnace, iron, lava, skeleton, stone, table, tree, water, zombie, plant]

Output as many pre-conditions as you can think of in the following format: {output_format}

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the the inventory tool required for executing this action.

To discover the objects within immediate distance, consider the following aspects:

- List all the tools in the inventory for each state transition.
- List the tools that are common across all the state transitions.
- The common tools must be present in all state transitions' tools; if no tools are required, simply output as "None".
- List the most advanced tool required for executing the action within the common tools.

Tool list: [None, wood pickaxe, stone pickaxe, iron pickaxe, wood sword, stone sword, iron sword].

Output as many pre-conditions as you can think of in the following format: {output_format}

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the inventory and status increase after executing this action.

To discover the increase about the inventory and status, consider the following aspects:

- List all the increase of the inventory and status for each state transition after executing the action.
- List the increase of the inventory and status that are common across all the state transitions after executing the action.
- The common increase must be present in all state transition's increase; if no inventory and status can be found, simply output as "None";

Output as many pre-conditions as you can think of in the following format: {output_format}

Learning Basic Attribute

In the Crafter environment, the world dynamics are unique.

You are tasked with discovering the unique dynamics of actions within this environment.

Given the state transitions described as follows: {partial_description}, identify the facing object changes for executing this action.

To discover the change about the facing object, consider the following aspects:

- List all the change about the facing object after executing the action for each state transition.
- List the common change about the facing object across all the state transitions.
- The common change must be present in all state transitions' facing object change; if no change can be found, simply output as "None".

inventory: [wood, stone, coal, iron, diamond, sapling, wood_pickaxe, stone_pickaxe, iron_pickaxe, wood_sword, stone_sword, iron_sword]

status: [health, food, drink, energy], the maximum value of status is 9 and you can use 'increased_to_9' to represent the status increase to the maximum value.

Output as many pre-conditions as you can think of in the following format: {output_format}

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the pre-conditions for the action '{action}'. Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the inventory materials used pre-conditions for the action '{action}' based on the discovered dynamics.

Precondition: You need to use {precondition} to execute the action.

For the verification, you need to consider the followings:

- List all the inventory materials that are consumed not gained for each state transition and their corresponding quantity,
- List the common inventory materials used across all the state transitions.
- Determine if the inventory materials used mentioned in the pre-condition are within the common inventory materials used.
- If it is within the common inventory materials used, then it is valid; otherwise, it is invalid.

Finally, if there are more advanced tools exist, then it is invalid; otherwise, it is valid.

materials are: wood, stone, coal, iron, diamond, sapling.

Output in this format: {output_format}

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the pre-conditions for the action '{action}'. Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the immediate objects pre-conditions for the action '{action}' based on the discovered dynamics.

Precondition: You only need {precondition} within immediate distance.

For the verification, you need to consider the followings:

- List all the objects within immediate distance before executing the action for each state transition only from the object list.
- Identify and list the objects that are present in all state transitions within the immediate distance before executing the action as common immediate objects, only from the object list.
- The common object ****must**** be present in all state transitions' immediate objects.
- Determine if there are objects mentioned in the pre-condition that are not in the common immediate objects.
- If there are objects mentioned in the pre-condition that are not in the common immediate objects, then it is invalid; otherwise, it is valid.

Object list: [coal, cow, diamond, furnace, iron, lava, skeleton, stone, table, tree, water, zombie, plant]

Output in this format: {output_format}

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the pre-conditions for the action '{action}'.
Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the facing object pre-conditions for the action '{action}' based on the discovered dynamics.

Precondition: You need to face {precondition} before execute the action.

For the verification, you need to consider the followings:

- List the facing object for each state transition before this action executed.
- List the union of all the facing object across all the state transitions before this action executed.
- Determine if the facing object mentioned in the pre-condition is within the union of all the facing object.
- If the facing object mentioned in the pre-condition is within the union of all the facing objects, then it is valid; otherwise, it is invalid.

Finally, if there are more advanced tools exist, then it is invalid; otherwise, it is valid.

Object list: [coal, cow, diamond, furnace, iron, lava, skeleton, stone, table, tree, water, zombie, plant]

Output in this format: {output_format}

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the pre-conditions for the action '{action}'.
Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the inventory pre-conditions for the action '{action}' based on the discovered dynamics.

Precondition: You only need {preconditions} to execute the action.

For the verification, you need to consider the followings:

List all the tools in the inventory for each state transition.

- List the tools that are common across all the state transitions.
- The common tools must be present in all state transitions' tools; if no tools are required, simply output as "None".
- If the tools mentioned in the pre-condition are within the common tools, then it is valid; otherwise, it is invalid.

Tool list: [None, wood pickaxe, stone pickaxe, iron pickaxe, wood sword, stone sword, iron sword].

Output in this format: {output_format}

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the outcome for the action '{action}'.

Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the inventory and status increase after executing this action. '{action}' based on the discovered dynamics.

predicted_increases: You only increase {precondition} in the inventory and status after executing the action.

For the verification, you need to consider the followings:

- List all the increases of the inventory and status for each state transition after executing the action as the increases.
- List the increases of the inventory and status that are common across all the state transitions after executing the action as the common increases.
- The common increases must be present in all state transition's increases; if no inventory and status increases can be found, simply output as "None".
- Determine if there are increases mentioned in the predicted_increases that are not in the common increases.
- If there are increases mentioned in the predicted_increases that are not in the common increases, then it is invalid; otherwise, it is valid.

Output in this format: {output_format}

inventory: [wood, stone, coal, iron, diamond, sapling, wood_pickaxe, stone_pickaxe, iron_pickaxe, wood_sword, stone_sword, iron_sword]

status: [health, food, drink, energy], the maximum value of status is 9 and you can use 'increased_to_9' to represent the status increase to the maximum value from the non-maximum value.

Verifying Basic Attribute

In the crafter environment, you have discovered unique dynamics and need to verify the outcome for the action '{action}'.

Given the state transitions before and after taking the action '{action}', described as follows: {sampled_descriptions}

You are asked to verify the facing object changes after executing this action. '{action}' based on the discovered dynamics.

predicted_changes: The facing object changes is {precondition} after executing the action.

For the verification, you need to consider the followings:

- List all the change about the facing object after executing the action for each state transition.
- List the common change about the facing object across all the state transitions.
- The common change must be present in all state transitions' facing object change; if no change can be found, simply output as "None".
- Determine if the predicted changes are within the common changes.
- If the predicted changes are within the common changes, then it is valid; otherwise, it is invalid.

Verifying Basic Attribute

Given randomly selected frames featuring '{obj}', please complete the following:

1. Identify which objects are most closely related to '{obj}'.
 - Rate their relevance using these terms: 'Very related', 'Not related'.
 - Here are the selected frames concerning '{obj}': {description}
2. Identify the common time relationship for this object.
3. Format your response as follows: {object_relationship_format}

Only consider the relationship with the following objects: ['grass', 'coal', 'cow', 'diamond', 'iron', 'lava', 'skeleton', 'stone', 'tree', 'water', 'zombie', 'plant', 'path', 'sand', 'plant-ripe']

For the relationship definitions:

- Very related: They can always be found together.
- Not related: They cannot be found together.

Verifying Basic Attribute

Given randomly selected frames featuring '{obj}' described as: {description}, please complete the following tasks:

1. Identify the objects and time most closely related to '{obj}' in the provided description.
 - Rate their relevance using the following terms: 'Very related', 'Not related'.
2. Based on the discovered dynamics {dynamics}, verify the relationships between the objects and time.
3. Output all the valid relationships within the discovered dynamics that match all the relevance levels of the objects in the provided frames only about the object '{obj}'.

Consider relationships with the following objects only: ['grass', 'coal', 'cow', 'diamond', 'furnace', 'iron', 'lava', 'skeleton', 'stone', 'table', 'tree', 'water', 'zombie', 'plant', 'path', 'sand', 'plant-ripe'].

Relationship definitions:

- Very related: They are always found together.
- Not related: They are never found together.

Output in the following format: {object_relationship_format}

Discover Subtask Attribute

In the Crafter environment, the world dynamics are complex and there are many subtasks to be completed.

Here are the dynamics about the core action {action}:

The precondition:

- it need to face: {self.facing_object_preconditions[action]}
- it need to have {self.immediate_object_preconditions[action]} within immediate distance
- it need to have {self.inventory_materials_precondition[action]} in the inventory.
- it need to have {self.inventory_tool_precondition[action]} in the inventory.

The outcome:

- its facing object changes to {self.facing_object_change[action]}
- its inventory and status outcome increases on {self.inventory_outcome[action]}

Note: 'None' indicates that there are no specific preconditions or outcomes for the corresponding elements of the dynamics."

Given the subtask "{subtask}", I want you to write the step-plan for completing this subtask.

First, read and understand all the provided world dynamics.

Then, locate the world dynamics that are relevant to the subtask.

Next, write the subtask's requirements, the step-plan for completing the subtask based on the world dynamics and the outcomes in this format {subtask_plan_format }

-steps: the general steps required to complete the subtask.

-termination_condition: when should this subtask be terminated; you should consider the status, potential danger and the outcome of the action

Finally, only output the plan for the subtask.

Here are a few examples: {examples}

Here are the observations on how to completing the subtask: {partial_description}

Verify Subtask Attribute

In the Crafter environment, the world dynamics are complex and there are many subtasks to be completed.

Here are the dynamics about the core action {action}:

The precondition:

- it need to face: {self.facing_object_preconditions[action]}
- it need to have {self.immediate_object_preconditions[action]} within immediate distance
- it need to have {self.inventory_materials_precondition[action]} in the inventory.
- it need to have {self.inventory_tool_precondition[action]} in the inventory.

The outcome:

- its facing object changes to {self.facing_object_change[action]}
- its inventory and status outcome increases on {self.inventory_outcome[action]}

Note: 'None' indicates that there are no specific preconditions or outcomes for the corresponding elements of the dynamics."

Given the subtask "{subtask}", I want you to verify the plan for completing this subtask {discovered_plan}.

First, read and understand all the provided world dynamics.

Then, locate the world dynamics that are relevant to the subtask.

Next, examine the provided plan's requirements, the step-plan and the termination conditions.

Lastly, if it can be applied across different examples, then it is a valid plan.

Here are the observations on how to completing the subtask: {partial_description}

Output within this format{subtask_verification_format}

Discover Subgoal Attribute

Given the following considerations:

- The game's final goal: collect diamonds and survive.
- All the subtasks: [['collect_coal', 'collect_diamond', 'collect_drink', 'collect_iron', 'collect_sapling', 'collect_stone', 'collect_wood', 'defeat_skeleton', 'defeat_zombie', 'eat_cow', 'eat_plant', 'make_iron_pickaxe', 'make_iron_sword', 'make_stone_pickaxe', 'make_stone_sword', 'make_wood_pickaxe', 'make_wood_sword', 'move', 'place_furnace', 'place_plant', 'place_stone', 'place_table', 'sleep']]
- Environment dynamics: {env_dynamics}

- Human player completed subtasks in chronological order: {human_demo}

Note: action_x_y means this action is performed consecutively for y times

Your task is to generate a plan by reordering the subtasks to help the player achieve the final game goal with different subgoals.

For plan generation, consider:

- List all the subtasks that contributes to the final goal from the human player's trajectory.
- Re-order the subtasks that fits the game's dynamics.
- For each subtask, state it as a subgoal with a description and the exact repetition of the subtask if needed.

Lastly, output the plan in this format: {output_format};

For example: {output_example}

Verify Subgoal Attribute

Given the following considerations:

- The game's final goal: collect diamonds and survive.
- All the subtasks: [['collect_coal', 'collect_diamond', 'collect_drink', 'collect_iron', 'collect_sapling', 'collect_stone', 'collect_wood', 'defeat_skeleton', 'defeat_zombie', 'eat_cow', 'eat_plant', 'make_iron_pickaxe', 'make_iron_sword', 'make_stone_pickaxe', 'make_stone_sword', 'make_wood_pickaxe', 'make_wood_sword', 'move', 'place_furnace', 'place_plant', 'place_stone', 'place_table', 'sleep']]
- Environment dynamics: {env_dynamics}

- Human player completed subtasks in chronological order: {human_demo}

Note: action_x_y means this action is performed consecutively for y times

- Discovered plan: {discovered_plan}

Your task is to verify the provided plan.

For plan verification, consider:

- List all the subtasks that contributes to the final goal from the human player's trajectory.
- Re-order the subtasks that fits the game's dynamics.
- If the provided plan can be applicable to the different trajectories, then it is valid.

Lastly, output this format: {output_format};