

# ADAM: AN EMBODIED CAUSAL AGENT IN OPEN-WORLD ENVIRONMENTS


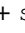

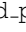
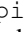
**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In open-world environments like Minecraft, existing agents face challenges in continuously learning structured knowledge, particularly causality. These challenges stem from the opacity inherent in black-box models and an excessive reliance on prior knowledge during training, which impair their interpretability and generalization capability. To this end, we introduce ADAM, An emboDied causal Agent in Minecraft, that can autonomously navigate the open world, perceive multimodal contexts, learn causal world knowledge, and tackle complex tasks through lifelong learning. ADAM is empowered by four key components: 1) an interaction module, enabling the agent to execute actions while documenting the interaction processes; 2) a causal model module, tasked with constructing an ever-growing causal graph from scratch, which enhances interpretability and diminishes reliance on prior knowledge; 3) a controller module, comprising a planner, an actor, and a memory pool, which uses the learned causal graph to accomplish tasks; 4) a perception module, powered by multimodal large language models, which enables ADAM to perceive like a human player. Extensive experiments show that ADAM constructs an almost perfect causal graph from scratch, enabling efficient task decomposition and execution with strong interpretability. Notably, in our modified Minecraft games where no prior knowledge is available, ADAM maintains its performance and shows remarkable robustness and generalization capability. ADAM pioneers a novel paradigm that integrates causal methods and embodied agents in a synergistic manner.

## 1 INTRODUCTION

Embodied agents exploring open-world environments mark a critical frontier in artificial intelligence (AI) research (Cassell, 2000; Xia et al., 2018; Savva et al., 2019). The ultimate goal is to build generally capable agents (GCAs) (Team et al., 2021) that can autonomously perform a broad range of tasks through perception, learning, and interaction (Mnih et al., 2015; Xi et al., 2023; Park et al., 2023). Minecraft (Nebel et al., 2016), a globally renowned 3D video game, serves as a representative open-world environment for these agents. It offers a randomly generated world of massive blocks, where players need to master complex crafting recipes (*e.g.*, planks  + sticks   $\rightarrow$  wood\_pickaxe ) and gather resources (*e.g.*, mining cobblestone  with wood\_pickaxe ) , progressively unlocking new items in the technology tree. The substantial freedom and precise simulation of physical laws in Minecraft render it an exceptional platform for researching GCAs.

In Minecraft, two primary approaches for developing GCAs have been extensively explored: reinforcement learning (RL)-based (Lin et al., 2021; Baker et al., 2022; Fan et al., 2022; Mao et al., 2022; Hafner et al., 2023) and large language model (LLM)-based (Wang et al., 2023a; Zhu et al., 2023; Qin et al., 2023a; Nottingham et al., 2023; Wang et al., 2023c;d). Specifically, RL agents learn through interactions and updating their black-box model weights, which poses challenges for interpretability, efficiency, and generalization. On the other hand, LLM-based agents possess and rely on rich prior knowledge of both virtual games and real worlds (Ouyang et al., 2022; Wei et al., 2022a; Achiam et al., 2023). Their reliance on omniscient data (*e.g.*, GPS coordinates, voxel blocks, biome, *etc.*, which are not explicitly observable by the player) (Wang et al., 2023a; Zhu et al., 2023; Wang et al., 2023d) presents challenges for generalization and human **gameplay observation** alignment.

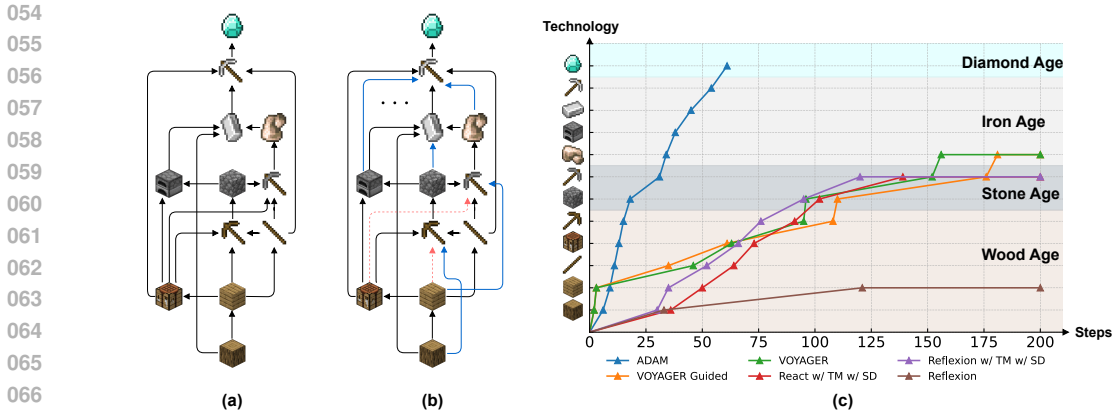


Figure 1: (a) The technology tree for acquiring diamonds in the Minecraft game. ADAM can precisely discover item dependencies from scratch. (b) Modified Minecraft technology tree, where the prior knowledge from the Internet or wiki does not align with the actual game dynamics. Red arrows denote removed dependencies, while blue arrows denote added dependencies. (c) In the game setting shown in (b), ADAM maintains the ability to learn the correct causal graph and successfully obtains diamonds, whereas other methods can only acquire raw\_iron within the step limit, and ADAM achieves a 4.6× speedup in obtaining raw\_iron compared to the SOTA.

To address these issues, we propose ADAM, An embodied causal Agent in Minecraft, that can autonomously navigate the open world, perceive multimodal contexts, learn causal world knowledge, and tackle complex tasks through lifelong learning. Specifically, ADAM is composed of four key modules as shown in Fig. 2: (1) **Interaction module**, which enables the agent to execute actions from the action space and processes the agent’s observable information into formatted records. (2) **Causal model module**, which includes two causal discovery (CD) methods. LLM-based CD utilizes interaction records to make causal assumptions. Intervention-based CD refines these assumptions to derive a causal subgraph. Multiple causal subgraphs are integrated into a comprehensive causal graph (i.e., technology tree). (3) **Controller module**, which includes a planner, an actor, and a memory pool. The planner can utilize the causal graph to perform task decomposition. The actor uses the subtasks for action choosing. The memory pool ensures the long-term context dependence. (4) **Perception module**, which is driven by multimodal LLMs (MLLMs), enabling ADAM to perceive its surroundings without relying on omniscient data, thereby achieving human-like **gameplay observation**.

Extensive experiments demonstrate that ADAM achieves a 2.2× speedup compared to the SOTA in the task of obtaining diamonds. In scenarios where crafting recipes are modified (Fig. 1b), only ADAM maintains the ability to obtain diamonds, while other methods can only acquire raw\_iron within the step limit, and ADAM achieves a 4.6× speedup in obtaining raw\_iron compared to the SOTA (Fig. 1c). ADAM demonstrates strong interpretability through constructing a nearly perfect technology tree (Fig. 1a) from scratch, whereas other methods exhibit at least 30% errors or omissions. Meanwhile, ADAM closely aligns with human **gameplay observation** without relying on omniscient metadata, while maintaining comparable environmental perception performance to methods that utilize such data.

Overall, our contributions are as follows:

- (1) **We introduce ADAM, an embodied causal agent** that autonomously navigate the open world, perceive multimodal contexts, learn causal world knowledge, and tackle complex tasks through lifelong learning.
- (2) **We tackle the limitations of existing embodied agents.** Our ADAM demonstrates strong **robustness** without relying on prior knowledge or omniscient metadata, unlike other LLM-based agents, while aligns human **gameplay observation** during exploration.
- (3) **We pioneer the integration of causal methods into open-world embodied agents**, allowing the agent to organize the learned knowledge in a rigorous causal graph, thereby demonstrating excellent interpretability.

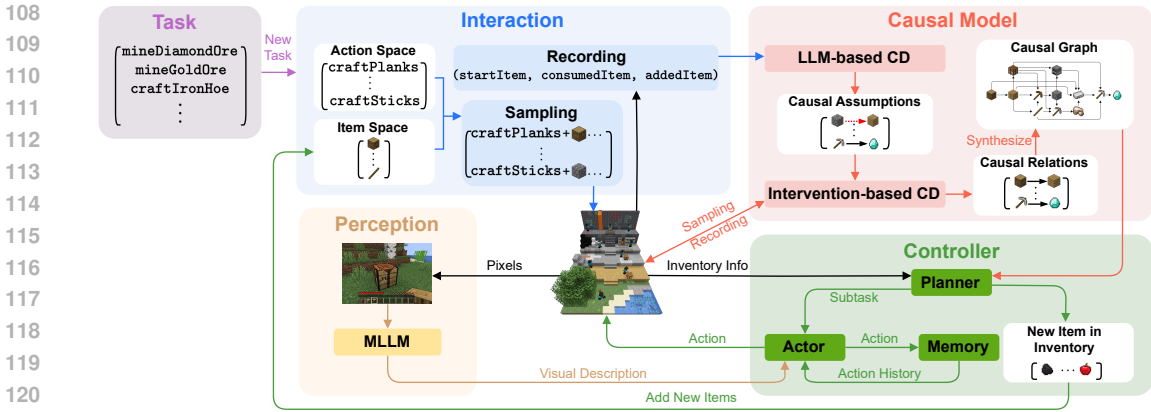


Figure 2: Four key modules of ADAM. The **interaction module** executes actions in the environment according to the task and records the processes. The **causal model module** identifies the causal relationship between items and actions to construct an ever-growing causal graph. The **controller module** implements task execution based on the learned causal graph. The **perception module** aligns the agent’s **gameplay observation** more closely with human.

- (4) We improve the CD performance by employing embodied agent-driven interventions, which enhances the accuracy and efficiency of CD compared to existing methods without interventions.

## 2 PRELIMINARIES

**Causal graphical models (CGMs).** A CGM represents the structure of causality within a system (Peters et al., 2017) by detailing the direct causal relationships among a set of variables  $X_1, \dots, X_n$ . It is characterized by a distribution over these variables and is associated with a directed acyclic graph (DAG), known as a causal graph. In this graph, each node corresponds to a variable, and each directed edge signifies a direct causal relation from  $X_i$  to  $X_j$ .

**Causal discovery from interventions.** Causal Discovery (CD) (Spirtes et al., 2001; Pearl, 2009; Peters et al., 2017; Glymour et al., 2019) is a fundamental process to infer causal relationship from data. The relationship is typically represented in the form of a causal graph. While observational data reveals correlations, interventions allow us to analyze causal dependencies between variables. Specifically, interventions alter the distribution of variables (e.g., types of initial items) during experimental sampling, which serves as a gold standard for CD (Eberhardt & Scheines, 2007). By observing their effects, we can identify causal relationships rather than merely correlations.

## 3 METHOD

In this section, we begin with the basic notations and definitions in our work (Sec. 3.1). Then, we give an overview of our ADAM framework (Sec. 3.2). Next, we detail the four modules of ADAM: interaction module (Sec. 3.3), causal model module (Sec. 3.4), controller module (Sec. 3.5), and perception module (Sec. 3.6).

### 3.1 NOTATIONS AND DEFINITIONS

We first introduce several key notations and definitions in our work. Sets are denoted by uppercase letters, and their elements by lowercase letters.

**Inventory:** The set  $I_t$  of items possessed by the agent at any step  $t$ . **Initialization:** At the initialization of Minecraft instances, the initial inventory  $I_0$  can be specified by the agent. **Action Space:** The set  $A$  of actions whose names (e.g., gatherWoodLog) are replaced by letters and invisible.

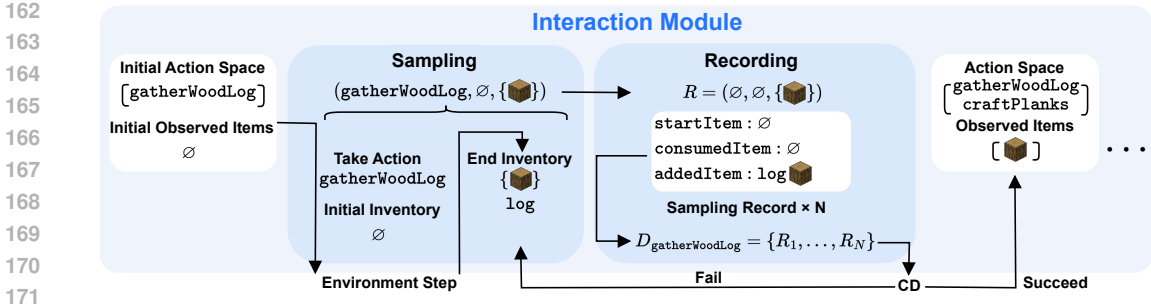


Figure 3: The interaction module has two core functionalities: **sampling** and **recording**. Sampling involves executing actions in the environment, and recording involves processing and documenting the observable information. For instance, the initial action space is  $\{\text{gatherWoodLog}\}$ , whose name is not exposed to ADAM and is denoted as  $\{a\}$  here (Note that, the original notation  $\{\text{gatherWoodLog}\}$  is retained in the figure for the illustrative purpose.). The initial observed item space is  $\emptyset$ . After executing  $a$  for one step, logs (log) are obtained. A sampling can be represented as  $(a, \emptyset, \{\text{log}\})$ , where  $\emptyset$  is the initial inventory and  $\{\text{log}\}$  is the inventory after this step. The result is recorded as  $R = (\emptyset, \emptyset, \{\text{log}\})$ , where the first  $\emptyset$  is the initial inventory and the second  $\emptyset$  indicates that no items are consumed, and  $\{\text{log}\}$  represents the items that are obtained. After sampling  $N$  times, data  $D_a = \{R_1, \dots, R_N\}$  is provided to the causal model module for CD. If the causal relation failed to be identified, resampling on  $a$  occurs; if successful, new actions like `craftPlanks` are enabled by the acquisition of log, and the observed item space is updated to  $\{\text{log}\}$ .

ADAM must independently discover the effects of these actions. **Movement Space:** The set  $M$  of basic movements whose names (e.g., `moveForward`, `moveBackward`) are visible to ADAM. **Step:** The agent takes an action  $a$  in the environment. A step ends either when action completes or when execution times out. **Observed Item Space:** The set  $S$  of all items that ADAM has encountered. Initially,  $S$  is empty. **Environmental Factors:** The set  $\mathcal{E}$  of environment conditions (e.g., biome, surrounding block types). **Task:** Denoted by the tuple  $(I_{\text{goal}}, \mathcal{E})$ , a task is accomplished at step  $t$  if  $I_{\text{goal}} \subseteq I_t$  and the factors  $\mathcal{E}$  are present within a certain distance around the agent.

### 3.2 OVERVIEW

ADAM comprises four modules as depicted in Fig. 2. Given a task  $(I_{\text{goal}}, \mathcal{E})$ , the **interaction module** enables the agent to execute each action  $a$  and records data  $D_a$ . This data is then utilized by the **causal model module**, which employs LLM-based CD to make causal assumptions and intervention-based CD to refine these assumptions into causal subgraphs. These subgraphs are integrated into a causal graph (i.e., technology tree). Once the causal graph  $\mathcal{G}$  contains all required items  $I_{\text{goal}}$  in the task, the **controller module** completes the task from an empty inventory, aided by visual descriptions from the **perception module**. Newly discovered items enable the execution of new unknown actions and the CD on these actions. This iterative process ensures the lifelong learning through continuous engagement and adaptation. ADAM is a general framework that can be extended to other open-world environments, as discussed in Appendix F.

### 3.3 INTERACTION MODULE

The interaction module (Fig. 3) enables the agent to execute actions for sampling and records observable information. Initially, the action space  $A$  contains one element `gatherWoodLog`, which is the most basic action in Minecraft and a common setup for Minecraft agents (Wang et al., 2023a). As the agent acquires certain new items, new actions are unlocked in  $A$ . For example, the action of mining becomes available only after the agent obtains a mining tool (e.g., a `woodenPickaxe`). By continuously collecting data on each action and coordinating with other modules, the nodes on the technology tree are progressively discovered by the agent.

**Sampling.** This involves initiating a Minecraft instance with all observed items  $S$  as the initial inventory  $I_0$ , and then executing action  $a$  to observe the agent’s inventory  $I_1$  after this step. The quantity of each item in  $S$  and  $I_1$  is recorded. This sampling, represented as  $(a, S, I_1)$ , is performed

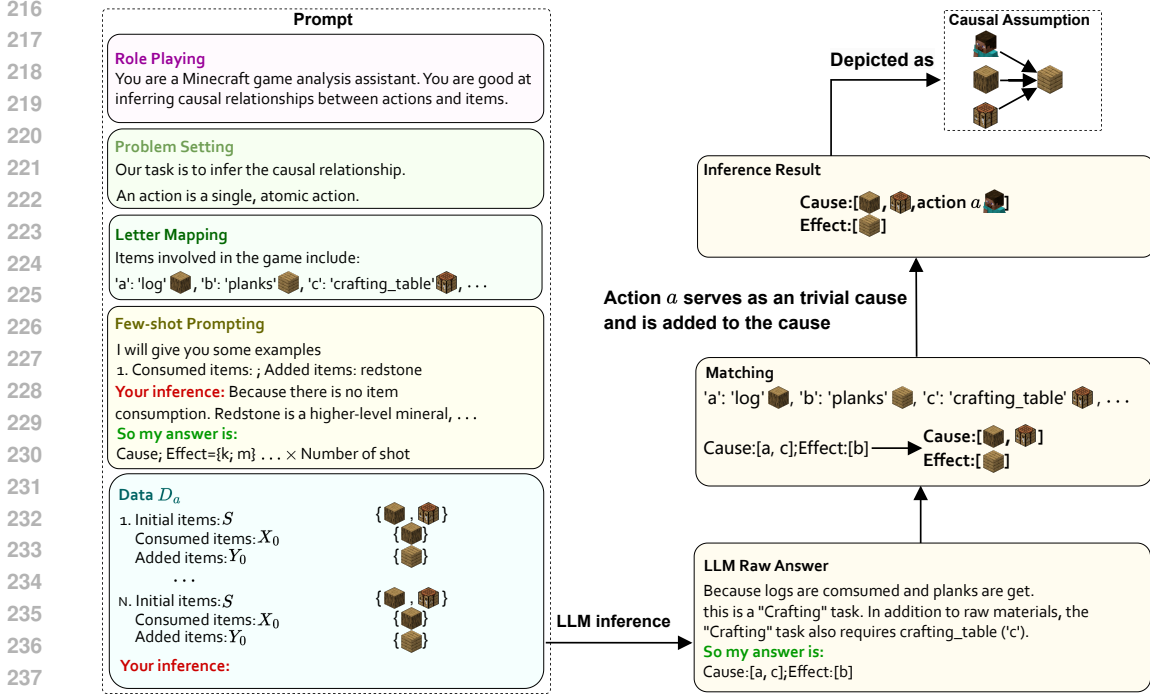


Figure 4: LLM-based CD performs causal reasoning under the guidance of the prompt. Role Playing assigns an analysis assistant role to the LLM. Problem Setting provides the reasoning task. Letter Mapping maps the item names to letters for the accurate output. Few-shot Prompting provides examples for chain-of-thought (Wei et al., 2022b) reasoning. Data  $D_a$  is presented in the same form as the few-shot examples. The output of LLM serves as the causal assumption.

$N$  times, focusing on one specific action at a time. Due to the same configuration of these  $N$  samplings, parallelization is available and accelerates the exploration.

**Recording.** This involves processing the results of samplings and documenting them. For the sampling  $(a, S, I_1)$ , consumed items are denoted by  $X$ , which includes items with decreased quantities and items that are present in  $S$  but absent in  $I_1$ . Conversely, obtained items are denoted by  $Y$ , which includes items with increased quantities and items that newly appear in  $I_1$ . A single record is denoted as  $R = (S, X, Y)$ . Such  $N$  records on action  $a$  collectively form data  $D_a = \{R_1, \dots, R_N\}$ .

### 3.4 CAUSAL MODEL MODULE

The causal model module uses data  $D_a$  to infer causal relationships and constructs causal subgraphs for each action  $a$ . These subgraphs are then integrated into a comprehensive causal graph.

A causal relationship is a stable and repeatable dependence in any step, where the agent takes an action  $a$  and the acquisition of items  $E$  (effect items) relies on the items  $C$  (cause items) possessed by the agent before the step. Since ADAM focuses on item acquisition, causal relations where no items are obtained do not contribute to the technology tree.

In the causal model module, **LLM-based CD** makes assumptions on the causal relationships, which effectively reduces the number of item nodes that need to be confirmed in the causal subgraph and achieves acceleration. Then, **intervention-based CD** refines these assumptions and accurately constructs causal subgraphs. We also detail our **optimization techniques** employed in this module.

**LLM-based CD.** The input to LLM-based CD (Fig. 4) is the data  $D_a$  containing  $N$  records, and the output is a causal assumption, which consists of cause items  $C$  and effect items  $E$ . The prompt is designed with five components: **(1) Role Playing**, which assigns a specific role to the LLM. In this

context, the LLM serves as a causal analysis assistant, dedicated to extracting causal relationships from data. **(2) Problem Setting**, which provides the reasoning task and the fundamental concepts of the environment. We avoid introducing specific environmental knowledge for generalization. **(3) Letter Mapping**, which involves mapping item names to letters, a simplification that facilitates the formatted output. **(4) Few-shot Examples**, which involves providing the LLM with several reasoning examples in a chain-of-thought (Wei et al., 2022b) style, including example questions, the reasoning processes, and the expected answering format. The examples are independent of the technology tree, thus preventing the introduction of prior knowledge. **(5) Data  $D_a$** , which consists of  $N$  records from the interaction module, and is formatted consistently with the few-shot examples.

**Intervention-based CD.** Intervention is a method to experimentally verify causal relationships among variables. Intervention-based CD (Fig. 5) can refine the causal assumptions and construct a highly accurate causal graph.

Before interventions, it has to be confirmed that  $C$  is a sufficient condition for  $E$ . If  $C$  has already lacked the necessary items to achieve  $E$ , then the vital edges are missing and the graph can not be corrected by excluding redundant edges. Specifically, by performing sampling  $(a, C, I_1)$  as described in Sec. 3.3, if  $E$  is consistently absent from  $I_1$ , the assumption is deemed incorrect, leading the LLM-based CD to re-infer the assumption. If these inferences continue to fail, the interaction module will resample data  $D_a$ .

Intervention-based CD performs samplings  $(a, C \setminus \{c\}, I_1)$  for each item  $c \in C$ . For each item  $e \in E$ , if there is always  $e \notin I_1$  within a maximum number of samplings, then  $c$  is the cause of  $e$ , and the edge  $c \rightarrow e$  is retained. If at least one sample includes  $e \in I_1$ , then  $c$  is not indispensable for  $e$ , and the edge  $c \rightarrow e$  is removed. Intervention-based CD yields accurate causal subgraphs, which are subsequently integrated into a complete causal graph.

**Optimization techniques.** (1) *Temporal Modeling (TM)*: Temporal information can help determine causal direction. Events occurring later cannot be the cause of events occurring earlier. This eliminates some edges in the causal graph. (2) *Subgraph Decomposition (SD)*: By focusing on the causal effects of individual actions, the causal model module processes a manageable number of items (only those in each cause-effect pair) at a time. This significantly reduces the complexity of the CD process.

### 3.5 CONTROLLER MODULE

During the execution of task  $(I_{\text{goal}}, \mathcal{E})$ , ADAM starts from an empty inventory (i.e.,  $I_0 = \emptyset$ ) to ensure the fair comparison with other methods. After all items in  $I_{\text{goal}}$  have been obtained in the causal graph, the controller module (Fig. 6) is responsible for executing the task. If there are newly discovered items, they will be added to the observed item space  $S$  and pend for a new cycle of CD, thus achieving lifelong learning. The controller module comprises three components as described below.

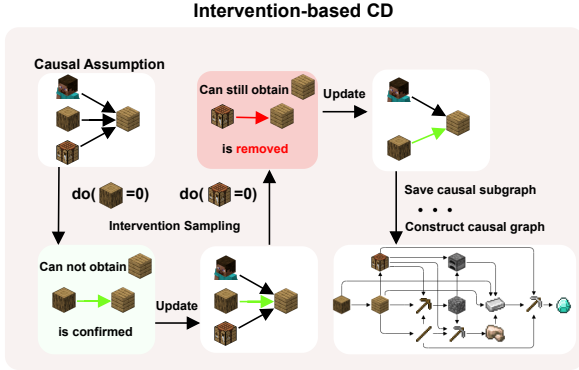


Figure 5: Intervention-based CD verifies causal assumptions. An example of causal assumption is that, under the action  $a$ ,  $\text{log}$  and  $\text{crafting\_table}$  contribute to the acquisition of  $\text{planks}$ . This assumption is denoted as  $(a, \{\text{log}, \text{crafting\_table}\}, \{\text{planks}\})$ . Intervention-based CD will verify each item in  $\{\text{log}, \text{crafting\_table}\}$ . When removing  $\text{log}$  from the inventory and executing action  $a$ ,  $\text{planks}$  cannot be obtained, proving that  $\text{log}$  is a dependency of  $\text{planks}$ , and this edge  $(\text{log} \rightarrow \text{planks})$  is retained in the causal graph (represented in green). When removing  $\text{crafting\_table}$  and executing action  $a$ ,  $\text{planks}$  still can be obtained, which shows that  $\text{crafting\_table}$  is not a dependency of  $\text{planks}$ , and this edge  $(\text{crafting\_table} \rightarrow \text{planks})$  is removed from the causal graph (represented in red). Intervention-based CD results in a corrected causal subgraph. Multiple subgraphs can be combined into the technology tree in Minecraft. The actions are not shown in the technology tree for the sake of simplicity.

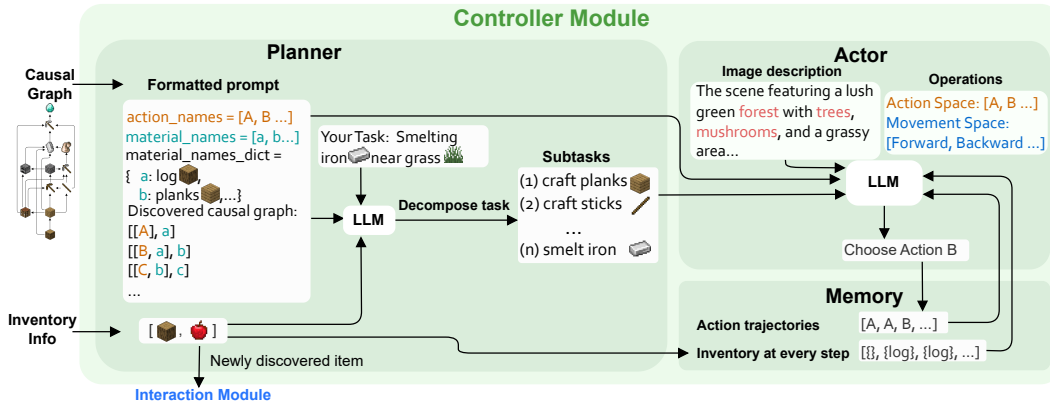


Figure 6: The controller module comprises three components. The Planner utilizes the reasoning capabilities of LLMs to decompose the task. It receives the current inventory and the learned causal graph as input. The Actor leverages LLMs to choose an action  $a$  in the the action space  $A$  or a movement  $m$  in the movement space  $M$ , and execute it in the environment. The Memory records the step information, including action trajectories and item changes every step.

**Planner.** The Planner utilizes LLMs to decompose the task with current inventory  $I_t$  at step  $t$  and the learned causal graph. Relying solely on the causal graph is suboptimal as actions may fail or have side effects. LLMs can fully utilize the inventory information and provide detailed thought process of the decomposition process, which is passed to the Actor for action choosing.

**Actor.** The Actor leverages LLMs to choose an action  $a \in A$  or a movement  $m \in M$  to execute. It receives the task decomposition from the Planner, the description of game image at this step, and the records from the Memory. In the task  $(I_{\text{goal}}, \mathcal{E})$ , the Actor prioritizes obtaining the items  $I_{\text{goal}}$ , because this process involves changing positions, which may affect the acquisition of  $\mathcal{E}$ .

**Memory.** The Memory records observable information during task execution, including action trajectories and inventory changes at each step. It plays a crucial role in tracking long-term dependencies and facilitates robust task execution.

### 3.6 PERCEPTION MODULE

The perception module utilizes MLLMs for environmental observation, enabling ADAM to perceive the world without relying on metadata such as the names of surrounding blocks, GPS coordinates, or the biome names, which are typically invisible to human players. This module captures first-person screenshots between steps, which are processed by MLLMs to generate descriptions. This text description is subsequently passed to the Actor in the controller module for action choosing.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETUP

In our study, we employ Mineflayer (PrismarineJS, 2023a), a JavaScript-based framework providing control APIs for the commercial Minecraft (version 1.19)<sup>1</sup>. The encapsulation of Mineflayer uses the implementation in VOYAGER (Wang et al., 2023a). For visual processing, we utilize prismarine-viewer (PrismarineJS, 2023b), an API for rendering game scenes from the agent’s perspective. ADAM and our baselines all use GPT-4-turbo (gpt-4-0125-preview) (Achiam et al., 2023) for LLM inference, with the temperature set to 0.3 based on our experiments in Appendix A. For visual description, we utilize LLaVA-v1.5-13B (Liu et al., 2024) in our perception module.


<sup>1</sup><https://www.minecraft.net>

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

Model	SHD	Model	SHD
ADAM	$2 \pm 2$	CDHRL	$10 \pm 4$
ADAM w/o TM,SD	$19 \pm 6$	CDHRL w/ SD	$6 \pm 2$
Reflexion	$24 \pm 9$	React w/ TM,SD	$5 \pm 2$
AutoGPT	$24 \pm 6$	Reflexion w/ TM,SD	$4 \pm 2$
Empty Graph	32	AutoGPT w/ TM,SD	$4 \pm 2$


Table 1: Structural Hamming Distance (SHD) between the learned causal graph and the target graph. For non-embodied agents without built-in interventions, even with our TM and SD (Section 3.4), the causal graph learned by these agents remains suboptimal. The CD used by CDHRL is based on SDI (Ke et al., 2019), which incorporates temporal modeling (TM) into its implementation, and still exhibits over 30% errors or omission, while ADAM can identify a nearly perfect causal graph.


## 4.2 BASELINES

In the absence of directly comparable works, we select representative methods as baselines and focus on their comparable aspects, including: (1) **ReAct** (Yao et al., 2023), which explicitly expresses the thought processes through chain-of-thought prompting (Wei et al., 2022b). (2) **Reflexion** (Shinn et al., 2024), derived from ReAct, which can reflect on its exploration history. (3) **AutoGPT** (Significant Gravitas), which can autonomously decompose tasks and execute subtasks in a ReAct-style. Baselines 1–3 can only perform text-based tasks and lack embodied components to interact with the environment. Hereafter, they are referred to as **non-embodied agents**, and we have adapted them with our interaction module for embodied exploration. (4) **VOYAGER** (Wang et al., 2023a) is an LLM-based embodied lifelong learning agent in Minecraft, featuring an automatic curriculum aiming to “discover as many diverse things as possible”. We also add our benchmarking tasks (e.g., obtaining diamonds ) to the curriculum for oriented explorations, denoted as **VOYAGER-Guided**. (5) **CDHRL** (Peng et al., 2022) introduces an RL agent that constructs hierarchical structures based on causal relationships. Given that RL agents have **disparate action spaces** and **magnitudes** of difference in episode length compared to LLM-based methods ( $10^5 \sim 10^8$  in **DreamerV3** (Hafner et al., 2023) and **DEPS** (Wang et al., 2023d) versus  $10^1 \sim 10^2$  in **VOYAGER** and our **ADAM**), we focus our comparison on the CD component of CDHRL. Detailed discussion of other Minecraft agents that are not directly comparable can be found in Appendix C.

## 4.3 MAIN RESULTS

**Interpretability.** We evaluate the interpretability of agents by assessing their ability to construct a causal graph. Structure Hamming Distance (SHD) (Zheng et al., 2024) can quantify the discrepancy between the learned causal graph and the target graph as presented in Tab. 1. Despite applying our TM and SD optimization (Section 3.4), non-embodied agents without built-in interventions fail to achieve the optimal accuracy. For CDHRL, we directly provide ADAM’s sampling data for its CD. CDHRL performs CD with all nodes in the causal graph, which hampers its performance as CDHRL shows improved performance when integrated with our SD optimization. Nevertheless, these competitive methods exhibit at least 30% errors or omissions, whilst ADAM is capable of learning a **nearly perfect causal graph**. **VOYAGER** does not organize knowledge in a causal graph.

**Efficiency.** Efficiency is evaluated in the original Minecraft as shown in Tab. 2. ADAM achieves a **2.2× speedup** compared to the SOTA in the task of obtaining diamonds . Its design facilitates parallel sampling (Section 3.3), which significantly boosts exploration efficiency. Due to the absence of intervention-based CD, non-embodied agents are unable to refine their causal assumptions, limiting their exploration speed and confining them to the lower levels of the technology tree. Additionally, ADAM achieves **higher success rate** across most tasks compared to other methods.

**Robustness.** We assess the robustness of agents in a modified Minecraft environment where crafting recipes are altered. The result is shown in Fig. 1c. In this scenario, there exists a misalignment between the LLM’s prior knowledge and the actual game dynamics. ADAM successfully obtains diamonds  in the modified Minecraft. The baselines lack a CD approach to learn and verify causal knowledge, and struggle with complex dependencies. The most advanced item that baseline



Framework	Wooden Tool	Stone Tool	Iron Tool	Diamond
React w/ TM w/ SD	51 ± 19(2/3)	96(1/3)	N/A (0/3)	N/A (0/3)
Reflexion w/ TM w/ SD	60 ± 27(3/3)	122 ± 56(2/3)	N/A (0/3)	N/A (0/3)
AutoGPT w/ TM w/ SD	49 ± 20(2/3)	103 ± 45(2/3)	N/A (0/3)	N/A (0/3)
VOYAGER	8 ± 2(3/3)	<b>10 ± 3(3/3)</b>	27 ± 11(3/3)	113 ± 41(2/3)
VOYAGER Guided	<b>7 ± 2(3/3)</b>	11 ± 2(3/3)	<b>24 ± 9(3/3)</b>	75 ± 20(2/3)
ADAM	23 ± 5(3/3)	33 ± 8(3/3)	53 ± 16(3/3)	68 ± 21(3/3)
ADAM Parallel	12 ± 2(3/3)	18 ± 3(3/3)	29 ± 5(3/3)	<b>34 ± 7(3/3)</b>

Table 2: Exploration steps in different tasks. Fewer steps indicates higher efficiency. Each method has three trials for a maximum length of 200 steps. The success rate is depicted in the parentheses. ADAM achieves a 2.2× speed compared to the SOTA in the task of obtaining diamonds, with a higher success rate.

	VOYAGER	VOYAGER w/o Meta	ADAM	ADAM w/o MLLM
Find a river	<b>16 ± 8(2/3)</b>	N/A (0/3)	21 ± 16(2/3)	N/A (0/3)
Gather log near river	<b>36(1/3)</b>	N/A (0/3)	40 ± 23(2/3)	N/A (0/3)
Smelting iron near grass	N/A (0/3)	N/A (0/3)	<b>95(1/3)</b>	N/A (0/3)

Table 3: Performance of ADAM and VOYAGER in tasks requiring environmental factors  $\mathcal{E}$ . Each method has three trials for a maximum length of 100 steps. The success rate is depicted in the parentheses. VOYAGER’s performance significantly declines when metadata is not accessible, whereas ADAM do not rely on metadata. MLLM contributes ADAM’s performance in this type of tasks.

agents manage to acquire is `raw_iron` 🪄. In terms of exploration speed, ADAM achieves a **4.6× speedup** in obtaining `raw_iron` 🪄. For more details, please refer to Tab. 7 in the Appendix.

**Lifelong learning.** ADAM utilizes CD methods to learn the effects of each action, obtaining causal subgraphs that include new items. These new items make it possible to perform more unknown actions, thereby continually expanding the knowledge of the game world in a bootstrapping manner and achieves lifelong learning. ADAM successfully learns a complex causal graph of all 41 actions we implement, as demonstrated in Fig. 7.

**Human gameplay alignment.** Avoiding the use of human-invisible metadata demonstrates ADAM’s alignment with human gameplay. We compare VOYAGER, a fully text-based agent that relies on metadata<sup>2</sup>. We test three tasks that requires environmental factors  $\mathcal{E}$ . The results are shown in Tab. 3. VOYAGER’s performance significantly declines when metadata is not accessible. ADAM performs well on tasks with  $\mathcal{E}$  relying solely on observable information. Ablation experiments demonstrate that MLLMs contribute to ADAM’s performance on this type of tasks.

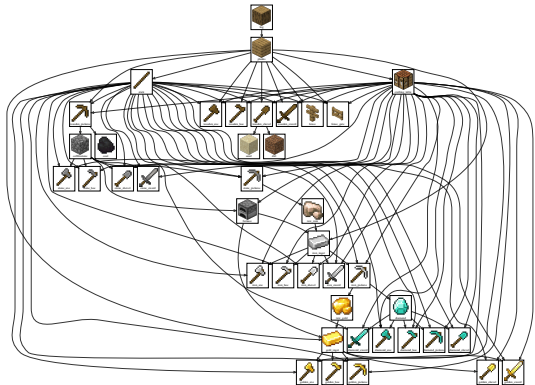


Figure 7: The causal graph learned in lifelong learning. ADAM successfully unlocks all 41 actions we implement and discovers accurate causal relationships.

#### 4.4 ABLATION STUDIES

**Ablation of LLM-based CD** Prior knowledge and inference capabilities are key factors in ablating LLM-based CD. Prior knowledge can be ablated in modified environments and enhanced through fine-tuning LLMs on MC-QA datasets. Inference capabilities can be ablated by replacing SOTA LLMs to smaller LLMs. Our result in Tab. 4 shows that ADAM primarily utilizes the reasoning abilities of LLMs rather than relying on prior knowledge.

<sup>2</sup>The information used by VOYAGER and ADAM is compared in Tab. 5 in the Appendix.

Model	R&E (success/all)	Model	R&E (success/all)
gpt-4-turbo-preview	<b>0.0 (35/35)</b>	gpt-4	0.1 (35/35)
gpt-4-turbo-preview <sup>†</sup>	0.1 (35/35)	gpt-4 <sup>†</sup>	0.1 (35/35)
gpt-3.5-turbo	0.2 (34/35)	Llama-2-70B	0.6 (23/35)
gpt-3.5-turbo <sup>†</sup>	0.3 (32/35)	Llama-2-70B <sup>†</sup>	0.9 (16/35)
Llama-2-70B-finetuned	0.4 (27/35)	Llama-2-13B-finetuned	1.8 (5/35)
Llama-2-70B-finetuned <sup>†</sup>	0.9 (15/35)	Llama-2-13B-finetuned <sup>†</sup>	N/A (0/35)

Table 4: Ablation of LLM-based CD. <sup>†</sup> means “in a modified environment”. We record the average number of redundant/error items (represented as R&E) in the causal assumption proposed by LLM-based CD, and the success rate after the intervention-based CD. LLMs with only strong prior knowledge but weak inference abilities cannot perform well.

**Ablation of intervention-based CD.** Without intervention-based CD, agents are forced to rely on exhaustive trials to learn the game knowledge, significantly impairing their efficiency and effectiveness. Our experimental results are shown in Fig. 8. Through interventions, ADAM achieves up to  $4.4 \times$  speed and a higher accuracy compared to the ablated group.

## 5 RELATED WORK

**Causality in Agent.** The integration of causality (Pearl, 2009; Peters et al., 2017; Schölkopf, 2022) into agents is primarily aimed at enhancing the learning efficiency (Méndez-Molina et al., 2020; Seitzer et al., 2021; Gasse et al., 2021; Sun et al., 2021; Peng et al., 2022). Peng et al. (2022) propose CDHRL to build high-quality hierarchical structures in complicated environments. Méndez-Molina et al. (2020) employ the causal models to restrict the search space. Zeng et al. (2023) distinguish agents with causality in two categories: ones relying on prior causal information and ones learn causality by causal discovery algorithms (Spirtes et al., 2000; Sun et al., 2007; Zhang & Hyvärinen, 2009; Zhang et al., 2011; Peters et al., 2014; Zhu et al., 2019). Our work aligns with the latter category and extends to a wider range of scenarios where prior knowledge is unknown.

**LLM/MLLM-Based Agent.** Leveraging the generalization capabilities of LLM (Brown et al., 2020; Touvron et al., 2023a;b) to empower agent systems with tools (Qin et al., 2023b; Schick et al., 2024; Shen et al., 2024) is an essential task (Xi et al., 2023; Wang et al., 2023b). Schick et al. (2024) design a framework to allow LLM to use external APIs to complete tasks. Qin et al. (2023b) build related benchmarks to evaluate performance in such tasks. In tasks involving interaction with the environment, Shinn et al. (2024) enhance the agent’s ability through language feedback signals. Wei et al. (2022b) employ Chain-of-Thought prompting method to optimize the reasoning capabilities of the LLM agent. However, the hallucination and interpretability challenges (Zhang et al., 2023) of LLMs also accompany these systems. In this work, we prove that the perspective of causal architecture can reduce the reliance on priors and enhance the robustness of inferences.

## 6 CONCLUSION

In this work, we introduce ADAM, an embodied causal agent in open-world environments. ADAM innovatively incorporates CD with embodied exploration, significantly improving the accuracy of CD while enhancing the efficiency and interpretability of embodied exploration. Not relying on prior knowledge, ADAM demonstrates strong robustness, and its multimodal perception aligns with human behavior. Our work sets a foundation for developing autonomous agents that can understand and manipulate environments in a causal manner.




Figure 8: Average number of steps and success rate to learn the causal subgraphs of “Collecting” actions (e.g., gatherIronOre), which are representative due to higher noise in their sampling data. We perform up to 20 steps for each action, and if the CD fails, it is counted as 20. Intervention contributes up to  $4.4 \times$  acceleration and higher success rate in the exploration.

540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593

## REPRODUCIBILITY STATEMENT

We provide detailed descriptions of our experimental setup in Section 4.1, and we use the same OpenAI API (gpt-4-0125-preview) for all LLM-based methods to ensure comparability and reproducibility. After the paper is published, we commit to releasing our code to support reproducibility.

## REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Bowen Baker, Ilge Akkaya, Peter Zhokov, Joost Huizinga, Jie Tang, Adrien Ecoffet, Brandon Houghton, Raul Sampedro, and Jeff Clune. Video pretraining (vpt): Learning to act by watching unlabeled online videos. *Advances in Neural Information Processing Systems*, 35:24639–24654, 2022.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Justine Cassell. Embodied conversational interface agents. *Communications of the ACM*, 43(4):70–78, 2000.
- Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of science*, 74(5):981–995, 2007.
- Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. Minedojo: Building open-ended embodied agents with internet-scale knowledge. *Advances in Neural Information Processing Systems*, 35:18343–18362, 2022.
- Maxime Gasse, Damien GRASSET, Guillaume Gaudron, and Pierre-Yves Oudeyer. Causal reinforcement learning using observational and interventional data. 2021.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- William H Guss, Cayden Codell, Katja Hofmann, Brandon Houghton, Noboru Kuno, Stephanie Milani, Sharada Mohanty, Diego Perez Liebana, Ruslan Salakhutdinov, Nicholay Topin, et al. The minerl 2019 competition on sample efficient reinforcement learning using human priors. *arXiv preprint arXiv:1904.10079*, 2019.
- William H Guss, Mario Ynocente Castro, Sam Devlin, Brandon Houghton, Noboru Sean Kuno, Crissman Loomis, Stephanie Milani, Sharada Mohanty, Keisuke Nakata, Ruslan Salakhutdinov, et al. The minerl 2020 competition on sample efficient reinforcement learning using human priors. *arXiv preprint arXiv:2101.11071*, 2021.
- Danijar Hafner. Benchmarking the spectrum of agent capabilities. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=1W0z96MFEoH>.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- Anssi Kanervisto, Stephanie Milani, Karolis Ramanauskas, Nicholay Topin, Zichuan Lin, Junyou Li, Jianing Shi, Deheng Ye, Qiang Fu, Wei Yang, et al. Minerl diamond 2021 competition: Overview, results, and lessons learned. *NeurIPS 2021 Competitions and Demonstrations Track*, pp. 13–28, 2022.
- Nan Rosemary Ke, Olexa Bilaniuk, Anirudh Goyal, Stefan Bauer, Hugo Larochelle, Chris Pal, and Yoshua Bengio. Learning neural causal models from unknown interventions. *CoRR*, abs/1910.01075, 2019. URL <http://arxiv.org/abs/1910.01075>.

- 594 Zaijing Li, Yuquan Xie, Rui Shao, Gongwei Chen, Dongmei Jiang, and Liqiang Nie. Optimus-  
595 1: Hybrid multimodal memory empowered agents excel in long-horizon tasks. *CoRR*,  
596 abs/2408.03615, 2024. doi: 10.48550/ARXIV.2408.03615. URL [https://doi.org/10.](https://doi.org/10.48550/arXiv.2408.03615)  
597 [48550/arXiv.2408.03615](https://doi.org/10.48550/arXiv.2408.03615).
- 598 Zichuan Lin, Junyou Li, Jianing Shi, Deheng Ye, Qiang Fu, and Wei Yang. Juewu-mc:  
599 Playing minecraft with sample-efficient hierarchical reinforcement learning. *arXiv preprint*  
600 *arXiv:2112.04907*, 2021.
- 601 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances*  
602 *in neural information processing systems*, 36, 2024.
- 603 Hangyu Mao, Chao Wang, Xiaotian Hao, Yihuan Mao, Yiming Lu, Chengjie Wu, Jianye Hao, Dong  
604 Li, and Pingzhong Tang. Seihai: A sample-efficient hierarchical ai for the minerl competition. In  
605 *Distributed Artificial Intelligence: Third International Conference, DAI 2021, Shanghai, China,*  
606 *December 17–18, 2021, Proceedings 3*, pp. 38–51. Springer, 2022.
- 607 Arquímides Méndez-Molina, Ivan Feliciano-Avelino, Eduardo F Morales, and Luis Enrique Sucar.  
608 Causal based q-learning. *Res. Comput. Sci.*, 149(3):95–104, 2020.
- 609 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-  
610 mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level  
611 control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- 612 Steve Nebel, Sascha Schneider, and Günter Daniel Rey. Mining learning and crafting scientific  
613 experiments: a literature review on the use of minecraft in education and research. *Journal of*  
614 *Educational Technology & Society*, 19(2):355–366, 2016.
- 615 Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi, Sameer  
616 Singh, and Roy Fox. Do embodied agents dream of pixelated sheep: Embodied decision mak-  
617 ing using language guided world modelling. In Andreas Krause, Emma Brunskill, Kyunghyun  
618 Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International Confer-*  
619 *ence on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume  
620 202 of *Proceedings of Machine Learning Research*, pp. 26311–26325. PMLR, 2023. URL  
621 <https://proceedings.mlr.press/v202/nottingham23a.html>.
- 622 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
623 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
624 low instructions with human feedback. *Advances in neural information processing systems*, 35:  
625 27730–27744, 2022.
- 626 Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and  
627 Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings*  
628 *of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22, 2023.
- 629 Judea Pearl. *Causality*. Cambridge university press, 2009.
- 630 Shaohui Peng, Xing Hu, Rui Zhang, Ke Tang, Jiaming Guo, Qi Yi, Ruizhi Chen, Xis-  
631 han Zhang, Zidong Du, Ling Li, Qi Guo, and Yunji Chen. Causality-driven hierar-  
632 chical structure discovery for reinforcement learning. In Sanmi Koyejo, S. Mohamed,  
633 A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural In-*  
634 *formation Processing Systems 35: Annual Conference on Neural Information Process-*  
635 *ing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9,*  
636 *2022*, 2022. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/](http://papers.nips.cc/paper_files/paper/2022/hash/7e9fbd01b3084956dd8a070c7bf30bad-Abstract-Conference.html)  
637 [7e9fbd01b3084956dd8a070c7bf30bad-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/7e9fbd01b3084956dd8a070c7bf30bad-Abstract-Conference.html).
- 638 Jonas Peters, Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with  
639 continuous additive noise models. 2014.
- 640 Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations*  
641 *and learning algorithms*. The MIT Press, 2017.
- 642 PrismarineJS. Prismarinejs/mineflayer, 2023a. URL [https://github.com/](https://github.com/PrismarineJS/mineflayer)  
643 [PrismarineJS/mineflayer](https://github.com/PrismarineJS/mineflayer). <https://github.com/PrismarineJS/mineflayer>.

- 648 PrismaticJS. Prismaticjs/prismatic-viewer, 2023b. URL [https://github.com/](https://github.com/PrismaticJS/prismatic-viewer)  
649 [PrismaticJS/prismatic-viewer](https://github.com/PrismaticJS/prismatic-viewer). URL <https://github.com/PrismaticJS/prismatic-viewer>.  
650 viewer.  
651
- 652 Yiran Qin, Enshen Zhou, Qichang Liu, Zhenfei Yin, Lu Sheng, Ruimao Zhang, Yu Qiao, and Jing  
653 Shao. Mp5: A multi-modal open-ended embodied system in minecraft via active perception.  
654 *arXiv preprint arXiv:2312.07472*, 2023a.
- 655 Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru  
656 Tang, Bill Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world  
657 apis. *arXiv preprint arXiv:2307.16789*, 2023b.  
658
- 659 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
660 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
661 models from natural language supervision. In *International conference on machine learning*, pp.  
662 8748–8763. PMLR, 2021.
- 663 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain,  
664 Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied  
665 ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.  
666 9339–9347, 2019.  
667
- 668 Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro,  
669 Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can  
670 teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36, 2024.  
671
- 672 Bernhard Schölkopf. Causality for machine learning. In *Probabilistic and Causal Inference: The*  
673 *Works of Judea Pearl*, pp. 765–804. 2022.
- 674 Maximilian Seitzer, Bernhard Schölkopf, and Georg Martius. Causal influence detection for improv-  
675 ing efficiency in reinforcement learning. *Advances in Neural Information Processing Systems*, 34:  
676 22905–22918, 2021.  
677
- 678 Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugging-  
679 gpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information*  
680 *Processing Systems*, 36, 2024.
- 681 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:  
682 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*  
683 *Systems*, 36, 2024.  
684
- 685 Significant Gravitas. AutoGPT. URL [https://github.com/Significant-Gravitas/](https://github.com/Significant-Gravitas/AutoGPT)  
686 [AutoGPT](https://github.com/Significant-Gravitas/AutoGPT).  
687
- 688 Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT  
689 press, 2000.
- 690 Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press,  
691 2001.  
692
- 693 Xiaohai Sun, Dominik Janzing, Bernhard Schölkopf, and Kenji Fukumizu. A kernel-based causal  
694 learning algorithm. In *Proceedings of the 24th international conference on Machine learning*, pp.  
695 855–862, 2007.
- 696 Yüwen Sun, Kun Zhang, and Changyin Sun. Model-based transfer reinforcement learning based on  
697 graphical model representations. *IEEE Transactions on Neural Networks and Learning Systems*,  
698 2021.  
699
- 700 Xiaojuan Tang, Jiaqi Li, Yitao Liang, Song-Chun Zhu, Muhan Zhang, and Zilong Zheng. Mars: Sit-  
701 uated inductive reasoning in an open-world environment. *CoRR*, abs/2410.08126, 2024. doi: 10.  
48550/ARXIV.2410.08126. URL <https://doi.org/10.48550/arXiv.2410.08126>.

- 702 Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob  
703 Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, et al. Open-ended  
704 learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808*, 2021.
- 705  
706 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
707 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Ar-  
708 mand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation  
709 language models. *CoRR*, abs/2302.13971, 2023a. doi: 10.48550/ARXIV.2302.13971. URL  
710 <https://doi.org/10.48550/arXiv.2302.13971>.
- 711 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-  
712 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher,  
713 Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy  
714 Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,  
715 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel  
716 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya  
717 Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar  
718 Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan  
719 Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen  
720 Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan  
721 Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez,  
722 Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-  
723 tuned chat models. *CoRR*, abs/2307.09288, 2023b. doi: 10.48550/ARXIV.2307.09288. URL  
724 <https://doi.org/10.48550/arXiv.2307.09288>.
- 725 Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan,  
726 and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models.  
727 *arXiv preprint arXiv:2305.16291*, 2023a.
- 728 Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai  
729 Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents.  
730 *arXiv preprint arXiv:2308.11432*, 2023b.
- 731 Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, Haowei Lin,  
732 Zhaofeng He, Zilong Zheng, Yaodong Yang, Xiaojian Ma, and Yitao Liang. JARVIS-1:  
733 open-world multi-task agents with memory-augmented multimodal language models. *CoRR*,  
734 abs/2311.05997, 2023c. doi: 10.48550/ARXIV.2311.05997. URL [https://doi.org/10.](https://doi.org/10.48550/arXiv.2311.05997)  
735 [48550/arXiv.2311.05997](https://doi.org/10.48550/arXiv.2311.05997).
- 736 Zihao Wang, Shaofei Cai, Anji Liu, Xiaojian Ma, and Yitao Liang. Describe, explain, plan  
737 and select: Interactive planning with large language models enables open-world multi-task  
738 agents. *CoRR*, abs/2302.01560, 2023d. doi: 10.48550/ARXIV.2302.01560. URL <https://doi.org/10.48550/arXiv.2302.01560>.
- 739  
740 Zihao Wang, Shaofei Cai, Zhancun Mu, Haowei Lin, Ceyao Zhang, Xuejie Liu, Qing Li, Anji Liu,  
741 Xiaojian Ma, and Yitao Liang. Omnijarvis: Unified vision-language-action tokenization enables  
742 open-world instruction following agents. *CoRR*, abs/2407.00114, 2024. doi: 10.48550/ARXIV.  
743 2407.00114. URL <https://doi.org/10.48550/arXiv.2407.00114>.
- 744  
745 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yo-  
746 gatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol  
747 Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large language mod-  
748 els. *Trans. Mach. Learn. Res.*, 2022, 2022a. URL [https://openreview.net/forum?](https://openreview.net/forum?id=yzkSU5zdwD)  
749 [id=yzkSU5zdwD](https://openreview.net/forum?id=yzkSU5zdwD).
- 750 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi,  
751 Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language  
752 models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.),  
753 *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Informa-*  
754 *tion Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December*  
755 *9, 2022*, 2022b. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/](http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html)  
[9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html).

- 756 Yue Wu, So Yeon Min, Shrimai Prabhunoye, Yonatan Bisk, Russ R Salakhutdinov, Amos Azaria,  
757 Tom M Mitchell, and Yuanzhi Li. Spring: Studying papers and reasoning to play games. *Advances*  
758 *in Neural Information Processing Systems*, 36, 2024.
- 759 Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe  
760 Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents:  
761 A survey. *arXiv preprint arXiv:2309.07864*, 2023.
- 762 Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson  
763 env: Real-world perception for embodied agents. In *Proceedings of the IEEE conference on*  
764 *computer vision and pattern recognition*, pp. 9068–9079, 2018.
- 765 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan  
766 Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International*  
767 *Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenRe-  
768 view.net, 2023. URL [https://openreview.net/pdf?id=WE\\_vluYUL-X](https://openreview.net/pdf?id=WE_vluYUL-X).
- 769 Haoqi Yuan, Chi Zhang, Hongcheng Wang, Feiyang Xie, Penglin Cai, Hao Dong, and Zongqing  
770 Lu. Plan4mc: Skill reinforcement learning and planning for open-world minecraft tasks. *arXiv*  
771 *preprint arXiv:2303.16563*, 2023.
- 772 Yan Zeng, Ruichu Cai, Fuchun Sun, Libo Huang, and Zhifeng Hao. A survey on causal reinforce-  
773 ment learning. *arXiv preprint arXiv:2302.05209*, 2023.
- 774 Kun Zhang and Aapo Hyvärinen. On the identifiability of the post-nonlinear causal model. In  
775 *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 647–  
776 655, 2009.
- 777 Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional  
778 independence test and application in causal discovery. In *Proceedings of the Twenty-Seventh*  
779 *Conference on Uncertainty in Artificial Intelligence*, pp. 804–813, 2011.
- 780 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,  
781 Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: a survey on hallucination in large  
782 language models. *arXiv preprint arXiv:2309.01219*, 2023.
- 783 Yujia Zheng, Biwei Huang, Wei Chen, Joseph Ramsey, Mingming Gong, Ruichu Cai, Shohei  
784 Shimizu, Peter Spirtes, and Kun Zhang. Causal-learn: Causal discovery in python. *Journal*  
785 *of Machine Learning Research*, 25(60):1–8, 2024.
- 786 Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. In  
787 *International Conference on Learning Representations*, 2019.
- 788 Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li,  
789 Lewei Lu, Xiaogang Wang, et al. Ghost in the minecraft: Generally capable agents for open-world  
790 environments via large language models with text-based knowledge and memory. *arXiv preprint*  
791 *arXiv:2305.17144*, 2023.
- 792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

## A LLM'S PRIOR ON MINECRAFT.

We utilize the Minecraft crafting recipes to construct an MC-QA dataset (introduced in Appendix B), aiming to evaluate the prior knowledge of various LLMs on Minecraft. We test and determine that 0.3 is the optimal temperature as shown in Fig. 9a. Then we test various LLMs on this dataset. The GPT series (Ouyang et al., 2022; Achiam et al., 2023) show significantly stronger Minecraft prior knowledge than other LLMs as shown in Fig. 9b.

Utilizing this dataset to fine-tune LLMs can improve their prior knowledge on Minecraft as shown in Fig. 9c. On the other hand, by modifying the crafting recipes in Minecraft, we can make the SOTA LLMs (e.g., GPT series) have no prior of this modified environment. This setup enables us to distinctively analyze the roles of *prior knowledge* and *inference capability* as shown in Fig. 9d, which serves as the basis of our ablation study.

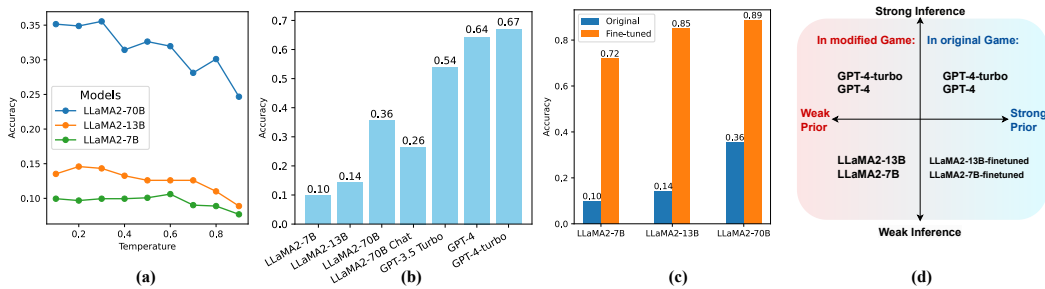


Figure 9: (a) LLaMA2 (Touvron et al., 2023b) demonstrates optimal accuracy in answering crafting recipes at a temperature of 0.3, measured as the ratio of correct answers to total questions. (b) Performance of open-source LLMs and GPT series models, showcasing their inherent prior knowledge of Minecraft. (c) Illustration of the improvement in performance for open-source LLMs fine-tuned with the crafting recipe dataset. (d) Categorizing the LLMs into four types based on their prior knowledge of Minecraft recipes and inference capabilities.

## B MC-QA DATASET

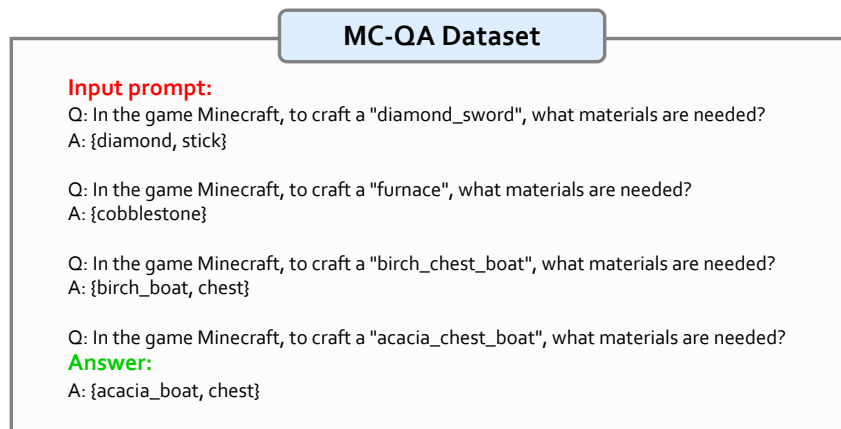


Figure 10: An example of the questions in MC-QA dataset.

Minecraft is an open-world game with a high degree of freedom. Players can freely gather materials (e.g., log 🪵), mine ore (e.g., iron\_ore ⬛), craft tools (e.g., wooden\_pickaxe 🪓), and more. Crafting recipes (e.g., planks 🪵 + sticks 🪵 → wood\_pickaxe 🪓) are the main knowledge in the Minecraft game, and serve as the basis for players to climb the technology tree and obtain more advanced items.



864 Given this, LLMs’ mastery of crafting recipes can reflect the strength of their prior knowledge in the  
 865 Minecraft game. We utilize the crafting recipes in Minecraft (version 1.19) to create the MC-QA  
 866 dataset. An example of the QA pairs in the dataset is shown in Fig. 10. The questions in this dataset  
 867 ask for the crafting ingredients required to obtain higher-level items in the technology tree, and the  
 868 answers are the ingredient items. LLMs need to give their answers in the specified format. The order  
 869 of the items in the answer is not required. For each question, we provide 3 examples to help LLMs  
 870 understand the QA task and the format of the answers. For situations where there are multiple ways  
 871 to craft the same item, we take them all into account to avoid the model being biased toward a fixed  
 872 understanding of the game. The dataset contains 754 QA pairs on the knowledge of obtaining items  
 873 in the Minecraft.

## 874 875 876 877 878 C AGENT IN MINECRAFT

880  
881  
882  
883 RL explorations in Minecraft agents focus on the efficient use of data (Baker et al., 2022; Fan et al.,  
 884 2022), hierarchical RL design (Lin et al., 2021; Mao et al., 2022), innovative architecture modeling  
 885 (Hafner et al., 2023), *etc.* Hafner et al. (2023) use world models to achieve a general and scalable RL  
 886 without human data or curricula. Much work (Guss et al., 2019; 2021; Kanervisto et al., 2022; Fan  
 887 et al., 2022) has made different simplifications for the Minecraft environment to facilitate the RL  
 888 agent systems. The MineDojo framework (Fan et al., 2022) provides an internet-scale knowledge  
 889 database and game environments for CLIP model (Radford et al., 2021) and RL training. These  
 890 efforts provide efficient optimization for agent sampling and interaction, but there is still a gap to  
 891 commercial Minecraft games with complete game features like in **VOYAGER** (Wang et al., 2023a)  
 892 and our work.

893 The reasoning capabilities and rich prior knowledge of LLMs have contributed to much work on  
 894 Minecraft agents (Yuan et al., 2023; Zhu et al., 2023; Wang et al., 2023a; Qin et al., 2023a; Notting-  
 895 ham et al., 2023; Wang et al., 2023d;c). **VOYAGER** (Wang et al., 2023a) and **GITM** (Zhu et al.,  
 896 2023) use LLMs’ prior knowledge of Minecraft and environment feedback to complete exploration  
 897 tasks in a text-based manner. Qin et al. (2023a) leverage MLLMs to introduce visual information as  
 898 the contextual basis for action execution. These methods more or less rely on prior knowledge of  
 899 Minecraft. Our ADAM shows effectiveness even when the game rules are modified.

900 There are also agents that operate in non-commercial Minecraft environments and utilize frame-level  
 901 control, including **Plan4MC** (Yuan et al., 2023), **DEPS** (Wang et al., 2023d), **JARVIS-1** (Wang  
 902 et al., 2023c), **OmniJARVIS** (Wang et al., 2024), and **Optimus-1** (Li et al., 2024). They need  
 903 approximately  $10^4$  steps for a task, in contrast to the action-level control in our system, which  
 904 involves around  $10^2$  steps for a task. Furthermore, these RL methods require training stages, whereas  
 905 our system does not involve weight updates. Notably, JARVIS-1 incorporates crafting recipes **as an**  
 906 **integral part** of the system, utilizing prior knowledge rather than learning from scratch.

907 Due to the complexity of the Minecraft environment, **Crafter** (Hafner, 2022) simplifies the game  
 908 space while maintaining the game rules of Minecraft. Crafter is a 2D game environment with craft-  
 909 ing routes similar to Minecraft. Agents such as **SPRING** (Wu et al., 2024) and **Mars** (Tang et al.,  
 910 2024) have been developed on this environment.

911 In the original implementation of **VOYAGER**, the environment feedback provides **detailed infor-**  
 912 **mation** such as crafting recipe errors (*e.g.*, “I cannot make an iron chestplate because I need: 7 more  
 913 iron ingots.”). We retain this informative feedback in all our experiments with **VOYAGER**, even in  
 914 environments with modified crafting recipes where the feedback may not align with the changes.

915 Tab. 5 shows the environmental information used by ADAM and **VOYAGER**. This setup is used in  
 916 our experiments at Sec. 4.3. **VOYAGER** does not have a visual input and needs omniscient metadata  
 917 (Meta) which is not explicitly exposed to human players, while ADAM utilizes visual input and other  
 observable information to make decisions.

	<b>VOYAGER</b>	<b>VOYAGER w/o Meta</b>	<b>ADAM</b>	<b>ADAM w/o MLLM</b>
Observation Space	Environment Feedback, Inventory, Meta	Environment Feedback, Inventory	Pixels, Inventory	Inventory
Action Space	Code	Code	Discrete	Discrete

Table 5: Comparison of environmental information used by ADAM and VOYAGER. VOYAGER does not have a visual input and needs to directly read the game information (Meta) which is not explicitly exposed to human players, while ADAM relies on the visual input and the information observable by human players to make decisions.




<b>Action Type</b>	<b>Action-Item Dependency</b>	<b>Data Quality</b>	<b>Skill Level</b>
Collecting 	Simple	Noise	Low
Crafting 	Complex	Clean	High
Smelting 	Complex	Noise	High

Table 6: Three action types in our experiment setting. "Smelting" actions have complex causal subgraphs, often leading to omissions in LLM-based CD. "Collecting" actions have noisy sampling data, and the results of the LLM-based CD are often redundant. The "Crafting" actions have complex causal subgraphs and clean sampling data.

## D IMPLEMENTATION DETAILS

### D.1 PROMPT

The prompt for LLM-based CD is shown in Fig. 11, which is composed of 5 components: (1) Role Playing, which assigns a specific role to the LLM; (2) Problem Setting, which provides specific details of the inference task; (3) Letter Mapping, which involves mapping item names to letters, a simplification that facilitates the formatted output; (4) Few-shot Prompting, which involves providing the LLM with several inference examples in chain-of-thought (Wei et al., 2022b) style; (5) Data  $D_a$ , which is the data collected by the interaction module for inference.

### D.2 ACTION SPACE AND MOVEMENT SPACE

We have implemented 41 discrete actions and 6 movements to ensure the agent can freely explore the Minecraft world through a diverse range of combinations. The actions can be divided into three categories: "Smelting", "Collecting", and "Crafting". "Smelting" actions have complex causal subgraphs, often leading to omissions in LLM-based CD. "Collecting" actions have noisy sampling data, and the results of the LLM-based CD are often redundant. The "Crafting" actions have complex causal subgraphs and clean sampling data.

The movement space corresponds to the low-level movement control visible to human players, including moving forward / backward, lowering / raising the agent's coordinates and turning left / right.

## E ROBUSTNESS

Tab. 7 shows our experiment result in the modified Minecraft environment where the crafting recipes are altered. ADAM can maintain its performance as it is equipped with CD methods, whereas agents that rely on prior knowledge struggle to explore efficiently. The result demonstrates the robustness and generalization capabilities of our ADAM architecture.

972 You are a Minecraft game analysis assistant. Our task is to infer the effect of an action and explore the causal  
 973 relationship by analyzing the items consumed and generated before and after an action. An action is a single, atomic  
 974 action consisting of  
 975 1. "Crafting" type actions, that is, combining raw materials into an item.  
 976 2. "Collecting" type actions, that is, collecting certain items. There may be some by-products in this process.  
 977 3. "Smelting" type actions, that is, consuming fuel and raw metal materials, and obtaining smelted items. A furnace is  
 978 needed.

979 It's possible that the action didn't have the expected effect, or that some additional items were collected, and we'll give  
 980 multiple sampling records to ensure robustness. Such records all correspond to the execution of the same action. Your  
 981 answer will only refer to the alphabetical codes of these items.  
 982 The answer should be in the format {Cause; Effect}

983 (Comment: The following is an example of the Observed Item Space. "m" and "n" are for few-shot learning and are not  
 984 included in the technology tree to avoid prior)  
 985 Items involved in the game include  
 986 'a': 'log'  
 987 'b': 'planks'  
 988 'c': 'crafting\_table'  
 989 'd': 'stick'  
 990 'e': 'wooden\_pickaxe'  
 991 'f': 'cobblestone'  
 992 'g': 'stone\_pickaxe'  
 993 'h': 'raw\_iron'  
 994 'i': 'furnace'  
 995 'j': 'iron\_ingot'  
 996 'k': 'iron\_pickaxe'  
 997 'l': 'diamond'  
 998 'm': 'redstone'  
 999 'n': 'redstone\_torch'

1000 log corresponds to a variety of log named xx\_log  
 1001 planks correspond to a variety of planks named xx\_planks

1002 1. Initial items: iron\_pickaxe, crafting\_table ; Consumed items: ; Added items: redstone, cobblestone,  
 1003 cobbled\_deepslate  
 1004 2. Initial items: iron\_pickaxe, crafting\_table ; Consumed items: ; Added items: cobblestone, redstone, andesite  
 1005 3. Initial items: iron\_pickaxe, crafting\_table ; Consumed items: ; Added items: granite  
 1006 4. Initial items: iron\_pickaxe, crafting\_table ; Consumed items: ; Added items: cobblestone, cobbled\_deepslate  
 1007 5. Initial items: iron\_pickaxe, crafting\_table ; Consumed items: ; Added items: redstone, cobblestone

1008 Your inference:  
 1009 Because there is no item consumption, this is a "Collecting" task. Redstone is a higher-level mineral and needs to be  
 1010 collected with an iron\_pickaxe ('k'). The rest of the items are considered additional items obtained during collection and  
 1011 will not be considered. So my answer is:  
 1012 Cause; Effect={k; m}

1013 1. Initial items: stick, redstone, crafting\_table ; Consumed items: stick, redstone; Added items: redstone\_torch  
 1014 2. Initial items: stick, redstone, crafting\_table ; Consumed items: ; Added items:  
 1015 3. Initial items: stick, redstone, crafting\_table ; Consumed items: redstone, stick; Added items: redstone\_torch  
 1016 4. Initial items: stick, redstone, crafting\_table ; Consumed items: stick, redstone; Added items: redstone\_torch  
 1017 5. Initial items: stick, redstone, crafting\_table ; Consumed items: redstone, stick; Added items: redstone\_torch

1018 Your inference:  
 1019 Because redstone and sticks are being used to craft redstone torches, this is a "Crafting" task. In addition to raw  
 1020 materials, the "Crafting" task also requires crafting\_table ('c'). So my answer is:  
 1021 Cause; Effect={c, d, m; n}

1022 (Our Data  $D$ )  
 1023 Your inference:

Figure 11: The prompt for LLM-based CD. The contents in red will be replaced in the inference process.

## 1017 F GENERALIZATION

1019 The ADAM architecture is a general framework for embodied agents operating in various open-  
 1020 world environments including Minecraft. When adapting ADAM to other application scenarios,  
 1021 some modifications may be necessary:

- 1023 (1) The world knowledge in Minecraft is the dependence between items and actions. Conse-  
 1024 quently, in this paper, items and actions are designed as causal graph nodes. When mi-  
 1025 grating to other environments, key elements related to the agent's task objectives can be  
 similarly designed as causal graph nodes.

Framework	Wooden Tool	Stone Tool	Iron Tool	Diamond
React w/ TM w/ SD	91 ± 34(2/3)	139(1/3)	N/A (0/3)	N/A (0/3)
Reflexion w/ TM w/ SD	76 ± 28(2/3)	120 ± 40(2/3)	N/A (0/3)	N/A (0/3)
AutoGPT w/ TM w/ SD	82 ± 25(2/3)	124(1/3)	N/A (0/3)	N/A (0/3)
VOYAGER	95 ± 33(2/3)	152 ± 43(2/3)	N/A (0/3)	N/A (0/3)
VOYAGER Guided	108 ± 35(2/3)	176(1/3)	N/A (0/3)	N/A (0/3)
ADAM	28 ± 4(3/3)	52 ± 14(3/3)	94 ± 27(3/3)	109 ± 34(2/3)
ADAM Parallel	<b>15 ± 2(3/3)</b>	<b>31 ± 7(3/3)</b>	<b>54 ± 14(3/3)</b>	<b>61 ± 18(2/3)</b>

Table 7: Performance in the modified Minecraft game. Each method has three trials for a maximum length of 200 steps. The success rate is depicted in the parentheses. ADAM can maintain its performance as it is equipped with CD methods, whereas agents that rely on prior knowledge struggle to explore efficiently. The result demonstrates the robustness and generalization capabilities of our ADAM architecture.

- (2) In ADAM, the perception module utilizes a vision-based MLLM and does not rely on omniscient metadata. This allows the module to adapt well to other visual tasks. If conditions permit (*e.g.*, a robot equipped with LiDAR), the perception module can provide more precise information, potentially further improving performance.
- (3) To model actions as a finite set of causal graph nodes, it is necessary to discretize the continuous action space. The granularity of this discretization should be determined based on the specific environment.