Causality-driven Hierarchical Structure Discovery for Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Hierarchical reinforcement learning (HRL) has been proven to be effective for tasks with sparse rewards, for it can improve the agent's exploration efficiency by discovering high-quality hierarchical structures (e.g., subgoals or options). However, automatically discovering high-quality hierarchical structures is still a great challenge. Previous HRL methods can only find the hierarchical structures in simple environments, as they are mainly achieved through the randomness of agent's policies during exploration. In complicated environments, such a randomness-driven exploration paradigm can hardly discover high-quality hierarchical structures because of the low exploration efficiency. In this paper, we propose CDHRL, a causality-driven hierarchical reinforcement learning framework, to build high-quality hierarchical structures efficiently in complicated environments. The key insight is that the causalities among environment variables are naturally fit for modeling reachable subgoals and their dependencies; thus, the causality is suitable to be the guidance in building high-quality hierarchical structures. Roughly, we build the hierarchy of subgoals based on causality autonomously, and utilize the subgoal-based policies to unfold further causality efficiently. Therefore, CDHRL leverages a causality-driven discovery instead of a randomness-driven exploration for high-quality hierarchical structure construction. The results in two complex environments, 2D-Minecraft and Eden, show that CDHRL can discover high-quality hierarchical structures and significantly enhance exploration efficiency.

1 Introduction

Reinforcement learning (RL) enables the intelligence of the agents by learning to take actions in the environments for maximum rewards [28, 20, 18]. For complex tasks with large state spaces and sparse delayed rewards [12], hierarchical reinforcement learning (HRL) extends RL methods' successes by discovering beneficial hierarchical structures, which helps the agent train multiple levels of policies and explore efficiently using high-level actions [23].

The major challenge of HRL methods is how to discover high-quality hierarchical structures (e.g., subgoals, skills, or options)[1]. A straightforward method is to manually define the hierarchical structures based on task-specific prior knowledge. However, in observing that the hand-crafted hierarchical structures require human expertise and yield in low generality, some researchers propose methods to discover subgoals automatically [21, 17, 34, 22]. These methods focus on how to discover subgoals from experience but follow an inefficient randomness-driven exploration paradigm to obtain experience. Specifically, the agent explores the environment through the randomness of the policy, which is implemented by raising the entropy of the policy or adding random noise to actions, and then

[1] In this paper, the hierarchical structure refers to the subgoal hierarchy in the environment.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

discovers subgoals from the collected experience. Randomness-driven exploration paradigm hardly discovers the high-quality hierarchical structure in complicated environments, such as Minecraft, because of low exploration efficiency, thus limiting the HRL method's performance.

To address the above issue, we propose a Causality-Driven Hierarchical Reinforcement Learning (CDHRL) framework which leverages the causality in the environment as the guidance to discover high-quality hierarchical structures. The key insight is that causality in a system, which explains the dependencies among objects, states, events, or processes, can be naturally used to describe a final goal with sub-goals structurally. Therefore, in CDHRL, we use the causalities among environment variables to guide the discovery of the subgoal hierarchy. Such a causality-driven paradigm avoids the inefficient random exploration process, and significantly enhances the efficiency of discovering hierarchical structures in complicated environments. Moreover, the generated high-quality hierarchical structure, which consists of subgoals and their dependency, is more instructive and efficient during training than the traditional fully-connected hierarchical structure.

In implementing CDHRL, we utilize an iterative boosting framework to progressively discover the causality and construct the hierarchical structure, since it is difficult to discover all causalities in complicated environments directly. The proposed CDHRL consists of two processes, causality discovery and subgoal hierarchy construction. Specifically, causality discovery refers discovery of causality between environment variables. In subgoal hierarchy construction, we find reachable subgoals, which are changes of controllable environment variables, and leverage causality to construct dependency of subgoals. The above two processes can iteratively boost each other: On one side, we can transform discovered causality between environment variables into the subgoal hierarchy. On the other side, we can utilize trained subgoal-based policies to enhance the efficiency of causality discovery. Taking a simplified example in Minecraft play, stonepickaxe (SP) and ironore(IO) are environment variables. After the agent learned subgoals of getting SP, CDHRL could control SP's distribution to discover causality from SP to IO, which indicates that the number of stonepickaxes affects the probability of acquiring ironore. Then, CDHRL adds subgoals about the effect variable IO to the subgoal hierarchy and trains corresponding subgoal-based policies. After that, we in turn discover further causality with IO as cause efficiently by utilizing subgoals about IO to change IO's distribution and collect intervention data. Such a causality-driven discovery paradigm makes the two processes, causality discovery and subgoal hierarchy construction, iteratively promote each other.

We verified our method in two typical complex tasks, including 2d-Minecraft [27] and simplified sandbox survival games Eden [4]. The results show that CDHRL discovers high-quality hierarchical structures and significantly enhances exploration efficiency. Compared to the state-of-the-art HRL method HAC [17] and HRL enhanced by curriculum learning [24], our method significantly outperforms existing studies in terms of both final performance and learning speed.

2 Related Works

2.1 Hierarchical Reinforcement Learning

Discovery of hierarchical structure plays an essential role in the Hierarchical Reinforcement Learning (HRL) algorithm. Some methods reduce the difficulty of discovering the hierarchical structure by adding artificial prior, such as sub-tasks dependency graph [27] and manually defined subgoals [9, 29]. Others may manually restrict the form of hierarchical structure to reduce the search space [11, 10, 6]. By introducing prior information, these methods perform well on specific tasks. However, it not only requires solid expert knowledge but also sacrifices versatility. Our work does not introduce downstream task-specific information, and thus has a broader application.

Some researchers propose methods to discover subgoals automatically. Vezhnevets et al. [30] and Nachum et al. [21] restrict the goal space as state space and train goal-based hierarchical policy end-to-end. Levy et al. [17] enable the multi-level hierarchical structure with hindsight corrections. Some methods discover hierarchical structure by optimizing additional objectives. Zhang et al. [34] assume that the subgoals should be near to the current state. Nachum et al. [22] claim that the optimal policy based on subgoals should have minimal loss than that based on primitive actions. Pitis et al. [24] consider the training progress for finding the next desired subgoal. Existing methods focus on discovering subgoals from data that is randomly collected. They are often applied in easy-to-explore environments like Maze [2] or Robot [8], but hard work in a complicated environment due to the

low exploration efficiency of their randomness-driven exploration paradigm. CDHRL follows the causality-driven exploration paradigm, which is more suitable for complicated environments.

2.2 Causal Reinforcement Learning

Causal reinforcement learning is a research direction that combines the causal effect with reinforcement learning. Pitis et al. [25] leverages local causal structures to improve the sample efficiency of off-policy RL. Méndez-Molina et al. [19] makes a trade-off between exploration and exploitation based given causal graph. Most of them utilize predefined causality graph as prior knowledge to enhance RL methods, However, none of them automatically discover causality in the environment to guide the exploration of hierarchical structure. Causal discovery methods [13, 14, 16] discover causality between variables through the data, which is an important issue in causal learning. Our method introduces a causal discovery-based exploration dramatically enhances the discovery efficiency of hierarchical structures in complex environments.

2.3 Relational Reinforcement Learning

Relational Reinforcement Learning (RRL) represents entities, actions, and policy by relational language to combine relation learning or reasoning with RL[7], which shares similar motivation with our method. Zambaldi et al. [33] adopt self-attention mechanism upon entities-centric representation to implements relation learning . Wang et al. [31] learns particular relations using a graph neural network to promote zero-shot performance on continuous control tasks. Chitnis et al. [5] and Kokel et al. [15] both learn relation model and planner in first-order language described environment. RRL aims to combine relation induction and reasoning with RL to promote policy generalization, which often works in structural environments with a clear definition of logical predicates. Our method adopts causality, which is more simple and common than logical relation, to guide the direction of subgoals discovery and thus enhance its efficiency in complicated environments.

3 Preliminary

3.1 Subgoal-based Markov Decision Process (MDP)

Hierarchical reinforcement learning that adopts subgoals as hierarchical structure extends the MDP framework by introducing the subgoal space $\mathcal{G}$. Subgoal-based MDP is a six-tuple, $(\mathcal{S}, \mathcal{G}, \mathcal{A}, T, r, \gamma)$, including a state set $\mathcal{S}$, a subgoal space $\mathcal{G}$, and an action set $\mathcal{A}$. The transition probability function $T(s, a, s')$ represents the probability of transition from s to s' with executing action a. $r(s, g, a, s')$ represents the goal-reaching reward function that indicates whether the agent achieves the subgoal g in transition (s, a, s'). And γ is the discount rate. The target of reinforcement learning is to find a subgoal-based policy $\pi : \mathcal{S}, \mathcal{G} \rightarrow \mathcal{A}$ that maximizes the accumulated rewards $R_\pi(s, g) = E[\sum_{n=0}^{\infty} \gamma^n R_{t+n+1} | s_t = s, g_t = g]$. Our proposed framework aims to pre-train the subgoal-based policy for downstream sparse reward tasks.

3.2 Causal Discovery

Causal discovery aims to infer causality through variables data under causal modeling and assumptions. There are two key parts of it: how to represent causality and how to obtain the data for discovering causality.

Structural Causal Model (SCM) is a kind of model that describes the system's causality. SCM over a finite number M of random variables X_i includes a directed acyclic graph (DAG) $C \in \{0, 1\}^{M \times M}$ as the causality graph, and a set of generating functions f_i of variables:

$$X_i := f_i(X_{pa(i,C)}, N_i), \ \ \forall i \in \{0, \ldots, M-1\}. \tag{1}$$

$c_{ij} = 1$ indicates that X_j is the parent of X_i, which means that X_j is the direct cause of X_i. N_i is the jointly-independent noise and $Pa(i, C)$ is the collection of the parent nodes (direct cause variables) of X_i in causality graph C.

Intervention sampling is a typical operation in causal discovery. Unlike standard sampling, it sets the distribution of certain variables in the system as fix value or uniform distribution, then samples to get intervention data. Intervention data has more causal information than observation data. In

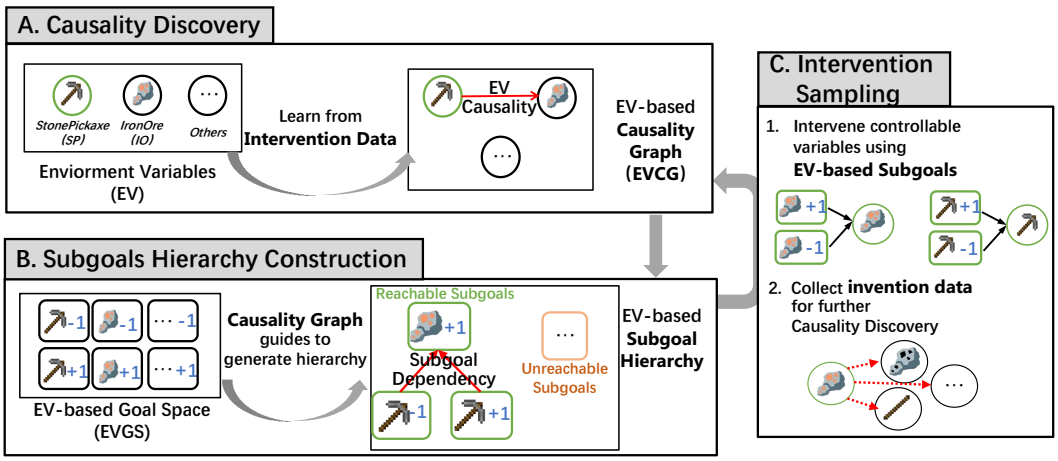

Figure 1: Overview of CDHRL framework. **A. Discover causality between Environment Variables (EVs):** Taking EV in Minecraft as a simplified example, we discover the causality between StonePickaxe (SP) and IronOre (IO) from SP's intervention data. **B. Build EV-based subgoal hierarchy for HRL subgoal training.** Guided by causality from SP to IO, we add reachable subgoals of increasing/decreasing IO values to the subgoal hierarchy and train corresponding subgoal-based policies. **C. Obtaining intervention data for causality discovery in the new round.** After getting subgoal-based policies about IO, the agent can further change IO's distribution and collect intervention data required in causality discovery (Step A) of the next iteration. (The pseudo-code of the framework is in Appendix B.1)

RL environments, we cannot do intervention directly because the variable generation mechanisms in the environment are unavailable. However, since the agent's behavior can affect the state of the environment, we train subgoal-based policies in CDHRL that can increase or decrease the values of variables to conduct the intervention.

4 Causality-Driven Hierarchical Reinforcement Learning

In this section, we present our Causality-Driven Hierarchical Reinforcement Learning (CDHRL) framework that leverages the causality of Environment Variables for high-efficient hierarchical structure discovery. Figure 1 shows the architecture of CDHRL. To enable the causality discovery and the subgoal hierarchy construction to promote each other, we first adopt environment variables to build a causality-based hierarchical structure that bridges causality discovery and hierarchical reinforcement learning (Section 4.1). Then, we provide the causality discovery methodology in the context of MDP (Section 4.2 & Section 4.4). Additionally, we propose the subgoal policy training methodology based on the causality-based hierarchical structure (Section 4.3).

4.1 Causality-based Hierarchical Structure

To bridge causality discovery and subgoal hierarchy construction, and to make them collaboratively promote each other, we introduce the main components of these two processes, causality graph and goal space, both based on environment variables. The **Environment Variables** (EV) are key factors of the environment state and can be seen as a kind of disentangled representation. Noted, to focus on the hierarchical structure discovery guided by the causality, we assume that EV can be obtained from state by an oracle mapping, $O : \mathcal{S} \rightarrow \mathcal{X}$. In fact, there are representation learning methods explicitly designed to learn the environment variables, such as CausalVAE[32] and DEAR[26]. We also conduct a sensitivity analysis on EV in Appendix A to verify the reasonableness of this assumption.

Environment Variable-based Causality Graph (EVCG). The causality graph describes the system's causality through a directed acyclic graph (DAG), where the source nodes of directed edges represent cause variables and destination nodes represent effect variables. The causality graph is shown in Figure 1.A. The process learns the environment variable-based causality graph (EVCG) of the RL environment, where nodes consist of environment variables and edges represent causality.

Environment Variables are key factors of environment state and the fundamental element of causality. Taking Minecraft as an example, environment variables consist of the number of tools (like StonePickaxe), materials (like IronOre), and surrounding objects (like Pigs). The causality represented by the arrow from StonePickaxe to IronOre means that the number of stone pickaxes the agent-owned affects the probability of acquiring iron ore.

Environment Variable-based Goal Space (EVGS). To leverage the discovered causality between the environment variables, we introduce environment variable-based goal space (EVGS) and construct the subgoal hierarchy upon EVGS, as shown in Figure 1.B. Different from prior studies that adopt the state or latent spaces as subgoal space [17, 22], EVGS consists of goals which are to change the value of environment variables. Thus the subgoal-based policies can be used to change the distribution of environment variables or complete downstream tasks. Formally, we introduce two fundamental changes, Increase (F_i) and Decrease (F_d), for variable values:

$$\begin{aligned} F_i(x_{i,t}, x_{i,t+1}) &= \mathbf{1}_{x_{i,t+1} > x_{i,t}} \\ F_d(x_{i,t}, x_{i,t+1}) &= \mathbf{1}_{x_{i,t+1} < x_{i,t}}. \end{aligned} \tag{2}$$

where $X_{i;t}$ is the value of X_i in state s_t. The EVGS is the Cartesian product of the environment variable set $\mathcal{X}$ and the changes set $\mathcal{F}_{change} = \{F_i, F_d\}$:

$$\mathcal{G}_{EVGS} = \{(X_i, F) | X_i \in \mathcal{X}, F \in \mathcal{F}_{change}\}. \tag{3}$$

The corresponding goal-reaching function $r(s_t, g, a, s_{t+1})$ is:

$$r(s_t, (X_i, F), a, s_{t+1}) = F(O(s_t)_i, O(s_{t+1})_i), \tag{4}$$

where O is the mapping from state to variable values.

Collaboration between Causality Discovery and Subgoal Hierarchy Construction. With the EV-based causality graph and EV-based goal space, we can naturally transform the causality graph into a subgoal hierarchy. Step A and B in Figure 1 show an example. Specifically, we divide the environment variables into *controllable* and unknown ones for the agent. The controllable variables are variables whose related subgoals have been added into the subgoal hierarchy (green ones, e.g., StonePickaxe) so that the agent can change their distribution to obtain their intervention data. Hence, in each iteration, the agent can discover causality from controllable variables (e.g., StonePickaxe) to the unknown variables (e.g., IronOre). Then, some unknown variables, which as the effect of the causality, are considered as new controllable variables for the agent, because the agent's behavior can affect them and thus recognize their cause variables. The subgoals related to new controllable environment variables are considered to be *reachable*. We then train reachable subgoal-based policies and add them to the subgoal hierarchy. In the next iteration, the agent can change the distribution of new controllable environment variables through subgoal-based policies and discover new causality which takes them as the cause variable. The iteration converges when there is no new causality, and the remained isolated variables are uncontrollable for the agent. Subgoals of uncontrollable variables will not be added to the subgoal hierarchy.

In this way, causality discovery and subgoal hierarchy construction can promote each other. On the one hand, the discovered causality between environment variables guides the agent to find available subgoals accurately and construct the subgoal hierarchy upon EVGS. On the other hand, the agent can utilize trained policies of subgoals to do intervention on environment variables. More complex causality between variables can be discovered along with the sampling data from poor to rich. The alternating iterative optimization above can address the challenge of discovering hierarchical structure.

4.2 Causality Learning

Causality in the Context of Markov Decision Process (MDP). We first introduce two key points of causality discovery in the context of MDP: what causality to discover and how to get learning data. 1) ***Causality is discovered within adjacent steps.*** According to the properties of MDP, the state s_{t+1} is only determined by state s_t and action a_t. Hence, we redefine the causality equation 1 as $X_{i,t+1} := f_i(X_{pa(i,C),t}, N_i)$, which means the variable values at step $t+1$ are only affected by that at step t. In other words, we aim to learn the causality from transition data of adjacent steps in the agent's trajectory. 2) ***Intervention data is obtained through subgoal-based policy.*** As mentioned in the preliminary, it is hard to do intervention in the RL environment directly since the generation mechanisms are unavailable. Therefore, we control the variables' distribution through the

agent's behavior to sample intervention data for causal discovery. Specifically, with the application of subgoals in EVGS, the agent can change the distribution of controllable environment variables and collect the intervention data subsequently. More details about the intervention process based on subgoal-based policy are introduced in Section 4.4.

Structural Causal Model Learning. The structural Causal Model (SCM) is a kind of model that describes the system's causality. We adopt a causal discovery methodology similarly to SDI [14] and intervention data collected from the agent's trajectories to learn the environment variables' SCM. As the equation 1 shown, SCM over M variables can be described by two sets of parameters: **structural parameters** $\eta \in R^{M \times M}$ model the causality graph C and **functional parameters** θs model the generating functions f.

The parameters η and θ are optimized through a two-phase iterative and continuous process. In the first function learning phase, we want to fix the input variables of generating function and optimize the functional parameters θ. So we sample hypothesis configuration C of the causality graph from Bernoulli distribution $Ber(\sigma(\eta))$, where σ represents the Sigmoid function. And then, we optimize θ by maximizing the likelihood of the intervention data. In the second structure learning phase, we want to learn the causality graph under fixed generating functions f. We estimate the gradient of η through fixed θ and a REINFORCE-like predictor proposed by Bengio et al. [3]. For the causality from variable X_j to X_i, the gradient g_{ij} of η_{ij} is estimated from X_j's intervention data set $\{X\}$ by the following formula:

$$g_{ij} = \sum_k (\sigma(\eta_{ij}) - c_{ij}^{(k)}) \left(\frac{L_{C,i}^{(k)}(X)}{\sum_{k'} L_{C,i}^{(k')}(X)} \right), \ \forall i \in \{0, \ldots, M-1\}, \tag{5}$$

where k represents the k-th draw of hypothesis configuration C. $L_{C,i}^{(k)}(X)$ represents the X_i's likelihood under the intervention data. The two optimization phases cycle until convergence. (Algorithm details are in Appendix B.2).

4.3 Subgoal-based Policy Training

As described in section 4.1, when new causalities are discovered, the subgoals related to new effect variables in the EVGS are added to the subgoal hierarchy by the guidance of new causalities. Figure 1.B shows the effect variable IO's subgoals are added to the subgoal hierarchy. Specifically, we first assumed the new effect variable X_i is controllable, and subgoals of X_i are reachable. We then train X_i's subgoal-based policies $\pi(a|s, (X_i, F))$. To induce an efficient training, the action space of $\pi(a|s, (X_i, F))$ consists of subgoals related to X_i's cause variables and primitive action set $\mathcal{A}$:

$$a \in \{(X_j, F') | X_j \in X_{pa(i,C)}, F' \in \mathcal{F}_{change}\} \bigcup \mathcal{A}. \tag{6}$$

Besides, to verify whether newly associated variables are controllable as supposed, we compare the final training success rates of the new subgoals with a preset threshold ϕ_{causal} before adding them to the subgoal hierarchy. When the corresponding subgoal-based policies are successfully trained, variables will be confirmed as controllable variables and can be intervened to obtain data in the next iteration. Along with the causality discovered from shallow to deep, the subgoal-based policies are trained layer by layer and finally consist of a multi-level hierarchy. After the above process converges to no new causality, subgoals related to remaining isolated variables are taken as unreachable ones. For details on subgoal-based policy training, please refer to the training algorithm in Appendix B.3.

Distinguishing reachable subgoals related to controllable environment variables and filtering unreachable ones during policy training can reduce the goal space reasonably. Meanwhile, through the dependency of subgoals suggested by causality between environment variables, we construct the action space for the subgoal-based policies, which is more efficient than primitive action space. The above two improvements in policy training lead to a more data-efficient learning process, verified in experiment results in Section 6.2).

4.4 Intervention Sampling

Considering that environment variables in RL environments cannot be directly intervened, we propose a methodology that adopts subgoal-based policies to obtain the intervention data. Benefit

from variable-based goal space, the subgoal-based policy could effectively change the distribution of controllable variables.

The detailed intervention sampling includes the following steps: **(1)** For controllable variable X_i, sample the desired value x_i from the uniform distribution, and then use subgoals (X_i, F) to drive the agent to change the value of X_i to x_i; **(2)** Collect the follow-up trajectory until X_i's value changes, and then obtain intervention data of X_i from all variables' values of adjacent steps. At the beginning of the optimization, there are no trained subgoal-based policies. Therefore, we set the agent's action as one of the variables first and do intervention on it to initiate the iterative process.

5 Experiment Setup

5.1 Environment

To verify the effectiveness of CDHRL, we choose two environments with discrete action space, namely 2d-Minecraft [27] and Eden [4] (more details of environments are introduced in Appendix C.1. They have complex hierarchical structures and very sparse rewards.

2D-Minecraft [27] is a simplified 2D version of the famous Minecraft [12]. In one episode with limited steps, the agent must navigate the map, pick up various materials, and craft tools until obtaining the diamond. The agent only can get a positive reward after getting the diamond.

Eden [4] is a survival game which is similar to "Don't Starve." To make a living in a 40×40 grid world, the agent with 8×8 limited vision range must chase animals to obtain food for maintaining satiety, find rivers to get water for preventing being thirsty, and collect materials to craft tools for surviving in the night. The agent takes survival as the unique target and cannot get any positive rewards except a negative reward after it dies and ends the episode.

5.2 CDHRL Implementation

We implement the agent based on multi-level DQN with HER [1], considering the action space is discrete. The training is divided into two stages, pre-training and adaptation. In the pre-training stage, we train multi-level subgoal-based policies until no new causality is discovered. In the adaptation stage, we train an upper controller on the pre-trained subgoals to maximize the task reward. The SCM is implemented like SDI [14]. The environment supports the oracle mapping O from environment state to variable values.

5.3 Baselines

We compare CDHRL with the following three methods for validation. For a fair comparison, we use the environment variable-based goal space $\mathcal{G}_{EVGS}$ to replace the goal space in these algorithms to ensure that the performance promotion is because of the difference in hierarchical structure discovery methods.

Oracle HRL (OHRL) is implemented as a two-level DQN with HER and an **oracle goal space**. The goal space is a subset of $\mathcal{G}_{EVGS}$ after artificial eliminating unreachable and useless subgoals.

HAC [17] is a state-of-the-art goal-conditioned HRL that discovers subgoals with a randomness-driven exploration paradigm. We implement a two-layer HAC **with subgoal space** $\mathcal{G}_{EVGS}$.

MEGA [24] is a kind of goal-conditioned HRL enhanced by curriculum learning, which pre-train subgoals in the order of their training progress. We set the initial subgoal distribution as the uniform distribution on $\mathcal{G}_{EVGS}$ and pre-trains subgoals with enough steps. After that, like the adaption stage in CDHRL, we train an upper controller on the trained subgoals.

6 Results

We first carry out extensive experiments to compare the performance of CDHRL and the state-of-the-arts in complex environments with a hard-to-explore hierarchical structure. Then, we dig deep ablation studies to show why CDHRL triumphs over the existing HRL paradigm. At last, we showcase the discovered causality-based hierarchical structure, which can be reasonably interpreted.

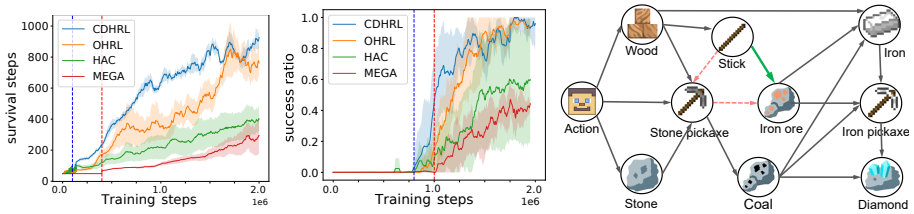

(a) Eden (b) 2D-Minecraft (c) Causality Graph in Minecraft

Figure 2: (a) Agent's survival time in Eden. (b) Success ratio in 2D-Minecraft. The vertical dotted lines indicate the end of pre-training of CDHRL and MEGA. Results are derived from average data in 8 trials. (c) The causal graph of 2D-Minecraft discovered by the agent. Some uncontrollable variables that unlinked are ignored here.

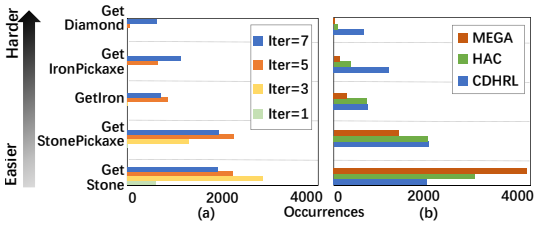

Figure 3: Exploration capability comparison. We select five explore milestones in 2D-Minecraft from easy-to-explore to hard-to-explore and record the occurrences in 10K test episodes to compare the agent's exploration capability. More occurrences on the hard-to-explore milestone represent higher exploration capability. (a) CDHRL's exploration capability iteratively increases along with the construction of hierarchical structures. (b) CDHRL shows much better exploration capability than HAC and MEGA. All compared methods are tested after trained 800K steps.

6.1 Comparative Analysis

We compare the performance on Eden and 2D-Minecraft of CDHRL with other methods. Figure 2 shows that CDHRL learns faster and achieves the best performance even though all the evaluated methods have the same subgoal space. Concretely, OHRL and HAC do not have the pre-training stage and start learning directly. However, the success ratios of OHRL and HAC in the first 1e6 steps of 2(b) keep nearly zero since they seldomly obtain the success trajectory data of getting diamonds. CDHRL learns much faster even if there is pre-training time overhead. MEGA has the pre-training stage for curriculum learning to identify the near-term and long-term subgoals. However, it fails to identify the effective hierarchical structures in these two environments with relatively large sub-goal space. The hierarchical structure discovered by CDHRL not only significantly enhances the agent's exploration capability in sparse reward, but also identifies the reachable subgoals that promotes the learning speed of subgoal-based policies. We will explain these phenomenons further in Section 6.2.

6.2 Insights

Further, we evaluate the detailed exploration capability and causal discovery efficiency and show how CDHRL outperforms existing HRL paradigms. We have two observations from the experiment results:

1) The causality-driven exploration paradigm enhances the exploration capability significantly. We experimentally evaluate the agent's exploration capability in 2D-Minecraft based on the achieved times of exploration milestones in the environment. Figure 3(b) shows the exploration capability comparison of different methods. The CDHRL's achieved times of hard-to-explore milestones are much higher than HAC and MEGA, indicating a much better exploration capability of CDHRL. Figure 3(a) shows the exploration capability increases along with iterations of causality discovery in CDHRL. The reason is that the agent's causality-driven exploration becomes more efficient as more subgoal-based policies and controllable variables, which other HRL methods cannot achieve. In conclusion, with the help of the causality-based exploration paradigm, the agent can achieve better exploration efficiency.

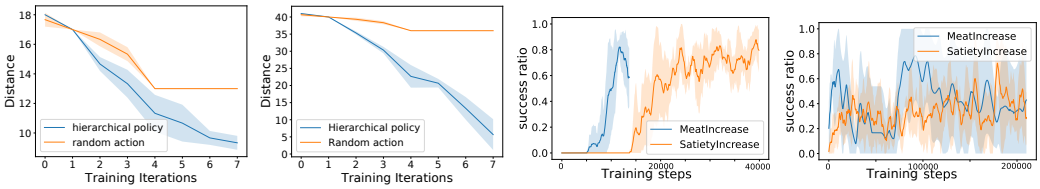

(a) SHD (b) SID (c) CDHRL (d) MEGA

Figure 4: (a) Structure Hamming Distance (SHD) and (b) Structural Interventional Distance (SID) between learned and ground-truth causality graph. The intervention data sampled by the assistance of hierarchical policy make the causality graph more accurate. The ground-truth causality graphs of Eden and 2D-Minecraft are in Appendix C.2.(c) and (d) are learning curve of two subgoals.

2) The subgoal hierarchy construction and causality discovery iteratively promote each other. On the one hand, the hierarchical structure makes the causality graph more accurate. We train the causal model in 2D-Minecraft using intervention data collected by hierarchical policy and random action, respectively. We record the causality graph of each iteration and compute their *Structure Hamming Distance* (SHD) and *Structural Interventional Distance* (SID) relative to the ground truth, where SHD and SID are typical metrics for the accuracy of causality graphs. The lower distance means the more accurate causality graph. Figure 4(a) and 4(b) show the result. The accuracy of the causality graph learned using the hierarchical policy is much better than using random action, especially for SID. Such results indicate that utilizing the hierarchical policy to obtain intervention data is essential for causality discovery.

On the other hand, the construction of the subgoal hierarchy is more efficient and reasonable. We compare the convergence speed and stableness of CDHRL and MEGA to achieve two subgoals in Eden. For these two subgoals, $MeatIncrease = (X_{Meat}, F_i)$ and $SatietyIncrese = (X_{Satiety}, F_i)$, the former is the prerequisite goal of the latter one and should be learnt by agent first. Figure 4(c) and 4(d) show that CDHRL achieves a faster and more stable learning process than MEGA. We attribute the poor performance of MEGA to the fact that it cannot identify the dependency of the subgoals and is unable to learn subgoals in order, resulting in much lower performance, especially when the environment is complex with a complicated dependency of subgoals.

In contrast, since X_{Meat} is the cause variable of $x_{Sariety}$, CDHRL infers the learning order of subgoals based on discovered causality and achieves stable and fast learning progress. Additionally, the reduced action space of multi-level policy transformed from causality also improves training efficiency.

6.3 An Showcase of Causality-based Hierarchical Structure

We showcase the established hierarchical structure based on discovered causality in 2D-MineCraft from the causality graph view. The learned causality graph can be reasonably interpreted. As shown in Figure 2(c), most of the causality in the graph meets with human cognition. We also find an interesting phenomenon that CDHRL sometimes may discover long-term causality and ignore corresponding short-term causality, e.g. “Stick” to “Iron re” (green edge) instead of “Stick” to “Stone pickaxe” to “Iron ore” (red edges) in Figure 2(c).CDHRL may benefit from such phenomenons because cutting unimportant causality details can improve the training speed and the generalization of hierarchical structure. Although there is some deviation between the discovered and the ground-truth causality graph, the causality-based hierarchical structure still improves the exploration efficiency.

7 Conclusion

Hierarchical reinforcement learning methods can hardly work well in complicated environments because of the low exploration efficiency of the randomness-driven exploration paradigm when discovering the hierarchical structure. We propose a causality-driven hierarchical reinforcement learning framework (CDHRL) to address the challenge, which autonomously builds high-quality hierarchical structures in an iterative boosting way. In two complicated environments with very sparse rewards, 2d-Minecraft and Eden, our method discovers a high-quality subgoal hierarchy and significantly enhances exploration efficiency. Limited in causal discovery algorithm capability and the complexity of variable changes set, our method can only be applied in environments with discrete environment variables. We leave broadening the application scope of CDHRL as our future work.

References

[1] Marcin Andrychowicz, Dwight Crow, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *NIPS*, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/453fadbd8a1a3af50a9df4df899537b5-Abstract.html.

[2] Akhil Bagaria and George Dimitri Konidaris. Option discovery using deep skill chaining. In *ICLR*, 2020.

[3] Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Nan Rosemary Ke, Sébastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher J. Pal. A meta-transfer objective for learning to disentangle causal mechanisms. In *ICLR*, 2020. URL https://openreview.net/forum?id=ryxWIgBFPS.

[4] Ruizhi Chen, Xiaoyu Wu, Yansong Pan, Kaizhao Yuan, Ling Li, TianYun Ma, JiYuan Liang, Rui Zhang, Kai Wang, Chen Zhang, Shaohui Peng, Xishan Zhang, Zidong Du, Qi Guo, and Yunji Chen. Eden: A unified environment framework for booming reinforcement learning algorithms, 2021.

[5] Rohan Chitnis, Tom Silver, Joshua B. Tenenbaum, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Glib: Efficient exploration for relational model-based reinforcement learning via goal-literal babbling. In *AAAI*, 2021.

[6] Caleb Chuck, Supawit Chockchowwat, and Scott Niekum. Hypothesis-driven skill discovery for hierarchical deep reinforcement learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5572–5579. IEEE, 2020. doi: 10.1109/IROS45743.2020.9340891. URL https://doi.org/10.1109/IROS45743.2020.9340891.

[7] Saso Dzeroski, Luc De Raedt, and Hendrik Blockeel. Relational reinforcement learning. In Jude W. Shavlik, editor, *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998), Madison, Wisconsin, USA, July 24-27, 1998*, pages 136–143. Morgan Kaufmann, 1998.

[8] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *ArXiv*, abs/1802.06070, 2019.

[9] Carlos Florensa, Yan Duan, and Pieter Abbeel. Stochastic neural networks for hierarchical reinforcement learning. In *ICLR*, 2017. URL https://openreview.net/forum?id=B1oK8aoxe.

[10] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*. PMLR, 2018. URL http://proceedings.mlr.press/v80/florensa18a.html.

[11] Sébastien Forestier, Yoan Mollard, and Pierre-Yves Oudeyer. Intrinsically motivated goal exploration processes with automatic curriculum learning. *CoRR*, abs/1708.02190, 2017. URL http://arxiv.org/abs/1708.02190.

[12] William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. In *IJCAI*. ijcai.org, 2019. doi: 10.24963/ijcai.2019/339. URL https://doi.org/10.24963/ijcai.2019/339.

[13] Amin Jaber, Murat Kocaoglu, Karthikeyan Shanmugam, and Elias Bareinboim. Causal discovery from soft interventions with unknown targets: Characterization and learning. In *NIPS*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/6cd9313ed34ef58bad3fdd504355e72c-Abstract.html.

[14] 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.

[15] Harsha Kokel, Arjun Manoharan, Sriraam Natarajan, Balaraman Ravindran, and Prasad Tadepalli. Reprel: Integrating relational planning and reinforcement learning for effective abstraction. In *ICAPS*, 2021.

[16] Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural DAG learning. In *ICLR*, 2020. URL https://openreview.net/forum?id=rklbKA4YDS.

[17] Andrew Levy, George Dimitri Konidaris, Robert Platt Jr., and Kate Saenko. Learning multi-level hierarchies with hindsight. In *ICLR*, 2019. URL https://openreview.net/forum?id=ryzECoAcY7.

[18] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *ICLR*, 2016. URL http://arxiv.org/abs/1509.02971.

[19] Arquímides Méndez-Molina, Ivan Feliciano-Avelino, Eduardo F. Morales, and Luis Enrique Sucar. Causal based q-learning. *Res. Comput. Sci.*, 149(3):95–104, 2020. URL http://rcs.cic.ipn.mx/2020_149_3/Causal%20Based%20Q-Learning.pdf.

[20] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, and et al. Human-level control through deep reinforcement learning. *Nat.*, 518(7540):529–533, 2015. doi: 10.1038/nature14236. URL https://doi.org/10.1038/nature14236.

[21] Ofir Nachum, Shixiang Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *NIPS*, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/e6384711491713d29bc63fc5eeb5ba4f-Abstract.html.

[22] Ofir Nachum, Shixiang Gu, Honglak Lee, and Sergey Levine. Near-optimal representation learning for hierarchical reinforcement learning. In *ICLR*, 2019. URL https://openreview.net/forum?id=H1emus0qF7.

[23] Ofir Nachum, Haoran Tang, Xingyu Lu, Shixiang Gu, Honglak Lee, and Sergey Levine. Why does hierarchy (sometimes) work so well in reinforcement learning? *CoRR*, abs/1909.10618, 2019. URL http://arxiv.org/abs/1909.10618.

[24] Silviu Pitis, Harris Chan, Stephen Zhao, Bradly C. Stadie, and Jimmy Ba. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *ICML*. PMLR, 2020. URL http://proceedings.mlr.press/v119/pitis20a.html.

[25] Silviu Pitis, Elliot Creager, and Animesh Garg. Counterfactual data augmentation using locally factored dynamics. *ArXiv*, abs/2007.02863, 2020.

[26] Xinwei Shen, Furui Liu, Hanze Dong, Qing Lian, Zhitang Chen, and Tong Zhang. Disentangled generative causal representation learning. *ArXiv*, abs/2010.02637, 2020.

[27] Sungryull Sohn, Junhyuk Oh, and Honglak Lee. Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. In *NIPS*, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/018dd1e07a2de4a08e6612341bf2323e-Abstract.html.

[28] Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*. The MIT Press, 1999. URL http://papers.nips.cc/paper/1713-policy-gradient-methods-for-reinforcement-learning-with-function-approximation.

[29] Chen Tessler, Shahar Givony, Tom Zahavy, Daniel J. Mankowitz, and Shie Mannor. A deep hierarchical approach to lifelong learning in minecraft. In *AAAI*. AAAI Press, 2017. URL http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14630.

[30] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In *ICML*. PMLR, 2017. URL http://proceedings.mlr.press/v70/vezhnevets17a.html.

[31] Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. Nervenet: Learning structured policy with graph neural networks. In *ICLR*, 2018.

[32] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jijie Wang. Causalvae: Disentangled representation learning via neural structural causal models. *CVPR*, 2021.

[33] Vinícius Flores Zambaldi, David Raposo, Adam Santoro, Victor Bapst, Yujia Li, Igor Babuschkin, Karl Tuyls, David P. Reichert, Timothy P. Lillicrap, Edward Lockhart, Murray Shanahan, Victoria Langston, Razvan Pascanu, Matthew M. Botvinick, Oriol Vinyals, and Peter W. Battaglia. Deep reinforcement learning with relational inductive biases. In *ICLR*, 2019.

[34] Tianren Zhang, Shangqi Guo, Tian Tan, Xiaolin Hu, and Feng Chen. Generating adjacency-constrained subgoals in hierarchical reinforcement learning. In *NIPS*, 2020. URL `https://proceedings.neurips.cc/paper/2020/hash/f5f3b8d720f34ebebceb7765e447268b-Abstract.html`.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 7
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No] We adopt open-source experimental environments.

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We offer the code of our method in the supplemental material.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] There are no personally identifiable information or offensive content.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]