

Boosting Efficiency in Task-Agnostic Exploration through Causal Knowledge

Yupei Yang¹, Biwei Huang^{2*}, Shikui Tu^{1*} and Lei Xu^{1,3}

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

²Halicioğlu Data Science Institute (HDSI), University of California San Diego, USA

³Guangdong Institute of Intelligence Science and Technology, Zhuhai, China

{yupei_yang, tushikui, leixu}@sjtu.edu.cn, bih007@ucsd.edu

Abstract

The effectiveness of model training heavily relies on the quality of available training resources. However, budget constraints often impose limitations on data collection efforts. To tackle this challenge, we introduce *causal exploration* in this paper, a strategy that leverages the underlying causal knowledge for both data collection and model training. We, in particular, focus on enhancing the sample efficiency and reliability of the world model learning within the domain of task-agnostic reinforcement learning. During the exploration phase, the agent actively selects actions expected to yield causal insights most beneficial for world model training. Concurrently, the causal knowledge is acquired and incrementally refined with the ongoing collection of data. We demonstrate that causal exploration aids in learning accurate world models using fewer data and provide theoretical guarantees for its convergence. Empirical experiments, on both synthetic data and real-world applications, further validate the benefits of causal exploration. The source code is available at <https://github.com/CMACH508/CausalExploration>.

1 Introduction

Deep neural networks have been incredibly successful in various domains, such as milestone achievements in Go games and control tasks [Silver *et al.*, 2016; Tassa *et al.*, 2018]. One key factor contributing to such remarkable performance is the availability of high-quality data for model training. However, in many practical applications, it remains data-hungry due to limited data collection efforts imposed by budget constraints [Fang *et al.*, 2017; Yoo and Kweon, 2019; Robine *et al.*, 2023]. This highlights the essential need to enhance both the sampling and model learning efficiency.

In this paper, we introduce *causal exploration* to tackle this challenge, a novel framework that makes use of the underlying causal knowledge to boost the data collection and model training processes. On one hand, acquiring and understanding causal knowledge unveils the fundamental mechanisms

behind the data generation process, thereby reducing the exploration space. In contrast to random data collection, causal exploration allows systematic action planning based on the identified causal structures. On the other hand, causal knowledge reflects the cause-and-effect dependencies among variables. Through the incorporation of causal structural constraints into the model, we can acquire causal dynamics models that eliminate redundant dependencies, as opposed to non-causal dense models, which have shown to provide more accurate estimations [Seitzer *et al.*, 2021; Huang *et al.*, 2022; Wang *et al.*, 2022].

Specifically, we focus on boosting the sample efficiency and reliability of the world model in the realm of task-agnostic reinforcement learning (RL). Different from methods that learn a fixed task from scratch, task-agnostic RL agent first learns a global world model that gathers information about the true environment on the data collected during exploration. Then based on the predictions of the learned model, the agent could make quick adaptations to various downstream tasks in a zero-shot manner given task-specific reward functions. This learning setup exhibits excellent generalization performance but also imposes high requirements on the accuracy of the model [Pathak *et al.*, 2019]. However, the data collection and world model learning processes are usually expensive due to extensive environment interactions, especially in large state spaces where discovering the optimal policy can be highly challenging [Burda *et al.*, 2018].

To address this problem, our causal exploration-based approach revolves around three primary aspects. First, we employ constraint-based methods to discover causal relationships among environment variables. Second, we formulate the dynamics model under causal structural constraints to enhance its reliability. Third, we propose several ways based on causal knowledge to improve the sample efficiency during exploration. In particular, for causal discovery, we present an efficient online method that selectively eliminates noisy samples and strategically gathers informative data points in an incremental way. Moreover, we learn to actively explore towards novel states that are expected to contribute most to model training. The learning process of the exploration policy is driven by intrinsic rewards, which measure both the agent’s level of surprise at the outcome and the quality of the training caused by the selected data. During exploration, causal knowledge and the world model are continuously re-

*corresponding author

fined with the ongoing collection of data. Our key contributions are summarized below.

- In order to enhance the sample efficiency and reliability of model training with causal knowledge, we introduce a novel concept: causal exploration, and focus particularly on the domain of task-agnostic reinforcement learning.
- To efficiently learn and use causal structural constraints, we develop an online method for causal discovery and formulate the world model with explicit structural embeddings. During exploration, we train the dynamics model under a novel weight-sharing-decomposition schema that can avoid additional computational burden.
- Theoretically, we show that, given strong convexity and smoothness assumptions, our approach attains a superior convergence rate compared to non-causal methods. Empirical experiments further demonstrate the robustness of our online causal discovery method and validate the effectiveness of causal exploration across a range of demanding reinforcement learning environments.

2 Preliminary

We consider task-agnostic RL within the framework of a Markov decision process characterized by state space \mathcal{S} and action space \mathcal{A} . In addition, to integrate causal information, we make the following assumptions throughout this paper.

Assumption 1. (Causal Factorization). *Both the state space and action space can be factorized. That is, $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n \in \mathbb{R}^n$ and $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_c \in \mathbb{R}^c$.*

Assumption 2. (Causal Sufficiency). *The state variables s_t are fully observable without any hidden confounders.*

Assumption 3. (Faithfulness Condition). *For a causal graph \mathcal{G} and the associated probability distribution P , every true conditional independence relation in P is entailed by the Causal Markov Condition applied to \mathcal{G} .*

Given these commonly made assumptions for causal discovery methods, we next define *transition causality* over the transition variables from $\{s_{t-1}, a_{t-1}\}$ to s_t .

Definition 1. (Transition Causality). *Under the Markov condition, the causal structures are over the state-action variables $\mathcal{U} = \{\mathcal{S}_{i,t-1}\}_{i=1}^n \cup \{\mathcal{A}_{j,t-1}\}_{j=1}^c$ and $\mathcal{V} = \{\mathcal{S}_{i,t}\}_{i=1}^n$, which can be represented by a directed acyclic graph $\mathcal{G} = (\{\mathcal{U}, \mathcal{V}\}, \mathcal{E})$ and its adjacency matrix D . Here, \mathcal{E} denotes the edge set and $D \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{V}|}$.*

Note that all edges are from \mathcal{U} to \mathcal{V} . If $s_{i,t-1} \in \mathcal{U}$ has a causal edge to $s_{j,t} \in \mathcal{V}$, then we call $s_{i,t-1}$ a parent of $s_{j,t}$ and have $D(i, j) = 1$. Take the example in Figure 1: we have $D(1, 1) = 1$ because $s_{1,t-1}$ is a parent of $s_{1,t}$, while $D(2, 1) = 0$ because $s_{2,t-1}$ does not have a causal edge to $s_{1,t}$. Here, we assume that the structural constraints are invariant over time t . The causal identification theory under these appropriate definitions and assumptions has been given in existing work.

Proposition 1. *Under the aforementioned assumptions, the causal adjacency matrix D is identifiable (see [Huang et al., 2021; Wang et al., 2022; Ding et al., 2022]).*




Figure 1: An illustration of causal relationships from $\mathcal{U} = \{s_{1,t-1}, s_{2,t-1}, a_{t-1}\}$ to $\mathcal{V} = \{s_{1,t}, s_{2,t}\}$ in the RL system.

3 Related Work

Task-agnostic RL. Over recent decades, task-agnostic exploration strategies have been an active research area to attain generalization [Aubret et al., 2019]. Existing methods focus on designing appropriate forms of intrinsic rewards, which can be broadly categorized into three types: (1) the number of times a state has been visited, which helps to guide the agent towards unexplored regions [Bellemare et al., 2016; Machado et al., 2020]; (2) curiosity about the environment dynamics, which is usually formalized as prediction errors of future states [Pathak et al., 2017; Kim et al., 2018]; and (3) information gain, which aims to improve the agent’s knowledge about the environment by maximizing the mutual information [Duan et al., 2016; Shyam et al., 2019]. However, existing works usually pay little attention to the support that causal structures can offer in improving exploration efficiency.

RL with causal discovery. The intersection of causal discovery and RL has become a popular trend in recent years. [Zhu et al., 2019] uses RL to search for the causal graph with the best score. [Hu et al., 2022] proposes to learn a hierarchical causal structure for subgoal-based policy learning. [Ding et al., 2022] and [Mutti et al., 2023] focus on providing tractable formulations of systematic generalization in RL tasks by employing a causal viewpoint. [Yu et al., 2023] learns a causal world model to generate explanations for the decision-making process. Theoretical evidence about the advantages of using a causal world-model in offline RL is given in [Zhu et al., 2022]. However, most of these works focus on either extracting underlying causal graphs from given data in a particular environment or using random exploration strategies for data collection, instead of directly utilizing causal structure as guidance to improve exploration efficiency in task-agnostic RL [Kosoy et al., 2022]. Experimental comparisons with some existing works that are close to us have been conducted in Section 6.

4 Discovering and Utilizing Causality for Learning World Models

After establishing proper assumptions and definitions, we proceed to introduce the methodology part for causal exploration. In this section, we initially assume that the causal structures are known and show how to explicitly incorporate causal knowledge into the world model and utilize it for model training. Next, we give an estimation procedure for causal structures.




Figure 2: An illustration of model architectures under (a) naive causal factorization; (b) our sharing-decomposition schema.

4.1 Causal Constraints for Forward Model

During task-agnostic exploration, the agent learns a world model f_{w^c} with parameter w^c , serving as an abstraction of the ground truth transitions in the environment. In other words, the world model f_{w^c} is designed to enable agents to predict future state \hat{s}_t based on current state s_{t-1} and action \mathbf{a}_{t-1} , represented by $\hat{s}_t = f_{w^c}(s_{t-1}, \mathbf{a}_{t-1}, e_t)$, where e_t is the corresponding random noise. However, based on the understanding that causal structures within environmental variables are typically sparse rather than dense, as suggested by [Huang *et al.*, 2022], such a framework could contain unnecessary dependencies. For instance, in the context of Figure 1, the variable $s_{2,t-1}$ does not causally affect $s_{1,t}$, and is thus identified as a non-parent node. Consequently, we only need to take a subset of $(s_{t-1}, \mathbf{a}_{t-1})$ as inputs of the model.

To reflect these constraints, we explicitly consider the causal structures over state and action variables to model the one-step transition dynamics, formulated as:

$$\hat{s}_t = \prod_{i=1}^n f_{w^c}^i(D_i \odot (s_{t-1}, \mathbf{a}_{t-1}), e_{i,t}), \quad (1)$$

where \odot denotes element-wise product, D_i and $e_{i,t}$ are the i -th column of causal matrix D and noise term e_t , respectively. However, a naive implementation requires training n world models since each factored dimension has its unique parents. It is likely to result in an explosive growth in computational complexity as the state dimension and network size increase.

We propose a sharing-decomposition schema to address such a problem. It is unnecessary for all of these n networks to be totally different. Instead, each of these models could share the first several layers as a common embedding. Then following the sharing module, we design predictive networks for each dimension. Suppose w_i^c is the network parameter for the i -th dimension, it is a combination of the shared parameter σ and decomposed parameter θ_i , written as

$$w_i^c = \sigma \cup \theta_i. \quad (2)$$

During the training time, each model focuses on a different aspect of the state in the decomposition part θ_i but shares a common knowledge σ . The number of shared layers is a hyperparameter that allows for a trade-off between the sharing and decomposition parts. By training forward models under this schema, our approach can both utilize causal information of the ground environment dynamics to generate accurate predictions and achieve a significant reduction in model




Figure 3: Average F1-score of the identified causal graph compared to the ground truth graph across various causal discovery methods. More details are given in Section 6.

parameters and computation time compared with naive decomposition. Figure 2 illustrates our sharing-decomposition schema during causal exploration. Moreover, we use the prediction error as the loss function for the optimization of w^c :

$$L_f(w^c) = \frac{1}{2} \|\hat{s}_t - s_t\|_2^2. \quad (3)$$

4.2 Efficient Online Causal Relationship Discovery

In this section, we show how to identify the causal adjacency matrix D . According to Definition 1, this can be transformed into determining whether there exists an edge between each pair of nodes in the causal graph \mathcal{G} . To achieve this goal, we start from a complete graph and then iteratively remove unnecessary edges using Conditional Independence Tests (CIT). Given that the edges follow the temporal order without instantaneous connection, we extend the PC algorithm [Spirtes *et al.*, 2000] to handle time-lagged causal relationships based on Kernel-based Conditional Independence (KCI) test [Zhang *et al.*, 2012] to identify the causal adjacency matrix D .

While causal discovery algorithms typically necessitate the collection of substantial causal information through data, it's important to note that accumulating more samples does not always confer an advantage: as the sample size increases, the time cost of causal algorithms also rises. Figure 9 in the appendix illustrates that the execution time of the PC algorithm based on KCI tests experiences exponential growth as sample size increases. Therefore, prioritizing the enhancement of data quality over quantity becomes paramount. In order to reduce the cost of the identification process, we design an efficient online causal relationship discovery method: instead of using all of the coming data for causal identification, we selectively collect representative data points during exploration in an incremental way. Specifically, we use the minibatch similarity and sample diversity criteria introduced in [Yoon *et al.*, 2021] as our selection strategies, which are defined as

$$\begin{aligned} \text{Similarity} &= \frac{\nabla f_{w^c}(b_t^i) \nabla f_{w^c}(B_t)^\top}{\|\nabla f_{w^c}(b_t^i)\| \cdot \|\nabla f_{w^c}(B_t)\|}, \\ \text{Diversity} &= \frac{-1}{t-1} \sum_{p \neq i}^{t-1} \frac{\nabla f_{w^c}(b_t^i) \nabla f_{w^c}(b_t^p)^\top}{\|\nabla f_{w^c}(b_t^i)\| \cdot \|\nabla f_{w^c}(b_t^p)\|}. \end{aligned} \quad (4)$$

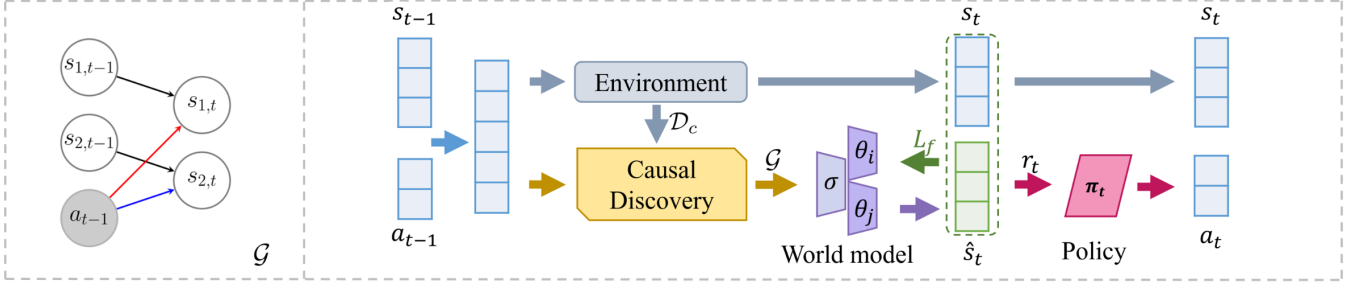


Figure 4: An overview of Causal Exploration framework. Throughout the process, the agent, guided by policy π_t , engages in exploration to gather data that are most beneficial for model training. Meanwhile, causal knowledge and the world model are continuously refined with the ongoing data collection.

Here b_t^i is the i -th data of the whole arrival batch \mathcal{B}_t . $\nabla f_{w^c}(b_t^i)$ and $\bar{\nabla} f_{w^c}(\mathcal{B}_t)$ are the gradient and average gradient of the sample and batch, respectively. A combination of these two criteria is used to select the top- κ data among \mathcal{B}_t for causal discovery during exploration. Experiments in Section 6 show that this online method significantly accelerates causal discovery without sacrificing overall accuracy.

It is noteworthy that our approach allows for flexibility in the choice of causal discovery methods. However, as has been discussed in [Ding *et al.*, 2022], constraint-based approaches usually exhibit better robustness compared with score-based methods under our specific configuration. Furthermore, we have also conducted comparative experiments among different causal discovery methods to demonstrate the correctness of the causal structure identified by the time-lagged PC algorithm in Section 6.

5 Boosting Efficiency through Causal Exploration

We now return to the fundamental question: how to enhance the data collection efficiency during causal exploration, thereby improving the performance of both causal discovery and model learning. To attain this goal, a commonly applied concept from active learning is the selection of samples that make the largest contributions to the model’s training loss. These samples are typically considered as a subset that the model is least familiar with. Hence, the prediction loss is used here as the intrinsic reward to guide exploration with a scaling weight η :

$$r_{t-1}^i = \frac{\eta}{2} \|\hat{s}_t - s_t\|_2^2. \quad (5)$$

This prediction loss can also be viewed as a validation of the agent’s causal beliefs. The larger the prediction error, the more surprised the agent is by the actual outcome, implying a greater deviation from the estimated values based on the causal structure and the world model. The faster the error rate drops, the more learning progress signals we acquire.

However, not all the novel states have a positive impact on the model. On the contrary, some noisy data may contribute significantly to prediction errors but can lead the model to an awful direction, which necessitates the agent to also pay attention to the inherent quality of the data during exploration.

To reflect this, we introduce active reward [Fang *et al.*, 2017] that measures the data quality as another intrinsic motivation. Once a new sample is collected at time step t , active reward is then calculated as the change of the model’s prediction ability before and after training. We use the prediction accuracy on a test set \mathcal{D}_h generated from episodes unseen before training to reflect the world model’s performance and formulate active reward as

$$r_{t-1}^a = \overline{L_f(w_{t-2}^c)} - \overline{L_f(w_{t-1}^c)}, \quad (6)$$

where $\overline{L_f(\cdot)}$ is the mean prediction error on \mathcal{D}_h and w_t^c denotes parameters of the trained world model at time t . The value of active reward reflects beneficial or detrimental training caused by the selected data. If the reward is always positive, it indicates that the agent has been selecting beneficial samples for training the world model. We combine prediction loss and active reward with a regularization weight β :

$$r_{t-1} = r_{t-1}^i + \beta r_{t-1}^a. \quad (7)$$

During causal exploration, the agent keeps searching for causal informative data by maximizing the expected rewards, which is

$$\mathbf{a}_t^* = \arg \max_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\tau \sim \pi} \left[\sum_t \gamma^t r_t \right], \quad (8)$$

where τ represents the trajectory generated by the exploration policy π and γ is the discount factor. Meanwhile, the world model minimizes the prediction loss. Since both of them contain the prediction error in equation (3) and (5), we can draw a conclusion that the learning of world models and causal exploration facilitate each other. Figure 4 shows an overview of causal exploration, details are shown in Algorithm 1.

5.1 Theoretical Analysis on Causal Exploration

In this subsection, we first present a mathematical criterion for evaluating the impact of causal exploration on sampling efficiency, and then we theoretically demonstrate the benefits of causal exploration in learning world models, especially when the causal graph is sparse.

Causal efficiency ratio. Let k denote the optimization step in world model training, f represent the world model, and

Algorithm 1 Task-agnostic Causal Exploration

```

1: Initialize: Forward world Model  $f$  with parameter  $w^c$ 
   Causal discovery period  $N$ , Dataset  $\mathcal{B}_t$ 
   Exploration policy  $\pi$  with memory buffer  $\mathcal{M}$ 
2: Output: Forward world model  $f$ 
3: for Episode = 1, 2, ... do
4:   Collect test set  $\mathcal{D}_h$  for the calculation of active reward
5:   Set current time  $t = 1$ 
6:   while  $t < T$  do
7:     Choose action  $\mathbf{a}_{t-1}$  and predict next state  $\hat{\mathbf{s}}_t$ 
8:     Obtain  $\mathbf{s}_t$  from environment
9:     Calculate intrinsic reward  $r_{t-1}$ 
10:    Store  $(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}, \mathbf{s}_t)$  into  $\mathcal{B}_t$ 
11:    Store  $(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}, r_{t-1}, \mathbf{s}_t)$  into  $\mathcal{M}$ 
12:    if  $t = 0 \pmod{N}$  then
13:      Do online causal discovery on top- $\kappa$  of  $\mathcal{B}_t$  and get
      causal graph  $\mathcal{G}$ 
14:    end if
15:    Train world model under  $\mathcal{G}$  on  $\mathcal{B}_t$  and update  $w^c$ 
16:    Train exploration policy on  $\mathcal{M}$  and update  $\pi$ 
17:  end while
18: end for
19: return Latest forward world model  $f$ 
    
```

$L_f(w)$ be the corresponding loss function. The model parameters incorporating causal knowledge after the k -th optimization are denoted as $w^c(k)$, while $w(k)$ represents the parameters without causal information¹. We use the convergence bounds ratio, denoted by ξ , calculated as the relative improvement in the model’s loss compared to the optimal value L_f^* , to assess the enhanced exploration efficiency resulting from the incorporation of causal knowledge:

$$\xi = \frac{L_f(w^c(k)) - L_f^*}{L_f(w(k)) - L_f^*}. \quad (9)$$

Note that this ratio measures how much closer the model with causal knowledge gets to the optimal value compared to those without causal information. If $\xi < 1$ consistently holds, it indicates an improvement in efficiency attributable to causal exploration. Next, we provide a theoretical analysis of the performance of causal exploration in linear cases, and empirical experiments are further conducted to show that causal exploration is still efficient in learning deep neural networks.

Theoretical guarantees. We consider learning a linear forward model with gradient descent, formulated as

$$\min_w L_f(w) = \|w^\top \cdot (\mathbf{s}_{t-1}, \mathbf{a}_{t-1}) - \mathbf{s}_t\|_2^2, \quad (10)$$

where $\mathbf{s}_t = w^{\star\top} \cdot (\mathbf{s}_{t-1}, \mathbf{a}_{t-1})$ is the ground truth value. Moreover, we make the following assumption:

Assumption 4. L_f is strong convex and smooth such that $\exists m > 0, M > 0$, for any $w \in \text{dom } L_f$, we have

$$MI \succeq \nabla^2 L_f(w) \succeq mI.$$

¹For each exploration time t , the model undergoes several optimization steps, influenced by the data volume and batch size.

The following theorem shows reduced error bound with causal exploration.

Theorem 1. Suppose Assumption 4 holds, and suppose the density of causal matrix D is δ and the model is initialized with w_0 . Then for every optimization step k , we have

$$\begin{aligned} L_f(w^c(k)) - L_f^* &\leq \delta^k (L_f(w(k)) - L_f^*) \\ &\leq \frac{M}{2} \left[\delta \left(1 - \frac{m}{M} \right) \right]^k \|w_0 - w^*\|_2^2. \end{aligned} \quad (11)$$

Remark. The first inequality in Equation (11) establishes an upper bound for ξ at δ^k . Given that $0 \leq \delta^k \leq 1$, this confirms the effectiveness of causal exploration in enhancing efficiency. The subsequent inequality establishes that the training error of causal exploration gradually converges to the optimal value, exhibiting an *exponential* convergence rate characterized by the decay of the factor $[\delta(1 - \frac{m}{M})]^k$. Moreover, this theorem implies that the advantages of causal exploration are relevant to the sparseness of causal structure. The sparser the causal structure, the faster our method learns. When the causal matrix D is a complete matrix ($\delta = 1$), causal exploration degenerates into non-causal prediction-based exploration. The proof for Theorem 1 is provided in Appendix A.

6 Experiments

To evaluate the effectiveness of causal exploration in complex scenarios, we conduct a series of experiments on both synthetic datasets and real-world applications, including the traffic light control task and MuJoCo control suites [Todorov *et al.*, 2012]. Meanwhile, we compare our proposed causal exploration with the following baseline methods:

1. **Curiosity** [Burda *et al.*, 2018] and **Plan2Explore** [Sekar *et al.*, 2020]: non-causal exploration methods that learn dense models with different forms of intrinsic rewards.
2. **CID** [Seitzer *et al.*, 2021]: focusing on detecting causal influences between actions \mathbf{a}_{t-1} and future states \mathbf{s}_t with random data collection. Causal exploration extends to discover dependencies among state transitions as well.
3. **ASR** [Huang *et al.*, 2022]: a differential approach that learns the causal matrix D as free parameters together with the world model f . Instead, the identification of D is completed by causal discovery methods in our approach.
4. **CRL** [Wang *et al.*, 2022]: learning causal dynamics models to approximate conditional mutual information rather than directly implementing CIT on the observed data.

Table 1 summarizes the attributes of these baselines and our proposed causal exploration from four perspectives. Moreover, following [Ke *et al.*, 2021], we use the following metrics to evaluate the performance of the learned models:

1. **Prediction Ability:** The error in predicting future states is a direct reflection of the model learning quality. In other words, lower prediction errors achieved with the same number of samples during exploration imply a

Approach	Causal Knowledge	Explicit Structural Embedding	Online Causal discovery	Non-random Data Collection
Curiosity [Pathak <i>et al.</i> , 2017]	✗	✗	✗	✓
Plan2Explore [Sekar <i>et al.</i> , 2020]	✗	✗	✗	✓
CID [Seitzer <i>et al.</i> , 2021]	✓	✗	✗	✗
ASR [Huang <i>et al.</i> , 2022]	✓	✓	✗	✗
CDL [Wang <i>et al.</i> , 2022]	✓	✗	✗	✓
Causal Exploration (ours)	✓	✓	✓	✓

Table 1: Attributes of existing baselines and our causal exploration. ✓ denotes that a method has an attribute, whereas ✗ denotes the opposite. "Causal Knowledge" here indicates whether causal information is utilized. Methods with "Explicit Structural Embedding" formulate dynamics models embedded with the causal matrix D . "Online Causal Discovery" and "Non-random Data Collection" measure the contribution of the method to boosting sampling and model learning efficiency from different perspectives, corresponding to Sections 4.2 and 5, respectively.

higher level of accuracy, signifying increased sampling efficiency.

- Structural Difference:** Evaluation of causal discovery methods is based on the difference between the identified causal matrix and the ground truth matrix.
- Downstream Performance:** For task-agnostic exploration, performance on downstream tasks also serves as a good metric to measure the model's understanding of the underlying causal structures of the environment.

6.1 Synthetic Datasets

Environments. We build our simulated environment following the state space model with controls. When the agent takes an action \mathbf{a}_t based on the current state, the environment provides feedback \mathbf{s}_{t+1} at the next time. We denote the generative environment as

$$\mathbf{s}_1 \sim \mathcal{N}(\mathbf{0}, I), \quad \mathbf{s}_t \sim \mathcal{N}(h(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}), \Sigma), \quad (12)$$

where Σ is the covariance matrix and h is the mean value as the ground truth transition function implemented by deep neural networks under causal graph \mathcal{G} . Specifically, the linear condition consists of a single-layer network, and the nonlinear function is three-layer MLPs with sigmoid activation.

Given that the sparsity of the causal graph is an important factor affecting the performance of our method, we choose to demonstrate the superiority of our approach on relatively low-density causal structures. That is, \mathcal{G} is generated by randomly connecting edges with a probability of p . Such sparse causal structures allow us to evaluate the ability of our method to accurately learn world models with limited data, which is a common scenario in many real-world applications. Below, we provide discussions on the benefits of causal exploration based on the empirical results.

Correctness of the identified causal structure. We compare our proposed method with different causal discovery methods. To be specific, both CID and CDL apply score-based causal discovery while ASR is a constraint-based method that combines causal structure with neural networks for differentiable training. However, CID only considers causal action influence and lacks an explicit causal structure, so we do not use it for comparison. Figure 5 illustrates an example of the causal matrix discovered by different methods, where the depth of the color indicates the number of times each edge is detected as a causal connection. Figure 3 further




Figure 5: A comparison of the discovered causal matrix with different methods when $|\mathcal{U}| = 12$ and $|\mathcal{V}| = 10$.

shows the averaged F1-score of these methods. The results reflect that the causal matrix learned through causal exploration obtains a smaller bias from the ground-truth matrix than others, an explanation for better performances.

Reliability of the learned world models. To show the advantage of causal exploration for improving the quality of model learning, we compare it with all of the baselines. Figure 6 are prediction errors of the world models during exploration under various state and action spaces. For each of these settings, we conduct 10 experiments and take the average value to reduce the impact of randomness. In all of these environments, causal exploration achieves lower prediction errors with fewer data. Compared to Curiosity and Plan2Explore, these advantages stem from the introduction of causality. In addition, the larger the dimension and the sparser the causal structure, the more prominent the advantages of our proposed method over non-causal method will be. The improvements over ASR and CID can be attributed to the effective design of exploration strategies. Furthermore, the experimental results compared with CDL underscore the importance of the stability of causal discovery methods and the necessity of explicit causal embeddings.

Advantage of causal-based models at the initial iteration. We also observe that many of the prediction error curves in Figure 6 have a different y intercept, suggesting that causal exploration is better even at the initial iteration. This is because the integration of causal matrix D into the model exclusively incorporates the parent nodes and eliminates extraneous input information when predicting future states, a departure from prior methodologies where these nodes were treated and initialized on par with parent nodes. We note that this advantage is also tied to the correctness of the identified causal matrix and the sparsity of the causal graph as in Theorem 1.




Figure 6: Prediction ability of different learned models on the synthetic environment under various state and action spaces.

Model	Reward	Duration	Length	Vehicles
Curiosity	-3.51	10.92	1.55	304
CID	-4.36	12.45	3.62	288
ASR	-1.69	8.23	0.92	312
CDL	-2.55	9.39	1.32	308
Ours	-1.37	7.21	0.31	314
<i>True Environment</i>	<i>-1.01</i>	<i>7.61</i>	<i>0.11</i>	<i>316</i>

Table 2: Performances of downstream policy learning in traffic signal control task.

Ablation study and generalization to some challenging scenarios. We further perform ablation studies: For online causal discovery, we provide results with various sampling methods; for the optimization of exploration strategy, we first investigate the impact of active reward and then perform causal exploration with some other forms of intrinsic reward. Besides, since the reliability of identified causal structures is an essential guarantee for the performance of causal exploration, we also consider challenging scenarios where the agent starts with wrong causal beliefs and encounters sudden structural changes. All of these corresponding experimental results, along with the implementation details, are given from Appendix B to Appendix D.

6.2 Real World Applications

Application to Traffic Signal Control Task. We further conduct experiments in the challenging traffic signal control task to evaluate our proposed method. Table 2 lists the performance of different learned world models on downstream policy learning. The reward here is a combination of several terms including the sum of queue length and so on, defined in Equation (3) in IntelliLight ([Wei *et al.*, 2018]). Duration refers to averaged travel time vehicles spent on approaching lanes (in seconds). Length is the sum of the length of waiting vehicles over all approaching lanes and Vehicles is the total number of vehicles that passed the intersection. We also train an agent in the true environment to verify the reliability of the learned models through task-agnostic exploration. The performance of the policy learned on the world models trained during causal exploration is comparable to, or even in some metrics better than, that of the true environment. Appendix C provides the detailed descriptions, including the causal graph on state variables in traffic signal control task, and comparison of the prediction errors during exploration.

Application to MuJoCo Tasks. We also evaluate causal exploration on the challenging MuJoCo tasks, where the

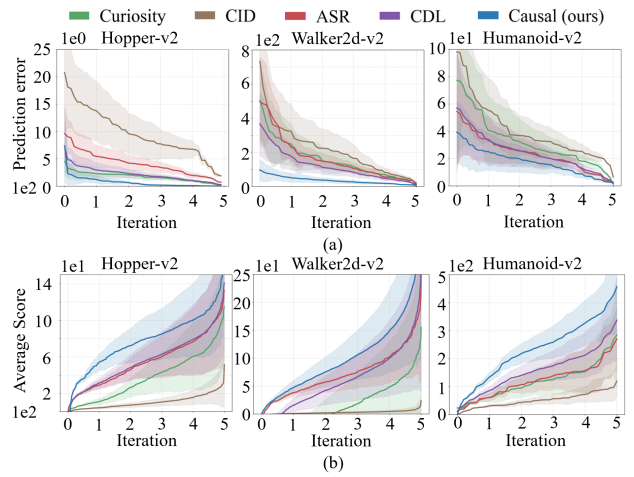


Figure 7: (a) Prediction errors in world model learning and (b) downstream policy learning performance on MuJoCo tasks.

state-action dimensions range from tens (Hopper-v2) to hundreds (Humanoid-v2). Implementation details and more experimental results including the identified causal structures are given in Appendix D. As shown in Figure 7(a), the world model learned during causal exploration keeps performing lower prediction errors than others. This ability of our causal method stems from its focus on causal structural constraints. By selectively incorporating parent nodes during state prediction, our approach maintains informed exploration even when exploratory signals decline, which prevents the agent from getting stuck in sub-optimal behaviors. Such a good performance is also achieved in a shorter time. Task performance in Figure 7(b) also reflects the reliability of these world models. We see that our proposed method still works well when the state space is large. As long as there is a strong causal relationship between the observed state variables, and the causal discovery algorithm can accurately identify the corresponding causal structure, our causal exploration method is believed to be well applied to high-dimensional situations.

7 Conclusion and Future Work

In this paper, we introduce causal exploration, a methodology designed to incorporate causal information from data for the purpose of learning world models efficiently. In particular, we employ causal exploration within the domain of task-agnostic reinforcement learning and design a sharing-decomposition schema to leverage causal structural knowledge for the world model. A series of experiments in both simulated environments and real-world tasks demonstrate the superiority of causal exploration, which highlights the importance of rich causal prior knowledge for efficient data collection and model learning. We would like to point out that this study primarily focuses on scenarios where states are fully observed. Future research directions involve addressing more complex scenarios including unobserved latent state variables and developing improved fault-tolerant mechanisms to enhance robustness. This work provides a promising direction for future endeavors in efficient exploration.

Acknowledgments

Yupei Yang, Shikui Tu and Lei Xu would like to acknowledge the support by the Shanghai Municipal Science and Technology Major Project, China (Grant No. 2021SHZDZX0102), and by the National Natural Science Foundation of China (62172273).

References

- [Aubret *et al.*, 2019] Arthur Aubret, Laetitia Matignon, and Salima Hassas. A survey on intrinsic motivation in reinforcement learning. *arXiv preprint arXiv:1908.06976*, 2019.
- [Bellemare *et al.*, 2016] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.
- [Burda *et al.*, 2018] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018.
- [Ding *et al.*, 2022] Wenhao Ding, Haohong Lin, Bo Li, and Ding Zhao. Generalizing goal-conditioned reinforcement learning with variational causal reasoning. *Advances in Neural Information Processing Systems*, 35:26532–26548, 2022.
- [Duan *et al.*, 2016] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, pages 1329–1338. PMLR, 2016.
- [Fang *et al.*, 2017] Meng Fang, Yuan Li, and Trevor Cohn. Learning how to active learn: A deep reinforcement learning approach. *arXiv preprint arXiv:1708.02383*, 2017.
- [Hu *et al.*, 2022] Xing Hu, Rui Zhang, Ke Tang, Jiaming Guo, Qi Yi, Ruizhi Chen, Zidong Du, Ling Li, Qi Guo, Yunji Chen, et al. Causality-driven hierarchical structure discovery for reinforcement learning. *Advances in Neural Information Processing Systems*, 35:20064–20076, 2022.
- [Huang *et al.*, 2021] Biwei Huang, Fan Feng, Chaochao Lu, Sara Magliacane, and Kun Zhang. Adarl: What, where, and how to adapt in transfer reinforcement learning. *arXiv preprint arXiv:2107.02729*, 2021.
- [Huang *et al.*, 2022] Biwei Huang, Chaochao Lu, Liu Leqi, José Miguel Hernández-Lobato, Clark Glymour, Bernhard Schölkopf, and Kun Zhang. Action-sufficient state representation learning for control with structural constraints. In *International Conference on Machine Learning*, pages 9260–9279. PMLR, 2022.
- [Ke *et al.*, 2021] Nan Rosemary Ke, Aniket Didolkar, Sarthak Mittal, Anirudh Goyal, Guillaume Lajoie, Stefan Bauer, Danilo Rezende, Yoshua Bengio, Michael Mozer, and Christopher Pal. Systematic evaluation of causal discovery in visual model based reinforcement learning. *arXiv preprint arXiv:2107.00848*, 2021.
- [Kim *et al.*, 2018] Hyoungseok Kim, Jaekyeom Kim, Yeonwoo Jeong, Sergey Levine, and Hyun Oh Song. Emi: Exploration with mutual information. *arXiv preprint arXiv:1810.01176*, 2018.
- [Kosoy *et al.*, 2022] Eliza Kosoy, Adrian Liu, Jasmine L Collins, David Chan, Jessica B Hamrick, Nan Rosemary Ke, Sandy Huang, Bryanna Kaufmann, John Canny, and Alison Gopnik. Learning causal overhypotheses through exploration in children and computational models. In *Conference on Causal Learning and Reasoning*, pages 390–406. PMLR, 2022.
- [Machado *et al.*, 2020] Marlos C Machado, Marc G Belle-mare, and Michael Bowling. Count-based exploration with the successor representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5125–5133, 2020.
- [Mutti *et al.*, 2023] Mirco Mutti, Riccardo De Santi, Emanuele Rossi, Juan Felipe Calderon, Michael Bronstein, and Marcello Restelli. Provably efficient causal model-based reinforcement learning for systematic generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9251–9259, 2023.
- [Nguyen *et al.*, 2017] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017.
- [Pathak *et al.*, 2017] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR, 2017.
- [Pathak *et al.*, 2019] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. In *International conference on machine learning*, pages 5062–5071. PMLR, 2019.
- [Rahaf and Lucas, 2019] Aljundi Rahaf and Caccia Lucas. Online continual learning with maximally interfered retrieval. In *NIPS*, 2019.
- [Robine *et al.*, 2023] Jan Robine, Marc Höftmann, Tobias Uelwer, and Stefan Harmeling. Transformer-based world models are happy with 100k interactions. *arXiv preprint arXiv:2303.07109*, 2023.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Seitzer *et al.*, 2021] Maximilian Seitzer, Bernhard Schölkopf, and Georg Martius. Causal influence detection for improving efficiency in reinforcement learning. *Advances in Neural Information Processing Systems*, 34:22905–22918, 2021.
- [Sekar *et al.*, 2020] Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak Pathak. Planning to explore via self-supervised world models. In *International Conference on Machine Learning*, pages 8583–8592. PMLR, 2020.

- [Shyam *et al.*, 2019] Pranav Shyam, Wojciech Jaśkowski, and Faustino Gomez. Model-based active exploration. In *International conference on machine learning*, pages 5779–5788. PMLR, 2019.
- [Silver *et al.*, 2016] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [Spirtes *et al.*, 2000] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT press, 2000.
- [Tassa *et al.*, 2018] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- [Todorov *et al.*, 2012] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- [Van Hasselt *et al.*, 2016] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- [Wang *et al.*, 2022] Zizhao Wang, Xuesu Xiao, Zifan Xu, Yuke Zhu, and Peter Stone. Causal dynamics learning for task-independent state abstraction. *arXiv preprint arXiv:2206.13452*, 2022.
- [Wei *et al.*, 2018] Hua Wei, Guanjie Zheng, Huaxiu Yao, and Zhenhui Li. Intellilight: A reinforcement learning approach for intelligent traffic light control. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2496–2505, 2018.
- [Yoo and Kweon, 2019] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 93–102, 2019.
- [Yoon *et al.*, 2021] Jaehong Yoon, Divyam Madaan, Eunho Yang, and Sung Ju Hwang. Online coreset selection for rehearsal-based continual learning. *arXiv preprint arXiv:2106.01085*, 2021.
- [Yu *et al.*, 2023] Zhongwei Yu, Jingqing Ruan, and Dengpeng Xing. Explainable reinforcement learning via a causal world model. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 4540–4548. IJCAI, 2023.
- [Zhang *et al.*, 2012] Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. *arXiv preprint arXiv:1202.3775*, 2012.
- [Zhu *et al.*, 2019] Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. *arXiv preprint arXiv:1906.04477*, 2019.
- [Zhu *et al.*, 2022] Zheng-Mao Zhu, Xiong-Hui Chen, Hong-Long Tian, Kun Zhang, and Yang Yu. Offline reinforcement learning with causal structured world models. *arXiv preprint arXiv:2206.01474*, 2022.