

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

6-3-2022

Offline Reinforcement Learning with Causal Structured World Models

Zheng-Mao Zhu

National Key Laboratory for Novel Software Technology, Nanjing University, China

Xiong-Hui Chen

National Key Laboratory for Novel Software Technology, Nanjing University, China & Polixir.ai, China

Hong-Long Tian

National Key Laboratory for Novel Software Technology, Nanjing University, China

Kun Zhang

Carnegie Mellon University & Mohamed bin Zayed University of Artificial Intelligence

Yang Yu

National Key Laboratory for Novel Software Technology, Nanjing University, China & Polixir.ai, China & Peng Cheng Laboratory, China

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC by 4.0

Uploaded 14 July 2022

Recommended Citation

Z.M. Zhu, X.H. Chen, H.L. Tian, K. Zhang and Y. Yu, "Offline Reinforcement Learning with Causal Structured World Models", 2022, arXiv:2206.01474

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

OFFLINE REINFORCEMENT LEARNING WITH CAUSAL STRUCTURED WORLD MODELS

A PREPRINT

Zheng-Mao Zhu¹, Xiong-Hui Chen^{1,3}, Hong-Long Tian¹, Kun Zhang^{2,5}, Yang Yu^{1,3,4*}

¹ National Key Laboratory for Novel Software Technology, Nanjing University

² Department of Philosophy, Carnegie Mellon University

³ Polixir.ai, ⁴ Peng Cheng Laboratory

⁵ Department of Machine Learning, Mohamed bin Zayed University of Artificial Intelligence
{zhuzm, chenxh, yuy}@lamda.nju.edu.cn, franktian424@gmail.com, kunz1@cmu.edu

ABSTRACT

Model-based methods have recently shown promising for offline reinforcement learning (RL), aiming to learn good policies from historical data without interacting with the environment. Previous model-based offline RL methods learn fully connected nets as world-models to map the states and actions to the next-step states. However, it is sensible that a world-model should adhere to the underlying causal effect such that it will support learning an effective policy generalizing well in unseen states. In this paper, We first provide theoretical results that causal world-models can outperform plain world-models for offline RL by incorporating the causal structure into the generalization error bound. We then propose a practical algorithm, oFfline mOdel-based reinforcement learning with CaUsal Structure (FOCUS), to illustrate the feasibility of learning and leveraging causal structure in offline RL. Experimental results on two benchmarks show that FOCUS reconstructs the underlying causal structure accurately and robustly. Consequently, it performs better than the plain model-based offline RL algorithms and other causal model-based RL algorithms.

1 Introduction

Offline RL has gained great interest since it enables RL algorithms to scale to many real-world applications, e.g., recommender systems [1, 2], autonomous driving [3], and healthcare [4], waiving costly online trial-and-error. In the offline setting, Model-Based Reinforcement Learning (MBRL) is an important direction that already produces significant offline learning performance [5, 6]. Moreover, learning models is also useful for training transferable policies [7, 6]. Therefore, there is increasing studies on learning world-models, from supervised prediction methods [8] to adversarial learning methods [9].

However, commonly there exists an underlying causal structure among the states and actions in various tasks. The causal structure can support learning a policy with better generalization ability. For example, in driving a car where the speed depends on the gas and brake pedals but not the wiper, a plain world-model trained on rainy days always predicts deceleration when the wiper is turned on and thus can not generalize to other weather situations. By contrast, a causal world-model can avoid the spurious dependence between wiper and deceleration (because of the rain) and hence help generalize well in unseen weather. In fact, empirical evidence also indicates that inducing the causal structure is important to improve the generalization of RL [10, 11, 12, 13, 14], but little attention was paid on causal world-model learning.

In this paper, we first provide theoretical support for the statement above: we show that a causal world-model can outperform a plain world-model for offline RL. From the causal perspective, we divide the variables in states and actions into two categories, namely, causal variables and spurious variables, and then formalize the procedure that learns a world-model from raw measured variables. Based on the formalization, we quantize the influence of the spurious

*Yang Yu is the corresponding author

dependence on the generalization error bound and thus prove that incorporating the causal structure can help reduce this bound.

We then propose a practical offline MBRL algorithm with causal structure, FOCUS, to illustrate the feasibility of learning causal structure in offline RL. An essential step of FOCUS is to learn the causal structure from data and then use it properly. Learning causal structure from data has been known as causal discovery [15]. There are some challenges in utilizing causal discovery methods in RL, and there are specific properties in the data that causal discovery may benefit from. Specifically, we extended the PC algorithm, which aims to find causal discovery based on the inferred conditional independence relations, to incorporate the constraint that the future cannot cause the past. Consequently, we can reduce the number of conditional independence tests and determine the causal direction. We further adopt kernel-based independence testing [16], which can be applied in continuous variables without assuming a functional form between the variables as well as the data distribution.

In summary, this paper makes the following key contributions:

- It shows theoretically that a causal world-model outperforms a plain world-model in offline RL, in terms of the generalization error bound.
- It proposes a practical algorithm, FOCUS, and illustrates the feasibility of learning and using a causal world-model for offline MBRL.
- Our experimental results verify the theoretical claims, showing that FOCUS outperforms baseline models and other online causal MBRL algorithms in the offline setting.

2 Related Work

Causal Structure Learning in Online MBRL. Despite that some methods have been proposed for learning causal structure in online RL, such methods all depend on interactions and do not have a mechanism to transfer into the offline setting. The core of online causal structure learning is to evaluate the performance or other metrics of one structure by interactions and choose the best one as the causal structure. [13] parameterizes the causal structure in the model and learns policies for each possible causal structure by minimizing the log-likelihood of dynamics. Given learned policies, it makes regression between the policy returns and the causal structure and then chooses the structure corresponding to the highest policy return. [17] (LNCM) samples causal structures from Multivariate Bernoulli distribution and scores those structures according to the log-likelihood on interventional data. Based on the scores, it calculates the gradients for the parameters of the Multivariate Bernoulli distribution and updates the parameters iteratively. [12] utilizes the speed of adaptation to learn the causal direction, which does not provide a complete algorithm for learning causal structure. By contrast, FOCUS utilizes a causal discovery method for causal structure learning, which only relies on the collected data to obtain the causal structure.

Causal Discovery Methods. The data in RL is more complex than it in traditional causal discovery, where the data is often discrete and the causal relationship is under a simple linear assumption. In recent years, practical methods have been proposed for causal discovery for continuous variables, which is the case we are concerned with in this paper. [18] is based on explicit estimation of the conditional densities or their variants, which exploit the difference between the characteristic functions of these conditional densities. The estimation of the conditional densities or related quantities is difficult, which deteriorates the testing performance especially when the conditioning set is not small enough. [19] discretizes the conditioning set to a set of bins and transforms conditional independence (CI) to the unconditional one in each bin. Inevitably, due to the curse of dimensionality, as the conditioning set becomes larger, the required sample size increases dramatically. By contrast, the KCI test [16] is a popular and widely-used causal discovery method, in which the test statistic can be easily calculated from the kernel matrices and the distribution can be estimated conveniently.

3 Preliminaries

Markov Decision Process (MDP). We describe the RL environment as an MDP with five-tuple $\langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ [20], where \mathcal{S} is a finite set of states; \mathcal{A} is a finite set of actions; P is the transition function with $P(\mathbf{s}'|\mathbf{s}, \mathbf{a})$ denoting the next-state distribution after taking action \mathbf{a} in state \mathbf{s} ; R is a reward function with $R(\mathbf{s}, \mathbf{a})$ denoting the expected immediate reward gained by taking action \mathbf{a} in state \mathbf{s} ; and $\gamma \in [0, 1]$ is a discount factor. An agent chooses actions \mathbf{a} according to a policy $\mathbf{a} \sim \pi(\mathbf{s})$, which updates the system state $\mathbf{s}' \sim P(\mathbf{s}, \mathbf{a})$, yielding a reward $r \sim R(\mathbf{s}, \mathbf{a})$. The agent's goal is to maximize the expected cumulative return by learning a good policy $\max_{\pi, P} \mathbb{E}[\gamma^t R(\mathbf{s}_t, \mathbf{a}_t)]$. The state-action value Q_π of a policy π is the expected discounted reward of executing action a from state \mathbf{s} and subsequently following policy π : $Q_\pi(\mathbf{s}, \mathbf{a}) := R(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\mathbf{s}' \sim P, \mathbf{a}' \sim \pi} [Q_\pi(\mathbf{s}', \mathbf{a}')]$.

Offline Model-based Reinforcement Learning. Model-based reinforcement learning typically involves learning a dynamics model of the environment by fitting it using a maximum-likelihood estimate of the trajectory-based data collected by running some exploratory policy [21, 22]. In the offline RL setting, where we only have access to the data collected by multiple policies, recent techniques build on the idea of pessimism (regularizing the original problem based on how confident the agent is about the learned model) and have resulted in better sample complexity over model-free methods on benchmark domains [23, 5].

4 Theory

In this section, we provide theoretical evidence for the advantages of a causal world-model over a plain world-model, which shows that utilizing a good causal structure can reduce the generalization error bounds in offline RL. Specifically, We incorporate the causal structure into the generalization error bounds, which include the model prediction error bound and policy evaluation error bound. The full proof can be found in Appendix A. In this paper, we focus on the linear case.

4.1 Model Prediction Error Bound

In this subsection, we formulize the procedure that learns a plain world-model and then connect the model prediction error bound with the number of spurious variables in the plain world-model. Specifically, we point out that the spurious variables lead the model learning problem to an ill-posed problem that has multiple optimal solutions, which consequently results in the increment of the model prediction error bound. Since model learning can be viewed as a supervised learning problem, we provide the model prediction error bound in a supervised learning framework.

Preliminary. Let \mathcal{D} denote the data distribution where we have samples $(\mathbf{X}, Y) \sim \mathcal{D}$, $\mathbf{X} \in \mathbb{R}^n$. The goal is to learn a linear function f to predict Y given \mathbf{X} . From the causal perspective, Y is generated from only its causal parent variables rather than all the variables in \mathbf{X} . Therefore we can split the variables in \mathbf{X} into two categories, $\mathbf{X} = (\mathbf{X}_{causal}, \mathbf{X}_{spurious})$:

- \mathbf{X}_{causal} represents the causal parent variables of Y , that is, $Y = f^*(\mathbf{X}_{causal}) + \epsilon_{causal}$, where f^* is the ground truth and ϵ_{causal} is a zero mean noise variable that $\mathbf{X}_{causal} \perp \epsilon_{causal}$.
- $\mathbf{X}_{spurious}$ represents the spurious variables that $\mathbf{X}_{spurious} \perp \mathbf{X}_{causal}$, but in some biased data sets $\mathbf{X}_{spurious}$ and \mathbf{X}_{causal} have strong relatedness. In other words, $\mathbf{X}_{spurious}$ can be predicted by \mathbf{X}_{causal} with small error, i.e., $\mathbf{X}_{spurious} = \mathbf{X}_{causal} \gamma_{spurious} + \epsilon_{spurious}$, where $\epsilon_{spurious}$ is the regression error with zero mean and small variance.

For clearly representation, we use $\mathbf{X}_{cau} \triangleq \mathbf{X} \circ \omega_{cau}$ (\circ represents element-wise multiplication) to replace \mathbf{X}_{causal} , where cau records the indices of \mathbf{X}_{causal} in \mathbf{X} and $(\omega_{cau})_i = \mathbb{I}(i \in cau)$. Correspondingly, we also use $\mathbf{X}_{spu} \triangleq \mathbf{X} \circ \omega_{spu}$ to replace $\mathbf{X}_{spurious}$. According to the definition of \mathbf{X}_{cau} , we have $Y = (\mathbf{X} \circ \omega_{cau})\beta^* + \epsilon_{cau}$, where $\omega_{cau} \circ \beta^*$ is the global optimal solution of the optimization problem

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}} [\mathbf{X}\beta - Y]^2. \quad (1)$$

Above problem is easy if the data is uniformly sampled from \mathcal{D} . However, in the offline setting, we only have biased data \mathcal{D}_{train} sampled by given policy π_{train} , where the optimization objective is

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} [\mathbf{X}\beta - Y]^2. \quad (2)$$

The Problem 2 has multiple optimal solutions due to the strong linear relatedness of \mathbf{X}_{spu} and \mathbf{X}_{cau} in \mathcal{D}_{train} , which is proved in Lemma 4.1.

Lemma 4.1. *Given that $\omega_{cau} \circ \beta^*$ is the optimal solution of Problem 1, suppose that in \mathcal{D}_{train} , $\mathbf{X}_{spu} = (\mathbf{X} \circ \omega_{cau})\gamma_{spu} + \epsilon_{spu}$ where $\mathbb{E}_{\mathcal{D}_{train}}[\epsilon_{spu}] = 0$ and $\gamma_{spu} \neq \mathbf{0}$, we have that $\hat{\beta}_{spu} \triangleq \omega_{cau} \circ (\beta^* - \lambda\gamma_{spu}) + \lambda\omega_{spu}$ is also an optimal solution of Problem 2 for any λ :*

$$\mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left[(|\mathbf{X}(\omega_{cau} \circ \beta^*) - Y|_2) \mid \mathbf{X} \right] = \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left[(|\mathbf{X}\hat{\beta}_{spu} - Y|_2) \mid \mathbf{X} \right]$$

The most popular method for solving such ill-posed problem is to add a regularization term for parameters β [24]:

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} [\mathbf{X}\beta - Y]^2 + k\|\beta\|^2, \quad (3)$$

where k is a coefficient. The form of Problem 3 corresponds to the form of the ridge regression, which provides a closed-form solution of k by Hoerl-Kennard formula [25].

In the following, we will first introduce the solution of λ under Problem 3 in Lemma 4.2, and then introduce the model prediction error bound with λ in Theorem 4.4. For ease of understanding, we provide a simple version where the dimensions of \mathbf{X}_{cau} and \mathbf{X}_{spu} are both one ($|\mathbf{X}_{cau}| = |\mathbf{X}_{spu}| = 1$).

Lemma 4.2 (λ Lemma). *Given λ as the coefficient in Lemma 4.1, and k in Problem 3 chosen by Hoerl-Kennard formula, we have the solution of λ in Problem 3 that:*

$$\lambda = \frac{\beta^* \gamma_{spu}}{\beta^{*2} + \gamma_{spu}^2 + 1 + \frac{\sigma_{spu}^2}{\sigma_{cau}^2} (1 + \frac{1}{(\beta^*)^2})} \quad (4)$$

Based on Lemma 4.2, we can find that the smaller σ_{spu}^2 (it means that \mathbf{X}_{spu} and \mathbf{X}_{cau} have stronger relatedness in the training dataset D_{train}), the larger λ . And we also have its bound:

Proposition 4.3. *Given λ as Formula 4, the bound of λ is that $-\frac{1}{2} \leq \lambda \leq \frac{1}{2}$.*

Theorem 4.4 (Spurious Theorem). *Let $\mathcal{D} = \{(X, Y)\}$ denote the data distribution, $\hat{\beta}_{spu}$ denote the solution in Lemma 4.1 with λ in Lemma 4.2, and $\hat{Y}_{spu} = X \hat{\beta}_{spu}$ denote the prediction. Suppose that the data value is bounded: $|X_i|_1 \leq X_{max}, i = 1, \dots, n$ and the error of optimal solution ϵ_{cau} is also bounded: $|\epsilon_{cau}|_1 \leq \epsilon_c$, we have the model prediction error bound:*

$$\mathbb{E}_{(X,Y) \sim D}[(|\hat{Y}_{spu} - Y|_1) | \mathbf{X}] \leq X_{max} |\lambda|_1 (|\gamma_{spu}|_1 + 1) + \epsilon_c. \quad (5)$$

Theorem 4.4 shows that

- The upper bound of the model prediction error $|\hat{Y}_{spu} - Y|_1$ increases by $X_{max} |\lambda|_1 (|\gamma_{spu}|_1 + 1)$ for each induced spurious variable X_{spu} in the model.
- When X_{spu} and X_{cau} have stronger relatedness (which means a bigger λ), the increment of the prediction model error bound led by X_{spu} is bigger.

4.2 Policy Evaluation Error Bound

Although in most cases, an accurate model ensures a good performance in MBRL, the model error bound is still an indirect evaluation compared to the policy evaluation error bound for MBRL. In this subsection, we apply the spurious theorem (Theorem 4.4) to offline MBRL and provide the policy evaluation error bound with the number of spurious variables.

Suppose that the state value and reward are bounded that $|S_{t,i}|_1 \leq S_{max}, R_t \leq R_{max}$, let λ_{max} denote the maximum of λ and γ_{max} denote the maximum of $|\gamma_{spu}|_1$, we have the policy evaluation error bound in Theorem 4.5.

Theorem 4.5 (RL Spurious Theorem). *Given an MDP with the state dimension n_s and the action dimension n_a , a data-collecting policy π_D , let M^* denote the true transition model, M_θ denote the learned model that M_θ^i predicts the i^{th} dimension with spurious variable sets spu_i and causal variables cau_i , i.e., $\hat{S}_{t+1,i} = M_\theta^i((S_t, A_t) \circ \omega_{cau_i \cup spu_i})$. Let $V_\pi^{M_\theta}$ denote the policy value of the policy π in model M_θ and correspondingly $V_\pi^{M^*}$. For an arbitrary bounded divergence policy π , i.e. $\max_S D_{KL}(\pi(\cdot|S), \pi_D(\cdot|S)) \leq \epsilon_\pi$, we have the policy evaluation error bound:*

$$|V_\pi^{M_\theta} - V_\pi^{M^*}| \leq \frac{2\sqrt{2}R_{max}}{(1-\gamma)^2} \sqrt{\epsilon_\pi} + \frac{R_{max}\gamma}{2(1-\gamma)^2} S_{max} [n_s \epsilon_c + (1 + \gamma_{max}) \lambda_{max} n_s (n_s + n_a) R_{spu}]$$

where $R_{spu} = \frac{\sum_{i=1}^{n_s} |spu_i|}{n_s(n_s + n_a)}$, which represents the spurious variable density, that is, the ratio of spurious variables in all input variables.

Theorem 4.5 shows the relation between the policy evaluation error bound and the spurious variable density, which indicates that:

- When we use a non-causal model that all the spurious variables are input, R_{spu} reaches its maximum value $\bar{R}_{spu} < 1$. By contrast, in the optimal causal structure, R_{spu} reaches its minimum value of 0.
- The density of spurious variables R_{spu} and the correlation strength of spurious variables λ_{max} both influence the policy evaluation error bound. However, if we exclude all the spurious variables, i.e., $R_{spu} = 0$, the correlation strength of spurious variables will have no effect.

5 Algorithm

In the theory section, we have provided the theoretical evidence about the advantages of a causal world-model over a plain world-model. Besides lower prediction errors, a causal world-model also matters for better decision-making in RL. In the condition that spurious variables do not increase prediction errors (e.g., spurious variables disturb only in

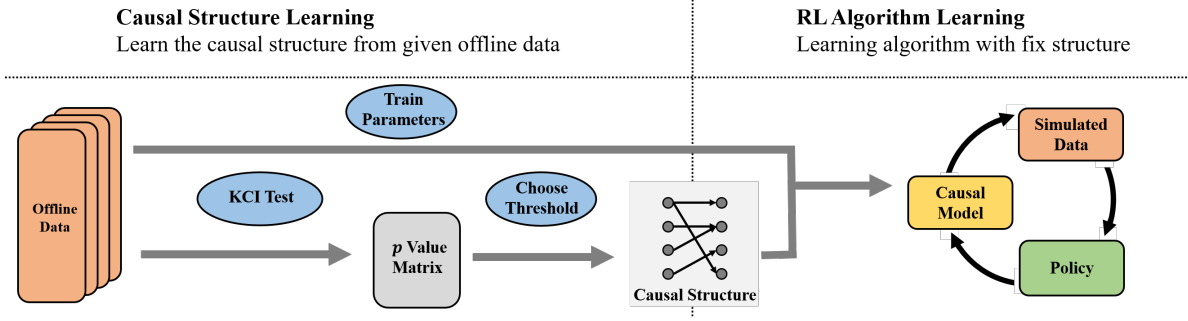


Figure 1: The architecture of FOCUS. Given offline data, FOCUS learns a p value matrix by KCI test and then gets the causal structure by choosing a p threshold. After combining the learned causal structure with the neural network, FOCUS learns the policy through an offline MBRL algorithm.

unreachable states), a wrong causal relation also leads to terrible decision-making. For example, rooster crowing can predict the rise of the sun, but forcing a rooster to crow for a sunny day is a natural decision if we have a wrong causal relation that rooster crowing causes the rise of the sun. In the above example, predicting the rise of the sun by rooster crowing is a zero-error world-model since rooster crowing on a rainy day is an unreachable state, but such world-model leads to terrible decision-making.

After demonstrating the necessity of a causal world-model in offline RL, in this section we propose a practical offline MBRL algorithm, FOCUS, to illustrate the feasibility of learning causal structure in offline RL. The main idea of FOCUS is to take the advantage of causal discovery methods and extend it to offline MBRL. Compared to previous online causal structure learning methods, the causal discovery method brings the following advantages in the offline setting:

- **Robust:** The structure is not influenced by the performance of the test data, which is artificially selected and biased.
- **Efficient:** The causal discovery method directly returns the structure by independence testing without any network training procedure and thus saves the samples for network training.

5.1 Preliminary

Conditional Independence Test. Independence and conditional independence (CI) play a central role in causal discovery [26, ?, 27]. Generally speaking, the CI relationship $X \perp\!\!\!\perp Y \mid Z$ allows us to drop Y when constructing a probabilistic model for X with (Y, Z) . There are multiple CI testing methods for various conditions, which provide the correct conclusion only given the corresponding condition. The kernel-based Conditional Independence test (KCI-test) [16] is proposed for continuous variables without assuming a functional form between the variables as well as the data distributions, which is the case we are concerned with in this paper. Generally, the hypothesis H_0 that variables are conditionally independent is rejected when p is smaller than the pre-assigned significance level, say, 0,05. In practice, we can design the significance level instead of a fixed value.

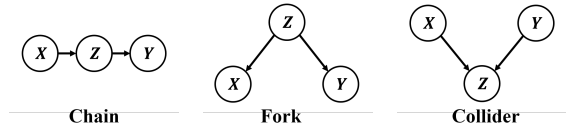


Figure 2: The three basic structure for three variables X, Y and Z , where Z plays an important role in causal discovery.

which is the case we are concerned with in this paper. Generally, the hypothesis H_0 that variables are conditionally independent is rejected when p is smaller than the pre-assigned significance level, say, 0,05. In practice, we can design the significance level instead of a fixed value.

Conditional Variables. Besides the specific CI test method, the conclusion of conditional independence testing also depends on the conditional variable Z , that is, different conditional variables can lead to different conclusions. Taking the triple (X, Y, Z) as an example, there are three typical structures, namely, *Chain*, *Fork*, and *Collider* as shown in Fig 2, where whether conditioning on Z significantly influences the testing conclusion.

- *Chain:* There exists causation between X and Y but conditioning on Z leads to independence.
- *Fork:* There does not exist causation between X and Y but not conditioning on Z leads to non-independence.
- *Collider:* There does not exist causation between X and Y but conditioning on Z leads to non-independence.

5.2 Building the Causal Structure from the KCI test

Applying the Independence Test in RL. Based on the preliminaries, given the two target variables X, Y and the condition variable Z , the KCI test returns a probability value $p = f_{KCI}(X, Y, Z) \in [0, 1]$, which measures the

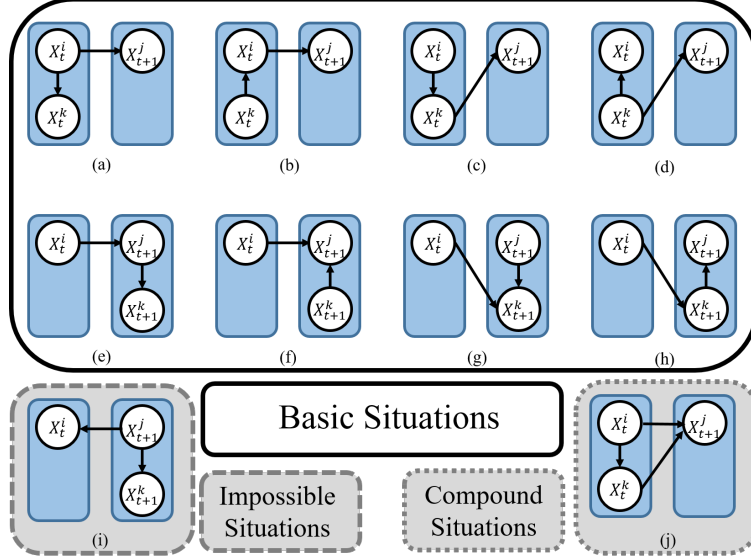


Figure 3: The basic, impossible and compound situations of the causation between target variables and condition variables. In the basic situations, **Top Line:** (a)-(d) list the situations that the condition variable X^k is in the t time step. **Bottom Line:** Similarly, (e)-(h) list the situations that the condition variable X^k is in the $t + 1$ time step.

probability that X and Y are conditionally independent given the condition Z . In other words, a small p implies that X and Y have causation given Z . To transform an implicit probability value into an explicit conclusion of whether the causation exists, we design a threshold p^* that:

$$Causation(X, Y) = \mathbb{I}(f_{KCI}(X, Y, Z) \leq p^*) \in \{0, 1\},$$

where $Causation(X, Y) = 1$ represents independence and 0 represents that causation exists. Details of choosing p^* can be found in Appendix B.1.

In model learning of RL, variables are composed of states and actions of the current and next timesteps and the causal structure refers to whether a variable in t timestep (e.g., the i^{th} dimension, X_t^i) causes another variable in $t + 1$ timestep (e.g., the j^{th} dimension, X_{t+1}^j). With the KCI test, we get the causal relation through the function $Causation(\cdot, \cdot)$ for each variable pair (X_t^i, X_{t+1}^j) and then form the causal structure matrix \mathcal{G} :

$$\mathcal{G}_{i,j} = Causation(X_t^i, X_{t+1}^j),$$

where $\mathcal{G}_{i,j}$ is the element in row i and column j of \mathcal{G} .

Choosing the Conditional Variable in RL. As we said in preliminaries, improper conditional variables can reverse the conclusion of independence testing. Therefore we have to carefully design the conditional variable set, which should include the intermediate variable of *Chain*, the common parent variable of *Fork*, but not the common son variable of *Collider*. Traditionally, the independence test has to traverse all possibilities of the conditional variable set and gives the conclusion, which is too time-consuming. However, in RL we can reduce the number of conditional independence tests by incorporating the constraint that the future cannot cause the past. Actually, this constraint limits the number of possible conditional variable set to a small value. Therefore we can even take a classified discussion for each possible conditional variable set. Before the discussion, we exclude two kinds of situations for simplicity:

- **Impossible situations.** We exclude some impossible situations as Fig 3 (i) by the temporal property of data in RL. Specifically, the direction of the causation cannot be from the variable of $t + 1$ time step to that of t time step because the effect cannot happen before the cause.
- **Compound situations.** We only discuss the basic situations and exclude the compound situations, e.g., Fig 3 (j), which is a compound of (a) and (c). It is because in such compound situations, the target variables X_t^i and X_{t+1}^j have direct causation (or it can not be a compound situation) and the independence testing only misjudges independence as non-independence but not non-independence as independence.

We list all possible situations of target variables X_t^i, X_{t+1}^j and condition variable $X_{t/t+1}^k$ in the world-model as shown in Fig 3. Based on the causal discovery knowledge in the preliminaries, we analyze basic situations in the following:

Algorithm 1 Causal Model Framework for Offline MBRL

Input: offline data set $\mathcal{D} = \{(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, r_t)\}$; model $\mathcal{M}(\cdot; \theta)$;

Stage 1: Causal Structure Learning

Get p value matrix G_p by KCI testing.

Get the threshold p^* by G_p .

Get causal structure mask matrix G by the threshold p^* .

Stage 2: Offline Reinforcement Learning

Choose an offline model-based reinforcement learning algorithm $\mathbf{Algo}(\cdot)$ and replace its model $\mathcal{M}(\cdot)$ by $\mathcal{M}_{Causal}(\cdot, G, \mathcal{M})$ (Algorithm 2 in Appendix).

Obtain the optimal policy $\pi^* = \mathbf{Algo}(\mathcal{D})$.

Return π^*

- **Top Line:** In (a)(b), whether conditioning on X_t^k does not influence the conclusion of causation; In (c), although X_t^k plays an intermediate variable in a *Chain* and conditioning on X_t^k leads to the conclusion of independence of X_t^i and X_{t+1}^j , the causal parent set of X_{t+1}^j will include X_t^k when testing the causal relation between X_t^k and X_{t+1}^j , which can offset the influence of excluding X_t^i . In (d), conditioning on Z is necessary for getting the correct conclusion of causation since X_t^k is the common causal parent in a *Fork* structure.
- **Bottom Line:** In (e)(f), whether conditioning on X_{t+1}^k does not influence the conclusion of causation; In (g), not conditioning on X_{t+1}^k is necessary to get the correct conclusion of causation since X_{t+1}^k is the common son in a *Collider* structure; In (h), although X_{t+1}^k plays an intermediate variable in a *Chain* and not conditioning on X_{t+1}^k leads to the conclusion of non-independence of X_t^i and X_{t+1}^j , including X_t^i in the causal parent set of X_{t+1}^j will not induce any problem since X_t^i does indirectly cause X_{t+1}^j .

Based on the above classified discussion, we can conclude our principle for choosing conditional variables in RL that: (1) Condition on the other variables in t time step; (2) Do not condition on the other variables in $t + 1$ time step.

5.3 Combining Learned Causal Structure with An Offline MBRL Algorithm

We combine the learned causal structure with an offline MBRL algorithm, MOPO [5], to form a causal offline MBRL algorithm as in Fig 1. The complete learning procedure is shown in Algorithm 1, where Algorithm 2 can be found in Appendix B.2. Notice that our causal model learning method could be combined with any offline MBRL algorithm in principle. More implementation details and hyperparameter values are summarized in Appendix B.1.

6 Experiments

To demonstrate that (1) Learning a causal world-model is feasible in offline RL and (2) a causal world-model can outperform a plain world-model and other related online methods in offline RL, we evaluate (1) **causal structure learning** and (2) **policy learning** on the Toy Car Driving and MuJoCo benchmark. Toy Car Driving is a simple and typical environment that is convenient to evaluate the accuracy of learned causal structure because We can design the causation between variables in it. The MuJoCo is the most common benchmark to investigate the performance in continuous controlling, where each dimension of the state has a specific meaning and is highly abstract. We evaluate FOCUS on the following indexes: (1) The *accuracy*, *efficiency* and *robustness* of causal structure learning. (2) The *policy return* and *generalization ability* in offline MBRL.

Baselines. We compare FOCUS with the sota offline MBRL algorithm, MOPO, and other online RL algorithms that also learn causal structure in two aspects, causal structure learning and policy learning. (1) MOPO [5] is a well-known and widely-used offline MBRL algorithm, which outperforms standard model-based RL algorithms and prior sota model-free offline RL algorithms on existing offline RL benchmarks. The main idea of MOPO is to artificially penalize rewards by the uncertainty of the dynamics, which can avoid the distributional shift issue. MOPO can be seen as the blank control with a plain world-model. (2) Learning Neural Causal Models from Unknown Interventions (LNCM) [17] is an online MBRL, in which the causal structure learning method can be transformed to the offline setting with a simple adjustment. We take LNCM as an example to show that an online method cannot be directly transferred into offline RL.

Environment. Toy Car Driving. Toy Car driving is a typical RL environment where the agent can control its direction and velocity to finish various tasks including avoiding obstacles and navigating. The information of the car, e.g., position, velocity, direction, and acceleration, can form the state and action in an MDP. In this paper, we use a 2D Toy

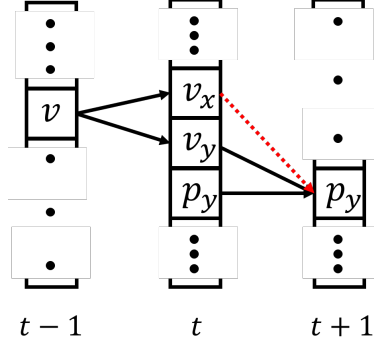


Figure 4: The visualization of the example. The red dotted arrow presents that $(v_x)_t$ is a spurious variable for $(p_y)_{t+1}$.

Car driving as the RL environment where the task of the car is to arrive at the destination (The visualization of states and a detailed description can be found in Appendix C.1). The state includes the direction d , the velocity (scalar) v , the velocity on the x -axis (one dimensional vector) v_x , the velocity on the y -axis v_y and the position (p_x, p_y) . The action is the steering angle a . The visualization of the causal graph can be found in Appendix C.1. This causal structure is designed to demonstrate how a variable become spurious for others and highlight their influence in model learning.

For example, when the velocity v_{t-1} maintains stationary due to an imperfect sample policy, $(v_x)_t$ and $(v_y)_t$ have strong relatedness that $(v_x)_t^2 + (v_y)_t^2 = v_{t-1}^2$ and one can represent the other. Since we design that $(p_y)_{t+1} - (p_y)_t = (v_y)_t$, $(v_x)_t$ and $(p_y)_{t+1} - (p_y)_t$ also have strong relatedness, which leads to that $(v_x)_t$ becomes a spurious variable of $(p_y)_{t+1}$ given $(p_y)_t$, despite that $(v_x)_t$ is not the causal parent of y_{t+1} . By contrast, when the data is uniformly sampled with various velocities, this spuriousness will not exist. **MuJoCo.** MuJoCo [28] is a general-purpose physics engine, which is also a well-known RL environment. MuJoCo includes multijoint dynamics with contact, where the variables of the state represent the positions, angles, and velocity of the agent. The dimensions of the state are from 3 to dozens. The limited dimensions and the clear meaning of each variable provide the convenience of causal structure learning.

Offline Data. We prepare three offline data sets, *Random*, *Medium*, and *Replay* for the Car Driving and MuJoCo. *Random* represents that data is collected by random policies. *Medium* represents that data is collected by a fixed but not well-trained policy, which is the least diverse. *Medium-Replay* is a collection of data that is sampled during training the *Medium* policy, which is the most diverse. The heat map of the data diversity is shown in Appendix C.1.

Table 1: The results on causal structure learning of our model and the baselines. Both the accuracy and the variance are calculated by five times experiments. *FOCUS (-KCI)* represents FOCUS with a linear independence test. *FOCUS (-CONDITION)* represents FOCUS with choosing all other variables as conditional variables.

INDEX	FOCUS	LNCM	FOCUS(-KCI)	FOCUS(-CONDITION)
ACCURACY	0.993	0.52	0.62	0.65
ROBUSTNESS	0.001	0.025	0.173	0.212
EFFICIENCY(SAMPLES)	1×10^6	1×10^7	1×10^6	1×10^6

6.1 Causal Structure Learning

We compare FOCUS with baselines on the causal structure learning with the indexes of the *accuracy*, *efficiency*, and *robustness*. The accuracy is evaluated by viewing the structure learning as a classification problem, where causation represents the positive example and independence represents the negative example. The efficiency is evaluated by measuring the samples for getting a stable structure. The robustness is evaluated by calculating the variance in multiple experiments. The results in Table 1 show that FOCUS surpasses LNCM on accuracy, robustness, and efficiency in causal structure learning. Noticed that LNCM also has a low variance because it predicts the probability of existing causation between any variable pairs with around 50%, which means that its robustness is meaningless.

6.2 Policy Learning

Policy Return. We evaluate the performance of FOCUS and baselines in the two benchmarks on three different and typical offline data sets. The results in Table 2 show that FOCUS outperforms baselines by a significant margin in

Table 2: The comparison on converged policy return in the two benchmarks. The detailed training curves are in Appendix C.1.

ENV	CAR DRIVING			MUJoCo(INVERTED PENDULUM)		
	RANDOM	MEDIUM	REPLAY	RANDOM	MEDIUM	REPLAY
FOCUS	68.1 ± 20.9	-58.9 ± 41.3	86.2 ± 18.2	23.5 ± 17.9	24.9 ± 14.1	49.2 ± 19.0
MOPO	-30.3 ± 49.9	-50.1 ± 34.2	46.2 ± 28.1	8.5 ± 6.2	2.5 ± 0.08	43.4 ± 7.7
LNCM	9.9 ± 42.5	-5.4 ± 32.5	11.4 ± 24.0	13.3 ± 0.9	3.1 ± 0.7	16.3 ± 6.4

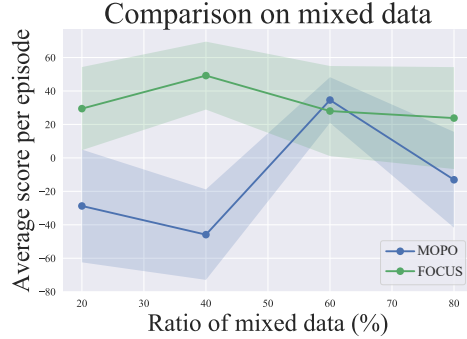


Figure 5: The comparison on generalization in mixed data.

most data sets. In *Random*, FOCUS has the most significant performance gains to the baselines in both benchmarks because of the accuracy of causal structure learning in FOCUS. By contrast, in *Medium-Replay*, the performance gains of FOCUS are least since the high data diversity in *Medium-Replay* leads to weak relatedness of spurious variables (corresponds to small λ), which verifies our theory. In *Medium*, the results in the two benchmarks are different. In Car Driving, the relatively high score of LNCM does not mean that LNCM is the best but all three fail. The failure indicates that extremely biased data makes even the causal model fail to generalize. However, the success of FOCUS in the Inverted Pendulum indicates that causal world-models depend less on the data diversity since FOCUS can still reach high scores in such a biased dataset where the baselines fail. Here we only provide the results in *Inverted Pendulum* but not all the environments in MuJoCo due to the characteristics of the robot control, specifically the frequency of observations, which we present a detailed description in Appendix C.1.

Generalization Ability. We compare the performance on different offline data sets, which is produced by mixing up *Medium-Replay* and *Medium* with different ratios. The $X\%$ in the x-axis represents that the data is mixed by 100% of the *Medium* and $X\%$ of the *Medium-Replay*. The results in Fig 5 show that FOCUS still performs well with a small ratio of *Medium-Replay* data while the baseline performs well only with a big ratio, which indicates that FOCUS is less dependent on the diversity of data.

6.3 Ablation Study

To evaluate the contribution of each component, we perform an ablation study for FOCUS. The results in Table 1 show that the KCI test and our principle of choosing conditional variables contribute to the causal structure learning of both accuracy and robustness.

7 Conclusion

In this paper, we provide theoretical evidence about the advantages of using a causal world-model in offline RL. We present error bounds of model prediction and policy evaluation in offline MBRL with causal and plain world-models. We also propose a practical algorithm, FOCUS, to address the problem of learning causal structure in offline RL. The main idea of FOCUS is to utilize causal discovery methods for offline causal structure learning. We design a general mechanism to solve problems in extending causal discovery methods in RL, which includes conditional variables choosing. Extensive experiments on the typical benchmark demonstrate that FOCUS achieves accurate and robust causal structure learning and thus significantly surpasses baselines in offline RL.

The limitation of FOCUS lies in: In our theory, we assume that the true causal structure is given. However, in practice, one needs to learn it from data and then use it. Quantifying the uncertainty in the learned causal structure from data is known to be a hard problem, and as one line of our future research, we will derive the generalization error bound with the causal structure learned from data.

References

- [1] Jing-Cheng Shi, Yang Yu, Qing Da, Shi-Yong Chen, and An-Xiang Zeng. Virtual-Taobao: Virtualizing real-world online retail environment for reinforcement learning. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI’19)*, Honolulu, HI, 2019.
- [2] Wenjie Shang, Qingyang Li, Zhiwei Qin, Yang Yu, Yiping Meng, and Jieping Ye. Partially observable environment estimation with uplift inference for reinforcement learning based recommendation. *Machine Learning*, (9):2603–2640, 2021.
- [3] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving video database with scalable annotation tooling. *CoRR*, 1805.04687, 2018.
- [4] O. Gottesman, F. Johansson, M. Komorowski, A. Faisal, David Sontag, Finale Doshi-Velez, and L. A. Celi. Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1):16–18, 2019.
- [5] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y. Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. MOPO: model-based offline policy optimization. In *Advances in Neural Information Processing Systems 33 (NeurIPS’20)*, 2020.
- [6] Xiong-Hui Chen, Yang Yu, Qingyang Li, Fan-Ming Luo, Zhiwei Tony Qin, Shang Wenjie, and Jieping Ye. Offline model-based adaptable policy learning. In *Advances in Neural Information Processing Systems 34 (NeurIPS’21)*, Virtual Conference, 2021.
- [7] Fan-Ming Luo, Shengyi Jiang, Yang Yu, Zongzhang Zhang, and Yi-Feng Zhang. Adapting environment sudden changes by learning context sensitive policy. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI’22)*, Virtual Conference, 2022.
- [8] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems 31 (NeurIPS’18)*, pages 2455–2467, Montréal, Canada, 2018.
- [9] Tian Xu, Ziniu Li, and Yang Yu. Error bounds of imitating policies and environments. In *Advances in Neural Information Processing Systems 33 (NeurIPS’20)*, pages 15737–15749, 2020.
- [10] Mark Edmonds, James Kubricht, Colin Summers, Yixin Zhu, Brandon Rothrock, Song-Chun Zhu, and Hongjing Lu. Human causal transfer: Challenges for deep reinforcement learning. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society (CogSci’18)*, 2018.
- [11] Josh Tenenbaum. Building machines that learn and think like people. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’18)*, page 5, 2018.
- [12] 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 *Proceedings of the 8th International Conference on Learning Representations (ICLR’20)*, 2020.
- [13] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems 32 (NeurIPS’19)*, pages 11693–11704, 2019.
- [14] Zheng-Mao Zhu, Shengyi Jiang, Yu-Ren Liu, Yang Yu, and Kun Zhang. Invariant action effect model for reinforcement learning. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI’22)*, Virtual Conference, 2022.
- [15] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2000.
- [16] Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI’11)*, pages 804–813, 2011.
- [17] Nan Rosemary Ke, Olexa Bilaniuk, Anirudh Goyal, Stefan Bauer, Hugo Larochelle, Chris Pal, and Yoshua Bengio. Learning neural causal models from unknown interventions. *CoRR*, 1910.01075, 2019.
- [18] Xiaohai Sun, Dominik Janzing, Bernhard Schölkopf, and Kenji Fukumizu. A kernel-based causal learning algorithm. In *Proceedings of the 24th International Conference on Machine Learning (ICML’07)*, pages 855–862, 2007.

- [19] Dimitris Margaritis. Distribution-free learning of bayesian network structure in continuous domains. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence (AAAI'05)*, pages 825–830, 2005.
- [20] Richard Bellman. A markovian decision process. *Journal of Mathematics and Mechanics*, 6(5):679–684, 1957.
- [21] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA'17)*, pages 1714–1721, 2017.
- [22] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. In *Proceedings of the 6th International Conference on Learning Representations, (ICLR'18)*, 2018.
- [23] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, 2020.
- [24] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik’s cube with a robot hand. *CoRR*, 1910.07113, 2019.
- [25] Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 42(1):80–86, 2000.
- [26] Judea Pearl et al. Models, reasoning and inference. *Cambridge University Press*, 19:2, 2000.
- [27] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009.
- [28] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5026–5033, 2012.

Appendix

A Theory

Definition A.1 (Optimization objective in data distribution \mathcal{D}):

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}} [\mathbf{X}\beta - Y]^2. \quad (6)$$

Definition A.2 (Optimization objective in data \mathcal{D}_{train}):

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} [\mathbf{X}\beta - Y]^2. \quad (7)$$

Definition A.3 (Optimization objective in data \mathcal{D}_{train} with regularization):

$$\min_{\beta} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} [\mathbf{X}\beta - Y]^2 + k \|\beta\|^2, \quad (8)$$

Lemma A.4. *Given that $\omega_{cau} \circ \beta^*$ is the optimal solution of Problem 1, suppose that in \mathcal{D}_{train} , $\mathbf{X}_{spu} = (\mathbf{X} \circ \omega_{cau})\gamma_{spu} + \epsilon_{spu}$ where $\mathbb{E}_{\mathcal{D}_{train}}[\epsilon_{spu}] = 0$ and $\gamma_{spu} \neq \mathbf{0}$, we have that $\hat{\beta}_{spu} \triangleq \omega_{cau} \circ (\beta^* - \lambda\gamma_{spu}) + \lambda\omega_{spu}$ is also an optimal solution of Problem 2 for any λ :*

$$\mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left[\left(|\mathbf{X}(\omega_{cau} \circ \beta^*) - Y|_2 \right) \mid \mathbf{X} \right] = \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left[\left(|\mathbf{X}\hat{\beta}_{spu} - Y|_2 \right) \mid \mathbf{X} \right]$$

Proof.

$$\begin{aligned} & \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left[\left(|\mathbf{X}(\omega_{cau} \circ \beta^*) - Y|_2 \right) \mid \mathbf{X} \right] \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}(\omega_{cau} \circ (\beta^* - \lambda\gamma_{spu} + \lambda\gamma_{spu})) - Y|_2 \right] \mid \mathbf{X} \right\} \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}(\omega_{cau} \circ (\beta^* - \lambda\gamma_{spu})) + \mathbf{X}(\omega_{cau} \circ \lambda\gamma_{spu}) - Y|_2 \right] \mid \mathbf{X} \right\} \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}(\omega_{cau} \circ (\beta^* - \lambda\gamma_{spu})) + \lambda(\mathbf{X}_{spu} - \epsilon_{spu}) - Y|_2 \right] \mid \mathbf{X} \right\} \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}(\omega_{cau} \circ (\beta^* - \lambda\gamma_{spu})) + \lambda(\mathbf{X} \circ \omega_{spu}) - Y|_2 \right] \mid \mathbf{X} \right\} \\ & \quad \left(\text{Since } \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}}[\epsilon_{spu}] = 0 \right) \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}(\omega_{cau} \circ (\beta^* - \lambda\gamma_{spu})) + \lambda\omega_{spu} - Y|_2 \right] \mid \mathbf{X} \right\} \\ &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{D}_{train}} \left\{ \left[|\mathbf{X}\hat{\beta}_{spu} - Y|_2 \right] \mid \mathbf{X} \right\} \\ & \quad \left(\text{Let } \hat{\beta}_{spu} \text{ denote } \omega_{cau} \circ (\beta^* - \lambda\gamma_{spu}) + \lambda\omega_{spu} \right) \end{aligned}$$

□

Lemma A.5 (λ Lemma). *Given λ as the coefficient in Lemma 4.1, and k in Problem 3 chosen by Hoerl-Kennard formula, we have the solution of λ in Problem 3 that:*

$$\lambda = \frac{\beta^* \gamma_{spu}}{\beta^{*2} + \gamma_{spu}^2 + 1 + \frac{\sigma_{spu}^2}{\sigma_{cau}^2} \left(1 + \frac{1}{(\beta^*)^2} \right)} \quad (9)$$

Proof. Since the solution of the ridge regression is

$$\beta(k) = (\mathbf{X}^T \mathbf{X} + k\mathbf{I})^{-1} \mathbf{X}^T \mathbf{Y},$$

we take $\hat{\beta}_{spu}$ into this solution and get:

$$\lambda = \frac{\sigma_{cau}^2 \beta^* \gamma_{spu} k}{\sigma_{cau}^2 \sigma_{spu}^2 + \sigma_{cau}^2 \gamma_{spu}^2 k + \sigma_{cau}^2 k + \sigma_{spu}^2 k + k^2} \quad (10)$$

Since k is chosen by Hoerl-Kennard formula that $k = \frac{\sigma_{spu}^2}{(\beta^*)^2}$, we have:

$$\begin{aligned}
 \lambda &= \frac{\sigma_{cau}^2 \beta^* \gamma_{spu}}{\sigma_{cau}^2 \sigma_{spu}^2 / k + \sigma_{cau}^2 \gamma_{spu}^2 + \sigma_{cau}^2 + \sigma_{spu}^2 + k} \\
 &= \frac{\sigma_{cau}^2 \beta^* \gamma_{spu}}{\sigma_{cau}^2 \sigma_{spu}^2 / \left(\frac{\sigma_{spu}^2}{(\beta^*)^2} \right) + \sigma_{cau}^2 \gamma_{spu}^2 + \sigma_{cau}^2 + \sigma_{spu}^2 + \frac{\sigma_{spu}^2}{(\beta^*)^2}} \\
 &= \frac{\sigma_{cau}^2 \beta^* \gamma_{spu}}{\sigma_{cau}^2 \beta^{*2} + \sigma_{cau}^2 \gamma_{spu}^2 + \sigma_{cau}^2 + \sigma_{spu}^2 + \frac{\sigma_{spu}^2}{(\beta^*)^2}} \\
 &= \frac{\sigma_{cau}^2 \beta^* \gamma_{spu}}{\sigma_{cau}^2 (\beta^{*2} + \gamma_{spu}^2 + 1) + \sigma_{spu}^2 + \frac{\sigma_{spu}^2}{(\beta^*)^2}} \\
 &= \frac{\beta^* \gamma_{spu}}{\beta^{*2} + \gamma_{spu}^2 + 1 + \frac{\sigma_{spu}^2}{\sigma_{cau}^2} \left(1 + \frac{1}{(\beta^*)^2} \right)}
 \end{aligned}$$

□

Proposition A.6. Given λ as Formula 4, we have

$$-\frac{1}{2} \leq \lambda \leq \frac{1}{2}.$$

Proof.

$$\begin{aligned}
 |\lambda| &= \frac{|\beta^* \gamma_{spu}|}{|\beta^{*2} + \gamma_{spu}^2 + 1 + \frac{\sigma_{spu}^2}{\sigma_{cau}^2} \left(1 + \frac{1}{(\beta^*)^2} \right)|} \\
 &\leq \frac{|\beta^* \gamma_{spu}|}{|\beta^{*2} + \gamma_{spu}^2 + 1| + \left| \frac{\sigma_{spu}^2}{\sigma_{cau}^2} \left(1 + \frac{1}{(\beta^*)^2} \right) \right|} \\
 &\leq \frac{|\beta^* \gamma_{spu}|}{|\beta^{*2} + \gamma_{spu}^2 + 1|} \\
 &\leq \frac{|\beta^* \gamma_{spu}|}{|2\beta^* \gamma_{spu} + 1|} \\
 &\leq \frac{1}{2}
 \end{aligned}$$

So we have : $-\frac{1}{2} \leq \lambda \leq \frac{1}{2}$.

□

Theorem A.7 (Spurious Theorem). Let $\mathcal{D} = \{(X, Y)\}$ denote the data distribution, $\hat{\beta}_{spu}$ denote the solution in Lemma 4.1 with λ in Lemma 4.2, and $\hat{Y}_{spu} = X\hat{\beta}_{spu}$ denote the prediction. Suppose that the data value is bounded: $|X_i|_1 \leq X_{max}, i = 1, \dots, n$ and the error of optimal solution ϵ_{cau} is also bounded: $|\epsilon_{cau}|_1 \leq \epsilon_c$, we have the model prediction error bound:

$$\mathbb{E}_{(X,Y) \sim \mathcal{D}}[(|\hat{Y}_{spu} - Y|_1) | X] \leq X_{max} |\lambda|_1 (|\gamma_{spu}|_1 + 1) + \epsilon_c. \quad (11)$$

Proof. Let \hat{Y}_{cau} denote $(\mathbf{X} \circ \omega_{cau})\beta^*$, we have

$$\begin{aligned}
 & \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|\hat{Y}_{spu} - Y|_1 \right) \mid \mathbf{X} \right] \\
 &= \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|(\hat{Y}_{spu} - \hat{Y}_{cau}) + (\hat{Y}_{cau} - Y)|_1 \right) \mid \mathbf{X} \right] \\
 &\leq \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|(\hat{Y}_{spu} - \hat{Y}_{cau})|_1 \right) \mid \mathbf{X} \right] + \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|(\hat{Y}_{cau} - Y)|_1 \right) \mid \mathbf{X} \right] \\
 &\leq \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|\mathbf{X}\lambda(-\omega_{cau} \circ \gamma_{spu} + \omega_{spu})|_1 \right) \mid \mathbf{X} \right] + \epsilon_c \\
 &= \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(|\mathbf{X}\lambda(-\gamma_{spu} + \omega_{spu})|_1 \right) \mid \mathbf{X} \right] + \epsilon_c \\
 &\leq \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(X_{max}|\lambda|_1 * |-\gamma_{spu} + \omega_{spu}|_1 \right) \mid \mathbf{X} \right] + \epsilon_c \\
 &\leq \mathbb{E}_{(\mathbf{X}, Y) \sim D_{test}} \left[\left(X_{max}|\lambda|_1 * (|\gamma_{spu}|_1 + 1) \right) \mid \mathbf{X} \right] + \epsilon_c \\
 &= X_{max}|\lambda|_1(|\gamma_{spu}|_1 + 1) + \epsilon_c
 \end{aligned}$$

□

Theorem A.8 (RL Spurious Theorem). *Given an MDP with the state dimension n_s and the action dimension n_a , a data-collecting policy π_D , let M^* denote the true transition model, M_θ denote the learned model that M_θ^i predicts the i^{th} dimension with spurious variable sets spu_i and causal variables cau_i , i.e., $\hat{S}_{t+1,i} = M_\theta^i((\mathbf{S}_t, \mathbf{A}_t) \circ \omega_{cau_i \cup spu_i})$. Let $V_\pi^{M_\theta}$ denote the policy value of the policy π in model M_θ and correspondingly $V_\pi^{M^*}$. For any bounded divergence policy π , i.e. $\max_S D_{KL}(\pi(\cdot|S), \pi_D(\cdot|S)) \leq \epsilon_\pi$, we have the policy evaluation error bound:*

$$|V_\pi^{M_\theta} - V_\pi^{M^*}| \leq \frac{2\sqrt{2}R_{max}}{(1-\gamma)^2} \sqrt{\epsilon_\pi} + \frac{R_{max}\gamma}{2(1-\gamma)^2} S_{max} [n_s \epsilon_c + (1 + \gamma_{max}) \lambda_{max} n_s (n_s + n_a) R_{spu}] \quad (12)$$

where $R_{spu} = \frac{\sum_{i=1}^{n_s} |spu_i|}{n_s(n_s + n_a)}$, which represents the spurious variable density, that is, the ratio of spurious variables in all input variables.

Proof. Before proving, we first introduce three lemmas:

Lemma A.9.

$$\begin{aligned}
 |V_\pi^{M_\theta} - V_\pi^{M^*}| &\leq |V_\pi^{M^*} - V_{\pi_D}^{M^*}| + |V_{\pi_D}^{M_\theta} - V_{\pi_D}^{M^*}| + |V_{\pi_D}^{M_\theta} - V_\pi^{M_\theta}| \\
 &\leq \frac{2\sqrt{2}R_{max}}{(1-\gamma)^2} \sqrt{\epsilon_\pi} + |V_{\pi_D}^{M_\theta} - V_{\pi_D}^{M^*}|
 \end{aligned}$$

Lemma A.10.

$$|V_{\pi_D}^{M_\theta} - V_{\pi_D}^{M^*}| \leq \frac{R_{max}}{1-\gamma} \sum_s |d_{\pi_D}^{M_\theta}(s) - d_{\pi_D}^{M^*}(s)| \sum_a \pi_D(a|s)$$

Lemma A.11.

$$|d_{\pi_D}^{M_\theta}(s) - d_{\pi_D}^{M^*}(s)| \leq \frac{\gamma}{(1-\gamma)} \sum_{s,a,s'} |M_\theta(S_t, A_t) - M^*(S_t, A_t)| \pi_D(a|s) d_{\pi_D}^{M^*}(s)$$

The detailed proof of these lemmas can be found in [9], which is omitted in this paper. Based on the model prediction error bound in Theorem 4.4, we have:

$$\begin{aligned}
 |M_\theta(S_t, A_t) - M^*(S_t, A_t)| &= \sum_{i=1}^{n_s} |M_\theta^i(S_t, A_t) - M^{*,i}(S_t, A_t)| \\
 &\leq \sum_{i=1}^{n_s} S_{max} [\epsilon_c + (\gamma_{max} + 1) \lambda_{max} |spu_i|] \\
 &= S_{max} [n_s \epsilon_c + (\gamma_{max} + 1) \lambda_{max} \sum_{i=1}^{n_s} |spu_i|] \\
 &= S_{max} [n_s \epsilon_c + (\gamma_{max} + 1) \lambda_{max} n_s (n_s + n_a) R_{spu}]
 \end{aligned}$$

With above lemmas, we have:

$$|V_{\pi}^{M_{\theta}} - V_{\pi}^{M^*}| \leq \frac{2\sqrt{2}R_{max}}{(1-\gamma)^2} \sqrt{\epsilon_{\pi}} + \frac{R_{max}\gamma}{2(1-\gamma)^2} S_{max} [n_s \epsilon_c + (\gamma_{max} + 1) \lambda_{max} n_s (n_s + n_a) R_{spu}]$$

□

B Algorithm

B.1 Choosing the Threshold of p-value

To be fair, we share a common p^* for the testing between any two variables. The choice of p^* significantly influences the accuracy of causal discovery that too small and too big both lead to causal misspecification. The intuition behind our choosing principle is that there is a significant gap in the p value between the causal relation and non-causal relation. Based on this intuition, we partition the probability range $[0, 1]$ into several intervals $[0, p_1), [p_1, p_2), \dots, [p_n, 1]$ according to the sorted p values $\{p_i\}_{i=1}^n$ and design p^* by the formula:

$$p^* = \arg \max_{p_i} \frac{p_{i+1}}{i+1} - \frac{p_i}{i}. \quad (13)$$

If we only consider the biggest gap between p_i , then we will easily choose a big but improper p^* due to the distribution of p_i in some intervals (e.g., $[0.5, 1]$) may be very sparse and thus leads to a big gap.

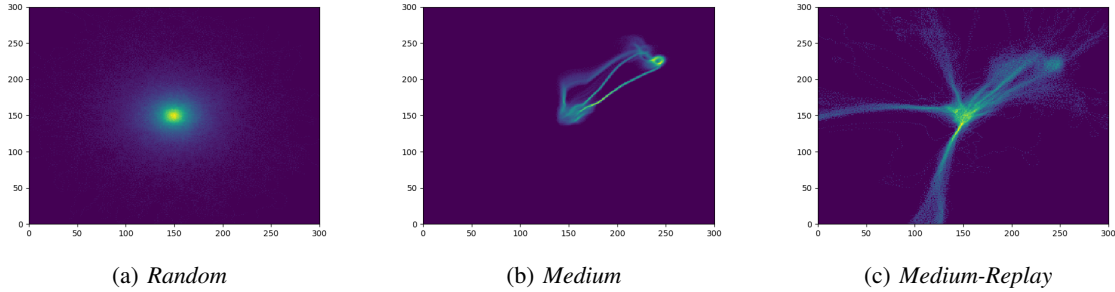


Figure 6: The heat map of the three offline data sets. The high brightness represents high data density.

B.2 Causal Structure Network

The details of Causal Structure Network is shown in Algorithm 2.

Algorithm 2 Causal Structure Network $\mathcal{M}_{Causal}(\cdot)$

Input: state $\mathbf{s}_t \in \mathbb{R}^{n_s}$, action $\mathbf{a}_t \in \mathbb{R}^{n_a}$,
 causal structure mask matrix $G \in \{0, 1\}^{(n_s, n_a) \times n_s}$,
 Make $\mathcal{M}_i(\cdot; \theta_i)$ as the copy of the basic model $\mathcal{M}(\cdot; \theta)$, where $i = 1, \dots, n_s$.
for $i = 1$ **to** n_s **do**
 Let $G_{\cdot, i}$ denote the i^{th} column of G
 Get the masked input $X = (\mathbf{s}_t, \mathbf{a}_t) \circ G_{\cdot, i}$
 Get prediction $\tilde{Y} = \mathcal{M}_i(X; \theta_i) \in \mathbb{R}^{n_s}$
 Let Y_i denote the i^{th} element of \tilde{Y} .
end for
Return $Y = (Y_i)_{i=1}^{n_s}$.

C Experiments

C.1 Environment Details

The heat map of the data diversity is shown in Fig 6. In *Random*, the data is clustered around the origin. In *Medium*, the data is gathered on a fixed trajectory from the origin to the destination. In *Medium-Replay*, the data is much more diverse where a lot of unseen data in above data sets is also sampled.

The visualization of the state in Car Driving and the ground truth of its causal graph are shown in Fig 8.

MuJoCo formulates robot control into MDPs with discrete timestep via equal interval sampling of the continuous-time. Therefore, for each timestep t , s_{t+1} is the result of numerous times of simulation based on s_t with repeated action a_t . Even if spurious variables are existed in one time of simulation, after numerous simulations, the causal effect will be propagated to almost variables, which leads to a full-connection causal graph ($R_{spu} = 0$). Therefore FOCUS degrades into vanilla MOPO in this scenario, which is meaningless to test. Fortunately, after analyzing the propagate progress of the dynamics, we found that the *Inverted Pendulum* is a special case where the causal graph will keep sparse after numerous simulations.

C.2 Experiment Result Details

The detailed training curves are shown in Fig 7.

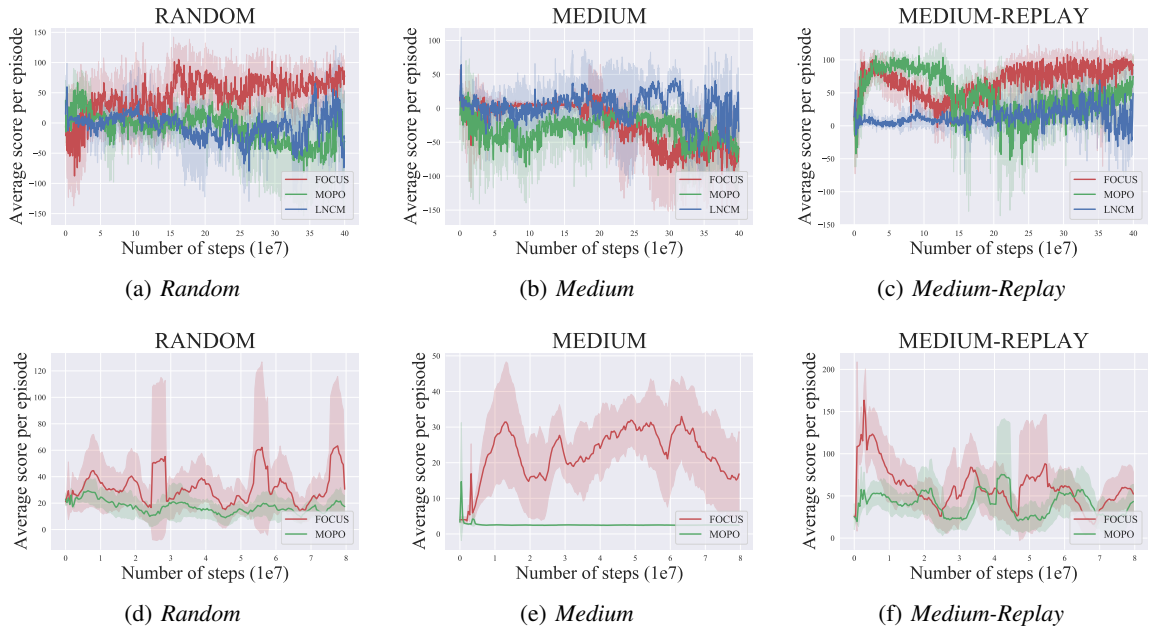


Figure 7: Comparison of FOCUS and the baselines in the two benchmarks. **(a)-(c):** The comparison in the Car Driving on the three datasets. **(d)-(f):** The comparison in the Inverted Pendulum of MuJoCo on the three datasets.

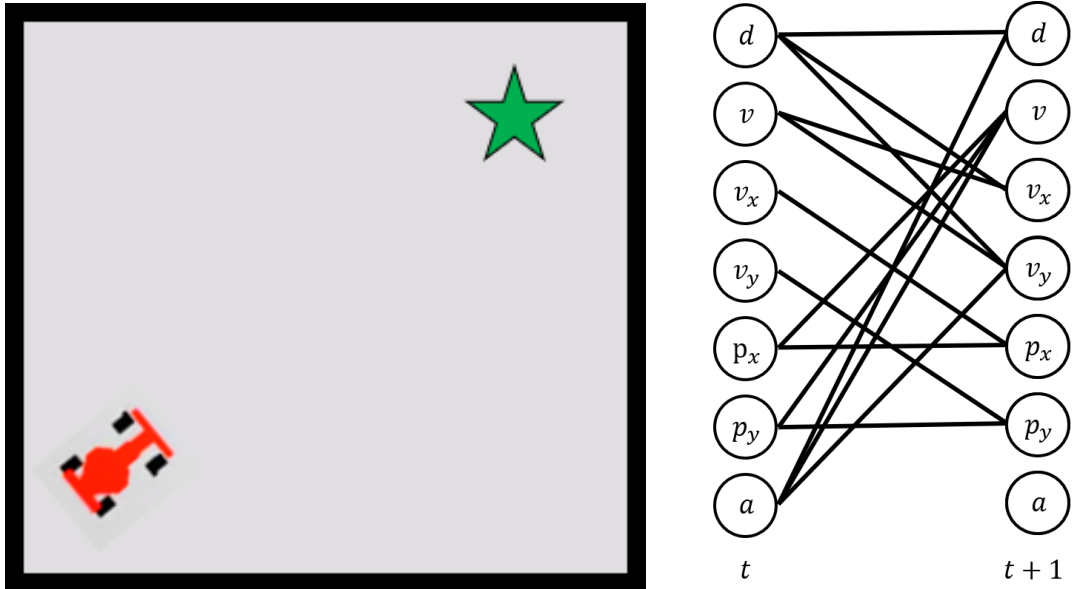


Figure 8: The visualization of the state and the causal structure for the Car Driving benchmark. **Left:** the Toy Car Driving. The goal of the agent is to arrive at the star-shape destination. **Right:** The ground truth of the causal structure in Toy Car Driving. The state is vector-based and its value is continuous.