

Learning by doing: an online causal reinforcement learning framework with causal-aware policy

Ruichu CAI^{1,2*}, Siyang HUANG¹, Jie QIAO¹, Wei CHEN¹, Yan ZENG³, Keli ZHANG⁴,
Fuchun SUN⁵, Yang YU⁶ & Zhifeng HAO⁷

¹*School of Computer Science, Guangdong University of Technology, Guangzhou 510006, China*

²*Pazhou Laboratory (Huangpu), Guangzhou 510555, China*

³*School of Mathematics and Statistics, Beijing Technology and Business University, Beijing 102401, China*

⁴*Huawei Noah's Ark Lab, Shenzhen 518116, China*

⁵*Department of Computer Science and Technology, Tsinghua University, Beijing 100190, China*

⁶*National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210093, China*

⁷*College of Science, Shantou University, Shantou 515063, China*

Received 2 August 2024/Revised 7 February 2025/Accepted 23 April 2025/Published online 29 October 2025

Abstract As a key component to intuitive cognition and reasoning solutions in human intelligence, causal knowledge provides great potential for reinforcement learning (RL) agents' interpretability towards decision-making by helping reduce the searching space. However, there is still a considerable gap in discovering and incorporating causality into RL, which hinders the rapid development of causal RL. In this paper, we consider explicitly modeling the generation process of states with the causal graphical model, based on which we augment the policy. We formulate the causal structure updating into the RL interaction process with active intervention learning of the environment. To optimize the derived objective, we propose a framework with theoretical performance guarantees that alternates between two steps: using interventions for causal structure learning during exploration and using the learned causal structure for policy guidance during exploitation. Due to the lack of public benchmarks that allow direct intervention in the state space, we design the root cause localization task in our simulated fault alarm environment and then empirically show the effectiveness and robustness of the proposed method against state-of-the-art baselines. Theoretical analysis shows that our performance improvement is attributed to the virtuous cycle of causal-guided policy learning and causal structure learning, which aligns with our experimental results. Codes are available at https://github.com/DMIRLAB-Group/FaultAlarm_RL.

Keywords causal reinforcement learning, reinforcement learning, causality, online reinforcement learning, causal structure learning

Citation Cai R C, Huang S Y, Qiao J, et al. Learning by doing: an online causal reinforcement learning framework with causal-aware policy. *Sci China Inf Sci*, 2026, 69(2): 122104, <https://doi.org/10.1007/s11432-024-4403-7>

1 Introduction

How to decide the next action in repairing the cascading failure under a complex dynamic online system? Such a question refers to multifarious decision-making problems in which reinforcement learning (RL) has achieved notable success [1–4]. However, most off-the-shelf RL methods contain a massive decision space and a black-box decision-making policy, thus usually suffering from low sampling efficiency, poor generalization, and lack of interpretability. As such, current efforts [5, 6] incorporate domain knowledge and causal structural information into RL to help reduce the searching space as well as improve the interpretability; e.g., a causal structure enables to locate the root cause guiding the policy decision. With the causal knowledge, recent RL approaches are mainly categorized as implicit and explicit modeling-based.

Implicit modeling-based approaches mostly ignore the detailed causal structure and only focus on extracting the task-invariant representations to improve the generalizability in unseen environments [7–12]. For instance, Ref. [8] proposed a method that extracted the reward-relevant representations while eliminating redundant information. In contrast, explicit modeling-based approaches seek to model the causal structure of the transition of the Markov decision process (MDP) [13–20]. For instance, Ref. [16] proposed a method to learn the causal structure among states and actions to reduce the redundancy in modeling while Ref. [13] utilized the causal structure of MDP through a planning-based method. However, these explicit modeling methods either rely on the causal knowledge from domain experts or might suffer from low efficiency in learning policy due to the indirect usage of causal structure in planning and the possible inefficient randomness-driven exploration paradigm.

* Corresponding author (email: cairuichu@gmail.com)




Figure 1 (Color online) Intervention-Inference-Guidance loop of online causal reinforcement learning.

Inspired by the intervention from causality and the decision nature of RL actions in online reinforcement learning: a random action is equivalent to producing an intervention on a certain state such that only its descendants will change while its ancestors will not; a decision could be made according to the causal influence of the action to a certain goal. As such, a causal structure can be learned through interventions by detecting the changing states, which in turn guides a policy with the causal knowledge from the learned causal structure. Although there has been recent interest in related subjects in causal reinforcement learning, most of them seek to learn a policy either with a fixed prior causal model or a learned but invariant one [13, 16, 19, 21], which does not naturally fit our case when the causal model is dynamically updated iteratively via interventions while learning policy learning (i.e., learning by doing), along with the theoretical identifiability and performance guarantees.

In this work, as shown in Figure 1, we propose an online causal reinforcement learning framework that reframes RL's exploration and exploitation trade-off scheme. In exploration, we devise an inference strategy using intervention to efficiently learn the causal structure between states and actions, modeling simultaneously causal dynamics of the environment; while in exploitation, we take the best of the learned structure to develop a causal-knowledge-triggered mask, which leads to a highly effective causal-aware policy. As such, the causal environment, the causal structure inference strategy, and the causal-aware policy construct a virtuous cycle to the online causal reinforcement learning framework.

In particular, our framework consists of causal structure learning and policy learning. For causal structure learning, we start by explicitly modeling the environmental causal structure from the observed data as initial knowledge. Then we formulate the causal structure updating into the RL interaction process with active intervention learning of the environment. This novel formulation naturally utilizes post-interaction environmental feedback to assess treatment effects after applying the intervention, thus enabling correction and identification of causality. For policy learning, we propose to construct the causal mask based on the learned causal structure, which helps directly reduce the decision space and thus improves sample efficiency. This leads to an optimization framework that alternates between causal discovery and policy learning to gain generalizability. Under some mild conditions, we prove the identifiability of the causal structure and the theoretical performance guarantee of the proposed framework.

To demonstrate the effectiveness of the proposed approach, we established a high-fidelity fault alarm simulation environment in the communication network in the operations and maintenance (O&M) scenario, which requires powerful reasoning capability to learn policies. We conduct comprehensive experiments in such an environment, and the experimental results demonstrate that the agent with causal learning capability can learn the optimal policy faster than the state-of-the-art model-free RL algorithms, reduce the exploration risk, and improve the sampling efficiency. Additionally, the interaction feedback from the environment can help learn treatment effects and thus update and optimize causal structure more completely. Furthermore, our framework with causality can also be unified to different backbones of policy optimization algorithms and be easily applied to other real-world scenarios.

The main contributions are summarized as follows.

- We propose an online causal reinforcement learning framework, including causal structure and policy learning. It interactively constructs compact and interpretable causal knowledge via intervention (doing), in order to facilitate policy optimization (learning).
- We propose a causal structure learning method that automatically updates local causal structures by evaluating the treatment effects of interventions during agent-environment interactions. Based on the learned causal model, we also develop a causal-aware policy optimization method triggered by a causal mask.
- We derive theoretical guarantees from aspects of both causality and RL: identifiability of the causal structure and performance guarantee of the iterative optimization on the convergence of policy that can be bounded by the causal structure.
- We experimentally demonstrate that introducing causal structure during policy training can greatly reduce the action space, decrease exploration risk, and accelerate policy convergence.

2 Related work

Reinforcement learning. RL solves sequential decision problems by trial and error, aiming to learn an optimal policy to maximize the expected cumulative rewards. RL algorithms can be conventionally divided into model-free and model-based methods. The key idea of the model-free method is that agents update the policy based on the experience gained from direct interactions with the environment. In practice, model-free methods are subdivided into value-based and policy-based ones. Value-based methods select the policy by estimating the value function, and representative algorithms include deep Q-network (DQN) [22], deep deterministic policy gradient (DDPG) [23], and dueling double DQN (D3QN) [24]. Policy-based methods directly learn the policy function without approximating the value function. The current mainstream algorithms are proximal policy optimization (PPO) [25], trust region policy optimization (TRPO) [26], A2C, A3C [27], and SAC [28]. The model-free approach reaches a more accurate solution at the cost of larger trajectory sampling, while the model-based approach achieves better performance with fewer interactions [29–33]. Despite the better performance of the model-based approach, it is still more difficult to train the environment model, and the model-free approach is more general for real-world applications. In this paper, we apply our approach to the model-free methods.

Causal reinforcement learning. Causal RL [34–36] is a research direction that combines causal learning with reinforcement learning. Ref. [16] proposed to extract relevant state representations based on the causal structure between partially observable variables to reduce the error of redundant information in decision-making. Refs. [7, 14] discovered simple causal influences to improve the efficiency of reinforcement learning. Refs. [37, 38] proposed counterfactual-based data augmentation to improve the sample efficiency of RL. Building dynamic models in model-based RL [5, 6, 39] based on causal graphs has also been widely studied recently. Ref. [5] leveraged the structural causal model as a compact way to encode the changeable modules across domains and applied them to model-based transfer learning. Ref. [6] proposed a causal world model for offline reinforcement learning that incorporated causal structure into neural network model learning. Most of them utilize pre-defined or pre-learned causal graphs as prior knowledge or detect single-step causality to enhance the RL policy learning. However, none of them used the intervention data of the interaction process with the environment to automatically discover or update the complex causal graph. Our method introduces a self-renewal interventional mechanism for the causal graph based on causal effects, which ensures the accuracy of causal knowledge and greatly improves the strategy's efficiency.

Causal discovery. Causal discovery aims to identify the causal relationships between variables. Typical causal discovery methods from observational data are constraint-based methods, score-based methods, and function-based methods. Constraint-based methods, such as PC and FCI algorithms [40], rely on conditional independence tests to uncover an underlying causal structure. Different from constraint-based methods, score-based methods use a score to determine the causal direction between variables of interest [41–43]. But both constraint-based methods and score-based methods suffer from the Markov equivalence class (MEC) problem, i.e., different causal structures imply the same conditional independence tests. By utilizing the data generation process assumptions, like the linear non-Gaussian assumption [44] and the additive noise assumption [45–47], function-based methods are able to solve the MEC problem and recover the entire causal structure.

Furthermore, leveraging additional interventional information can provide valuable guidance for the process of causal discovery [48, 49]. An intuitive concept involves observing changes in variables following an intervention on another variable. If intervening in one variable leads to changes in other variables, it suggests a potential causal relationship between the intervened variable and the variables that changed.

3 Problem formulation

In this section, we majorly give our model assumptions and relevant definitions to formalize the problem. We concern the RL environment with an MDP $\langle \mathcal{S}, \mathcal{A}, p, r, \gamma \rangle$, where $\mathcal{S}$ denotes the state space, $\mathcal{A}$ denotes the action space, $p(\mathbf{s}'|\mathbf{s}, a)$ denotes the dynamic transition from state $\mathbf{s} \in \mathcal{S}$ to the next state $\mathbf{s}'$ when performing action $a \in \mathcal{A}$ in state $\mathbf{s}$, r is a reward function with $r(\mathbf{s}, a)$ denoting the reward received by taking action a in state $\mathbf{s}$, and $\gamma \in [0, 1]$ is a discount factor.

To formally investigate the causality in online RL, we make the following factorization state space assumption.

Assumption 1 (Factorization state space). The state variables in the state space $\mathcal{S} = \{s_1 \times s_2 \times \cdots \times s_{|\mathcal{S}|}\}$ can be decomposed into disjoint components $\{s_i\}_{i=1}^{|\mathcal{S}|}$.

Assumption 1 implies that the factorization state space has explicit semantics on each state component and thus the causal relationship among states can be well defined. Such an assumption can be satisfied through an abstraction of states which has been extensively studied [9, 50].

Considering that causality implies the underlying physical mechanism, we can formulate the one-step Markov decision process with the causal graphical model¹⁾ [51] as follows.

Definition 1 (Causal graph on Markov decision process). Let $\mathcal{G} = (V_{\mathcal{S}}, E)$ denote the causal graph where $V_{\mathcal{S}}$ is the vertex set defined on the state space, and the edge set E represents the causal relationships among vertex. Given the total time span $[1, 2, \dots, T]$, the causal relationship on the one-step transition dynamics can be represented through the factored probability:

$$p(s_1^t, s_2^t, \dots, s_{|\mathcal{S}|}^t | s_1^{t-1}, s_2^{t-1}, \dots, s_{|\mathcal{S}|}^{t-1}) = \prod_{i=1}^{|\mathcal{S}|} p(s_i^t | \mathbf{s}_{\mathbf{Pa}_i}^{t-1}), \quad (1)$$

where $|\mathcal{S}|$ is the support of the state space, $\mathbf{Pa}_i := \{s_j | s_j \rightarrow s_i \in E\}$ denotes the parent set of s_i according to causal graph $\mathcal{G}$, and $\mathbf{s}_{\mathbf{Pa}_i}^{t-1}$ is the parent states from the last time step.

To establish a rigorous framework for causal reasoning in MDPs, we introduce the following assumptions, which generalize classical causal assumptions to the temporal domain.

Assumption 2 (Causal Markov assumption in MDP). A causal graph $\mathcal{G} = \{V_{\mathcal{S}}, E\}$ and a probability transition distribution $p(\mathbf{s}_i^t | \mathbf{s}_{\mathbf{Pa}_i}^{t-1})$ satisfy the Markov condition if and only if for every s_i^t state, s_i^t is independent of $\mathbf{s}_i^{1:T} \setminus \{\mathbf{s}_{\text{Des}_i^{\mathcal{G}}}^{t+1:T} \cup \mathbf{s}_{\mathbf{Pa}_i^{\mathcal{G}}}^{t-1}\}$ given $\mathbf{s}_{\mathbf{Pa}_i^{\mathcal{G}}}^{t-1}$ for all t in MDP, where $\mathbf{s}_i^{1:T}$ denote the set of state variables i from time 1 to T , and $\mathbf{s}_{\text{Des}_i^{\mathcal{G}}}^{t+1:T}$ denotes the descendant of s_i from time $t+1$ to T .

Assumption 3 (Causal faithfulness assumption in MDP). Let $\mathcal{G} = \{V_{\mathcal{S}}, E\}$ be a causal graph and $p(\mathbf{s}_i^t | \mathbf{s}_{\mathbf{Pa}_i}^{t-1})$ a transition distribution generated by $\mathcal{G}$. $\langle \mathcal{G}, p \rangle$ satisfies the faithfulness condition if and only if every conditional independence relation true in p is entailed by the causal Markov condition applied to $\mathcal{G}$ at any time in MDP.

Assumption 4 (Causal sufficiency assumption in MDP). A set of state variables $V_{\mathcal{S}}$ in $\mathcal{G}$ is causally sufficient if and only if there are no latent confounders of any two observed state variables at any time in MDP.

These assumptions are just the generalized version of the original one in the time domain such that causal structure is defined between the last time and the current time using independence. With these assumptions, we can develop the identifiability results for learning a causal graph in MDP.

An example of such a causal graph in MDP is given in Figure 2(a). In our framework, actions are modeled as interventions, which inherently influence the state. To capture this, we explicitly consider the impact of each action on the state. Without loss of generality, we can model the action on each state as a binary treatment $I_i \in \{0, 1\}$ for state s_i , where $I_i = 0$ indicates the state receives no intervention (natural evolution), and $I_i = 1$ indicates the state receives the treatment (treated) under which an intervention is performed. For example, $I_2 = 1$ at time t in Figure 2(a) means that there is an intervention $\text{do}(s_2)$ on s_2^t such that the effect of all parents on s_2^t is removed. Such action modeling is commonly encountered in many scenarios like network operation, robot control, etc. In such a case, we have $p(\text{do}(s_2^t) | \mathbf{s}_{\mathbf{Pa}_2}^{t-1}) = p(\text{do}(s_2^t))$ [52]. As such, the policy serves as the treatment assignment for each state, and the action space is structured such that each dimension corresponds to a binary intervention on a specific state variable (i.e., I_i for s_i). This design ensures that the action space spans the same dimensions as the state space: every state variable has an associated intervention “lever” in the action space so that we can intervene

¹⁾ Generally, in causality, a directed acyclic graph that represents a causal structure is termed a causal graph [40]. Here we generalize each state variable at a timestep t as one variable of interest.




Figure 2 (Color online) Illustration of online causal reinforcement learning framework. (a) A full-time causal graph in MDP and the action on the state can be viewed as an intervention. (b) The summary causal graph of (a) where each state would trigger the next state's occurrence, resulting in a cascade error. The actions on (c) s_1 and (d) s_2 from the policy depends on a given situation S^{t-1} as well as the causal mask.

the state and measure the effect of certain outcomes. This allows us to learn the causal influence within each state, which will further improve policy learning by selecting the most influenced action to the goal. While we assume full intervention capability across all state dimensions for simplicity, this framework readily extends to scenarios where certain states remain non-intervenable by omitting their corresponding action dimensions. Based on Definition 1, we can define the average treatment effect among states.

Definition 2 (Average treatment effect (ATE) on states). Let s_i and s_j denote two different state variables. Then the treatment effect of s_i on s_j is

$$\mathcal{C}_{s_i \rightarrow s_j} = \mathbb{E}[s_j(I_i = 1) - s_j(I_i = 0)], \quad (2)$$

where $s_j(I_i = 1)$ denotes the potential outcome of s_j if s_i were treated (intervened), $s_j(I_i = 0)$ denotes the potential outcome if s_i were not treated [53].

Intuitively, the potential outcome depicts the outcome of the state in performing different treatments and the ATE evaluates the treatment effect on the outcome. That is, ATE answers the question that when an agent performs an action $\text{do}(s_i)$, how is the average cause of an outcome of s_j [52]? Such a question suggests that an action applied to a state will solely influence its descendants and not its ancestors. This aspect is crucial for causal discovery, as it reveals the causal order among the states. Moreover, the treatment is not necessarily binary since our goal is to infer the causal order by the property of intervention in action, i.e., an intervention on the cause will influence its effect, which is also held in multi-treatment [54] or the continuous-treatment [55]. One can simply modify the corresponding ATE to adapt to the general treatment. For simplicity, we assume binary treatment in this work. To further accomplish the causal discovery, we assume that the states satisfy the causal sufficiency assumption [51], i.e., there are no hidden confounders and all variables are observable.

4 Framework

In this section, with proper definitions and assumptions, we first propose a general online causal reinforcement learning framework, which consists of two phases: policy learning and causal structure learning. Then, we describe these two phases in detail and provide a performance guarantee for them. The overall flow of our framework is eventually summarized in Algorithm 1.

4.1 Causal-aware policy learning

The general objective of RL is to maximize the expected cumulative reward by learning an optimal policy $\max_{\pi} \mathbb{E} \left[\sum_{t=0}^T \gamma^t r(s^t, a^t) \right]$. Inspired by viewing the action as the intervention on state variables, we use the fact that the causal structure $\mathcal{G}$ among state variables is effective in improving the policy decision space, proposing the causal-aware policy $\pi_{\mathcal{G}}(\cdot|s)$ with the following objective function for optimization:

$$\max_{\pi_{\mathcal{G}}} \mathbb{E} \left[\sum_{t=0}^T \gamma^t r(s^t, a^t) \right]. \quad (3)$$

Let us consider a simple case where we have already obtained a causal graph $\mathcal{G}$ of the state-action space. We now define a causal policy and associate it with the state-space causal structure $\mathcal{G}$.

Algorithm 1 Online causal reinforcement learning training process.

Input: Policy network θ ; replay buffer $\mathcal{B}$; causal structure $\mathcal{G}$;
while θ not converged **do**
 // Causal-aware policy learning
 while $t < T$ **do**
 $a^t \leftarrow$ Causal policy $\pi_{\mathcal{G}}(\cdot | \mathbf{s}^t)$ with causal mask $M_{\mathbf{s}^t}(\mathcal{G})$;
 $\mathbf{s}^{t+1}, r^t \leftarrow \text{Env}(\mathbf{s}^t, a^t)$;
 $\mathcal{B} \leftarrow \mathcal{B} \cup \{a^t, \mathbf{s}^t, r^t, \mathbf{s}^{t+1}\}$;
 // Causal structure learning
 for $i \leq |\mathcal{S}|$ **do**
 for $j \leq |\mathcal{S}|$ **do**
 Estimate $\hat{\mathcal{C}}_{s_i \rightarrow s_j}^{Att}$ from $\mathcal{B}$;
 Infer the causal relation between s_i, s_j based on $\hat{\mathcal{C}}_{s_i \rightarrow s_j}^{Att}$ (Theorem 1);
 Prune redundant edges of $\mathcal{G}$;
 Update θ with $\mathcal{B}$.

Definition 3 (Causal policy). Given a causal graph $\mathcal{G}$ on the state space, we define the causal policy $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$ under the causal graph $\mathcal{G}$ as follows:

$$\pi_{\mathcal{G}}(\cdot | \mathbf{s}) = M_{\mathbf{s}}(\mathcal{G}) \circ \pi(\cdot | \mathbf{s}), \quad (4)$$

where $M_{\mathbf{s}}(\mathcal{G})$ is the causal mask vector at state $\mathbf{s}$ w.r.t. $\mathcal{G}$, $\pi(\cdot | \mathbf{s})$ is the action probability distribution, and $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$ is the distribution of causal policy where each action is masked according to $M_{\mathbf{s}}(\mathcal{G})$.

The causal mask $M_{\mathbf{s}}(\mathcal{G}) = \{m_{\mathbf{s},a}^{\mathcal{G}}\}_{a=1}^{|\mathcal{A}|}$ is induced by the causal structure and the current state, aiming to pick out causes of the state and refine the searching space of policy. In other words, it ensures that all irrelevant actions can be masked out. For example, in a cascade error scenario of communication in Figure 2(b), where each state (e.g., system fault alarm) would trigger the next state's occurrence, resulting in cascade and catastrophic errors in communication networks, the goal here is to learn a policy that can quickly eliminate system fault alarms. The most effective and reasonable solution is to intervene on the root cause of the state, to prevent possible cascade errors. In Figure 2(b), we should intervene on s_2 since s_1 is not an error and s_2 is the root cause of the system on its current state.

For more general cases, based on the causal structure of errors, we can obtain the *TopK* causal order representing K possible root-cause errors and construct the causal mask vector to refine the decision space to a subset of potential root-cause errors. This is, the i -th element in $M_{\mathbf{s}}(\mathcal{G})$ is not masked ($m_{\mathbf{s},i}^{\mathcal{G}} = 0$) only if $s_i \in \text{Top}K_{\tilde{\mathcal{G}}}$ where $\text{Top}K_{\tilde{\mathcal{G}}}$ is the *TopK* causal order of $\tilde{\mathcal{G}}$, and $\tilde{\mathcal{G}} := \mathcal{G} \setminus \{s_i | s_i^t = 0\}$, K denotes the number of candidate causal actions. It is worth mentioning that different tasks correspond to different causal masks, but the essential role of the causal mask is to use causal knowledge to retain task-related actions and remove task-irrelevant actions, thus helping the policy to reduce unnecessary sampling. For example, for some goal Y , the causal mask can be set to $m_{\mathbf{s},i}^{\mathcal{G}} \propto |\mathcal{C}_{i \rightarrow y}|$ which is proportional to the causal effect where $\mathcal{C}_{i \rightarrow y} = \mathbb{E}[Y(I_i = 1) - Y(I_i = 0)]$ so that the causal mask can be task-specific for different goal. Note that some relevant causal imitation learning algorithms exist that utilize similar mask strategies [35,56]. However, they focus on imitation learning settings other than reinforcement learning. And they use the causal structure accurately while we take the best of causal order information, allowing the presence of transitory incomplete causal structures in iterations and improving computational efficiency.

In practice, we use an actor-critic algorithm PPO [25] as the original policy, which selects the best action via maximizing the Q value function $Q(\mathbf{s}^t, a^t)$. Notice that our method is general enough to be integrated with any other RL algorithms.

4.2 Causal structure learning

In this phase, we relax the assumption of giving $\mathcal{G}$ as a prior and aim to learn the causal structure through the online RL interaction process. As discussed before, an action is to impose a treatment and perform an intervention on the state affecting only its descendants while not its ancestors. As such, we develop a two-stage approach for learning causal structure with orientation and pruning stages.

In the orientation stage, we aim to estimate the treatment effect for each pair to identify the causal order of each state. However, due to the counterfactual characteristics in the potential outcome [52], i.e., we can not observe both control and treatment happen at the same time, and thus a proper approximation must be developed. In this work, instead of estimating ATE, we propose to estimate the average treatment effect for the treated sample (ATT) [57]:

$$\hat{\mathcal{C}}_{s_i \rightarrow s_j}^{Att} = \frac{1}{n} \sum_{\{k: I_i=1\}} [s_j^{(k)}(I_i = 1) - \hat{s}_j^{(k)}(I_i = 0)], \quad (5)$$

where n denotes the number of treated samples when $I_i = 1$, $s_j^{(k)}(I_i = 1)$ is the k -th observed sample, and $\hat{s}_j^{(k)}(I_i = 0)$ is an estimation that can be estimated from the transition in 1.

Theorem 1. Given a causal graph $\mathcal{G} = (V_S, E)$, for each pair of states s_i, s_j with $i \neq j$, s_i is the ancestor of s_j , i.e., s_i has a direct path to s_j if and only if $|\mathcal{C}_{s_i \rightarrow s_j}^{Att}| > 0$.

Please see Appendix B for detailed proofs of all theorems and lemmas. Theorem 1 ensures that ATT can be used to identify the causal order. However, redundant edges might still exist even when accounting for the causal order. To address this, we introduce a pruning stage and formulate a pruning method using a score-based approach to refine the causal discovery results. Specifically, the aim of causal structure learning can be formalized as maximizing the score of log-likelihood with an ℓ_0 -norm penalty:

$$\max_{\mathbf{G}} \sum_{t=1}^T \sum_{i=1}^{|S|} \log p(s_i^t | \mathbf{s}_{\mathbf{Pa}_i}^{t-1}) - \alpha \|\mathbf{G}\|_0, \quad (6)$$

where $\mathbf{G}$ is the adjacency matrix of the causal graph [41]. Note that such ℓ_0 -norm can be relaxed to a quadratic penalty practically for optimization [58] but we stick to the ℓ_0 -norm here for the theoretical plausibility. Then by utilizing the score in (6), we can prune the redundant edges by checking whether the removed edge can increase the score above. We continue the optimization until no edge can be removed. By combining the orientation and the pruning stage, the causal structure is identifiable, which is illustrated theoretically in Theorem 2.

Theorem 2 (Identifiability). Under the causal faithfulness and causal sufficiency assumptions, given the correct causal order and large enough data, the causal structure among states is identifiable from observational data.

4.3 Performance guarantees

To analyze the performance of the optimization of the causal policy, we first list the important Lemma 1 where the differences between two different causal policies are highly correlated with their causal graphs, and then show that policy learning can be well supported by the causal learning.

Lemma 1. Let $\pi_{\mathcal{G}^*}(\cdot | \mathbf{s})$ be the policy under the true causal graph $\mathcal{G}^* = (V_S, E^*)$. For any causal graph $\mathcal{G} = (V_S, E)$, when the defined causal policy $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$ converges, the following inequality holds:

$$D_{TV}(\pi_{\mathcal{G}^*}, \pi_{\mathcal{G}}) \leq \frac{1}{2} (\|M_{\mathbf{s}}(\mathcal{G}) - M_{\mathbf{s}}(\mathcal{G}^*)\|_1 + \|\mathbf{1}_{\{a: m_{\mathbf{s},a}^{\mathcal{G}^*}=1 \wedge m_{\mathbf{s},a}^{\mathcal{G}}=1\}}\|_1), \quad (7)$$

where $\|M_{\mathbf{s}}(\mathcal{G}) - M_{\mathbf{s}}(\mathcal{G}^*)\|_1$ is the ℓ_1 -norm of the masks measuring the differences of two policies, $\mathbf{1}$ is an indicator function and $\|\mathbf{1}_{\{a: m_{\mathbf{s},a}^{\mathcal{G}^*}=1 \wedge m_{\mathbf{s},a}^{\mathcal{G}}=1\}}\|_1$ measures the number of actions that are not masked on both policies.

Lemma 1 shows that the total variation distance between two policies $\pi_{\mathcal{G}^*}(\cdot | \mathbf{s})$ and $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$, is upper bounded by two terms that depend on the divergence between the estimated causal structure (causal masks) and the true one. It bridges the gap between causality and reinforcement learning, which also verifies that causal knowledge matters in policy optimization. In turn, this lemma facilitates the improvement of the value function's performance, as shown in Theorem 3.

Theorem 3. Given a causal policy $\pi_{\mathcal{G}^*}(\cdot | \mathbf{s})$ under the true causal graph $\mathcal{G}^* = (V_S, E^*)$ and a policy $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$ under the causal graph $\mathcal{G} = (V_S, E)$, recalling $R_{\max}$ is the upper bound of the reward function, we have the performance difference of $\pi_{\mathcal{G}^*}(\cdot | \mathbf{s})$ and $\pi_{\mathcal{G}}(\cdot | \mathbf{s})$ being bounded as

$$V_{\pi_{\mathcal{G}^*}} - V_{\pi_{\mathcal{G}}} \leq \frac{R_{\max}}{(1-\gamma)^2} (\|M_{\mathbf{s}}(\mathcal{G}) - M_{\mathbf{s}}(\mathcal{G}^*)\|_1 + \|\mathbf{1}_{\{a: m_{\mathbf{s},a}^{\mathcal{G}^*}=1 \wedge m_{\mathbf{s},a}^{\mathcal{G}}=1\}}\|_1). \quad (8)$$

An intuition of performance guarantees is that policy exploration helps to learn better causal structures through intervention, while better causal structures indicate better policy improvements. The detailed proofs of the above lemma and theorems are in Appendix B.

5 Experiments

In this section, we first discuss the basic setting of our designed environment as well as the baselines used in the experiments. Then, to evaluate the proposed approach, we conducted comparative experiments on the environment and provided the numerical results and detailed analysis.

5.1 Environment design

Since most commonly used RL benchmarks do not explicitly allow causal reasoning, we constructed FaultAlarmRL, a simulated fault alarm environment based on the real alarm data in the real-world application of wireless communication networks [59].

FaultAlarmRL environment is designed to mimic the operation process in a large communication network within an MDP framework. In the O&M process of such networks, efficiently and accurately locating the root cause of alarms within a given time period is crucial. Timely fault elimination improves O&M efficiency and ensures communication quality. In real wireless networks, the alarm event sequences of different nodes influence each other through the node topology, and the causal mechanisms between different types of alarm events are also affected by the underlying topology.

The simulation environment contains 50 device nodes and 18 alarm types, with the true causal relationships between alarm types and the meaning of each alarm type shown in Appendix E, Table E2. Alarm events are generated by root cause events based on the alarm causal graph and device topology graph propagation. There also exist spontaneous noise alarms in the environment. To mimic the operation in a large communication network, we designed an MDP transition environment modified from the topological Hawkes process. For example, the number of alarm events that occur in X_{t+1} is determined by the number of alarms in the previous time interval X_t without decay. This means that alarms persist until they are “fixed” and this type of transition constructs an MDP environment where the alarm propagation process can be expressed as

$$\begin{aligned} p(s_{t+1}|s_t, a_t; G_V, G_N) &= P(\mathbf{X}_{t+1}|\mathbf{X}_t; G_V, G_N) \\ &= \prod_{n \in N, v \in V} P(X_{n,v,t+1}|X_{n,PA_v,t}) \\ &= \prod_{n \in N, v \in V} \text{Pois}(X_{n,v,t+1}; \lambda_v(n, t+1)), \end{aligned}$$

where $X_{n,v,t+1}$ is the count of occurrence events of event type v at node n in the time interval $[t+1-\Delta t, t+1]$, Pois is the Poisson distribution, and $\lambda_v(n, t)$ is the Hawkes process intensity function. Specifically, $\lambda_v(n, t)$ is defined as

$$\lambda_v(n, t) = \mu_v + \sum_{v' \in PA_v} \sum_{n' \in N} \sum_{k=0}^K \alpha_{v',v,k} \hat{A}_{n',n}^K \kappa X_{n',v',t-1},$$

where $X_{n,v,t-1}$ is the count of occurrence alarms of type v at node n in the time interval $[t-1-\Delta t, t-1]$, κ is the exponential kernel function, k is the maximum hop, $\alpha_{v',v,k}$ is the propagation intensity function of the alarm, $\hat{A} := D^{-1/2}AD^{-1/2}$ is the normalized adjacency matrix of the topological graph, A is the adjacency matrix, D is the diagonal degree matrix, $\hat{A}_{n',n}^K$ denotes the n', n -th entries of the K -hop topological graph, and μ_v is the spontaneous intensity function of the alarm v .

The state in FaultAlarmRL is the current observed alarm information, which includes the time of the fault alarm, the fault alarm device, and the fault alarm type. The state space has $50 \times 18 \times 2 = 1800$ dimensions. The action space contains 900 discrete actions, each of which represents a specific alarm type on a specific device. We define the reward function as

$$r = \frac{N_t - N_{t+1}}{N_t} - \frac{t}{\text{step}_{\max}},$$

where N_t represents the number of alarms at time t , and $\text{step}_{\max}$ is the maximum number of steps in an episode, which is set to 100. Please see Appendix E for further details on the hyper-parameters of the environment. Additionally, we further evaluate our method in *cart-pole* environment from the OpenAI Gym toolkit (see Appendix D).

5.2 Experimental setups

We evaluate the performance of our methods in terms of both causal structure learning and policy learning. We first sampled 2000 alarm observations from the environment for the pre-causal structure learning. We learn the initial causal structure leveraging the causal discovery method topological Hawkes process (THP) [59] that considers the topological information behind the event sequence. In policy learning, we take the SOTA model-free algorithms PPO [25], SAC [28], D3QN [24], and DQN [22] which are suitable for discrete cases as the baselines, and call the algorithms after applying our method Causal PPO, Causal SAC, Causal D3QN, and Causal DQN. For a fair




Figure 3 (Color online) (a) Cumulative rewards of Causal PPO, Causal SAC, Causal D3QN, Causal DQN with THP-initialized structures and random-initialized structures, respectively, and baselines; (b) intervention steps of our proposed approach compared to the baselines; (c) average number of alarms per episode for our methods compared to the baselines; (d) the F1 score of causal structure learning from different methods.

comparison, we use the same network structure, optimizer, learning rate, and batch size when comparing the native methods with our causal methods. We measure the performance of policy learning in terms of cumulative reward, number of interactions, and average number of alerts per episode. In causal structure learning, Recall, Precision, F1, Accuracy and SHD are used as the evaluation metrics. All results were averaged across four random seeds, with standard deviations shown in the shaded area.

5.3 Analysis of policy learning

To evaluate the performance of our methods, the cumulative reward, the number of interventions, and the average number of alarms are used as evaluation metrics. As shown in Figure 3(a), our methods significantly outperform the native algorithms after introducing our framework. It can be found that our algorithms only need to learn fewer rounds to reach higher cumulative rewards, which proves that the learned causal structure indeed helps to narrow the action space, and greatly speed up the convergence of the policy.

We also show the results of different algorithms on the number of intervention steps in Figure 3(b). Impressively, our method requires fewer interventions to eliminate all the environmental alarms and does not require excessive exploration in the training process compared to the baselines. This is very important in real-world O&M processes, because too many explorations may pose a huge risk. The above result also reflects that policies with causal structure learning capabilities have a more efficient and effective training process and sampling efficiency.


As shown in Figure 3(c), our method has a much smaller average number of alarms compared to the baselines. This indicates that our methods can detect root cause alarms in time, and thus avoid the cascade alarms generated from the environment. It is worth noting that the huge performance difference between our methods and baselines shows that the learned causal mechanisms of the environment play a pivotal role in RL.

5.4 Analysis of causal structure learning

To better demonstrate the effectiveness of our method, we only provide a small amount of observational data in the early causal structure learning. As shown in Table 1, the causal structure learned by THP in the initial stage has

Table 1 Results of causal structure learning. The best results are in bold.

Method	F1 score	Precision	Recall	Accuracy	SHD
THP	0.638 $\pm$ 0.017	0.775 $\pm$ 0.020	0.543 $\pm$ 0.015	0.824 $\pm$ 0.007	57.00 $\pm$ 2.121
Causal PPO (THP)	0.861 $\pm$ 0.018	0.865 $\pm$ 0.007	0.856 $\pm$ 0.029	0.921 $\pm$ 0.009	26.00 $\pm$ 2.915
Causal SAC (THP)	0.858 $\pm$ 0.013	0.871 $\pm$ 0.007	0.846 $\pm$ 0.024	0.919 $\pm$ 0.007	26.25 $\pm$ 2.165
Causal D3QN (THP)	0.836 $\pm$ 0.015	0.849 $\pm$ 0.021	0.823 $\pm$ 0.014	0.904 $\pm$ 0.009	31.00 $\pm$ 3.000
Causal DQN (THP)	0.832 $\pm$ 0.020	0.848 $\pm$ 0.020	0.817 $\pm$ 0.025	0.904 $\pm$ 0.013	31.00 $\pm$ 4.062
Random initiation	0.188 $\pm$ 0.013	0.130 $\pm$ 0.009	0.130 $\pm$ 0.009	0.669 $\pm$ 0.017	107.5 $\pm$ 5.362
Causal PPO (Random)	0.840 $\pm$ 0.019	0.847 $\pm$ 0.015	0.834 $\pm$ 0.025	0.909 $\pm$ 0.011	29.50 $\pm$ 3.640
Causal SAC (Random)	0.837 $\pm$ 0.019	0.864 $\pm$ 0.015	0.811 $\pm$ 0.022	0.908 $\pm$ 0.010	29.75 $\pm$ 3.269
Causal D3QN (Random)	0.839 $\pm$ 0.016	0.847 $\pm$ 0.022	0.832 $\pm$ 0.017	0.907 $\pm$ 0.011	30.25 $\pm$ 3.491
Causal DQN (Random)	0.830 $\pm$ 0.019	0.849 $\pm$ 0.025	0.813 $\pm$ 0.020	0.904 $\pm$ 0.013	31.25 $\pm$ 4.085

**Figure 4** (Color online) (a) Ground truth; discovered causal graphs by (b) Causal PPO, (c) Causal SAC, (d) Causal DQN, (e) Causal D3QN with THP-initialized causal structure. Rows and columns correspond to the variables in FaultAlarmRL environment, where a cell (i, j) indicates a causal influence from variable i to variable j .

a large distance from the ground truth. However, as we continue to interact with the environment, our methods gradually update the causal graph, bringing the learned causal structure closer to the ground truth. From Table 1 we can see that the F1 score values of our causal method are all over 0.8, which is significant compared to the initial THP result. The learned causal structures are given in Figure 4. We can see that the proposed method can indeed identify the correct structure and interestingly all the root cause variables are mostly identified due to the identification of the causal order. In order to verify the robustness of our causal graph updating mechanism, we also conducted experiments on the initial random graph. As shown in Table 1, even if the initial random graph is far from the ground truth, through continuous interactive updating, we can eventually learn a more accurate causal structure compared with the THP algorithm. In addition, as shown in Figure 3(d), our methods converge to the optimal value early in the pre-training period for the learning of causal structure, regardless of whether it is given a random graph or a prior graph, which indicates that a small amount of intervention up front is enough to learn the causal structure. Taking Causal PPO as an example, its F1 score has reached 0.7 after only 20 episodes. This shows that even in the case of a random initial causal structure, our method can still achieve a correct causal graph by calculating the treatment effects and performing the pruning step, which is more robust in the application.




Figure 5 (Color online) (a) Cumulative rewards, (b) intervention steps, and (c) average number of alarms per episode for Causal PPO based on THP initialization structures at different K values; (d) the F1 scores of causal structure learning based on Causal PPO with THP initialization structure for different K values.

5.5 Sensitivity analysis

The parameter K represents the number of potential root-cause errors considered in the causal order. We further conduct sensitivity experiments to evaluate the sensitivity of the hyperparameter K , which controls the TopK causal order in policy learning. We conduct a sensitivity analysis using Causal PPO as a case study. The results are given in Figures 5(a)–(d), which show the variations in the accuracy and robustness of policy learning and causal structure learning for different values of K . Specifically, when the K is too large (e.g., $K > 11$), the candidate action under the causal mask would also be large, increasing the redundancy of the action space which decreases the policy’s performance. Similarly, when the K value is small (e.g., $K < 5$), the policy’s performance worsens because the overly constrained action space may limit the exploration of optimal actions. Thus, the K controls the trade-off between the exploration and the exposition in our method.

6 Conclusion

This paper proposes an online causal reinforcement learning framework with a causal-aware policy that injects the causal structure into policy learning while devising a causal structure learning method by connecting the intervention and the action of the policy. We theoretically prove that our causal structure learning can identify the correct causal structure. To evaluate the performance of the proposed method, we constructed a FaultAlarmRL environment. Experiment results show that our method achieves accurate and robust causal structure learning as well as superior performance compared with SOTA baselines for policy learning.

Acknowledgements This work was supported in part by National Science and Technology Major Project (Grant No. 2021ZD01-11501), National Science Fund for Excellent Young Scholars (Grant No. 62122022), National Natural Science Foundation of China (Grant Nos. U24A20233, 62206064, 62206061, 62476163, 62406078, 62406080), and Guangdong Basic and Applied Basic Research Foundation (Grant No. 2023B1515120020).

Supporting information Appendixes A–E. The supporting information is available online at info.scichina.com and link.springer.com. The supporting materials are published as submitted, without typesetting or editing. The responsibility for scientific accuracy and content remains entirely with the authors.

References

- 1 Sutton R S, Barto A G. Reinforcement learning: an introduction. *Robotica*, 1999, 17:229–235
- 2 Kober J, Bagnell J A, Peters J. Reinforcement learning in robotics: a survey. *Int J Robot Res*, 2013, 32: 1238–1274
- 3 Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 2016, 529: 484–489
- 4 Shalev-Shwartz S, Shammah S, Shashua A. Safe, multi-agent, reinforcement learning for autonomous driving. 2016. ArXiv:161003295
- 5 Sun Y, Zhang K, Sun C. Model-based transfer reinforcement learning based on graphical model representations. *IEEE Trans Neural Netw Learn Syst*, 2023, 34: 1035–1048
- 6 Zhu Z M, Chen X H, Tian H L, et al. Offline reinforcement learning with causal structured world models. 2022. ArXiv:220601474
- 7 Sontakke S A, Mehrjou A, Itti L, et al. Causal curiosity: RL agents discovering self-supervised experiments for causal representation learning. In: *Proceedings of International Conference on Machine Learning*, 2021. 139: 9848–9858
- 8 Zhang A, McAllister R T, Calandra R, et al. Learning invariant representations for reinforcement learning without reconstruction. In: *Proceedings of the 9th International Conference on Learning Representations*, 2021
- 9 Tomar M, Zhang A, Calandra R, et al. Model-invariant state abstractions for model-based reinforcement learning. 2021. ArXiv:210209850
- 10 Bica I, Jarrett D, van der Schaar M. Invariant causal imitation learning for generalizable policies. In: *Proceedings of Advances in Neural Information Processing Systems*, 2021. 3952–3964
- 11 Sodhani S, Levine S, Zhang A. Improving generalization with approximate factored value functions. In: *Proceedings of ICLR2022 Workshop on the Elements of Reasoning: Objects, Structure and Causality*, 2022
- 12 Wang Z, Xiao X, Zhu Y, et al. Task-independent causal state abstraction. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems, Robot Learning workshop*. 2021
- 13 Ding W, Lin H, Li B, et al. Generalizing goal-conditioned reinforcement learning with variational causal reasoning. In: *Proceedings of Advances in Neural Information Processing Systems*, 2022. 26532–26548
- 14 Seitzer M, Schölkopf B, Martius G. Causal influence detection for improving efficiency in reinforcement learning. In: *Proceedings of Advances in Neural Information Processing Systems*, 2021. 34: 22905–22918
- 15 Huang B, Feng F, Lu C, et al. Adarl: what, where, and how to adapt in transfer reinforcement learning. In: *Proceedings of the 10th International Conference on Learning Representations*, 2022
- 16 Huang B, Lu C, Legi L, et al. Action-sufficient state representation learning for control with structural constraints. In: *Proceedings of International Conference on Machine Learning*, 2022. 9260–9279
- 17 Wang L, Yang Z, Wang Z. Provably efficient causal reinforcement learning with confounded observational data. In: *Proceedings of Advances in Neural Information Processing Systems*, 2021. 21164–21175
- 18 Liao L, Fu Z, Yang Z, et al. Instrumental variable value iteration for causal offline reinforcement learning. 2021. ArXiv:2102.09907
- 19 Volodin S, Wichers N, Nixon J. Resolving spurious correlations in causal models of environments via interventions. 2020. ArXiv:2002.05217
- 20 Zhang A, Lipton Z C, Pineda L, et al. Learning causal state representations of partially observable environments. 2019. ArXiv:1906.10437
- 21 Lee T E, Zhao J A, Sawhney A S, et al. Causal reasoning in simulation for structure and transfer learning of robot manipulation policies. In: *Proceedings of 2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 4776–4782
- 22 Mnih V, Kavukcuoglu K, Silver D, et al. Playing Atari with deep reinforcement learning. 2013. ArXiv:13125602
- 23 Lillicrap T P, Hunt J J, Pritzel A, et al. Continuous control with deep reinforcement learning. In: *Proceedings of the 4th International Conference on Learning Representations, San Juan*, 2016
- 24 Wang Z, Schaul T, Hessel M, et al. Dueling network architectures for deep reinforcement learning. In: *Proceedings of International Conference on Machine Learning*, 2016. 1995–2003
- 25 Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms. 2017. ArXiv:170706347
- 26 Schulman J, Levine S, Abbeel P, et al. Trust region policy optimization. In: *Proceedings of International Conference on Machine Learning*, 2015. 1889–1897
- 27 Mnih V, Badia A P, Mirza M, et al. Asynchronous methods for deep reinforcement learning. In: *Proceedings of the 33rd International Conference on Machine Learning*, 2016. 1928–1937
- 28 Haarnoja T, Zhou A, Abbeel P, et al. Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: *Proceedings of International Conference on Machine Learning*, 2018. 1861–1870
- 29 Kaiser L, Babaeizadeh M, Milos P, et al. Model-based reinforcement learning for Atari. 2019. ArXiv:190300374
- 30 Sutton R S. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bull*, 1991, 2: 160–163
- 31 Janner M, Fu J, Zhang M, et al. When to trust your model: model-based policy optimization. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019
- 32 García C E, Prett D M, Morari M. Model predictive control: theory and practice-a survey. *Automatica*, 1989, 25: 335–348
- 33 Luo F M, Xu T, Lai H, et al. A survey on model-based reinforcement learning. *Sci China Inf Sci*, 2024, 67: 121101
- 34 Zeng Y, Cai R, Sun F, et al. A survey on causal reinforcement learning. 2023. ArXiv:2302.05209
- 35 de Haan P, Jayaraman D, Levine S. Causal confusion in imitation learning. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019
- 36 Sonar A, Pacelli V, Majumdar A. Invariant policy optimization: towards stronger generalization in reinforcement learning. In: *Proceedings of the 3rd Annual Conference on Learning for Dynamics and Control*, 2021. 21–33

- 37 Lu C, Huang B, Wang K, et al. Sample-efficient reinforcement learning via counterfactual-based data augmentation. 2020. ArXiv:201209092
- 38 Pitis S, Creager E, Garg A. Counterfactual data augmentation using locally factored dynamics. In: Proceedings of Advances in Neural Information Processing Systems, 2020. 33: 3976–3990
- 39 Wang Z, Xiao X, Xu Z, et al. Causal dynamics learning for task-independent state abstraction. In: Proceedings of International Conference on Machine Learning, 2022. 23151–23180
- 40 Spirtes P, Glymour C, Scheines R. Causation, Prediction, and Search. Cambridge: MIT Press, 2001
- 41 Chickering D M. Optimal structure identification with greedy search. *J Mach Learn Res*, 2002, 3: 507–554
- 42 Ramsey J, Glymour M, Sanchez-Romero R, et al. A million variables and more: the Fast Greedy Equivalence Search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *Int J Data Sci Anal*, 2017, 3: 121–129
- 43 Huang B, Zhang K, Lin Y, et al. Generalized score functions for causal discovery. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018. 1551–1560
- 44 Shimizu S, Hoyer P O, Hyvärinen A, et al. A linear non-Gaussian acyclic model for causal discovery. *J Mach Learn Res*, 2006. 7: 2003–2030
- 45 Hoyer P, Janzing D, Mooij J M, et al. Nonlinear causal discovery with additive noise models. In: Proceedings of Advances in Neural Information Processing Systems, 2008
- 46 Peters J, Mooij J M, Janzing D, et al. Causal discovery with continuous additive noise models. *J Mach Learn Res*, 2014, 15: 2009–2053
- 47 Cai R, Qiao J, Zhang Z, et al. Self: structural equational likelihood framework for causal discovery. In: Proceedings of the AAAI Conference on Artificial Intelligence, 2018. 1787–1794
- 48 Brouillard P, Lachapelle S, Lacoste A, et al. Differentiable causal discovery from interventional data. In: Proceedings of Advances in Neural Information Processing Systems, 2020. 33: 21865–21877
- 49 Tigas P, Annadani Y, Jesson A, et al. Interventions, where and how? Experimental design for causal models at scale. In: Proceedings of Advances in Neural Information Processing Systems, 2022. 35: 24130–24143
- 50 Abel D. A theory of abstraction in reinforcement learning. 2022. ArXiv:2203.00397
- 51 Peters J, Janzing D, Schölkopf B. Elements of Causal Inference: Foundations and Learning Algorithms. Cambridge: MIT Press, 2017
- 52 Pearl J. Causality. Cambridge: Cambridge University Press, 2009
- 53 Rosenbaum P R, Rubin D B. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 1983, 70: 41–55
- 54 Lopez M J, Gutman R. Estimation of causal effects with multiple treatments: a review and new ideas. *Stat Sci*, 2017, 32: 432–454
- 55 Callaway B, Goodman-Bacon A, Sant’Anna P H. Difference-in-differences with a continuous treatment. Technical Report, National Bureau of Economic Research, 2024
- 56 Samsami M R, Bahari M, Salehkaleybar S, et al. Causal imitative model for autonomous driving. 2021. ArXiv:211203908
- 57 Athey S, Imbens G W, Wager S. Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *J R Stat Soc Ser B-Stat Methodol*, 2018, 80: 597–623
- 58 Zheng X, Aragam B, Ravikumar P K, et al. Dags with no tears: continuous optimization for structure learning. In: Proceedings of Advances in Neural Information Processing Systems, 2018. 31
- 59 Cai R, Wu S, Qiao J, et al. THPs: topological Hawkes processes for learning causal structure on event sequences. *IEEE Trans Neural Netw Learn Syst*, 2022, 35: 479–493