

Explainable Reinforcement Learning via a Causal World Model

Zhongwei Yu, Jingqing Ruan and Dengpeng Xing*

Institute of Automation, Chinese Academy of Sciences

{yuzhongwei2021, ruanjingqing2019, dengpeng.xing}@ia.ac.cn

Abstract

Generating explanations for reinforcement learning (RL) is challenging as actions may produce long-term effects on the future. In this paper, we develop a novel framework for explainable RL by learning a causal world model without prior knowledge of the causal structure of the environment. The model captures the influence of actions, allowing us to interpret the long-term effects of actions through causal chains, which present how actions influence environmental variables and finally lead to rewards. Different from most explanatory models which suffer from low accuracy, our model remains accurate while improving explainability, making it applicable in model-based learning. As a result, we demonstrate that our causal model can serve as the bridge between explainability and learning.

1 Introduction

Many real-world applications like finance and healthcare require AI systems to be well understood by users due to the demand for safety, security, and legality [Gunning *et al.*, 2019]. Aiming to help people better understand and work with AI systems, the field of Explainable AI (XAI) has recently attracted increasing interest from researchers. For example, a number of explanatory tools have been developed to pry into the black box of deep neural networks [Bach *et al.*, 2015; Selvaraju *et al.*, 2020; Wang *et al.*, 2021b].

However, the domain of explainable reinforcement learning (XRL) has been neglected for a long time. Many XRL studies adopt classic tools of XAI such as saliency maps [Nikulin *et al.*, 2019; Joo and Kim, 2019; Shi *et al.*, 2021]. These tools are not designed for sequential decision-making and are weak in interpreting the temporal dependencies of RL environments. Therefore, some studies investigate explaining specific components of the decision process, e.g., observations [Koul *et al.*, 2018; Raffin *et al.*, 2019], actions [Fukuchi *et al.*, 2017; Yau *et al.*, 2020], policies [Amir and Amir, 2018; Coppens *et al.*, 2019], and rewards [Juozapaitis *et al.*, 2019].

However, these studies rarely combine explanations with the dynamics of environments, which is important for understanding the long-term effects produced by agents' actions. In addition, real-world environments usually contain dynamics unknown to users, making it crucial to interpret these dynamics using explanatory models. Model-based RL (MBRL) uses predictive world models [Nagabandi *et al.*, 2018; Kaiser *et al.*, 2020; Janner *et al.*, 2021] to capture such dynamics. However, these models are usually densely-connected neural networks and cannot be used for the purpose of explanation.

Psychological research suggests that people explain the world through causality [Sloman, 2005]. In this paper, we propose a novel framework that uses an interpretable world model to generate explanations. Rather than using a dense and fully-connected model, we perform causal discovery to construct a sparse model that is aware of the causal relationships within the dynamics of environments. In order to explain agents' decisions, the proposed causal model allows us to construct causal chains which present the variables causally influenced by the agent's actions. The proposed model advances the existing work that uses causality for explainable RL [Madumal *et al.*, 2020b; Madumal *et al.*, 2020a], as it does not require a causal structure provided by domain experts, and is applicable to continuous action space.

Apart from interpreting the world, humans also use causality to guide their learning process [Cohen *et al.*, 2020]. However, the trade-off between interpretability and performance [Gunning *et al.*, 2019; Longo *et al.*, 2020; Puiutta and Veith, 2020] indicates that explainable models are usually inaccurate and can hardly benefit learning. On the contrary, our model is sufficiently accurate, leading to a performance close to dense models in MBRL. Therefore, we can train the agent and explain its decisions through exactly the same model, making explanations more faithful to the agent's intention. This is significant for overcoming the issue that post-hoc explanations like saliency maps can sometimes fail to faithfully unravel the decision-making process [Atrey *et al.*, 2020].

Our main contributions are as follows: 1) We learn a causal model that captures the environmental dynamics without prior knowledge of the causal structure. 2) We design a novel approach to effectively extract the causal influence of actions, allowing us to derive causal chains for explaining the agent's decisions. 3) We show that our explanatory model is accurate enough to guide policy learning in MBRL.

*Corresponding author.

¹The extended version of this paper with appendices is available at <https://arxiv.org/pdf/2305.02749.pdf>.

2 Background

2.1 Causality in Reinforcement Learning

RL integrated with causality has recently become noticed by RL researchers. For example, Lu [2018], Wang [2021a] and Yang [2022] *et al.* use causal inference to improve the robustness against confounders or adversarial intervention; Dietterich *et al.* [2018] improve learning efficiency by removing the variables unrelated to the agent’s action; Nair *et al.* [2019] construct a causal policy model; Seitzer *et al.* [2021] improve exploration by detecting the causal influence of actions; Wang [2022] and Ding [2022] *et al.* investigate causal world models as we do. However, they mainly use causality to improve generalization rather than generate explanations.

Only a few studies consider improving explainability using a causal model. Madumal *et al.* [2020b] propose the Action Influence Model (AIM), a causal model specialized for RL, to generate explanations about the agent’s actions. Another study [Madumal *et al.*, 2020a] further combines the AIM with decision trees to improve the quality of explanations. However, these approaches require finite action space, and the causal structure is given beforehand by human experts. In addition, they only consider a low-accuracy model, which cannot be used for policy learning. Volodin [2021] proposes a method to learn a sparse causal graph of hidden variables abstracted from high-dimensional observations, which improves the understanding of the environment. However, the approach provides no insight into the agent’s behavior.

2.2 Structural Causal Model

A *Structural Causal Model* (SCM) [Pearl *et al.*, 2016], denoted as a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{F})$, formalizes the causal relationships between multiple variables. \mathcal{U} and \mathcal{V} contain the variables of the SCM, where $\mathcal{U} = \{u_1, \dots, u_p\}$ is the set of *exogenous variables* and $\mathcal{V} = \{v_1, \dots, v_q\}$ is the set of *endogenous variables*. $\mathcal{F} = \{f^1, \dots, f^q\}$ is the set of *structural equations*, where f^j formulates how v_j is quantitatively decided by other variables. We call f^j a “structural” equation as it defines the subset of variables denoted as $PA(v_j) \subseteq \mathcal{U} \cup \mathcal{V} \setminus \{v_j\}$ (i.e., the parent variables) that directly decide the value of v_j .

An SCM is usually represented as a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{U} \cup \mathcal{V}, \mathcal{E})$ called the *causal graph*. The node set of \mathcal{G} is exactly $\mathcal{U} \cup \mathcal{V}$ and the edge set \mathcal{E} is given by the structural equations: $(x_i, v_j) \in \mathcal{E} \Leftrightarrow x_i \in PA(v_j)$, where $x_i \in \mathcal{U} \cup \mathcal{V}$. For simplicity, symbols like u_i , v_j , and x_i may denote the names or the values of the variables according to the context. In addition, we consider stochastic structural equations, where f^j outputs the posterior distribution of v_j conditioned on its direct parents $PA(v_j)$:

$$\Pr(v_j|PA(v_j)) \sim f^j(PA(v_j)). \quad (1)$$

2.3 Action Influence Model

An AIM, denoted as the tuple $(\mathcal{U}, \mathcal{V}, \mathcal{F}, A)$, is a causal model specialized for RL. Here, \mathcal{U} and \mathcal{V} follow the same definition in SCM. \mathcal{F} is the set of structural equations, and A is the action space of the agent. Different from SCM, each structural equation is related to not only an endogenous variable but also a unique action in A . In other words, there exists

a structural equation $f_a^j \in \mathcal{F}$ for any action $a \in A$ and any endogenous variable $v_j \in \mathcal{V}$ to describe how v_j is causally determined under action a . We use $PA_a(v_j) \subseteq \mathcal{U} \cup \mathcal{V} \setminus \{v_j\}$ to denote the causal parents of v_j under action a . Then, the posterior distribution of v_j under action a is given by

$$\Pr(v_j|PA_a(v_j), a) \sim f_a^j(PA_a(v_j)). \quad (2)$$

As a result, there exist overall $|A| \times |\mathcal{V}|$ causal equations in the AIM. In fact, we may also reckon the AIM as an ensemble of $|A|$ SCMs, where each SCM accounts for the influence of a unique action in A .

3 Factorized MDP

We are interested in tasks where the action and state can be factorized into multiple variables, and formalize such a task as a Factorized MDP (FMDP) denoted by the tuple $\langle S, A, O, R, P, T, \gamma \rangle$. Here, S , A , and O respectively denote the *state space*, *action space*, and *outcome space*. Each state $s \in S$ is factorized into n_s state variables such that $s = (s_1, \dots, s_{n_s})$, where s_i is the i -th state variable. Similarly, we have $a = (a_1, \dots, a_{n_a})$ for each action and $o = (o_1, \dots, o_{n_o})$ for each outcome. Figure 1(a) illustrates an example of the factorization for a simple 2-grid environment called the Vacuum world [Russell *et al.*, 2010]. Its state variables include the *position* of the vacuum and whether the places are clean ($clean_1$ and $clean_2$); it contains only one action variable a that is chosen from Left, Right, and Suck (making the place clean). The action of Left or Right leads to an outcome of failure when blocked by the world boundary.

On each step, the agent observes the current state s and takes an action a , then the state transits and the outcome is produced according to the transition probability $P(o', s'|s, a)$, leading to a *transition tuple* denoted as $\delta = (s, a, o, s')$. Meanwhile, the reward is given by the overall reward function $R(\delta)$. Following the reward decomposition [Juozapaitis *et al.*, 2019], we factorize R as the summation of n_r *reward variables*, given by $R(\delta) = \sum_{i=1}^{n_r} r_i(\delta)$. γ is the discount factor for computing returns. T is the termination condition deciding whether the episode terminates based on the transition δ . How reward variables $\{r_i\}_{i=1}^{n_r}$ and the termination condition T depend on the transition δ is defined by users according to their demands. In those cases where R and T contain components unknown to users, we may put these components into the outcome variables. That is, we use an outcome variable o_i to indicate the unknown reward, and define the corresponding reward variable as $r_i \equiv o_i$.

4 The Proposed Framework

There exist two perspectives on the causal model of the dynamics of the environment: 1) In a unified SCM, action variables are merely nodes of the causal structure and are treated evenly as state variables. 2) In an AIM, each action specifies a peculiar causal structure, leading to a good understanding of both the environment and the agent’s decisions [Madumal *et al.*, 2020b]. To make it clearer, Figure 1 illustrates how a unified SCM and an AIM pertain to our setting in the above-mentioned Vacuum world. However, directly learning an AIM is intractable as we must divide data into $|A|$ subsets to respectively learn $|A|$ SCMs.




Figure 1: The illustration of causal models in the Vacuum world. (a) illustrates the Vacuum world, where $position = 1$, $clean_1 = True$, and $clean_2 = False$. (b) and (c) respectively illustrate the causal graphs of the SCM and the AIM of the Vacuum world.

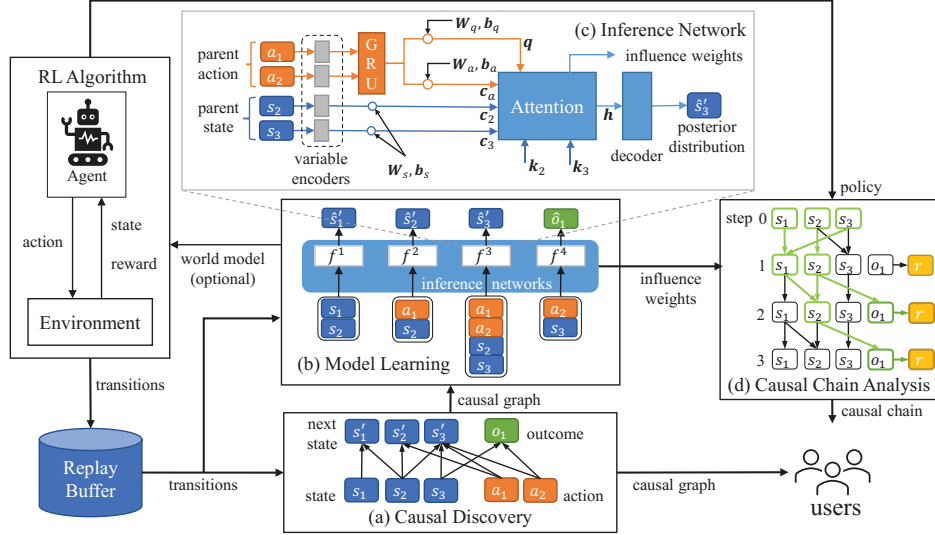


Figure 2: The illustration of the proposed framework. (a) shows an example of the causal graph identified by causal discovery. (b) illustrates the structure of the proposed model. (c) shows the inference network that approximates the structural equation of s'_3 . (d) illustrates the causal chain analysis, where the causal chain is highlighted in bold and green.

This reduces sample efficiency, produces redundant parameters, and cannot be applied to infinite action space. Therefore, we seek to build a unified SCM that can be converted to an AIM based on specially-designed structural equations. As illustrated in Figure 2(a), the exogenous variables in this SCM are the state variables s and action variables a ; the endogenous variables are the next-state variables s' and outcome variables o . In the following description, we define $\mathbf{u} := (s, a) = (s_1, \dots, s_{n_s}, a_1, \dots, a_{n_a})$ as the input (exogenous) variables and $\mathbf{v} := (s', o) = (s'_1, \dots, s'_{n_s}, o_1, \dots, o_{n_o})$ as the output (endogenous) variables of our model. In this way, a transition tuple can also be written as $\delta = (\mathbf{u}, \mathbf{v})$.

The workflow of the proposed framework is illustrated in Figure 2. When the agent interacts with the environment, we store the transition data into a replay buffer \mathcal{D} . Using the stored transitions, we perform *causal discovery* to identify the causal graph of the above-mentioned SCM. Then, we fit the causal equations for the variables of the next state s' and the outcome o using *Inference Networks*. Together with the known reward function R and termination condition T , we construct a causal world model that captures the dynamics of the environment. The attention weights (influence weights) in the inference networks capture the action influence, which

allows us to perform *causal chain analysis* to reveal the variables that are causally influenced by the agent's action. The discovered causal graph interprets the environmental dynamics, and the causal chain analysis provides explanations about the agent's decisions. Moreover, this causal world model can be used by MBRL algorithms to facilitate learning.

4.1 Causal Discovery

We assume that output variables v are produced independently conditioned on the input variables u , as independence underlies the intuition of humans to segment the world into components. Under this assumption, it is proven that the causal graph is bipartite, where no lateral edge exists in v , and each edge starts in u and ends in v . Therefore, we only need to determine whether there exists a causal edge for each variable pair (u_i, v_j) , where $1 \leq i \leq n_s + n_a$ and $1 \leq j \leq n_s + n_o$. Studies have shown that *conditional independent tests* (CITs) can be used to perform efficient causal discovery [Wang *et al.*, 2022; Ding *et al.*, 2022]. In this work, we implement CITs through Fast CIT [Chalupka *et al.*, 2018] and determine each edge using the following rule:

$$u_i \in PA(v_j) \iff (u_i \perp\!\!\!\perp v_j | \mathbf{u}_{-i}), \quad (3)$$

where \mathbf{u}_{-i} denotes all variables in \mathbf{u} other than u_i . In Appendix A, we provide the theoretical basics of causal discovery and prove a theorem showing that Equation 3 leads to sound causal graphs.

4.2 Attention-based Inference Networks

To perform causal inference on the discovered causal graph, we fit the structural equation of each output variable v_j using an inference network denoted as f^j , which takes the causal parents $PA(v_j)$ as inputs and predicts the posterior distribution $\Pr(v_j|PA(v_j))$. These inference networks should adapt to the structural changes of the causal graph, as the agent’s exploratory behaviors may reveal undiscovered causal relationships and lead to new causal structures. To achieve this, Ding *et al.* [2022] use Gated Recurrent Unit (GRU) networks that sequentially input all parent variables without discriminating the state and the action. To model the action influence, we design the attention-based inference networks as illustrated in Figure 2(c).

To handle heterogenous input variables (which may be scalars or vectors of different lengths), we first use *variable encoders* (each uniquely belongs to an input variable) to individually map the input variables to vectors of the same length. These encoders are shared by the inference networks of all output variables. In particular, we use $\tilde{\mathbf{u}} = \{\tilde{s}_1, \dots, \tilde{s}_n, \tilde{a}_1, \dots, \tilde{a}_m\}$ to denote these encoding vectors of input variables.

Then, for each inference network f^j , we compute the *contribution vectors* of the parent state variables through linear transforms:

$$\mathbf{c}_i^j = \mathbf{W}_s^j \tilde{s}_i + \mathbf{b}_s^j, \quad s_i \in Pa(v_j). \quad (4)$$

The usage of these contribution vectors is equal to the “value vectors” in key-value attention. We use the term “contribution vectors” since the word “value” is ambiguous in the context of RL.

Each inference network f^j contains a GRU network g^j , which receives the action variables in $PA(v_j)$ and outputs the action embedding \mathbf{e}^j . Then, we feed this action embedding into linear transforms to respectively obtain the query vector \mathbf{q}^j and the action contribution vector \mathbf{c}_a^j :

$$\mathbf{e}^j = GRU^j(\{\tilde{a}_i\}_{a_i \in PA(v_j)}), \quad (5)$$

$$\mathbf{q}^j = \mathbf{W}_q^j \mathbf{e}^j + \mathbf{b}_q^j, \quad (6)$$

$$\mathbf{c}_a^j = \mathbf{W}_a^j \mathbf{e}^j + \mathbf{b}_a^j. \quad (7)$$

The projection matrices $\mathbf{W}_s^j, \mathbf{W}_q^j, \mathbf{W}_a^j$ and bias vectors $\mathbf{b}_s^j, \mathbf{b}_q^j, \mathbf{b}_a^j$ are all trainable parameters of f^j . We use the superscript j to indicate that these parameters belong to f^j .

Each state variable s_i is allocated a key vector \mathbf{k}_i , which is a trainable parameter learned by gradient descent. We do not use the superscript j for key vectors as they are shared by inference networks of all output variables. The *influence weights* (i.e., attention weights) of the state variables in $PA(v_j)$ and the action are then computed by

$$\alpha_i^j = \frac{\exp(\mathbf{k}_i^T \mathbf{q}^j)}{1 + \sum_{s'_i \in PA(v_j)} \exp(\mathbf{k}_{s'_i}^T \mathbf{q}^j)}, \quad (8)$$

$$\alpha_a^j = \frac{1}{1 + \sum_{s'_i \in PA(v_j)} \exp(\mathbf{k}_{s'_i}^T \mathbf{q}^j)}. \quad (9)$$

We then compute the hidden representation of the posterior distribution using the weighted sum of the value vectors:

$$\mathbf{h}^j = \sum_{s_i \in PA(v_j)} \alpha_i^j \cdot \mathbf{c}_i^j + \alpha_a^j \cdot \mathbf{c}_a^j. \quad (10)$$

Finally, the *distribution decoder* D^j maps \mathbf{h}^j to the predicted posterior distribution:

$$\Pr(v_j|PA(v_j)) \sim D^j(\mathbf{h}^j). \quad (11)$$

We assume the type of this posterior distribution is previously known and D^j only outputs the parameters of the distribution. In our implementation, we use normal distribution (parameterized by the mean and variance) for real-number variables and use categorical distribution (parameterized by the probability of each class) for discrete variables.

The inference networks $\{f_j\}_{j=1}^{n_s+n_o}$ are trained by maximizing the log-likelihood of the transition data stored in \mathcal{D} , written as

$$\mathcal{L}_{infer} = \sum_{j=1}^{n_s+n_o} \frac{1}{|\mathcal{D}|} \sum_{\delta \in \mathcal{D}} \log \Pr(v_j|PA(v_j)). \quad (12)$$

4.3 Causal Chain Analysis

In order to generate explanations, we first describe how our model can be converted to an AIM. Noticing that the key vectors $\{\mathbf{k}_i\}_{i=1}^n$ are trainable parameters, the influence weights only depend on the numeric value of action variables. Therefore, the influence weight α_i^j captures how much the output variable v_j depends on state variable s_i under the given action \mathbf{a} . In order to generate laconic explanations, we define $PA_{\mathbf{a}}(v_j) := \{s_i \in PA(v_j) \mid \alpha_i^j > \tau\}$ as the parent set of v_j with salient dependencies under the action \mathbf{a} , where $\tau \in [0, 1]$ is a given threshold. In this way, we convert the SCM to an AIM, where the structural equation for v_j under action \mathbf{a} is written as

$$f_{\mathbf{a}}^j(PA_{\mathbf{a}}(v_j)) = D^j\left(\sum_{s_i \in PA_{\mathbf{a}}(v_j)} \alpha_i^j \cdot \mathbf{c}_i^j(s_i) + \alpha_a^j \cdot \mathbf{c}_a^j\right). \quad (13)$$

Since we use the AIM for the purpose of explanation, it is tolerable to set a larger threshold, which allows us to ignore parent variables that are not influential enough. Madumal *et al.* have introduced methods to generate good explanations using an AIM. The key is to build a causal chain containing the variables that (i) are causally affected by the actions, and (ii) causally lead to rewards. A single causal chain starting from state s and action \mathbf{a} leads to the explanation for “why the agent took \mathbf{a} on s ”. Contrastive explanations for “why the agent took \mathbf{a} instead of \mathbf{b} on s ” can be obtained by comparing the causal chains produced by the factual action \mathbf{a} and the counterfactual action \mathbf{b} . Details can be found in Appendix B and the AIM paper [Madumal *et al.*, 2020a].

The rest of this section introduces how to derive a causal chain starting from the state s^t at step t and an action \mathbf{a}^t (can be factual or counterfactual) using our model. First of all, we use our model and the agent’s policy to simulate the

most-likely trajectory $\delta^t, \delta^{t+1}, \dots, \delta^{t+H-1}$, where H denotes the number of simulation steps. The symbol δ^{t+k} denotes the transition tuple on step $t+k$, where the actions \mathbf{a}^{t+k} for $k \geq 1$ are produced by the agent’s policy. For a factual causal chain, this simulation is not necessary if factual data of these future states and actions is available.

Then, we build an extended graph containing the state, outcome, and reward variables of these H steps. The edges of this graph accord to the structure of the AIM derived above. That is, if $s_i \in PA_{\mathbf{a}^{t+k}}(v_j)$ then there exists an edge from s_i^{t+k} to v_j^{t+k} for all $k = 1, \dots, H$. It is worth mentioning that we treat the first transition δ^t differently since \mathbf{a}^t is exactly the action being explained: If $PA(v_j) \cap \mathbf{a}^t = \emptyset$, then v_j^t is not affected by the choice of \mathbf{a}^t . In this case, no edge will be established from any state s_i^t to v_j^t .

Afterward, the explainee may specify the target variables (a subset of reward variables) he/she is interested in. Otherwise, all reward variables will be considered. We perform a graph search from the starting state variables \mathbf{s}^t and highlight all paths from \mathbf{s}^t to the target rewards. These paths together form the *causal chain* of action \mathbf{a}^t starting from \mathbf{s}^t . Based on this causal chain, the explanation can be presented as a picture or a natural-language description.

4.4 Model-based RL

XAI literature has widely discussed the trade-off between interpretability and performance, which is also reflected in our model. A sparser causal graph (discovered using a smaller threshold η) is usually easier to read and produces clearer explanations. However, it also enforces the model to infer posterior distributions using less information from input variables \mathbf{u} , leading to inferior accuracy. In Appendix C, we provide a theorem that formally shows that decreasing the threshold η leads to a denser causal graph (i.e., lower interpretability) and also higher predicting accuracy.

In order to show our model is accurate enough to do more than generate post-hoc explanations, we consider applying our world model to MBRL to facilitate policy learning. We use a bootstrap ensemble containing 5 models to alleviate the effect of the epistemic uncertainty [Chua *et al.*, 2018]. For each iteration, we first collect real transition data into the model buffer \mathcal{D} . Then, we update the world model by causal discovery and fitting structural equations using the data in \mathcal{D} . Afterward, we perform k -step model-rollouts [Janner *et al.*, 2021] to generate simulated data for updating the policy. In our implementation, the policy is trained using Proximal Policy Optimization [Schulman *et al.*, 2017]. The pseudo-code of the learning procedure is given in Appendix D.

5 Experiments

We present examples of causal chains in two representative environments: Lunarlander-Continuous for the continuous action space, and Build-Maine for the discrete action space. To verify whether our approach can produce correct causal chains, we design an environment to measure the accuracy of recovering causal dependencies of the ground-truth AIM. To evaluate the performance of our model in MBRL,

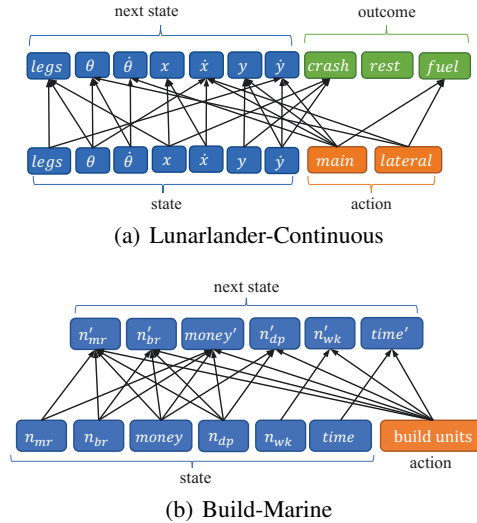


Figure 3: The discovered causal graphs of two environments.

we perform experiments in two extra environments: Cartpole and Lunarlander-Discrete. The Build-Marine environment is adapted from one of the StartCraftII mini-games in SC2LE [Samvelyan *et al.*, 2019]; the Cartpole and Lunarlander environments are classic control problems provided by OpenAI Gym [Brockman *et al.*, 2016]. Our source code is available at <https://github.com/EaseOnWay/Explainable-Causal-Reinforcement-Learning>.

5.1 Explanation Results

Lunarlander-Continuous

We factorize the state into 7 variables ($x, y, \dot{x}, \dot{y}, \theta, \dot{\theta}, legs$) indicating 1) the horizontal position, 2) the vertical position, 3) the horizontal velocity, 4) the vertical velocity, 5) the angle, 6) the angle velocity, and 7) whether the two legs are in contact with the ground. The action includes 2 continuous variables ranged in $(-1, 1)$, respectively controlling the throttles of the main and the lateral engines. The environment contains 3 outcome variables, including 1) the fuel cost due to firing the engine, 2) whether the lander crashed, and 3) whether the rocket is resting.

In this experiment, we learn a post-hoc model to generate explanations for a previously trained policy. We first use the policy (with noise) to collect 150k samples into the buffer \mathcal{D} . Then, we use these samples to discover the causal graph (with the threshold $\eta = 0.05$) and train the inference networks. The resulting causal graph is presented in Figure 3(a). The environment contains many kinds of rewards, leading to complicated causal chains if we consider them all. To make our explanation clearer, we present a causal chain in Figure 4 considering only two kinds of rewards: 1) the reduction of the distance to the target location, and 2) the reduction of the angle (i.e., balancing the rocket). This causal chain shows that the agent’s action \mathbf{a}^t first influences the velocities ($\dot{\theta}, \dot{x}$, and \dot{y}) and thereby reduces (or increases) the angle (θ) and the distance ($\sqrt{x^2 + y^2}$). In addition, we observe that no parent of the outcome variable *rest* is discovered in the causal graph.

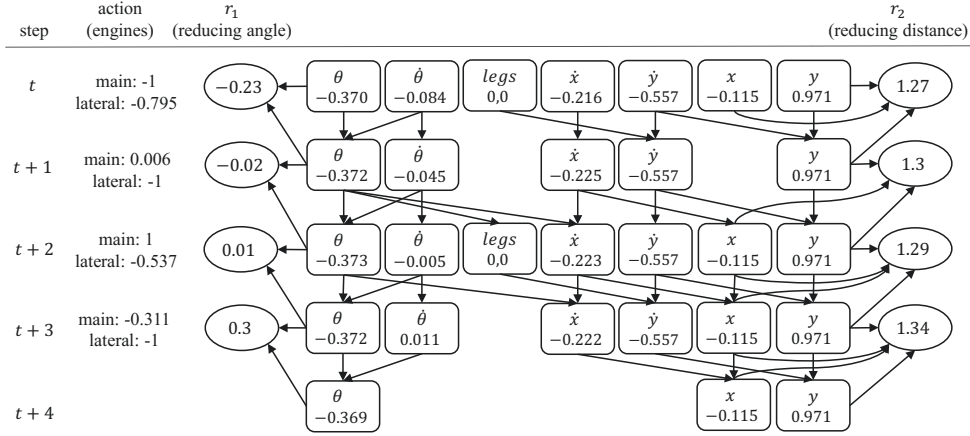


Figure 4: An example of a 4-step causal chain on Lunarlander-Continuous

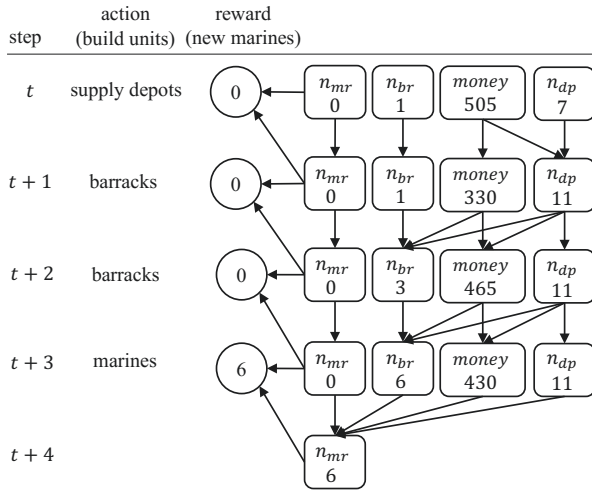


Figure 5: An example of a 4-step causal chain on Build-Marine

This means the policy provides insufficient opportunities to reveal its causality. As a result, the variable *rest* is excluded from all causal chains, showing that it is not an important consideration in the agent’s policy.

Build-Marine

The original observation space provided by the SC2LE interface contains hundreds of variables, which is intractable for causal discovery. In our implementation, we define the state as the tuple containing only 6 variables denoted as $(n_{wk}, n_{mr}, n_{br}, n_{dp}, money, time)$, namely 1) the number of workers, 2) the number of marines, 3) the number of barracks, 4) the number of supply depots, 5) the amount of money, and 6) the game time. We are interested in the macro-level decision-making and therefore define the action as one discrete variable indicating which unit (workers, marines, barracks, supply depots, or none) to be built. The micro-level control of building these units (e.g., determining where to place the new barracks) is implemented by simple rules. The goal is to produce as many marines as possible within 15 minutes. Therefore, the player is rewarded with 1 for every

	Direct	Full+Attn	Caus+Attn (ours)
AIM accuracy	97.0 %	90.0%	99.3%

Table 1: The accuracy of recovering the causal dependencies of the AIM. “Direct” means the direct approach mentioned in Section 5.2; “Full” means using a full graph; “Caus” means using a causal graph; and “Attn” means using attention.

newly produced marine. In addition, this environment contains no outcome variable.

In this experiment, both the policy and the causal model are obtained by model-based learning (see Section 4.4). We present the final causal graph in Figure 3(b) (discovered using the threshold $\eta = 0.15$) and an example of the causal chain in Figure 5. The causal chain shows that our attention-based inference networks successfully reason the causal dependencies under different actions, which reflect the following rules of the StarCraftII game: 1) Building new barracks requires at least one supply depot; 2) marines are built from barracks; 3) the number of marines is limited by the number of supply depots; and 4) building more units requires sufficient money in hand. This causal chain explains the reason why the agent builds supply depots: to gain permission to build barracks and provide enough supplies for building marines. Interestingly, we discover no causal relationship between n_{wk} and $money'$. For human players, it is common sense that more workers increase the efficiency of collecting minerals and thus lead to a higher income. Since the causal model is learned using the transition data produced along with policy training, this missing edge indicates that the agent explored inadequately for building more workers, providing insufficient evidence to reveal this causal relation.

5.2 Accuracy of Recovering Action Influence

Good explanations are generated from correct causal chains, which require us to accurately recover the AIMs of environments. In Section 4, we have mentioned a *Direct* approach that learns the AIM by splitting the data buffer \mathcal{D} into $|\mathcal{A}|$ sub-buffers and performing causal discovery for each action $a \in \mathcal{A}$. Though this direct approach is theoretically sound, it

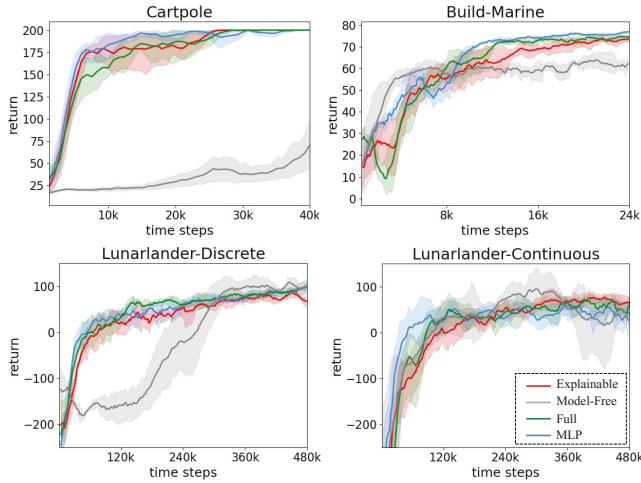


Figure 6: The training curves. Our explainable model (red) is compared to non-explainable dense models (green and blue) to show the performance cost of using a sparse causal graph. The grey curves show the performance without models.

suffers from poor sample efficiency and high computational complexity. Noticing that the causal dependencies under different actions usually share similar structures, our approach takes 2 stages: 1) In the **causal stage**, we learn a unified SCM, whose causal graph summarizes causal dependencies for all actions; 2) in the **attention stage**, we then transfer this SCM into an AIM based on attention weights (influence weights).

To verify whether our approach can accurately recover the AIM, we design a simple environment that contains spurious correlations to confuse neural networks (see Appendix E.2 for details). We compare our approach with two baselines: 1) the Direct approach mentioned above, and 2) a non-causal approach that uses a full causal graph and only relies on attention. The accuracy of recovering the ground-truth causal dependencies of the AIM using non-i.i.d data is shown in Table 1. These results show that: 1) the causal graph discovered in the causal stage precludes most spurious correlations, making our approach more effective than the Direct approach in practice; 2) and attention alone is insufficient to accurately extract the causal influence of actions.

Further, we examined the causal chains derived solely from attention (where full causal graphs are used). In these chains, we found plenty of spurious correlations, which lead to unreasonable explanations (e.g., “the number of supply depots naturally grows with time” in Build-Marine). An example of such a causal chain and the related discussion are provided in Appendix F. This result shows that causal discovery is an indispensable process for producing reasonable explanations.

5.3 Performance in Model-Based RL

We evaluate the performance of model-based policy learning using our explanatory model. We compare the learning performance with the **Model-Free** approach that learns policies without models. In addition, we consider two dense models as baselines: 1) the model that concatenates all exogenous variables u as inputs and infers endogenous variables using a Multi-Layer Perceptron (MLP), and 2) the **Full** model that

adopts the same networks as ours whereas uses a full causal graph.

The MLP model is the most commonly used in MBRL, and the Full model follows the state-of-the-art modular architecture [Ke *et al.*, 2021]. These dense models are not suitable for generating explanations. However, they are more accurate (if well-trained) than our explainable model as they are allowed to predict each output variable based on the complete inputs. Existing studies show that causal models generalize better than dense models [Wang *et al.*, 2022; Ding *et al.*, 2022]. However, we focus on ordinary learning problems and do not consider using our model for generalization. Therefore, we stress that the goal of this experiment is not to obtain a higher performance than dense baselines. On the contrary, We aim to figure out: 1) whether the proposed model can be of help to the learning process, and 2) how much performance our sparse model sacrifices for improved explainability.

The learning curves are shown in Figure 6. In all environments, the performance of our explanatory model is very close to the dense baselines. Compared to the model-free approach, the model-based approaches significantly learn faster in Cartpole and Lunarlander-Discrete and converge to higher returns in Build-Marine. The returns of all tested approaches are close in Lunarlander-Continuous, whereas the model-based approaches improve the stability of learning. These results show that our model improves explainability at an acceptable cost in performance and well balances the interpretability-accuracy trade-off. Therefore, our model can simultaneously guide policy learning and explain decisions, leading to better consistency between explanations and the agent’s cognition of the environment.

6 Conclusion and Future Work

This paper proposes a framework that learns a causal world model to generate explanations about agents’ actions. To achieve this, we perform causal discovery to identify the causal structure in the environment and fit causal equations through attention-based inference networks. These inference networks produce the influence weights that capture the influence of actions, which allow us to perform causal chain analysis in order to generate explanations. The proposed framework does not require the structural knowledge provided by human experts and is applicable to infinite action spaces. Apart from generating explanations, we successfully applied our model to model-based RL, showing that the model can be the bridge between learning and explainability.

A weakness of our approach is that it requires a known factorization of the environment, which limited its application scope. There exists a number of studies aiming to learn the causal feature set from raw observations [Zhang *et al.*, 2020; Volodin, 2021; Zhang *et al.*, 2021]. Future work will put representation learning into consideration for better applicability. In addition, we currently consider model-based policy learning as the usage of our model apart from generating explanations. However, this usage does not make full use of the advantage of a causal model. Future work will investigate better usage of our model to further improve learning.

Acknowledgments

This work was supported in part by National Key R&D Program of China (No.2022ZD0116405) and in part by the National Nature Science Foundation of China under Grant (62073324).

References

- [Amir and Amir, 2018] Dan Amir and Ofra Amir. HIGHLIGHTS: Summarizing Agent Behavior to People. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, Stockholm, Sweden, 2018.
- [Atrey *et al.*, 2020] Akanksha Atrey, Kaleigh Clary, and David Jensen. Exploratory Not Explanatory: Counterfactual Analysis of Saliency Maps for Deep Reinforcement Learning, February 2020. arXiv:1912.05743 [cs].
- [Bach *et al.*, 2015] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLOS ONE*, 10(7):e0130140, July 2015.
- [Brockman *et al.*, 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, June 2016. arXiv:1606.01540 [cs].
- [Chalupka *et al.*, 2018] Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Fast Conditional Independence Test for Vector Variables with Large Sample Sizes, April 2018. arXiv:1804.02747 [cs, stat].
- [Chua *et al.*, 2018] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models, November 2018. arXiv:1805.12114 [cs, stat].
- [Cohen *et al.*, 2020] Alexandra O. Cohen, Kate Nussenbaum, Hayley M. Dorfman, Samuel J. Gershman, and Catherine A. Hartley. The rational use of causal inference to guide reinforcement learning strengthens with age. *npj Science of Learning*, 5(1):16, December 2020.
- [Coppens *et al.*, 2019] Youri Coppens, Kyriakos Efthymiadis, Tom Lenaerts, and Ann Nowe. Distilling Deep Reinforcement Learning Policies in Soft Decision Trees. In *Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence*, pages 1–6, 2019.
- [Dietterich *et al.*, 2018] Thomas G. Dietterich, George Trimpontias, and Zhitang Chen. Discovering and Removing Exogenous State Variables and Rewards for Reinforcement Learning, June 2018. arXiv:1806.01584 [cs, stat].
- [Ding *et al.*, 2022] Wenhao Ding, Haohong Lin, Bo Li, and Ding Zhao. Generalizing Goal-Conditioned Reinforcement Learning with Variational Causal Reasoning, July 2022. arXiv:2207.09081 [cs, stat].
- [Fukuchi *et al.*, 2017] Yosuke Fukuchi, Masahiko Osawa, Hiroshi Yamakawa, and Michita Imai. Autonomous Self-Explanation of Behavior for Interactive Reinforcement Learning Agents. In *Proceedings of the 5th International Conference on Human Agent Interaction*, pages 97–101, Bielefeld Germany, October 2017. ACM.
- [Gunning *et al.*, 2019] David Gunning, Mark Stefik, Jaesik Choi, Timothy Miller, Simone Stumpf, and Guangzhong Yang. XAI—Explainable artificial intelligence. *SCIENCE ROBOTICS*, 4(7120):2, 2019.
- [Janner *et al.*, 2021] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to Trust Your Model: Model-Based Policy Optimization, November 2021. arXiv:1906.08253 [cs, stat].
- [Joo and Kim, 2019] Ho-Taek Joo and Kyung-Joong Kim. Visualization of Deep Reinforcement Learning using Grad-CAM: How AI Plays Atari Games? In *2019 IEEE Conference on Games (CoG)*, pages 1–2, 2019. ISSN: 2325-4289.
- [Juozapaitis *et al.*, 2019] Zoe Juozapaitis, Anurag Koul, Alan Fern, Martin Erwig, and Finale Doshi-Velez. Explainable Reinforcement Learning via Reward Decomposition. In *Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence*, pages 1–7, 2019.
- [Kaiser *et al.*, 2020] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H. Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Afroz Mohiuddin, Ryan Sepassi, George Tucker, and Henryk Michalewski. Model-Based Reinforcement Learning for Atari, February 2020. arXiv:1903.00374 [cs, stat].
- [Ke *et al.*, 2021] Nan Rosemary Ke, Aniket Didolkar, Sarthak Mittal, Anirudh Goyal, Guillaume Lajoie, Stefan Bauer, Danilo Rezende, Yoshua Bengio, Michael Mozer, and Christopher Pal. Systematic Evaluation of Causal Discovery in Visual Model Based Reinforcement Learning, July 2021. arXiv:2107.00848 [cs, stat].
- [Koul *et al.*, 2018] Anurag Koul, Sam Greysdanus, and Alan Fern. Learning Finite State Representations of Recurrent Policy Networks, November 2018. arXiv:1811.12530 [cs, stat].
- [Longo *et al.*, 2020] Luca Longo, Randy Goebel, Freddy Lecue, Peter Kieseberg, and Andreas Holzinger. Explainable Artificial Intelligence: Concepts, Applications, Research Challenges and Visions. In *Machine Learning and Knowledge Extraction*, volume 12279, pages 1–16. Springer International Publishing, Cham, 2020. Series Title: Lecture Notes in Computer Science.
- [Lu *et al.*, 2018] Chaochao Lu, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Deconfounding Reinforcement Learning in Observational Settings, December 2018. arXiv:1812.10576 [cs, stat].
- [Madumal *et al.*, 2020a] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. Distal Explanations for Model-free Explainable Reinforcement Learning, September 2020. arXiv:2001.10284 [cs].

- [Madumal *et al.*, 2020b] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. Explainable Reinforcement Learning through a Causal Lens. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2493–2500, April 2020.
- [Nagabandi *et al.*, 2018] Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7559–7566, 2018.
- [Nair *et al.*, 2019] Suraj Nair, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Causal Induction from Visual Observations for Goal Directed Tasks, October 2019. arXiv:1910.01751 [cs, stat].
- [Nikulin *et al.*, 2019] Dmitry Nikulin, Anastasia Ianina, Vladimir Aliev, and Sergey Nikolenko. Free-Lunch Saliency via Attention in Atari Agents. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 4240–4249, Seoul, Korea (South), October 2019. IEEE.
- [Pearl *et al.*, 2016] Judea Pearl, Madelyn Glymour, and Nicholas P. Jewell. *Causal inference in statistics: a primer*. Wiley, Chichester, West Sussex, 2016.
- [Puiutta and Veith, 2020] Erika Puiutta and Eric MSP Veith. Explainable Reinforcement Learning: A Survey, May 2020. arXiv:2005.06247 [cs, stat].
- [Raffin *et al.*, 2019] Antonin Raffin, Ashley Hill, René Traoré, Timothée Lesort, Natalia Díaz-Rodríguez, and David Filliat. Decoupling feature extraction from policy learning: assessing benefits of state representation learning in goal based robotics, June 2019. arXiv:1901.08651 [cs, stat].
- [Russell *et al.*, 2010] Stuart J. Russell, Peter Norvig, and Ernest Davis. *Artificial intelligence: a modern approach*. Prentice Hall series in artificial intelligence. Prentice Hall, Upper Saddle River, 3rd ed edition, 2010.
- [Samvelyan *et al.*, 2019] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob Foerster, and Shimon Whiteson. The StarCraft Multi-Agent Challenge, December 2019. arXiv:1902.04043 [cs, stat].
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, August 2017. arXiv:1707.06347 [cs].
- [Seitzer *et al.*, 2021] Maximilian Seitzer, Bernhard Schölkopf, and Georg Martius. Causal Influence Detection for Improving Efficiency in Reinforcement Learning, December 2021. arXiv:2106.03443 [cs].
- [Selvaraju *et al.*, 2020] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. *International Journal of Computer Vision*, 128(2):336–359, February 2020. arXiv:1610.02391 [cs].
- [Shi *et al.*, 2021] Wenjie Shi, Gao Huang, Shiji Song, Zhuoyuan Wang, Tingyu Lin, and Cheng Wu. Self-Supervised Discovering of Interpretable Features for Reinforcement Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021.
- [Sloman, 2005] Steven Sloman. *Causal Models: How People Think about the World and Its Alternatives*. Oxford University Press, 08 2005.
- [Volodin, 2021] Sergei Volodin. *CauseOccam : Learning Interpretable Abstract Representations in Reinforcement Learning Environments via Model Sparsity*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2021.
- [Wang *et al.*, 2021a] Lingxiao Wang, Zhuoran Yang, and Zhaoran Wang. Provably Efficient Causal Reinforcement Learning with Confounded Observational Data. In *Advances in Neural Information Processing Systems*, volume 34, pages 21164–21175. Curran Associates, Inc., 2021.
- [Wang *et al.*, 2021b] Rui Wang, Xiaoqian Wang, and David I. Inouye. Shapley Explanation Networks, April 2021. arXiv:2104.02297 [cs].
- [Wang *et al.*, 2022] Zizhao Wang, Xuesu Xiao, Zifan Xu, Yuke Zhu, and Peter Stone. Causal Dynamics Learning for Task-Independent State Abstraction, June 2022. arXiv:2206.13452 [cs].
- [Yang *et al.*, 2022] Chao-Han Huck Yang, I-Te Danny Hung, Yi Ouyang, and Pin-Yu Chen. Training a Resilient Q-Network against Observational Interference, January 2022. arXiv:2102.09677 [cs, eess] version: 3.
- [Yau *et al.*, 2020] Herman Yau, Chris Russell, and Simon Hadfield. What Did You Think Would Happen? Explaining Agent Behaviour through Intended Outcomes. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, Vancouver, Canada, 2020.
- [Zhang *et al.*, 2020] Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarin Gal, and Doina Precup. Invariant Causal Prediction for Block MDPs. In *Proceedings of the 37th International Conference on Machine Learning*, pages 11214–11224. PMLR, November 2020. ISSN: 2640-3498.
- [Zhang *et al.*, 2021] Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning Invariant Representations for Reinforcement Learning without Reconstruction, April 2021. arXiv:2006.10742 [cs, stat].