Invariance Discovery for Systematic Generalization in Reinforcement Learning

Anonymous Authors[1]

Abstract

In the sequential decision making setting, an agent aims to achieve *systematic generalization* over a large, possibly infinite, set of environments. Such environments are modeled as discrete Markov decision processes with both states and actions represented through a feature vector. The underlying structure of the environments allows the transition dynamics to be factored into two components: one that is environment-specific and another one that is shared. Consider a set of environments that share the laws of motion as an illustrative example. In this setting, the agent can take a finite amount of *reward-free* interactions from a subset of these environments. The agent then must be able to *approximately* solve any planning task defined over any environment in the original set, relying on the above interactions only. Can we design a provably efficient algorithm that achieves this ambitious goal of systematic generalization? In this paper, we give a partially positive answer to this question. First, we provide the first tractable formulation of systematic generalization by employing a *causal* viewpoint. Then, under specific structural assumptions, we provide a simple learning algorithm that allows us to guarantee any desired planning error up to an unavoidable sub-optimality term, while showcasing a polynomial sample complexity.

1. Introduction

Whereas recent breakthroughs have established Reinforcement Learning (RL) (Sutton and Barto, 2018) as a powerful tool to address a wide range of sequential decision making problems, the curse of generalization (Kirk et al., 2021) is still a main limitation of commonly used techniques. RL algorithms deployed on a given task are usually effective in discovering the correlation between an agent's behavior and the resulting performance from large amounts of labeled samples. However, those algorithms are usually unable to discover basic cause-effect relations between the agent's behavior and the environment dynamics. Crucially, the aforementioned correlations are oftentimes specific to the task, and they are unlikely to be of any use for addressing different tasks. Instead, some universal causal relations generalize over the environments, and once learned can be exploited for solving any task. Let us consider as an illustrative example an agent interacting with a large set of physical environments. While each of these environments can have its specific dynamics, we expect the basic laws of motion to hold across the environments, as they encode general causal relations. Once they are learned, there is no need to discover them again from scratch when facing a new task, or an unseen environment. Even if the dynamics over these relations can change, such as moving underwater is different than moving in the air, or the gravity can change from planet to planet, the underlying causal structure still holds. This knowledge alone often allows the agent to solve new tasks in unseen environments with a few, or zero, interactions.

We argue that we should pursue this kind of generalization in RL, which we call *systematic generalization*, where learning universal causal relations from interactions with a few environments allows us to approximately solve any task in any other environment without further interactions. Although this problem setting might seem overly ambitious or even far-fetched, in this document we provide the first tractable formulation of systematic generalization (Section 3), thanks to a set of structural assumptions that are motivated by a causal viewpoint. Especially, we consider a large, potentially infinite, set of reward-free environments, or a *universe*, the agent can freely interact with. Crucially, these environments share a common causal structure that explains a significant portion, but not all, of their transition dynamics. Can we design a provably efficient algorithm that guarantees an arbitrarily small planning error for any possible task that can be defined over the set of environments, by taking reward-free interactions with a generative model?

In this document, we provide a partially positive answer to this question by presenting a simple but principled causal model-based approach (see Figure 1). This algorithm interacts with a finite subset of the universe to learn the causal structure underlying the set of environments in the form of a causal dependency graph $\mathcal{G}$ (Wadhwa and Dong, 2021). The

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the SCIS workshop.

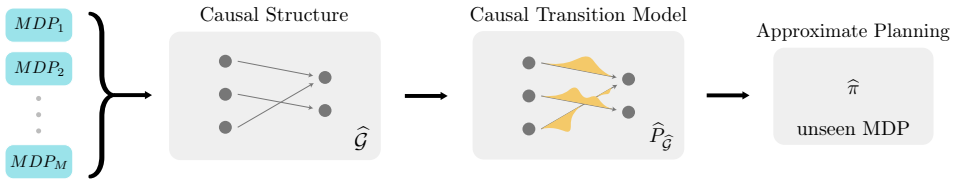

Figure 1: High-level illustration of causal model-based approach to systematic generalization.

causal transition model, which encodes the dynamics that is common across the environment, is obtained by estimating the Bayesian network $P_{\mathcal{G}}$ over $\mathcal{G}$ (Dasgupta, 1997) from a mixture of the environments. Then, the learned model is employed by a planning oracle to provide an approximately optimal policy for a latent environment and a given reward function. We can show that this simple recipe allows achieving any desired planning error up to an unavoidable error term, which is inherent to the setting. Especially, we provide an analysis of the sample complexity (Section 4) of the proposed approach, which is polynomial in all the relevant quantities of the problem.

Finally, with this work we aim to connect several active research areas on reward-free RL (Jin et al., 2020), multi-task RL (Brunskill and Li, 2013), model-based RL (Sutton and Barto, 2018), factored MDPs (Rosenberg and Mansour, 2021), causal RL (Zhang et al., 2020), experimental design (Ghassami et al., 2018), independence testing (Canonne et al., 2018), into a general framework where individual progresses can be enhanced beyond the sum of their parts.

2. Notation

We will denote a set of integers $\{1, \dots, a\}$ as $[a]$, and the probability simplex over the space $\mathcal{A}$ as $\Delta_{\mathcal{A}}$. For any $A \in \mathcal{A}$, we denote with $A[Z]$ the vector $(A_i)_{i \in Z}$. Given two probability measures P and Q over a discrete space $\mathcal{A}$, their L_1-distance is $\|P - Q\|_1 = \sum_{A \in \mathcal{A}} |P(A) - Q(A)|$. We will denote by $\mathcal{U}_{\mathcal{A}}$ the uniform distribution over $\mathcal{A}$.

Graphs We define a graph $\mathcal{G}$ as a pair $\mathcal{G} := (\mathcal{V}, E)$, where $\mathcal{V}$ is a set of nodes and $E \subseteq N \times N$ is a set of edges between them. We call $\mathcal{G}$ a *directed graph* if all of its edges E are directed. We also define the in-degree of a node to be its number of incoming edges: $\text{degree}_{\text{in}}(A) = |\{(B, A) : (B, A) \in E, \forall B\}|$. $\mathcal{G}$ is said to be a *Directed Acyclic Graph* (DAG) if it is a directed graph without cycles. We call $\mathcal{G}$ a *bipartite graph* if there exists a partition $X \cup Y = \mathcal{V}$ such that none of the nodes in X and Y are connected by an edge.

Causal Graphs and Bayesian Networks For a set $\mathcal{X}$ of random variables, we represent the causal structure over $\mathcal{X}$ with a DAG $\mathcal{G}_{\mathcal{X}} = (\mathcal{X}, E)$,[1] which we call the *causal graph* of $\mathcal{X}$. For each pair of variables $A, B \in \mathcal{X}$, a directed edge $(A, B) \in \mathcal{G}_{\mathcal{X}}$ denotes that B is conditionally dependent on A. For every variable $A \in \mathcal{X}$, we denote as $\text{Pa}(A)$ the *causal parents* of A, i.e., the set of all the variables $B \in \mathcal{X}$ on which A is conditionally dependent, $(B, A) \in \mathcal{G}_{\mathcal{X}}$. A Bayesian network (Dean and Kanazawa, 1989) over the set $\mathcal{X}$ is defined as $\mathcal{N} := (\mathcal{G}_{\mathcal{X}}, P)$, where $\mathcal{G}_{\mathcal{X}}$ specifies the *structure* of the network, i.e., the dependencies between the variables in $\mathcal{X}$, and the distribution $P : \mathcal{X} \to \Delta_{\mathcal{X}}$ specifies the conditional probabilities of the variables in $\mathcal{X}$, such that $P(\mathcal{X}) = \prod_{X_i \in \mathcal{X}} P_i(X_i | \text{Pa}(X_i))$.

Markov Decision Processes We define a *discrete* episodic Markov Decision Process (MDP) (Puterman, 2014) as $\mathcal{M} := ((\mathcal{S}, d_{\mathcal{S}}, n), (\mathcal{A}, d_{\mathcal{A}}, n), P, H, r)$, where $\mathcal{S}$ is a set of $|\mathcal{S}| = S$ states and $\mathcal{A}$ is a set of $|\mathcal{A}| = A$ actions, such that every $s \in \mathcal{S}$ can be represented through a d_S-dimensional vector of discrete features taking value in $[n]$, and $a \in \mathcal{A}$ through a d_A-dimensional vector of discrete features taking value in $[n]$.[2] P is a transition model such that $P(s'|s, a)$ gives the conditional probability of the next state s' having taken action a in state s, H is the horizon of an episode, and $r : \mathcal{S} \times \mathcal{A} \to [0, 1]$ is a deterministic reward function. A stochastic *policy* $\pi_h(a|s)$ denotes the conditional probability of taking action a in state s at step h. The *value function* $V_h^{\pi} : \mathcal{S} \to \mathbb{R}$ associated to π is defined as $V_h^{\pi}(s) := \mathbb{E}_{\pi}\left[\sum_{h'=h}^{H} r(s_{h'}, a_{h'}) \,\middle|\, s_h = s\right]$. We will write $V_{\mathcal{M},r}^{\pi}$ to denote V_1^{π} in the MDP $\mathcal{M}$ with reward function r.

3. Problem Formulation

In our setting, a learning agent aims to master a large, potentially infinite, set $\mathbb{U}$ of environments modeled as discrete MDPs without rewards,

$$\mathbb{U} := \left\{\mathcal{M}_i = ((\mathcal{S}, d_S, n), (\mathcal{A}, d_A, n), P_i, \mu)\right\}_{i=1}^{\infty},$$

which we call a *universe*. The agent can draw a finite amount of experience by interacting with the MDPs in $\mathbb{U}$. From these interactions alone, the agent aims to acquire sufficient knowledge to approximately solve any task that can be specified over the universe $\mathbb{U}$. Specifically, a *task* is defined as any pairing of an MDP $\mathcal{M} \in \mathbb{U}$ and a reward function r, whereas *solving it* refers to providing a slightly sub-optimal policy via planning, i.e., without taking additional interactions. We call this problem *systematic generalization*, which

[1] We will omit the subscript $\mathcal{X}$ whenever clear from the context.

[2] Note that any *tabular* MDP can be formulated under this alternative formalism by taking $n = 2$, $d_S = S$, and $d_A = A$.

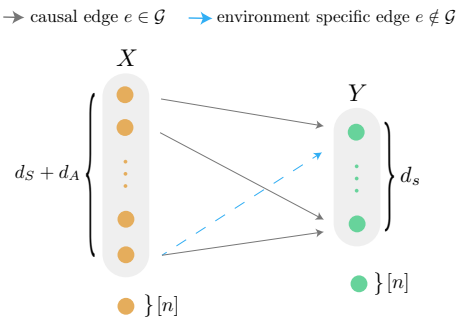

Figure 2: Illustration of the causal structure $\mathcal{G}$ of $\mathbb{U}$.

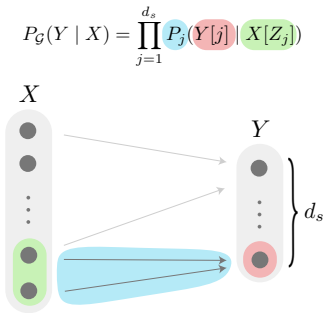

Figure 3: Illustration of the causal transition model $P_{\mathcal{G}}$.

we can formalize as follows.

Definition 1 (Systematic Generalization)**.** *For any latent MDP $\mathcal{M} \in \mathbb{U}$ and any given reward $r : \mathcal{S} \times \mathcal{A} \to [0, 1]$, the systematic generalization problem requires the agent to provide a policy π, such that $V^*_{\mathcal{M},r} - V^{\pi}_{\mathcal{M},r} \leq \epsilon$ up to any desired $\epsilon > 0$.*

Since the set $\mathbb{U}$ is infinite, we clearly require additional structure to make the problem feasible. On the one hand, the state space $(\mathcal{S}, d_S, n)$, action space $(\mathcal{A}, d_A, n)$, and initial state distribution μ are shared across $\mathcal{M} \in \mathbb{U}$. The transition dynamics P_i is instead specific to each MDP $\mathcal{M}_i \in \mathbb{U}$. However, we assume the presence of a *common causal structure* that underlies the transition dynamics of the universe, and relates the single transition models P_i.

3.1. The Causal Structure of the Transition Dynamics

To ease the notation, we denote the current state-action features with a random vector $X = (X_i)_{i \in [d_S + d_A]}$, and the next state features with a random vector $Y = (Y_i)_{i \in [d_S]}$. For each environment $\mathcal{M}_i \in \mathbb{U}$, the conditional dependencies between Y and X are represented through a bipartite dependency graph $\mathcal{G}_i$. Clearly, each environment can display its own dependencies, but we assume there is a set of dependencies that represent general causal relationships between the features, and that appear in any $\mathcal{M}_i \in \mathbb{U}$. In particular, we call the intersection $\mathcal{G} := \cap_{i=0}^{\infty} \mathcal{G}_i$ the *causal structure* of $\mathbb{U}$, which is the set of conditional dependencies that are common across the universe. In Figure 2, we show an illustration of such a causal structure. We assume the causal structure $\mathcal{G}$ is time-consistent, i.e., $\mathcal{G}^{(h)} = \mathcal{G}^{(1)}$ for any step $h \in [H]$, and sparse, which means that the number of features $X[z]$ on which a feature $Y[j]$ is dependent on is bounded from above.

Assumption 1 (Z-sparseness)**.** *Let $Z \in \mathbb{N}$. The causal structure $\mathcal{G}$ is Z-*sparse *if* $\max_{j \in [d_S]} \text{degree}_{\text{in}}(Y[j]) \leq Z$.

Given a causal structure $\mathcal{G}$, and without losing generality,[3] we can express each transition model P_i as

$$P_i(Y|X) = P_{\mathcal{G}}(Y|X) F_i(Y|X)$$
$$P_{\mathcal{G}}(Y|X) = \prod_{j=1}^{d_S} P_j(Y[j]|X[Z_j]) \tag{1}$$

in which $P_{\mathcal{G}}$ is the Bayesian network over the causal structure $\mathcal{G}$, whereas F_i includes environment-specific factors affecting the conditional probabilities,[4] the Z_j are the set of indices z such that $(X[z], Y[j]) \in \mathcal{G}$. Since it represents the conditional probabilities due to universal causal relations in $\mathbb{U}$, we call $P_{\mathcal{G}}$ the *causal transition model* of $\mathbb{U}$, for which we show an illustration in Figure 3. We assume the causal transition model $P_{\mathcal{G}}$ is also time-consistent, i.e., $P^{(h)}_{\mathcal{G}} = P^{(1)}_{\mathcal{G}}, \forall h \in [H]$, and that it explains a significant part of the transition dynamics of $\mathcal{M}_i \in \mathbb{U}$.

Assumption 2 (λ-sufficiency)**.** *Let $\lambda \in [0, 1]$ be a constant. The causal transition model $P_{\mathcal{G}}$ is* causally λ-sufficient *if*

$$\sup_X \|P_{\mathcal{G}}(\cdot|X) - P_i(\cdot|X)\|_1 \leq \lambda, \qquad \forall P_i \in \mathcal{M}_i \in \mathbb{U}.$$

Notably, the parameter λ controls the amount of the transition dynamics that is due to the universal causal relations $\mathcal{G}$ ($\lambda = 0$ means that $P_{\mathcal{G}}$ is sufficient to explain the transition dynamics of any $\mathcal{M}_i \in \mathbb{U}$, whereas $\lambda = 1$ implies no shared structure between the transition dynamics of the $\mathcal{M}_i \in \mathbb{U}$). In this paper, we argue that learning the causal transition model $P_{\mathcal{G}}$ is a good target for systematic generalization and we provide theoretical support for this claim in Section 4.

3.2. A Class of Training Environments

Even if the universe $\mathbb{U}$ admits the structure that we presented in the last section, it is still an infinite set. Instead, the agent can only interact with a finite subset of discrete MDPs

$$\mathbb{M} := \{\mathcal{M}_i = ((\mathcal{S}, d_S, n), (\mathcal{A}, d_A, n), P_i, \mu)\}_{i=1}^{M} \subset \mathbb{U},$$

[3]Note that one can always take $P_{\mathcal{G}}(Y|Z) = 1, \forall (X, Y)$ to avoid shared structure on the transition dynamics.

[4]The parameters in F_i are numerical values such that P_i remains a well-defined probability measure.

which we call a *class* of size M. Crucially, the causal structure $\mathcal{G}$ is a property of the full set $\mathbb{U}$, and if we aim to infer it from interactions with a finite class $\mathbb{M}$, we have to assume that $\mathbb{M}$ is informative enough.

Assumption 3 (Diversity)**.** *Let $\mathbb{M} \subset \mathbb{U}$ be class of size M. We say $\mathbb{M}$ is* causally diverse *if $\mathcal{G} = \cap_{i=1}^{M} \mathcal{G}_i = \cap_{i=1}^{\infty} \mathcal{G}_i$.*

Analogously, if we aim to infer the causal transition model $P_{\mathcal{G}}$ from interactions with the transition models P_i of the single MDPs $\mathcal{M}_i \in \mathbb{M}$, we have to assume that $\mathbb{M}$ is balanced in terms of the conditional probabilities displayed by its components, so that the factors that do not represent universal causal relations even out while learning.

Assumption 4 (Evenness)**.** *Let $\mathbb{M} \subset \mathbb{U}$. We say $\mathbb{M}$ is* causally even *if $\mathbb{E}_{i \sim \mathcal{U}_{[M]}} \left[F_i(Y[j]|X) \right] = 1, \forall j \in [d_S]$.*

Whereas in this paper we assume that $\mathbb{M}$ is *diverse* and *even* by design, we leave as future work the interesting problem of selecting such a class from active interactions with $\mathbb{U}$, which would add to our problem formulation flavors of active learning and experimental design (Hauser and Bühlmann, 2014; Kocaoglu et al., 2017; Ghassami et al., 2018).

4. Sample Complexity of Systematic Generalization with a Generative Model

We have access to a class $\mathbb{M}$ of discrete MDPs within a universe $\mathbb{U}$, from which we can draw interactions with a generative model $P(X)$. We would like to solve the systematic generalization problem as described in Definition 1. This problem requires to provide, for any combination of a (latent) MDP $\mathcal{M} \in \mathbb{U}$, and a given reward function r, a planning policy $\widehat{\pi}$ such that $V^*_{\mathcal{M},r} - V^{\widehat{\pi}}_{\mathcal{M},r} \leq \epsilon$. Especially, can we design an algorithm that guarantees this requirement with high probability by taking a number of samples K that is polynomial in ϵ and the relevant parameters of $\mathbb{M}$? Here we give a partially positive answer to this question, by providing a simple but provably efficient algorithm that guarantees systematic generalization over $\mathbb{U}$ up to an unavoidable sub-optimality term ϵ_λ that we will later specify.

The algorithm implements a model-based approach into two separated components. The first, for which we provide the pseudocode in Algorithm 1, is the procedure that actually interacts with the class $\mathbb{M}$ to obtain a principled estimation $\widehat{P}_{\widehat{\mathcal{G}}}$ of the causal transition model $P_{\mathcal{G}}$ of $\mathbb{U}$. The second, is a planning oracle that takes as input a reward function r and the estimated causal transition model, and returns an optimal policy $\widehat{\pi}$ operating on $\widehat{P}_{\widehat{\mathcal{G}}}$ as an approximation of the transition model P_i of the true MDP $\mathcal{M}_i$. We provide an upper bound to the sample complexity of the Algorithm 1.

Lemma 4.1. *Let $\mathbb{M} = \{\mathcal{M}_i\}_{i=1}^{M}$ be a class of M discrete MDPs, let $\delta \in (0,1)$, and let $\epsilon > 0$. The Algorithm 1 returns a causal transition model $\widehat{P}_{\widehat{\mathcal{G}}}$ such that $Pr(\|\widehat{P}_{\widehat{\mathcal{G}}} - P_{\mathcal{G}}\|_1 \geq \epsilon) \leq \delta$ with a sample complexity*

$$K = O\Big(M d_S^3 Z^2 n^{3Z+1} \log\Big(\frac{4 M d_S^2 d_A n^Z}{\delta} \Big) \Big/ \epsilon^2 \Big).$$

Algorithm 1 Causal Transition Model Estimation

Input: class of MDPs $\mathbb{M}$, error ϵ, confidence δ
let $K' = C'\big(d_S^2 Z^2 n \log(2 M d_S^2 d_A/\delta)/\epsilon^2\big)$
set the generative model $P(X) = \mathcal{U}_X$
for $i = 1, \dots, M$ **do**
  let $P_i(Y|X)$ be the transition model of $\mathcal{M}_i \in \mathbb{M}$
  $\widehat{\mathcal{G}}_i \leftarrow$ *Causal Structure Estimation* $(P_i, P(X), K')$
end for
let $\widehat{\mathcal{G}} = \cap_{i=1}^{M} \widehat{\mathcal{G}}_i$
let $K'' = C''\big(d_S^3 n^{3Z+1} \log(4 d_S n^Z/\delta)/\epsilon^2\big)$
let $P_{\mathbb{M}}(Y|X)$ be the mixture $\frac{1}{M}\sum_{i=1}^{M} P_i(Y|X)$
$\widehat{P}_{\widehat{\mathcal{G}}} \leftarrow$ *Bayesian Network Estimation* $(P_{\mathbb{M}}, \widehat{\mathcal{G}}, K'')$
Output: causal transition model $\widehat{P}_{\widehat{\mathcal{G}}}$

Having established the sample complexity of the causal transition model estimation, we can now show how the learned model $\widehat{P}_{\widehat{\mathcal{G}}}$ allows us to approximately solve, via a planning oracle, any task defined by a combination of a latent MDP $\mathcal{M}_i \in \mathbb{U}$ and a given reward function r.[5]

Theorem 4.2. *Let $\delta \in (0,1)$ and $\epsilon > 0$. For a latent discrete MDP $\mathcal{M} \in \mathbb{U}$, and a given reward function r, a planning oracle operating on the causal transition model $\widehat{P}_{\widehat{\mathcal{G}}}$ as an approximation of $\mathcal{M}$ returns a policy $\widehat{\pi}$ such that*

$$Pr\big(V^*_{\mathcal{M}_i,r} - V_{\mathcal{M}_i,r} \geq \epsilon_\lambda + \epsilon\big) \leq \delta,$$

where $\epsilon_\lambda = 2\lambda H^3 d_S n^{2Z+1}$, and $\widehat{P}_{\widehat{\mathcal{G}}}$ is obtained from Algorithm 1 with $\delta' = \delta$ and $\epsilon' = \epsilon/2H^3 n^{Z+1}$.

Theorem 4.2 establish the sample complexity of systematic generalization through Lemma 4.1. For the discrete MDP setting, we have that $\widetilde{O}(M H^6 d_S^3 Z^2 n^{5Z+3})$, which reduces to $\widetilde{O}(M H^6 S^4 A^2 Z^2)$ in the tabular setting. Unfortunately, we are only able to obtain systematic generalization up to an unavoidable sub-optimality term ϵ_λ. This error term is related to the λ-sufficiency of the causal transition model (Assumption 2), and it accounts for the fact that $P_{\mathcal{G}}$ cannot fully explain the transition dynamics of each $\mathcal{M} \in \mathbb{U}$, even when it is estimated exactly. This is inherent to the ambitious problem setting, and can be only overcome with additional interactions with the test MDP $\mathcal{M}$.

[5] To provide this result in the discrete MDP setting, we have to further assume that the transition dynamics P_i of the target MDP $\mathcal{M}_i$ admits factorization analogous to (1), such that we can write $P_i(Y|X) = \prod_{j=1}^{d_S} P_{i,j}(Y[j]|X[Z_j^{'}])$, where the scopes $Z_j^{'}$ are given by the environment-specific causal structure $\mathcal{G}_i$, which we assume to be $2Z$-sparse (Assumption 1).

References

Brunskill, E. and Li, L. (2013). Sample complexity of multi-task reinforcement learning. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.

Canonne, C. L., Diakonikolas, I., Kane, D. M., and Stewart, A. (2018). Testing conditional independence of discrete distributions. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–57. IEEE.

Dasgupta, S. (1997). The sample complexity of learning fixed-structure bayesian networks. *Machine Learning*, 29(2):165–180.

Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation. *Computational intelligence*, 5(2):142–150.

Ghassami, A., Salehkaleybar, S., Kiyavash, N., and Bareinboim, E. (2018). Budgeted experiment design for causal structure learning. In *Proceedings of the International Conference on Machine Learning*.

Hauser, A. and Bühlmann, P. (2014). Two optimal strategies for active learning of causal models from interventional data. *International Journal of Approximate Reasoning*, 55(4):926–939.

Jin, C., Krishnamurthy, A., Simchowitz, M., and Yu, T. (2020). Reward-free exploration for reinforcement learning. In *Proceedings of the International Conference on Machine Learning*.

Kirk, R., Zhang, A., Grefenstette, E., and Rocktäschel, T. (2021). A survey of generalisation in deep reinforcement learning. *arXiv preprint arXiv:2111.09794*.

Kocaoglu, M., Shanmugam, K., and Bareinboim, E. (2017). Experimental design for learning causal graphs with latent variables. In *Advances in Neural Information Processing Systems*.

Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Rosenberg, A. and Mansour, Y. (2021). Oracle-efficient regret minimization in factored mdps with unknown structure. In *Advances in Neural Information Processing Systems*.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Wadhwa, S. and Dong, R. (2021). On the sample complexity of causal discovery and the value of domain expertise. *arXiv preprint arXiv:2102.03274*.

Zhang, A., Lyle, C., Sodhani, S., Filos, A., Kwiatkowska, M., Pineau, J., Gal, Y., and Precup, D. (2020). Invariant causal prediction for block mdps. In *Proceedings of the International Conference on Machine Learning*.