

# Planning Under Distribution Shifts with Causal POMDPs

Matteo Ceriscioli, Karthika Mohan

Oregon State University, Corvallis, USA  
 {ceriscim, karthika.mohan}@oregonstate.edu

## Abstract

In the real world, planning is often challenged by distribution shifts. As such, a model of the environment obtained under one set of conditions may no longer remain valid as the distribution of states or the environment dynamics change, which in turn causes previously learned strategies to fail. In this work, we propose a theoretical framework for planning under partial observability using Partially Observable Markov Decision Processes (POMDPs) formulated using causal knowledge. By representing shifts in the environment as interventions on this causal POMDP, the framework enables evaluating plans under hypothesized changes and actively identifying which components of the environment have been altered. We show how to maintain and update a belief over both the latent state and the underlying domain, and we prove that the value function remains piecewise linear and convex (PWLC) in this augmented belief space. Preservation of PWLC under distribution shifts has the advantage of maintaining the tractability of planning via  $\alpha$ -vector-based POMDP methods.

## Introduction

Planning aims to compute an optimal way for an agent to act, assuming access to a complete and correct model of the environment. In stochastic settings, this often means identifying a policy that maps information states to actions and maximizes long-term reward (Kaelbling, Littman, and Cassandra 1998). Difficulties arise once the agent is deployed in a different environment and the underlying state distribution or transition dynamics change. Distribution shifts undermine the assumption that the transition and observation processes remain fixed and accurately describe how states evolve.

To illustrate this issue, consider a delivery rover whose policy and world model were obtained in an environment with specific conditions. When the same rover is deployed in a different setting, relevant factors, such as friction, visibility, or mechanical stress, may change. A conventional POMDP planner cannot identify which mechanisms have shifted or why performance has degraded. A causally informed model can represent such changes as interventions on specific components, reason about their effects on the environment and the sensors, and support effective planning under the resulting distribution shift.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This paper examines whether planning remains tractable in the presence of distribution shifts. The tractability of finite-horizon POMDPs relies on the fact that their value functions are PWLC in the belief state (Smallwood and Sondik 1973). When the environment is affected by an unknown shift, the resulting ambiguity in the transition dynamics introduces an additional source of uncertainty into the model. Although certain structured forms of uncertainty are known to preserve the PWLC structure (Osogami 2015), general forms of ambiguity can cause finite-horizon POMDPs to lose their PWLC properties (Saghafian 2018). We consider a causal formulation of POMDPs in which distribution shifts are represented as interventions. Under this formulation, we show that the value function retains its PWLC structure despite the presence of unknown shifts. Thus this ensures that standard  $\alpha$ -vector-based POMDP planning algorithms remain applicable.

Causal modeling is widely used to support robust decision making under perturbations of the data-generating process (Pearl and Bareinboim 2011; Pearl 2018; Schölkopf 2022), as it offers a principled framework for analyzing how distribution shifts affect outcomes. Such shifts can be naturally represented as interventions in a causal model (Richens and Everitt 2024; Ceriscioli and Mohan 2025), which allows systematic reasoning about their consequences.

The primary contribution of this paper is a causal framework for planning under partial observability in the presence of distribution shifts, which preserves the PWLC structure of the value function. The framework enables the evaluation of policies under specified shifts and supports the detection and identification of changes in the environment.

All proofs are included in the supplementary material.

## Preliminaries

Let  $X$  be a node in a graph  $G$ . In this paper we denote the set of parent nodes of  $X$  as  $Pa(X)$ . Random variables are represented with upper-case letters, e.g.  $X$ , and the corresponding instantiations with lower-case letters, e.g.  $x$ . The set of observable values of a random variable  $X$  is called  $dom(X)$ . Given a set  $\mathbf{X}$ ,  $\Pi(\mathbf{X})$  is the set of all probability distributions over  $\mathbf{X}$ .

**Factored POMDP.** Partially Observable Markov Decision Processes (POMDPs) (Åström 1965; Kaelbling,

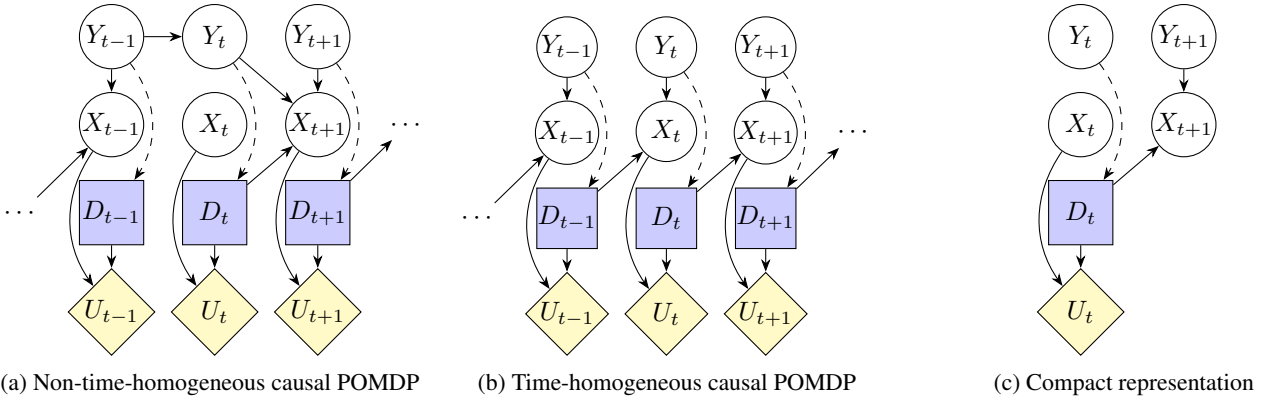


Figure 1: On the left (a), a non-time-homogeneous causal POMDP where the graph structure and variable dependencies can change across timesteps. If the model is faithful, meaning every solid edge represents a true direct causal dependency, any change in the graph implies a change in the transition function. In the center (b), a time-homogeneous causal POMDP with a consistent graph across timesteps. On the right (c), a simplified depiction of the time-homogeneous POMDP that omits repeated past nodes while fully capturing causal dependencies.

Littman, and Cassandra 1998) model sequential decision-making in which the agent lacks full observability of the environment. When each state in a POMDP is represented as a vector  $s = (v_1, \dots, v_n)$  of instantiations of random variables  $V_1, \dots, V_n$ , the model is called a factored POMDP (Boutilier and Poole 1996). This structure allows the transition function to be decomposed into sub-functions involving only subsets of the state variables.

**Causal Influence Diagrams.** In this paper, we study factored POMDPs whose intra- and inter-temporal causal structure is expressed using Causal Influence Diagrams (CIDs) (Heckerman 1995; Everitt et al. 2021). CIDs extend Causal Bayesian Networks (CBNs) (Pearl 2009) by incorporating decision-theoretic elements, offering a structured way to represent causal relationships while specifying what the agent observes and what it influences through its actions.

**Definition 1** (Causal Influence Diagram (Everitt et al. 2021)). A *Causal Influence Diagram* (CID) is a Causal Bayesian Network  $M = (G = (\mathbf{V}, \mathbf{E}), P)$ , where  $P$  is a joint probability distribution compatible with the conditional independences encoded in the DAG  $G$ . The nodes in  $\mathbf{V}$  are partitioned into decision ( $\mathbf{D}$ ), utility ( $\mathbf{U}$ ), and chance ( $\mathbf{C}$ ) nodes,  $\mathbf{V} = (\mathbf{D}, \mathbf{U}, \mathbf{C})$ . Each utility node  $U_i$  is assigned a utility function  $f_i : \text{dom}(Pa(U_i)) \rightarrow \mathbb{R}$ .

In a CID, the environment is described by the set of chance nodes  $\mathbf{C}$ . Each chance node  $C \in \mathbf{C}$  corresponds to a random variable, The set of decision nodes  $\mathbf{D}$  contains all the variables which value is set by an agent, for each  $D \in \mathbf{D}$  it is possible to assign a policy  $\pi : \text{dom}(Pa(D)) \rightarrow \mathcal{A}$ .

### Causal POMDPs

There exist various causal representations of POMDPs, including models that integrate causal structure into online planning or representation learning. For example, CAR-DESPOT incorporates causal information into the AR-DESPOT planner to reason about confounding in robotic

environments (Cannizzaro and Kunze 2023). Earlier work has proposed a causal POMDP model for causal representation learning and has shown its usefulness for zero-shot learning in complex tasks (Sontakke et al. 2021). Other approaches use causal modeling to recover hidden causal dynamics within partially observable environments (Liang and Boularias 2021), or employ causal abstractions to improve long-horizon reasoning (Gao et al. 2025).

In this work, we adapt the formulation of Ceriscioli and Mohan (2025), which models a POMDP as an infinite CID unrolled over time containing a decision node and a utility node at each timestep  $t$ , corresponding respectively to the action and reward at time  $t$ .

**Definition 2** (Causal POMDP). A causal POMDP is a tuple  $(\mathbf{V}, \mathcal{A}, \mathbf{T}, R, \mathbf{V}_o, O, \gamma)$  where:

1.  $\mathbf{V} = (V_1, \dots, V_n)$  is the ordered set of state variables.
2.  $\mathcal{A}$  is the finite set of actions.
3.  $T : \text{dom}(\mathbf{V}) \times \mathcal{A} \rightarrow \Pi(\text{dom}(\mathbf{V}))$  is the state-transition function. Given  $\mathbf{V}^{(t)}$  the set of state variables before the transition, and  $\mathbf{V}^{(t+1)}$  the set of state variables after the transition, it is possible to decompose  $T$  as follows:

$$T(v^{(t)}, a^{(t)}, v^{(t+1)}) = \prod_{i=1}^n T_i(v_i^{(t+1)} \mid a^{(t)}, pa(v_i^{(t+1)})) \quad (1)$$

with  $Pa(V_i^{(t+1)}) \subseteq \mathbf{V}^{(t)} \cup (\bigcup_{k < i} V_k^{(t+1)})$ . Each function  $T_i$  governs the transition of the state variable  $V_i$ .

4.  $R : \mathbf{V}_R \times \mathcal{A} \rightarrow \mathbb{R}$  is a reward function, where  $\mathbf{V}_R \subseteq \mathbf{V}$ .
5.  $\mathbf{V}_o = \{V_{o,1}, \dots, V_{o,p}\}$  is a set of observable variables, forming the agent’s observation.
6.  $\mathbf{O} = \{O_1, \dots, O_p\}$   $O : \text{dom}(\mathbf{V}) \times \mathcal{A} \times \text{dom}(\mathbf{V}_o) \rightarrow [0, 1]$  is the set of conditional observation probabilities.
7.  $\gamma \in [0, 1)$  is the discount factor.

The set of state variables induces the set of states  $\mathbf{S} := \text{dom}(\mathbf{V}) = \text{dom}(V_1) \times \dots \times \text{dom}(V_n)$ . Similarly, the set of observations is  $\mathbf{\Omega} := \text{dom}(\mathbf{V}_o)$ . The agent receives a

reward  $R(s, a)$  after taking an action  $a \in \mathcal{A}$  in state  $s \in \mathbf{S}$ .

Observable variables can be included in the state ( $\mathbf{V}_o \subseteq \mathbf{V}$ ) without loss of expressiveness: for any causal POMDP where  $\mathbf{V}_o \cap \mathbf{V} = \emptyset$ , there exists an equivalent causal POMDP with  $\mathbf{V}'_o \subseteq \mathbf{V}$  where transitions and rewards do not depend on  $\mathbf{V}_o$ .

**Assumption 1.** The observable variables are also state variables, i.e.  $\mathbf{V}_o \subseteq \mathbf{V}$ .

When Assumption 1 holds, the observation is fully determined by the state and generated through the state-transition function. Therefore, the conditional observation probability function simplifies to  $O(v', a, v'_o) = 1$  if  $v'_o$  is compatible with  $v'$  and 0 otherwise. Since  $v'$  completely determines  $v'_o$ , we can write  $O(v', v'_o)$  instead of  $O(v', a, v'_o)$ .

Under Assumption 1, a causal POMDP  $(\mathbf{V}, \mathcal{A}, \mathbf{T}, R, \mathbf{V}_o, O, \gamma)$  induces a CID with an infinite DAG  $G = (\mathbf{V}', E)$ , where  $\mathbf{V} = \bigcup_{t=0}^{\infty} \mathbf{V}^{(t)}$  and  $\mathbf{V}^{(t)} := \mathbf{V}$  (see Figure 1b). Time-homogeneity of the transition function allows a compact CID representation including only two consecutive timesteps  $t$  and  $t + 1$ , with decisions and utilities at  $t$ , and edges involving  $D$  and  $U$ , between timesteps, and within  $t + 1$  (see Figure 1c). If all state variables are observable, this reduces to a causal MDP.

Distribution shifts generally alter the state transition function, and if the distribution shift is unknown, then estimating it requires planning using a non-time-homogeneous POMDP, non-time-homogeneous causal POMDPs can also be represented with infinite CIDs, as illustrated in Figure 1a.

## Planning under Distribution Shifts

Being able to discern the causal relationships governing the environment allows us to infer the state dynamics under changing conditions. In a causal model, *interventions* are deliberate alterations of components or mechanisms of the model. A *distribution shift* is any change in the probability distribution of a random variable. In this paper, we focus on shifts that affect the variables describing the environment in which we plan, specifically, the variables used to factorize the POMDP state in the causal POMDP model. A *domain* denotes the probabilistic configuration of the environment under a particular shift, including the original one.

In this section, we develop the use of causal POMDPs for planning under potential distribution shifts. We first formalize distribution shifts as interventions within the underlying causal model. A key consequence of modeling distribution shifts as interventions is that they can be embedded directly into the agent’s belief and planning process. We proceed by describing how to evaluate a policy under a specified distribution shift. Finally, we address the problem of planning while concurrently identifying the distribution shift, showing how to update the agent’s belief over states and domains and that the resulting belief value function remains piecewise linear and convex.

## Modeling Distribution Shifts as Interventions

As shown in previous work (Richens and Everitt 2024; Ceriscioli and Mohan 2025), interventions are an effective

way to represent distribution shifts in a causal model. Applying an intervention  $\sigma$  to a set of variables may change the joint distribution  $P(\mathbf{V})$ . When the model and intervention are known, it is possible to compute the updated distribution  $P(\mathbf{V}; \sigma)$  which is called the *interventional distribution*, opposed to the *observational distribution*  $P(\mathbf{V})$ , which is the one we observe when no intervention is applied.

We introduce stochastic shifts, a family of soft interventions.

**Definition 3** (Stochastic Shift Intervention). Let  $M = (G, \Theta)$  be a CBN, and  $X$  be a discrete random variable with  $\text{dom}(X) = \{x_1, \dots, x_m\}$ . A *Stochastic Shift Intervention*  $\sigma$  is an intervention associated with a matrix:

$$A_\sigma = \begin{pmatrix} p_{11} & \cdots & p_{1m} \\ \vdots & \ddots & \vdots \\ p_{m1} & \cdots & p_{mm} \end{pmatrix} \quad \text{with } \sum_j p_{ij} = 1 \text{ for every } i. \quad (2)$$

such that when applied to a random variable  $X$ , its conditional distribution is updated as follows:

$$P(X = x_j \mid pa(X); \sigma) = \sum_{i=1}^m p_{ij} P(X = x_i \mid pa(X)) \quad (3)$$

The condition  $\sum_j p_{ij} = 1$  is equivalent to requiring that each row in  $A_\sigma$  sums up to one. Observe that if the probability distribution of  $X$  is represented with a probability vector, e.g.  $P(X \mid pa(X)) = (P(x_1 \mid pa(X)), \dots, P(x_m \mid pa(X)))^T$ , then we can apply the stochastic shift intervention  $\sigma$  by matrix multiplication, i.e.

$$P(X \mid pa(X); \sigma) = A_\sigma^T P(X \mid pa(X)) \quad (4)$$

$A_\sigma$  is the identity matrix iff  $\sigma$  is the identity intervention  $\sigma_{id}$ , which leaves the domain and distribution unchanged.

**Example.** Let  $X \sim \text{Unif}(\{1, 2, 3\})$  and  $\sigma$  be a stochastic shift s.t.  $p_{11} = p_{22} = 1$ , and  $p_{31} = p_{32} = \frac{1}{2}$ . Each time  $X$  takes the value 3,  $\sigma$  remaps it to 1 or 2 with equal probability. So  $P(X = x; \sigma) = \frac{1}{2}$  if  $x \in \{1, 2\}$  and 0 if  $x = 3$ .

A stochastic shift intervention  $\sigma$  maps probability distributions over a finite set to other distributions. The following result shows that, for any starting distribution, there exists a  $\sigma$  that maps it to any target distribution.

**Proposition 1.** *Given a conditional distribution  $P(X \mid pa(X))$  and an arbitrary target conditional distribution  $P'(X \mid pa(X))$ , it is possible to define a stochastic shift intervention  $\sigma$  s.t.  $P'(X \mid pa(X)) = P(X \mid pa(X); \sigma)$ .*

Note that even if the original distribution  $P(X \mid pa(X))$  and the shifted distribution  $P(X \mid pa(X); \sigma)$  are both observed, in the general case it is not always possible to uniquely determine the distribution shift  $\sigma$ , as it is possible that two different shifts applied to the same distribution generate the same distribution. Suppose  $X$  is the outcome of a fair coin flip, i.e.,  $X \sim \text{Bern}(\frac{1}{2})$ , also suppose that after altering the coin we observe  $\text{Bern}(\frac{3}{4})$  as the new distribution, then both

$$A_\sigma = \begin{pmatrix} 1 & 0 \\ 0.5 & 0.5 \end{pmatrix} \quad A_{\sigma'} = \begin{pmatrix} 0.5 & 0.5 \\ 1 & 0 \end{pmatrix} \quad (5)$$

correspond to shifts that map  $\text{Bern}(\frac{1}{2})$  to  $\text{Bern}(\frac{3}{4})$ .

## Evaluating Policies under Distribution Shifts

Let  $M = (\mathbf{V}, \mathcal{A}, \mathbf{T}, R, \mathbf{V}_o, O, \gamma)$  be a causal POMDP,  $\pi : \text{dom}(\mathbf{V}_o) \rightarrow \mathcal{A}$  a policy, and  $\sigma$  a stochastic shift intervention representing a distribution shift. When planning under partial observability, the state is generally not completely available to the agent and therefore it maintains a belief  $b_S$  about the state that updates whenever it receives an observation (Kaelbling, Littman, and Cassandra 1998). Evaluating a policy consists in computing its expected return, represented by the belief value function  $V^\pi(b)$  using the Bellman equation for causal POMDPs:

$$\begin{aligned} V^\pi(b_S) &= \sum_{s \in S} R(s, \pi(b_S)) b_S(s) + \\ &+ \gamma \sum_{s', o'} O(s', o') \sum_s \prod_{V_i \in \mathbf{V}^{(t+1)}} T_i(v_i | \pi(b_S), pa(v_i)) V^\pi(b_S^{a, o}) \end{aligned} \quad (6)$$

**Known shift  $\sigma$ .** We write  $V^\pi(b_S; \sigma)$  to denote the value function when the environment is affected by the shift  $\sigma$ . Its expression is identical to that in Equation 6, except that each transition term  $T_i(v_i | \pi(b_S), pa(v_i))$  is replaced by  $T_i(v_i | \pi(b_S), pa(v_i); \sigma)$ . Note that if  $\sigma = \sigma_{id}$  then  $V^\pi(b_S; \sigma) = V^\pi(b_S)$ . Once the shift  $\sigma$  is fixed and the transition functions are updated according to Equation 3, the problem of computing the value function of a causal POMDP reduces to that of a standard POMDP.

## Planning under Unknown Distribution Shifts

Now we consider the task of planning in an environment affected by an unknown distribution shift.

**State and Domain Estimation.** As part of planning with causal POMDP under an unknown distribution shift, it is no longer sufficient for the agent to keep track of its belief about the state, as it also needs to keep a belief about the unknown domain. It is possible to define a prior on the state and the domain separately as  $b_S(s)$  and  $b_\Sigma(\sigma)$ , however, since in the general case at each timestep the observation depends both on the previous state and the domain, the state belief and the domain belief become coupled and therefore it is appropriate to keep track of them using a joint belief  $b(s, \sigma)$ . If we have no prior knowledge we can express that as uniform priors on both the states and the domain  $b(s, \sigma) = b(\sigma|s)b(s) = b_\Sigma(\sigma)b_S(s)$  where  $b_S(s) \sim \text{Unif}(S)$ , and  $b_\Sigma(\sigma) \sim \text{Unif}(\Sigma)$ .

The following proposition shows that the joint belief over states and domains admits a Bayesian update analogous to the standard POMDP filter, but extended to account for domain-dependent transition dynamics.

**Proposition 2.** [State-Domain Joint Belief Update] *Let  $b$  be the current joint belief over states and domains,  $a$  be the action taken at time  $t$ , and  $o'$  the observation received after performing  $a$ . Then the updated state state-domain joint belief  $b'$  is:*

$$b'_{o', a}(s', \sigma) = \frac{O(s', o')}{P(o' | a, b)} \sum_s b(s, \sigma) \prod_{V_i \in \mathbf{V}^{(t+1)}} T(v_i | pa(V_i); \sigma) \quad (7)$$

such that  $s' = (v_1, \dots, v_n)$  for  $\{V_1, \dots, V_n\} = \mathbf{V}^{(t+1)}$ .

## Preservation of PWLC under Distribution Shift

We establish that, despite the presence of an unknown shift, the structural properties that support tractable POMDP planning remain intact.

**Value Function for Casual POMDPs under an Unknown Shift.** To demonstrate that the value function of a causal POMDP subject to a distribution shift and planning horizon  $n$  is piecewise linear and convex, and that it admits a finite vector representation analogous to that of standard POMDPs, we propose an approach similar to Porta, Spaan, and Vlassis (2005) and present a constructive proof.

The following lemma shows that the value function under an unknown distribution shift can still be expressed as a maximum over linear functionals of the joint belief.

**Lemma 1.** *Let  $\Sigma$  be a set of stochastic shift interventions, the value function of a causal POMDP with set of states  $S$  subject to an unknown distribution shift in the set  $\Sigma$  with planning horizon  $n$  can be expressed as:*

$$V_n(b) = \max_{\{\alpha_n^i\}_i} \sum_{s \in S} \int_{\Sigma} \alpha_n^i(s, \sigma) b(s, \sigma) d\sigma \quad (8)$$

for some functions  $\alpha_n^i : S \times \Sigma \rightarrow \mathbb{R}$ .

Using the characterization of  $V_n(b)$  provided by Lemma 1, the following theorem completes our analysis by proving that the value function under distribution shifts is still piecewise linear and convex, thereby preserving the structural form that underlies  $\alpha$ -vector-based planning.

**Theorem 1.** *Let  $\Sigma$  be a set of stochastic shift interventions, the value function of a causal POMDP subject to an unknown distribution shift in the set  $\Sigma$  with planning horizon  $n$  is piecewise linear and convex in the joint state-domain belief  $b$ .*

This confirms that distribution shifts do not compromise the representational tractability of the planning problem.

## Conclusions

Recognizing distribution shifts as a critical source of uncertainty in planning, this work introduces a framework for decision-making under such shifts using causal POMDPs. By representing the state in a factorized form and distribution shifts as interventions on the underlying causal model, causal POMDPs enable both the evaluation of policies under specified shifts and planning in an unknown domain by maintaining a joint belief over states and domains. We show that, even under the uncertainty in the transition dynamics that arises from an unknown distribution shift, the finite-horizon value function remains piecewise linear and convex with respect to the joint state-domain belief, thereby preserving the structural properties that support  $\alpha$ -vector-based planning in standard POMDPs.

## References

- Åström, K. J. 1965. Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications*, 10: 174–205.
- Boutilier, C.; and Poole, D. 1996. Computing optimal policies for partially observable decision processes using compact representations. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2*, AAAI'96, 1168–1175. AAAI Press. ISBN 026251091X.
- Cannizzaro, R.; and Kunze, L. 2023. CAR-DESPOT: Causally-Informed Online POMDP Planning for Robots in Confounded Environments. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018–2025.
- Ceriscioli, M.; and Mohan, K. 2025. Agents Robust to Distribution Shifts Learn Causal World Models Even Under Mediation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Everitt, T.; Carey, R.; Langlois, E. D.; Ortega, P. A.; and Legg, S. 2021. Agent Incentives: A Causal Perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(13): 11487–11495.
- Gao, H.; Xu, T.; Zhang, T.; Guo, Y.; Zhao, C.; Ren, J.; Jiang, Y.; Guo, S.; and Chen, F. 2025. Causal dreamer for partially observable model-based reinforcement learning. *Neurocomputing*, 652: 131012.
- Heckerman, D. 1995. A Bayesian approach to learning causal networks. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, UAI'95, 285–295. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1558603859.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1): 99–134.
- Liang, J.; and Boularias, A. 2021. Inferring Time-delayed Causal Relations in POMDPs from the Principle of Independence of Cause and Mechanism. In Zhou, Z.-H., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 1944–1950. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Osogami, T. 2015. Robust partially observable Markov decision process. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 106–115. Lille, France: PMLR.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. USA: Cambridge University Press, 2nd edition. ISBN 052189560X.
- Pearl, J. 2018. Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM '18*, 3. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355810.
- Pearl, J.; and Bareinboim, E. 2011. Transportability of Causal and Statistical Relations: A Formal Approach. *Proceedings of the AAAI Conference on Artificial Intelligence*, 25(1): 247–254.
- Porta, J. M.; Spaan, M. T. J.; and Vlassis, N. 2005. Robot Planning in Partially Observable Continuous Domains. In *Proceedings of Robotics: Science and Systems*. Cambridge, USA.
- Richens, J.; and Everitt, T. 2024. Robust agents learn causal world models. In *The Twelfth International Conference on Learning Representations*.
- Saghafian, S. 2018. Ambiguous partially observable Markov decision processes: Structural results and applications. *Journal of Economic Theory*, 178: 1–35.
- Schölkopf, B. 2022. *Causality for Machine Learning*, 765–804. New York, NY, USA: Association for Computing Machinery, 1 edition. ISBN 9781450395861.
- Smallwood, R. D.; and Sondik, E. J. 1973. The Optimal Control of Partially Observable Markov Processes over a Finite Horizon. *Operations Research*, 21(5): 1071–1088.
- Sontakke, S. A.; Mehrjou, A.; Itti, L.; and Schölkopf, B. 2021. Causal Curiosity: RL Agents Discovering Self-supervised Experiments for Causal Representation Learning. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 9848–9858. PMLR.