

CAR-DESPOT: Causally-Informed Online POMDP Planning for Robots in Confounded Environments

Ricardo Cannizzaro and Lars Kunze

Abstract—Robots operating in real-world environments must reason about possible outcomes of stochastic actions and make decisions based on partial observations of the true world state. A major challenge for making accurate and robust action predictions is the problem of confounding, which if left untreated can lead to prediction errors. The partially observable Markov decision process (POMDP) is a widely-used framework to model these stochastic and partially-observable decision-making problems. However, due to a lack of explicit causal semantics, POMDP planning methods are prone to confounding bias and thus in the presence of unobserved confounders may produce underperforming policies. This paper presents a novel causally-informed extension of “anytime regularized determinized sparse partially observable tree” (AR-DESPOT), a modern anytime online POMDP planner, using causal modelling and inference to eliminate errors caused by unmeasured confounder variables. We further propose a method to learn offline the partial parameterisation of the causal model for planning, from ground truth model data. We evaluate our methods on a toy problem with an unobserved confounder and show that the learned causal model is highly accurate, while our planning method is more robust to confounding and produces overall higher performing policies than AR-DESPOT.

I. INTRODUCTION

Robots must be capable of making decisions to achieve their tasks efficiently and safely, especially in safety-critical and high-risk applications. This must be done in a manner that is robust to the range of sources of uncertainty typically encountered in real-world environments with stochastic action outcomes, and partial and noisy sensor observations [1].

The partially observable Markov decision process (POMDP) is a popular framework that provides a principled approach for formulating sequential decision-making problems for agents operating in stochastic and partially-observable environments [2]. It has been successfully demonstrated in a wide range of robotics domains [3][4][5].

Despite this, standard POMDPs lack the causal semantics required to articulate system relationships in a causal manner and thus the application of causal analysis is not possible. This is problematic when unobserved confounding between the agent’s action selection and outcome is present in the system. Confounding occurs when there is a common cause of the target cause and effect variables, which introduces a spurious correlation between variables that is entangled with the true causation we are interested in [6]. Without the proper causal treatment, this can lead to prediction errors caused by




Fig. 1: (a) *GridWorldConfounded*. The robot moves in the four cardinal directions in a grid to reach the goal while avoiding occupied cells. Variable sensor interference (orange cell) from an electromagnet is an unobserved confounder. (b) The SCM representation of a UCPOMDP. The agent’s reactive action selection and state updates are both conditionally dependent on the unobserved confounder. Causal interventions (red crosses) are taken on actions to eliminate confounding bias in transition predictions. Observed variables are shown as white nodes, unobserved as grey.

confounding bias. For POMDP models where the confounder is unobservable (UCPOMDPs), it is not possible to control for it in analysis. In these cases standard POMDP planners may generate underperforming policies [7][8].

This has significant implications for robots operating in real-world environments. Sources of confounding can arise from the environment, which system designers may not have anticipated and/or are not observable to the robot. Examples include weather phenomena (e.g., wind), surface conditions (e.g., slippery surfaces), extraneous factors (e.g., magnetic fields, radiation sources), and adversarial agents whom may manipulate conditions to hinder the robot’s progression.

To address this problem, this paper proposes a novel causally-informed extension of the efficient anytime regularized determinized sparse partially observable tree (AR-DESPOT) online planner [9][10], that uses causal modelling and inference to eliminate policy bias errors caused by unmeasured confounding variables. To address the further problem that the system dynamics model may not be fully known for planning *a priori*, we also propose a method to learn offline the partial model parameterisation from data.

We experimentally evaluate our learning and planning methods on a toy problem with an unobserved confounder to demonstrate the benefits of combining casual modelling and inference with POMDP planning. While we are motivated by real-world robotics problems, we start here with a minimal

This work is supported by the Australian Defence Science & Technology Group and the EPSRC RAILS project (grant reference: EP/W011344/1).

The authors are with the Oxford Robotics Institute, Dept. Engineering Science, University of Oxford, UK. ricardo@robots.ox.ac.uk

toy problem representative of an autonomous inspection mission as an initial proof of concept.

The main contributions of this work are: 1) the novel formulation of the CAR-DESPOT algorithm; 2) a quantitative evaluation on a UCPOMDP toy problems; and 3) a demonstration of how components of the causal model can be learned offline from ground truth observational data.

II. RELATED WORK

A. Causality in Robotics

Recent work has identified a growing need for causality as a component of robotics and trustworthy autonomous systems [11][12][13]. While research into causality in robotics is still nascent, it has been proposed to hold great potential benefits such as addressing machine learning domain transfer issues by providing artificial agents an understanding of cause and effect [14]. Further, they can be used to achieve desirable properties of autonomous systems, including explainability, fairness, and accountability [11], which is typically not possible with “black-box” deep learning methods.

Existing works on causal discovery in robotics aim to discover action-outcome causal relationships, including from human demonstrations in table-setting tasks [15] and from simulated robot actions in block stacking tasks [16]. In the former, authors learn temporal co-occurrences to approximate causal relationships, while in the latter authors learn the structure and parameterisation of a causal Bayesian network. In our work, we address a sub-problem of causal discovery: causal parameterisation learning. Further, we represent the system as a structural causal model (SCM) [6], which is a more expressive type of causal model that also permits the more sophisticated counterfactual inference (see III-A).

Causal reasoning has been applied to planning problems in earlier works. One approach uses causal action models to create plans for office delivery jobs [17], while another uses causal reasoning for coordination of multiple household cleaning robots [18]. More recent work combines causal learning and reasoning to explain, predict, and prevent plan failures in robotic block stacking tasks [19]. None of these approaches adopt a principled POMDP/MDP formulation used in our work, nor do they incorporate any probabilistic causal machinery [6] as we do. To the best of our knowledge, we are not aware of any existing online POMDP-based robot planner that utilises causal modelling and inference.

B. POMDP-Based Robot Planning

In POMDP planning problems, the agent must find an optimal (or near-optimal) policy, which maps its belief over world states to the action to take to optimise its long-term reward. The true world state is unknown to the agent. Instead, it must rely on partial observations of the world to update its belief, which is a probability distribution over the possible states, and reason over this belief. POMDPs are extensively used in the probabilistic planning literature; a review of POMDPs in robotics is presented in [1].

Online policy search is a POMDP planning method that scales well to larger POMDPs [1]. Online planning alternates between policy planning and execution, with the agent searching for only the immediate next-best action at each time step. Several online planners have achieved notable success in robot POMDP planning problems, including POMDP-lite [20], Adaptive Belief Tree [21], POMCP [22], and AR-DESPOT [10]. However, none of these POMDP planners use explicitly causal models, thus when used for planning in environments with unobserved confounders, they are prone to confounding bias and thus may generate sub-optimally performing policies [8].

C. Causal Reinforcement Learning

Recent work has begun bridging the gap between reinforcement learning (RL) and causal inference research communities under the umbrella of causal RL [7][8][23][24].

While we adopt a SCM-based (PO)MDP formulation alike these works, they differ from ours in that they address model-based reinforcement learning problems, which involve two sub-problems: 1) learning the transition and reward distributions; and, 2) deciding on the optimal action sequence (i.e., planning). Thus, online RL-based methods inherently require the agent to take experiments in the world to learn beneficial state-observation sequences, which is not always feasible in mobile robotics applications due to safety constraints (e.g., autonomous driving), where robot recovery is not practical (e.g., underwater or planetary robots), or where the damage to the robot and infrastructure is too costly to allow the robot to fail while learning the value of poor action-observation sequences (e.g., industrial inspection). Thus, we focus on the application of causality to planning, in mobile robotics applications for which online experimentation is not suitable.

III. BACKGROUND KNOWLEDGE

A. Causal Modelling and Inference

Causal models encode the causal relationships that describe the data-generation process of a system. Critically, these are causal rather than merely association-based. As such, they encode relationships where the flow of information is strictly in the direction from cause to effect.

Judea Pearl’s *Ladder of Causality* defines a hierarchy of three causal inference types of increasing complexity [25]. The first level, *association*, is the most primitive and relates to the activity of seeing. Queries of this type involve reasoning about how taking an observation of one variable (X) should change our belief in another related variable (Y); e.g., $P(Y|X)$. The second level, *intervention*, relates to the activity of doing; that is taking an action to externally intervene on the system and interrupt the natural causal relationships that govern the data-generation process. This is represented with the *do*(·) operator [6], which defines an intervention to modify a set of variable outcomes in the causal model. Interventional inference queries involve reasoning about how taking an action to set a variable (X) to a fixed value (x) should change our belief in another; e.g., $P(Y|do(X = x))$. The third level, *counterfactual*, relates to

the activity of imagining other worlds with outcomes other than what we have observed (i.e., worlds that are counter to our factual evidence).

Causal Bayesian networks (CBNs) and **structural causal models (SCMs)** are two common generative causal probabilistic graphical model types used in causal modelling and causal inference [6]. For both, causal relationships are represented by an underlying causal directed acyclic graph (DAG). A causal DAG G is defined as $G = \langle V, E \rangle$, where V is the set of nodes, representing random variables in the causal model, and E is the set of directed edges, the presence of which between two variables denoting the existence of a conditional dependency of the child on the parent node.

In the case of CBNs, a causal DAG is combined with conditional probability distributions. SCMs are also constructed from an underlying causal DAG G , but are different to CBNs in that they separate the stochastic elements of a model from the deterministic ones [6]. This permits their use in answering counterfactual queries, which cannot be done with CBNs alone. A structural causal model M is defined as a 4-tuple $M = \langle U, V, F, P(U) \rangle$, where U is a set of exogenous random variables that are determined by factors external to the model, V is a set of endogenous variables, F is a set of deterministic functions used to assign the endogenous variable v_i its value as a function of its graph parents pa_i and corresponding exogenous variable u_i , and $P(U)$ is a probability distribution over the exogenous random variables:

$$u_i \sim P(U_i) \quad v_i := f_i(u_i, pa_i) \quad (1)$$

Through conditioning and intervention operations, CBNs and SCMs can be used to answer observational and interventional inference queries, which in general cannot be expected to coincide. For models with confounding variables, the **intervention** operation disentangles the spurious correlation from the true causation by blocking the influence of the confounding parent variables in the graph. Conversely, the **observational** distribution contains both the true causation and spurious correlation introduced by the confounder and thus it may provide a biased estimate of causal effects.

B. POMDP Planning

A POMDP is defined as a tuple $\langle S, A, T, Z, O, R \rangle$, where S is a set of world states, A is a set of actions available to the agent, T is the state-transition function ($T(s, a, s')$) specifying the probability distribution over arriving in successor state s' after taking action a from current state s , Z is a set of observations the agent can take, O is the observation function specifying the probability distribution ($O(s', a, z)$) over possible observations z made after taking a given action a and arriving in a given state s' , and R is the reward function specifying the immediate reward received by taking an action in a given state and arriving in a successor state.

During policy search and execution, the agent starts with an initial belief b_0 . After taking an action at $t - 1$ the agent receives an observation z_t and updates its belief b_t using Bayes' Rule. A policy $\pi : B \rightarrow A$ defines the action $a \in A$ to take when the agent holds the belief $b \in B$. The value V

of a policy π at a belief b can be defined as the expected total discounted reward for following π given the initial belief b : $V(b) = E[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(b_t)) | b_0 = b]$, given a discount factor $0 < \gamma < 1$ applied to reduce the value of rewards further into the future. Using the *argmax* decision rule, the optimal *greedy* policy π^* is the policy π for which the policy value V_π is maximised for a given initial belief b_0 .

Constructing a belief tree is one method for online planning. By rooting the tree at the current belief b_0 and performing a lookahead search, a policy π can be found that maximises $V_\pi(b_0)$. With each node in the tree representing a belief, a node branches into an edge for each action and branches again for each observation. Trials are conducted by simulating different possible sequences of alternating actions and observations. The policy with the highest expected value can be found by applying the Bellman optimality principle to recursively calculate the maximum value of action branches and average value of observation branches. Completing the belief tree search produces an approximately optimal policy.

C. AR-DESPOT Planner

AR-DESPOT [9][10] is an anytime online POMDP algorithm that addresses key computational challenges with planning under uncertainty, which can in the worst case render POMDP planning computationally intractable [26]. While a standard belief tree includes executions for all possible combinations of actions and observations, a DESPOT contains only a set of randomly sampled *scenarios*. Thus, a search through a DESPOT is more efficient, as it focuses only on a subset of the complete belief tree. By searching through a DESPOT instead of a standard belief tree, AR-DESPOT is able to solve difficult POMDP problems with extremely large observation or state spaces [10].

Since constructing a full DESPOT is not practical for very large POMDPs, the AR-DESPOT algorithm performs anytime heuristic search. By constructing the DESPOT incrementally following a heuristic, AR-DESPOT is able to scale up to very large POMDPs by focusing search on promising areas of the DESPOT. This balance of optimality and efficiency motivates its use in our work.

IV. PROPOSED METHOD FOR MODEL LEARNING AND CAUSAL PLANNING

Our proposed method consists of two components: 1) learning the parameterisation of the SCM model transition function; and, 2) using the learned model and interventional inference to improve POMDP planning performance in the presence of an unobserved confounder. We begin by describing our adopted causal representation of the POMDP model.

A. Causal POMDP Representation

We adopt a SCM-based formulation of POMDPs to permit interventional and counterfactual inference similar to the SCM-based MDP formulation in [8]. Although in this work we focus on interventional inference, this formulation sets up future work to apply counterfactual inference methods.

An **SCM-POMDP** is an SCM augmented with POMDP components, such that the set of endogenous variables V contains the world state $s_k \in S$ at time k ; the agent's choice of action $a_{k+1} \in A$; and successor state s_{k+1} , observation generated $z_{k+1} \in Z$, and immediate reward r_{k+1} received after taking the action. The system's transition and observation probability functions and reward assignment are split into stochastic and deterministic components, respectively encoded by the set of exogenous variables $P(U)$ and assignment functions F .

POMDPs with unobserved confounding (**UCPOMDPs**) are an extension of POMDPs, in which there is a set of unobserved confounding variables U_k that influence both the agent's choice of action A_{k+1} and the successor state S_{k+1} . **SCM-UCPOMDPs** are UCPOMDPs represented as SCMs. This is illustrated in Fig. 1b. Critically, the agent's choice of action is not an independent distribution, as it is in the SCM-POMDP formulation. Instead, it is conditionally dependent on one or more confounding variables. This dependence of the agent's action selection on unobserved confounder variables represents the *reactive* tendencies of the agent; that is, how the agent is likely to act reflexively in response to the external influences. This is in contrast to the *deliberative* decision-making process of the agent, in which actions are made independent of external influences. The exogenous variable distributions $P(U)$ and assignment functions F for a SCM-UCPOMDP (Eq 1) follow the structure of the SCM-UCPOMDP causal DAG shown in Fig 1b.

We represent the transition function under the agent's reactive behaviour using the observational distribution $P(S'|A, S)$. Here, the back-door path introduced by the confounder means that the observational distribution contains the spurious correlation through $A \leftarrow U \rightarrow S'$ in addition to causal influence directly through $A \rightarrow S'$. In the case of deliberative decision-making, we use causal semantics to represent the transition function using the interventional distribution $P(S'|do(A = a), S)$. Here, the *do*-operator is used to surgically remove the influence of the confounding variables on the agent's choice of action, thus retaining only the causal influence $A \rightarrow S'$.

B. Model Parameter Learning Method

Our proposed method aims to learn the parameterisation of an SCM-UCPOMDP model from a ground truth model in a privileged setting. We are motivated by real-world robotics applications where the full system dynamics model is not fully known *a priori*, but instead needs to be created using a hybrid method combining domain knowledge and observational data in an offline setting prior to robot deployment.

We define a ground truth SCM-UCPOMDP model M that serves as a proxy for the real-world system dynamics. During learning, we wish to obtain a model \hat{M} that learns the relative transition probabilities from a data set generated by observing the robot operate under the dynamics given by M . We assume sufficient domain knowledge is available to identify the existence of causal relationships between action, confounder, and state variables. Thus we assume the model

structure is known and seek to learn the conditional probabilities. We further assume the transition function can be represented as a combination of a deterministic assignment function $f_{s'}(s, \Delta s)$ of current state s and relative state transition Δs , and a relative transition probability $P(\Delta S|A, U)$. We further assume the transition function can be partitioned into two components, one for when the agent is in a state $s \in S_{UC}$ where it is under the influence of confounding (P_{UC}) and the other where the agent is not (P_0), such that:

$$P(S'|S, A, U) = \begin{cases} P_{UC}(S'|\Delta S, A, U) & \text{for } s \in S_{UC} \\ P_0(S'|\Delta S, A) & \text{otherwise,} \end{cases} \quad (2)$$

and we learn the relative transition functions for each case: $P_{UC}(\Delta S|A, U)$ and $P_0(\Delta S|A)$. In the context of the *Grid-WorldConfounded* problem, P_{UC} describes the transition probabilities when the robot is located in the orange cell and thus affected by the electromagnet (Figure 1a), while P_0 describes the transition probabilities when located in the other cells and thus not affected.

We further assume a two-phase data-collection approach, common in robot industrial inspection applications: an initial site survey with humans in the loop to collect observational robot data followed by autonomous robot deployment. In this setting we argue that it is reasonable to assume that the robot-and-human site survey team is equipped with sensors to observe all ground truth model variables in a privileged setting, including the confounder that is unobserved during autonomous deployment. Given this, we also seek to learn a parameterisation for $P(U)$, the independent distribution governing the probabilities of the confounding variable. We assume a sensor model is known, specifying $P(Z|S')$, and assume a known deterministic reward function specified by system designers. Thus, we do not learn these probabilities.

To begin the learning process, we sample a data set D of N number of observations from the ground truth model M , represented as a probabilistic program in the Python-based Pyro probabilistic programming language (PPL) [27]. Each observation is a record of $\langle UC, U, A, \Delta S \rangle$, where UC is a Boolean value indicating whether the agent was under the influence of the confounder when selecting its choice of action. Next, we define a learning model \hat{M} with the same causal structure as M , but with unknown parameterisation of $P(U)$, $P_{UC}(\Delta S|A, U)$, and $P_0(\Delta S|A)$ each represented as a categorical distribution. Finally, we learn the parameterisation of these three distributions by fitting \hat{M} to the data set D through automated stochastic variational inference (SVI) [28], [29] in Pyro PPL. We use the Adam optimizer [30] in PyTorch, and *TraceEnum_ELBO* which is a trace implementation of evidence lower bound (ELBO) based SVI in Pyro. We note that the confounder does not bias the learning method because the confounder is controlled for using the privileged observations of it.

C. Causally-Informed POMDP Planning Method

Here we describe CAR-DESPOT, a novel causally-informed extension of AR-DESPOT that combines the on-line planning performance of AR-DESPOT with the causal

expressivity of SCMs to eliminate policy errors due to confounding bias.

Since the SCM-UCPOMDPs subsume standard POMDP models, they retain properties of POMDPs and therefore can be used in their place. Thus, in CAR-DESPOT we formulate models as SCM-UCPOMDPs to take advantage of the existing POMDP-based search machinery in AR-DESPOT to find approximately optimal policies.

However, we make a novel and fundamental change to how the transition function is sampled during exploration rollouts. While AR-DESPOT samples from observational transition distribution $P(S'|S, A)$, which is prone to errors from unobserved confounders, CAR-DESPOT samples from the interventional distribution $P(S'|S, do(A = a))$ to avoid bias errors. We anticipate this will improve the planner performance in two ways: 1) it will avoid the heuristic incorrectly guiding the incremental DESPOT construction towards exploring less promising areas of the tree; and, 2) it will prevent the planner, given a constructed DESPOT, from incorrectly valuing discovered action-observation outcomes, thus leading to sub-optimal policies. The same interventional modification is applied to computing the POMDP observation function probability distribution. However, for models for which the observation is conditional only on the successor state, such as the toy model considered here, $P(Z|S', A)$ and $P(Z|S', do(A = a))$ are equivalent.

We represent SCM-UCPOMDP models as causal generative probabilistic programs in Pyro PPL and use Importance Sampling inference [31] to compute interventional transition posterior distributions, used to inform planner rollouts. We implement CAR-DESPOT by extending the C++ implementation of AR-DESPOT¹ provided by authors in [10].

V. EXPERIMENTS

We evaluate our proposed model learning and causally-informed POMDP planning methods on a minimal toy problem inspired by the autonomous robot inspection domain.

A. Problem Description

GridWorldConfounded (Fig. 1a) is a modified version of the *GridWorld* problem [32] with an unobserved confounder, the presence of which is known to exist *a priori*. The robot is located on a 2-dimensional grid and its aim is to navigate to a goal location to inspect an important asset as part of an industrial site inspection mission, while avoiding a collision with occupied cells. The robot always begins at the bottom-left grid cell (0, 0) and knows the position of the goal which is always (0, 3). The robot can move one grid cell in each of the four cardinal directions. Due to actuator errors, when executing its movement action there is a small chance the robot may drift and end up either to the left or right (with respect to the direction of travel) with probability of 5% in either direction. The robot receives a noiseless observation of its position after each action.

Critically, an electromagnet is located in cell (1, 2) and affects the robot’s orientation sensor measurements and action

selection when in the magnet’s region of influence, located to the left in (0, 2). The orientation error caused by the magnet $U_{OrientationError}$ fluctuates randomly over time and is not able to be perceived by the robot due to lack of a suitable sensor. At each time step, the orientation error is sampled from an independent categorical distribution over discrete values $\{-90^\circ, 0^\circ, +90^\circ\}$ with probabilities $[0.10, 0.80, 0.10]$.

The unobserved confounder influences the robot in two ways. First, when the robot is at (0, 2), the robot’s reaction is to take an action with the probabilities given shown in Table I. These probabilities reflect the reflexive rules internal to the robot’s reactive decision-making; here they are chosen to demonstrate the effects of confounding. The second effect of the confounding variable is on the stochastic execution of the movement action. If under the influence of the magnet, the robot will attempt to move in its intended direction but because of orientation error, it will actually move with an offset given by the orientation influence. For example, when taking the *RIGHT* action (0° heading in grid frame) given a $+90^\circ$ orientation error offset, the robot will instead move up. The drift dynamics are applied after the effects of confounding. When the robot is outside of the region of influence, only the drift transition probabilities apply.

TABLE I: The reactive action selection probability distribution for the *GridWorldConfounded* problem, conditional on the unobserved confounding variable $U_{OrientationError}$.

$P(A U_{OrientationError})$	-90°	0°	$+90^\circ$
<i>RIGHT</i>	0.05	0.45	0.05
<i>UP</i>	0.85	0.05	0.85
<i>LEFT</i>	0.05	0.45	0.05
<i>DOWN</i>	0.05	0.05	0.05

Each movement action gives a base reward of -1 . After an action, if the robot is at the goal it receives $+100$ reward and the scenario terminates, else if it is in collision it receives -50 reward and terminates; otherwise the problem continues.

While this is a simplified and highly abstract problem, the magnet acts as a demonstration of how unobserved confounders can feasibly arise in practice during robot deployments in complex and uncontrolled real-world environments.

B. Model Parameter Learning

We generate a data set of size $N = 800,000$ by sampling from the ground truth model M and use this as an input to SVI training. We use a learning rate $lr = 0.01$ and 5001 training steps to obtain the learned model \hat{M} , including the parameterisation of: $P_{UC}(\Delta Position, A, U_{OrientationError})$, $P_0(\Delta Position, A)$, and $P(U_{OrientationError})$. We learn 106 latent model parameters, using uniform Dirichlet priors for categorical distribution parameters in our learning model and the *AutoDelta* automatic guide for the SVI guide. We report the final ELBO-loss, which based on the Kullback–Leibler (KL) divergence, and the KL divergence of the learned full transition probability distribution $\hat{P}(S'|S, A, U)$.

¹AR-DESPOT - <https://github.com/AdaCompNUS/despot>

C. Planner Evaluation

We evaluate CAR-DESPOT on the GridWorldConfounded toy problem and compare it with a baseline method OAR-DESPOT, a version of AR-DESPOT modified to be compatible with SCM-UCPOMDP model representation but using observational transition quantities as done in the original algorithm. To assess the impact of our proposed learning method on planning performance, we evaluate each method in 2 cases, in which: 1) the learned model \hat{M} is used for planning but the ground truth model M is used for policy execution as a proxy to the real-world; and, 2) the ground truth model is used for both planning and policy execution.

The performance is assessed according to the average total discounted reward over 100 independent online planner executions. We define the maximum scenario length, maximum search depth, and policy simulation depth as 15 steps. A search budget of 150 seconds is given to the planners for each time step. This was chosen empirically to allow the planners to sample a sufficient number of trials during the incremental DESPOT creation to sample the low-probability transition outcomes leading to failure in the toy problem. The default discount factor $\gamma = 0.95$ and target gap ratio $\epsilon = 0.95$ is used, with regularisation pruning constant $\alpha = 0.01$. The number of scenarios is 500, a sufficient amount to achieve performance in similar problems in [9]. To avoid penalising the first step of online planning, all required transition inference queries are computed and memoized prior to planning. Observation function inference queries are computed as-needed during belief updates but cached for future re-use. 5000 particles are used in Importance Sampling for both inference types; this was chosen empirically to accurately infer probabilities for low-probability outcomes. All planner evaluations are performed on a desktop PC with Ubuntu 20.04 operating system, using CPU-only, on an AMD Ryzen 9 5900x 12-core processor with 128 GB of RAM.

VI. RESULTS AND DISCUSSION

A. Analysis of learned model

We begin by presenting and analysing the partially-learned SCM-UCPOMDP model of the GridWorldConfounded problem. SVI training yields an ELBO-loss of 3.300M (4.125 per data point) and produces interventional relative distribution probabilities each with an absolute error of less than 0.01 from ground truth values. Further, the probabilities for $U_{OrientationError}$, $[0.10, 0.80, 0.10]$, are accurately recovered during training: $[0.1004, 0.7997, 0.0999]$. The KL divergence of the learned full transition probability distribution $\hat{P}(S'|S, A, U)$ from ground truth distribution $P(S'|S, A, U)$ is $D_{KL} = 0.0021$, indicating a very close fit.

Fig 2 shows the learned parameters produce calculated **interventional** transition probabilities and inferred **observational** relative transition probabilities that are extremely close to the ground truth probabilities, when under the influence of confounding. Observational transition probabilities are inferred using Importance Sampling. These relative transition probabilities are provided as exemplar inference results; the




Fig. 2: Interventional and observational transition probabilities of ground truth and learned models, under the influence of confounding, as a function of action choice. Formulating actions as causal interventions removes confounding bias errors present in the observational transition predictions.

full (i.e., non-relative) transition inference results used for planning are computed immediately prior to plan time.

All non-zero probability outcomes are captured by the learned model, even very low probability events (e.g., $p < 0.001$), demonstrating that both the drift and magnet interference transitions dynamics are recovered accurately by the learning method. The non-confounded observation and interventional transition probabilities are also recovered accurately. Due to space constraints these results are not shown as they are a subset of the confounded transition dynamics.

Critically, the learned model is able to capture the difference between the observational and interventional transition dynamics when under confounding. Comparing the probabilities shown in Fig 2, the learned model correctly captures the probability mass shift due to confounding bias in the observational transition dynamics. This is most evident for the up action, for which the probability mass of outcome $(0, +1)$ is redistributed to outcomes $(-1, 0)$ and $(+1, 0)$; this causes the observational distribution to significantly underestimate the interventional probability: 0.1914 instead of 0.7229. We expect that when the learned model is used for planning with OAR-DESPOT, the up action will be undervalued when under the effects of confounding. We expect this when using the ground truth model also.

B. Analysis of planning performance

We present and discuss the planner evaluation results for the GridWorldConfounded problem. The distribution of policy performance over 100 runs of each planner and model combination is given in Fig 3. We report the total discounted reward of policies executed on the ground truth model. These

TABLE II: Mean total discounted rewards per model and algorithm combination. Standard errors are also given.

	Learned Model	Ground Truth Model
OAR-DESPOT	0.74 ± 7.73	10.81 ± 8.53
CAR-DESPOT	21.93 ± 9.21	15.54 ± 9.34




Fig. 3: Distribution of CAR-DESPOT and baseline planner performance, with learned and ground truth models, according to total discounted reward. Policies from CAR-DESPOT perform better overall due to eliminating confounding bias.

performances are summarised in Table II.

The performance distributions of all planner-model combinations are largely bimodal, with rewards either clustered around the lower end of the reward range or upper end, corresponding to task failure by colliding with a wall or task success by reaching the goal respectively. Notably, for OAR-DESPOT there is also a small third peak around -10 corresponding to task failure by means of exhausting the maximum allowable number of steps. This occurred in 4 runs for learned model and 1 run for ground truth, in which the OAR-DESPOT-generated policies caused oscillatory behaviour that prevented reaching the goal within the action budget. This did not occur in any runs for CAR-DESPOT.

For the **learned model**, the task-success peak is more concentrated for CAR-DESPOT than OAR-DESPOT and has a higher mode. The reason for this is that since there is no confounding bias in the inferred interventional transition dynamics when the robot is at $(0, 2)$ and thus under the influence of the confounder, CAR-DESPOT correctly estimates the value of taking actions to move along the left-hand side of the grid, which is shorter and therefore

results in a higher score due to fewer movement penalties and cumulative rewards being discounted over a shorter time horizon. Thus, in all 100 runs CAR-DESPOT chooses to take the left-hand path. There is also another benefit of taking the shorter path: the cumulative probability of accidentally drifting into a wall resulting in task failure is lower due to the fewer steps taken. Thus, more runs achieve task success: 53 for CAR-DESPOT and 42 for OAR-DESPOT; this further increases the mean performance. Additionally, 39 runs achieve the maximum score. The lower tail end of the task-success peak reflects a small number of runs which exhibited minor oscillatory behaviour before reaching the goal. The task-failure peak is again more concentrated but with a lower mode. These reflect runs in which the outcomes of the correctly chosen actions by chance resulted in unintended outcomes, due either to the drift dynamics or the confounding variable influencing the outcome, leading to task failure. Since the CAR-DESPOT policies choose the left-hand path, failures occur earlier in the run and thus are discounted less.

Conversely, OAR-DESPOT underestimates the value of taking the left-hand path due to the confounding bias in the inferred transition probabilities for taking the up action at $(0, 2)$ when under the influence of the confounder. Because of this, it has a strong preference to take the right-hand path to the goal, which is free from confounding but has a longer action cost and higher cumulative chance of failure due to drift. This yields a lower total discounted score. Interestingly, in 4 of the runs OAR-DESPOT takes the left-hand path and achieves the maximum possible score. This is likely due the incremental DESPOT construction missing either: 1) the failure outcomes at the confounder influence location; or, 2) the success outcome for taking the right-hand path. This demonstrates that CAR-DESPOT is able to generate overall better performing policies than OAR-DESPOT.

Similar results are observed for the **ground truth model**. The CAR-DESPOT task-success peak is again narrower and concentrated around a higher value than that of OAR-DESPOT, while the task-failure peak is narrower and concentrated around a lower value, for the same reasons identified previously. The number of task-success runs using the ground truth model with CAR-DESPOT is slightly lower compared to the learned model case: 48 compared to 53. This slight difference in frequency-based counts is expected due to the stochastic nature of the problem — even executing an optimal policy can occasionally and randomly result in task failure. However, we expect that these task-success counts would converge in the limit of performing more evaluation runs. Notwithstanding this, these results demonstrate that, when using the ground truth model, planning with CAR-DESPOT produces a better overall performance than OAR-DESPOT. Further, the similar results obtained from use of learned and ground truth models for planning further demonstrates the ability of our learning method to learn the relative transition dynamics accurately.

C. Limitations and Improvements

One limitation of our model learning method is the assumption that transition dynamics can be decomposed into a relative transition function and a known assignment function of the current state and relative state change. While this is suitable for the robot 2D position state representation in our toy problem, it may not be suitable for other state representations. In these cases, the full transition function of size $|S| \times |A| \times |S|$ will need to be learned, which will require a larger data set and an increased number of training steps.

We have evaluated our proposed method on a simple toy POMDP model with a relatively small state size (12), action space (4), and observation space (30). Despite this, the search time required to compute a good-performing policy online at each step is found to be 150s, which is prohibitively high for a mobile robot operating in real-time.

Finally, analysis in [8] has shown that for MDP-based problems with unobserved confounders, optimising policies based on the interventional transition function is not always guaranteed to produce optimal policies. Rather, policy optimisation using counterfactual inference produces a dominant policy. In this approach, the agent's reactive action selection is taken into account, which provides information about the state of the hidden confounder. We expect that this improvement will translate to our work in POMDP planning.

VII. CONCLUSION

In this paper we presented a novel causally-informed extension of the AR-DESPOT online POMDP planner, using causal modelling and inference to eliminate confounding bias from unmeasured confounding variables. Through experimental validation on a toy POMDP problem with unobserved confounding, we demonstrated that our method outperforms the baseline method. Using SCMs, we integrated causal reasoning into the POMDP formulation to enable robust POMDP-based planning for robots operating in real-world confounded environments.

REFERENCES

- [1] H. Kurniawati, "Partially observable markov decision processes and robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 253–277, 2022.
- [2] R. D. Smallwood and E. J. Sondik, "The optimal control of partially observable markov processes over a finite horizon," *Operations research*, vol. 21, pp. 1071–1088, 1973.
- [3] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online pomdp planning for autonomous driving in a crowd," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 454–460, 2015.
- [4] M. Budd, P. Duckworth, N. Hawes, and B. Lacerda, "Bayesian reinforcement learning for single-episode missions in partially unknown environments," in *6th Conference on Robot Learning (CoRL 2022)*, OpenReview, 2022.
- [5] A. Goldhoorn, A. Garrell, R. Alquézar, and A. Sanfeliu, "Searching and tracking people with cooperative mobile robots," *Autonomous robots*, vol. 42, pp. 739–759, 2018.
- [6] J. Pearl, *Causality: Models, reasoning, and inference, second edition*. Cambridge university press, 2009.
- [7] E. Bareinboim, A. Forney, and J. Pearl, "Bandits with unobserved confounders: A causal approach," in *Advances in Neural Information Processing Systems*, pp. 1342–1350, 2015.
- [8] J. Zhang and E. Bareinboim, "Markov decision processes with unobserved confounders: A causal approach," tech. rep., Technical Report R-23, Purdue AI Lab, 2016.
- [9] N. Ye, A. Somani, D. Hsu, and W. S. Lee, "Despot: Online pomdp planning with regularization," *Journal of Artificial Intelligence Research*, vol. 58, pp. 231–266, 2017.
- [10] A. Somani, N. Ye, D. Hsu, and W. S. Lee, "Despot: Online pomdp planning with regularization," in *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [11] N. Ganguly, D. Fazlija, M. Badar, M. Fischella, S. Sikdar, J. Schrader, J. Wallat, K. Rudra, M. Koubarakis, G. K. Patro, W. Z. E. Amri, and W. Nejdl, "A review of the role of causality in developing trustworthy ai systems," 2023.
- [12] T. Hellström, "The relevance of causation in robotics: A review, categorization, and analysis," *Paladyn, Journal of Behavioral Robotics*, vol. 12, pp. 238–255, 2021.
- [13] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and Brain Sciences*, vol. 40, p. e253, 11 2017.
- [14] J. Pearl, "The seven tools of causal inference, with reflections on machine learning," *Commun. ACM*, vol. 62, pp. 54–60, 2 2019.
- [15] C. Uhde, N. Berberich, K. Ramirez-Amaro, and G. Cheng, "The robot as scientist: Using mental simulation to test causal hypotheses extracted from human activities in virtual reality," in *IEEE International Conference on Intelligent Robots and Systems*, pp. 8081–8086, IEEE, 2020.
- [16] M. Diehl and K. Ramirez-Amaro, "Why did i fail? a causal-based method to find explanations for robot failures," *IEEE Robotics and Automation Letters*, vol. 7, pp. 8925–8932, 2022.
- [17] M. Beetz and H. Grosskreutz, "Causal Models of Mobile Service Robot Behavior," in *Proceedings of AIPS*, pp. 163–170, Linköping University Electronic Press, 1998.
- [18] E. Aker, A. Erdogan, E. Erdem, and V. Patoglu, "Causal reasoning for planning and coordination of multiple housekeeping robots," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6645 LNAI, pp. 311–316, Springer, 2011.
- [19] M. Diehl and K. Ramirez-Amaro, "A causal-based approach to explain, predict and prevent failures in robotic tasks," *Robotics and Autonomous Systems*, vol. 162, p. 104376, 2023.
- [20] M. Chen, E. Frazzoli, D. Hsu, and W. S. Lee, "Pomdp-lite for robust robot planning under uncertainty," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5427–5433, 5 2016.
- [21] H. Kurniawati and V. Yadav, *An Online POMDP Solver for Uncertainty Planning in Dynamic Environment*, pp. 611–629. Cham: Springer International Publishing, 2016.
- [22] D. Silver and J. Veness, "Monte-carlo planning in large pomdps," *Advances in neural information processing systems*, vol. 23, 2010.
- [23] P. A. Ortega, M. Kunesch, G. Delétang, T. Genewein, J. Grau-Moya, J. Veness, J. Buchli, J. Degraeve, B. Piot, J. Pérolat, T. Everitt, C. Tallec, E. Parisotto, T. Erez, Y. Chen, S. E. Reed, M. Hutter, N. de Freitas, and S. Legg, "Shaking the foundations: delusions in sequence models for interaction and control," *CoRR*, vol. abs/2110.10819, 2021.
- [24] M. Gasse, D. Grasset, G. Gaudron, and P.-Y. Oudeyer, "Causal reinforcement learning using observational and interventional data," 2021.
- [25] J. Pearl and D. Mackenzie, *The book of why: the new science of cause and effect*. Basic books, 2018.
- [26] O. Madani, S. Hanks, and A. Condon, "On the undecidability of probabilistic planning and infinite-horizon partially observable markov decision problems," in *AAAI/IAAI*, pp. 541–548, 1999.
- [27] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, "Pyro: Deep universal probabilistic programming," *Journal of Machine Learning Research*, vol. 20, 2019.
- [28] D. Wingate and T. Weber, "Automated variational inference in probabilistic programming," 1 2013.
- [29] R. Ranganath, S. Gerrish, and D. M. Blei, "Black box variational inference," *Journal of machine learning research*, vol. 33, pp. 814–822, 2014.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference for Learning Representations*, 2015.

- [31] T. Klock and H. K. van Dijk, "Bayesian estimates of equation system parameters: An application of integration by monte carlo," *Econometrica*, vol. 46, no. 1, pp. 1–19, 1978.
- [32] R. S. Sutton, "Integrated architectures for learning, planning, and reacting based on approximating dynamic programming," in *Proceedings of the 7th International Conference on Machine Learning (1990)*, pp. 216–224, Elsevier, 1990.