

Transfer Reinforcement Learning under Unobserved Contextual Information

Yan Zhang and Michael M. Zavlanos

Department of Mechanical Engineering and Materials Science, Duke University, USA
{yan.zhang2, michael.zavlanos}@duke.edu

Abstract—In this paper, we study a transfer reinforcement learning problem where the state transitions and rewards are affected by the environmental context. Specifically, we consider a demonstrator agent that has access to a context-aware policy and can generate transition and reward data based on that policy. These data constitute the experience of the demonstrator. Then, the goal is to transfer this experience, excluding the underlying contextual information, to a learner agent that does not have access to the environmental context, so that they can learn a control policy using fewer samples. It is well known that, disregarding the causal effect of the contextual information, can introduce bias in the transition and reward models estimated by the learner, resulting in a learned suboptimal policy. To address this challenge, in this paper, we develop a method to obtain causal bounds on the transition and reward functions using the demonstrator’s data, which we then use to obtain causal bounds on the value functions. Using these value function bounds, we propose new Q learning and UCB- Q learning algorithms that converge to the true value function without bias. We provide numerical experiments for robot motion planning problems that validate the proposed value function bounds and demonstrate that the proposed algorithms can effectively make use of the data from the demonstrator to accelerate the learning process of the learner.

Keywords—Causal inference; transfer learning; reinforcement learning; causal bounds

I. INTRODUCTION

Reinforcement learning (RL) methods have been widely used to solve sequential decision making problems in unknown stochastic environments [1]. Often, these methods require many samples from the environment to find an optimal control policy, which becomes impractical when they are used to control physical systems, such as vehicles or robots, for which sampling can be expensive. One solution to this problem is to use Transfer Learning (TL) [2]. The goal in TL is to use the experience gained from solving prior *source tasks* to learn a policy for a new related *target task* using much fewer samples than if this new policy was learned from the beginning.

TL methods can be classified depending on what specific knowledge is transferred from the source tasks to the target task. For example, the methods proposed in [3]–[6] focus on transferring value functions, transition functions, and policy functions between the source and target tasks. Specifically,

the key idea in [3] is to transfer the value functions from prior tasks to a new task to provide a good initialization for the learning algorithm. On the other hand, in [4] it is assumed that the source and target tasks share the same transition functions and only differ in the rewards, so that the transition features learnt from the source tasks can be transferred to the new task. In [5], a meta learning method is proposed to find a good initial policy for a new task, using experience from prior source tasks, which can be adapted to obtain the optimal policy using only a few samples. Along these lines, [6] assumes that the source and target tasks share the same hierarchical structure so that the skills learned to solve a low level problem can be directly transferred to the target task.

In contrast to the methods proposed in [3]–[6], the approaches discussed in [7]–[11] directly transfer the transition and reward samples from the source tasks to the target task. As discussed in [9], transferring samples directly is more general compared to transferring value functions, transition functions, and policy functions, since it does not impose any requirements on the specific RL algorithm used to solve the source tasks. The reason is that every RL algorithm relies on transition and reward samples collected from the environment. Specifically, the methods proposed in [7]–[9] augment the target task dataset with the samples from the source tasks so that fewer new samples are needed from the target task. Since the source and target tasks usually have different reward and transition distributions, transferring the source task data to the target task introduces bias in the estimated value functions or policies [8], [12]. The methods proposed in [10], [11] rely on importance sampling to correct the bias introduced by using samples from source tasks. Since unbiased importance weights can result in large variance, biased weights are usually employed in practice to reduce the variance. Nevertheless, the computation of importance weights is expensive and requires fitting regression models for the rewards and transitions, as shown in [10], [11]. This introduces additional bias when the chosen regression model does not perfectly fit the true reward and transition functions.

In this paper, we propose a new TL framework that, as in [7]–[11], relies on sample transfer between source and target tasks but, unlike [7]–[11], it does not introduce bias in the learnt value functions and policies. Specifically, we assume the presence of context in the environment, and consider a

contextual Markov Decision Process (MDP) [13], where the transition and reward functions are affected by this context that is subject to a fixed distribution. Moreover, we assume a demonstrator agent, who can observe the environmental context and can collect transition and reward samples using a context-aware policy. These samples, excluding the contextual information, are provided to a learner agent as the source task dataset. The target task is to let the learner agent, who cannot observe the context, find the optimal context-unaware policy in the same contextual MDP environment. In this problem, different contexts can be thought of as different source tasks. Compared to [7]–[11], since the context is hidden, the problem considered here is more challenging because the learner cannot identify which source task the transition and reward samples belong to. Therefore, the methods in [7], [8], [10], [11] cannot be directly applied. Furthermore, since the source datasets are provided without contextual information, transferring these data using model-based RL methods like [9] can cause significant bias. This is because hidden contexts make the reward and transition models non-identifiable, as it is well known in the causal inference literature [14]. To remove bias from the estimated transition and reward models and enable the learner agent to learn an optimal policy using fewer new samples, we extend the discretization method developed in Section 8.2 in [14] so that we can derive causal bounds on the transition and reward functions using the source dataset. Then, given these causal bounds, we formulate two optimization problems that can be efficiently solved to obtain upper and lower bounds on the value functions, respectively. Using these value function bounds, we develop Causal Bound Constrained Q learning [15] and UCB- Q learning [16]–[18] algorithms that converge to the true value function without bias, unlike the methods in [7]–[11]. This is because, the true value functions are shown to lie within the proposed causal bounds and, therefore, projection of the value function iterates on the intervals defined by these bounds does not affect convergence of the Q learning and UCB- Q learning methods which are known to be unbiased. The role of the causal bound constraints is to prevent these learning algorithms from returning poor value function estimates or exploring unnecessary states at the beginning of the learning process, reducing in this way the sampling complexity of our method.

To the best of our knowledge, the most relevant work to the method proposed here is [19]. Specifically, in [19] a multi-arm bandit problem is considered with binary action and reward spaces. Moreover, a similar method to the one proposed here is employed to compute causal bounds on the expected rewards of the bandit machines, which are used to improve on the performance of the UCB algorithm. Compared to [19], here we consider a sequential decision making problem on arbitrary discrete and finite action and reward spaces, which is more general. Moreover, we assume that the transition functions are also affected by the context

and propose additional causal bounds for them. Note that the learning problem for the target task considered in this paper can be also thought of as a Partial-Observable Markov Decision Process (POMDP) problem. Therefore, the learner can apply any RL algorithm used for POMDPs, e.g., [20], to solve the problem by itself. Nevertheless, these methods are sample inefficient. Instead, here we use the source dataset given by the demonstrator to reduce the number of samples needed by the learner to find the optimal policy.

The rest of this paper is organized as follows. In Section II, we define the proposed TL problem and discuss modeling bias due to the hidden context at the learner agent. In Section III, we develop our proposed method to compute the causal bounds on the reward, transition, and value functions. In Section IV, we propose a Q learning and a UCB- Q learning algorithm based on the value function bounds. In Section V, we numerically validate the proposed TL algorithms and causal bounds on robot motion planning problems. Finally, in Section VI, we conclude the paper.

II. PRELIMINARIES AND PROBLEM DEFINITION

In this section, we provide formal definitions of Markov Decision Processes (MDP) and contextual MDPs. We also formulate the problem of transferring experience from contextual MDPs to standard MDPs where the context is not observable, and discuss the challenges that arise in doing so.

A. MDPs and Contextual MDPs

Consider a MDP defined as a 4-tuple $(s_t, a_t, P(s_{t+1}|s_t, a_t), R(s_t, a_t, s_{t+1}))$, where $s_t \in \mathcal{S}$ and $a_t \in \mathcal{A}$ denote the state and action at time t , $P(s_{t+1}|s_t, a_t)$ is the transition probability from state s_t to s_{t+1} when taking action a_t , and $R(s_t, a_t, s_{t+1})$ is the reward received when action a_t is taken at state s_t to transition state s_{t+1} . In what follows, we assume that the state and action spaces, \mathcal{S} and \mathcal{A} , respectively, are finite and discrete. Moreover, we define a policy function $\pi(a_t|s_t) \rightarrow [0, 1]$ as the probability of choosing action a_t at state s_t . A run of the MDP of length T that is generated according to policy π is called an episode. The goal is to find the optimal policy π^* so that the discounted accumulated reward

$$\mathbb{E}_{\rho^0} \left[\sum_{t=0}^T \gamma^t R(s_t, a_t, s_{t+1}) \right]$$

is maximized, where $\gamma \in (0, 1)$ is the discount factor and ρ^0 is the distribution of the agent’s initial state. Furthermore, we define the state-based value function at state s , $V(s) := \mathbb{E} \left[\sum_{t=0}^T \gamma^t R(s_t, a_t, s_{t+1}) | s_0 = s \right]$, and the action-based value function at state-action pair (s, a) , $Q(s, a) := \mathbb{E} \left[\sum_{t=0}^T \gamma^t R(s_t, a_t, s_{t+1}) | s_0 = s, a_0 = a \right]$. Finding the optimal policy π^* is equivalent to computing the above two value functions under the optimal policy.

Given a MDP, a contextual MDP is defined as a 5-tuple $(s_t, a_t, P^u(s_{t+1}|s_t, a_t), R^u(s_t, a_t, s_{t+1}), \rho(u))$, with transition function $P^u(s_{t+1}|s_t, a_t)$ and reward function $R^u(s_t, a_t, s_{t+1})$ inherited from the original MDP but parameterized by a contextual variable $u \in \mathcal{U}$ [13]. The context variable u is sampled from a distribution $\rho(u)$ at the beginning of an episode and remains fixed for the duration of that episode. The contextual parametrization of the transition probabilities and reward functions in a contextual MDP is motivated by many practical applications. For example, the motion of a robot under the same control action is affected by contextual information related to the environmental conditions, e.g., lighting and wind conditions. Similarly, a patient’s health status under the same treatment is affected by contextual information related to the patient’s profile, e.g., their age, gender, or weight. Given a contextual MDP problem, the goal is to find a parameterized policy function $\pi^*(a_t|s_t, u)$ that maximizes the contextual accumulated reward

$$\mathbb{E}_{\rho^0} \left[\sum_{t=0}^{\infty} \gamma^t R^u(s_t, a_t, s_{t+1}) \right]$$

for every $u \in \mathcal{U}$.

B. The Experience Transfer Problem

In what follows, we consider the problem of transferring the experience from a demonstrator agent who is aware of the contextual information included in the environment to another learner agent who does not have access to this information. By transferring experience we are able to use much fewer data to learn optimal control policies in RL for the agent with no contextual information.

Specifically, consider a contextual MDP $(s_t, a_t, P^u(s_{t+1}|s_t, a_t), R^u(s_t, a_t, s_{t+1}), \rho(u))$ and assume that a demonstrator agent can observe the context u and has knowledge of the optimal contextual policy $\pi^*(a_t|s_t, u)$. We model the experience of the demonstrator as state transition and reward pairs (s_t, a_t, s_{t+1}, r_t) that can be obtained by executing the optimal policy $\pi^*(a_t|s_t, u)$ as follows. At the beginning of episode k , the environment randomly samples a contextual variable u_k from the distribution $\rho(u)$ and a starting state s_0 from the distribution ρ^0 . The context u_k is revealed to the demonstrator. Then, the demonstrator executes the optimal policy $\pi^*(a_t|s_t, u_k)$ for T time steps and records the state transitions and rewards (s_t, a_t, s_{t+1}, r_t) , excluding information about the context. Upon termination of this episode, a new episode is initialized and the process is repeated for K episodes. In the end, the experience of the demonstrator is collected in the set $\{(s_t, a_t, s_{t+1}, r_t)\}_{t=0:KT}$, which we call the observational data. Note that we do not require the samples of the demonstrator to be provided in the form of trajectories. Moreover, samples from different contexts need not be differentiated from each other. Compared to [7], [8], [10], [11] where the source samples should be

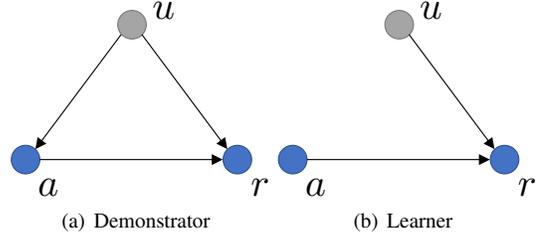


Figure 1. Causal graphs to represent the behavior of the demonstrator and the learner. In both plots, the grey nodes represent the unobserved context u , and the blue nodes represent the action taken and the reward received at certain state. The arrow represents the direction of the causal effect. At each arrow, the end node’s value is influenced by the source node’s value.

arranged according to different tasks or trajectories, we only need to know the joint distribution $P(s_t, a_t, s_{t+1}, r_t)$ in the source samples, which imposes minimal restrictions on the source data. This is important in applications where the demonstrators are not allowed to directly operate in the environment but only respond to independent samples of state and context pairs.

Together with the demonstrator, we also consider a learner agent who interacts with the same contextual environment but cannot observe u . In this case, the goal of the learner is to find the optimal policy $\pi^*(a_t|s_t)$ that maximizes the accumulated reward averaged over the unobservable context u , that is, the learner agent solves the problem

$$\max_{\pi(a_t|s_t)} \mathbb{E}_{\rho^0, \rho(u)} \left[\sum_{t=0}^{\infty} \gamma^t R^u(s_t, a_t, s_{t+1}) \right]. \quad (1)$$

Note that Problem (1) can be solved by the learner alone using any standard RL algorithm, e.g., Q learning. An optimal policy $\pi^*(a_t|s_t)$ that does not depend on the context u exists because the context variable u is subject to a stationary distribution $\rho(u)$ and its statistics can be marginalized to get new transition and reward functions that do not contain u . Nevertheless, this process typically requires a lot of data to learn the optimal policy. Our goal is to use experience from a context-aware demonstrator to help the learner find the optimal policy using much fewer data. We formally state the problem of interest in this paper as follows.

Problem 1: (Experience Transfer with Hidden Context) Given the experience data set $\{(s_t, a_t, s_{t+1}, r_t)\}_{t=0:KT}$ of a context-aware demonstrator that excludes contextual information, design learning algorithms for a context-unaware learner agent that use these data to find the optimal policy for problem (1) with only a small number of new data samples.

The key challenge in solving Problem 1 is that if the demonstrator’s experience is used by the learner without considering the causal effect of the hidden context, it can result in biased estimation of the transition and reward models and, therefore, in a suboptimal learner policy. We discuss this issue in detail in Section II-C. Note that there

are many practical situations where the context that is known to the demonstrator is not revealed to the learner. An example is transferring motion planning experience to a robot that is not equipped with the necessary sensors to observe contextual information in its environment. Another example is transferring experience of medical treatment from one patient to another, where contextual information about the patients' profiles needs to remain confidential.

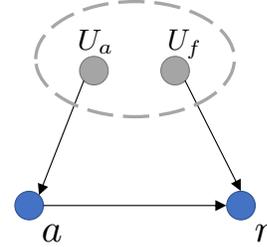


Figure 2. An auxiliary causal graph with auxiliary variables U_a and U_f .

C. Estimated Model Bias due to Hidden Context

A straightforward way to use the demonstrator's experience $\{(s_t, a_t, s_{t+1}, r_t)\}_{t=0:KT}$ is to directly estimate the model of the environment, that is, the transition and reward functions, from these samples. Then, any planning method for MDPs can be applied, e.g., Value Iteration [1], to find the optimal policy. This approach is called model-based RL and has a lower sampling complexity compared to model-free RL approaches, see [21]. However, when the data $\{(s_t, a_t, s_{t+1}, r_t)\}_{t=0:KT}$ is collected according to a hidden context, using this approach to estimate models of the transition and reward functions will result in bias and, therefore, a suboptimal policy for problem (1).

To see this, consider the estimation of the expected reward $\mathbb{E}[r_t]$ by taking action a_t at a certain state s_t . The causal relationship between the unobserved context u , the action a , and the reward r for the demonstrator and the learner is shown in Figure 1. Specifically, in the case of the demonstrator, both the action and the reward are affected by the context u . On the other hand, in the case of the learner, the reward depends on the context u but the action is independent of u since context information is not observable at the learner. Assume now that the demonstrator collects the observational data as described in Section II-B, which corresponds to the causal graph in Figure 1(a). Moreover, let $P(r, a)$ be the joint distribution of the action and reward pair at state s , and $P(s', a)$ be the joint distribution of the action and next state pair at state s , which are computed using the observational data. These distributions are called the expected reward $\mathbb{E}[r_t]$ marginalizing out the unobserved context u , the learner needs to estimate the causal effect of action a ,

$$\mathbb{E}[r|do(a)], \quad (2)$$

where $do(a)$ assigns a specific value a to the action regardless of what the context u is, [14]. Moreover, we have that $\mathbb{E}[r|do(a)] = \sum_r rP(r|do(a))$, where

$$P(r|do(a)) = \sum_u P(r|a, u)P(u). \quad (3)$$

Since the context u is missing from the observational distribution $P(r, a)$, we cannot compute (3) directly. Instead, using $P(r, a)$, we can compute $\mathbb{E}[r|a]$ as $\mathbb{E}[r|a] =$

Index	a = 1	a = 2	a = 3
1	1	1	1
2	1	1	2
3	1	2	1
4	1	2	2
5	2	1	1
6	2	1	2
7	2	2	1
8	2	2	2

Table I

ALL POSSIBLE MAPPINGS FROM THE ACTION SPACE $\mathcal{A} = \{1, 2, 3\}$ TO THE REWARD SPACE $\mathcal{R} = \{1, 2\}$.

$\sum_r rP(r|a) = \sum_r r \frac{P(r, a)}{P(a)}$. Furthermore, we have that

$$\begin{aligned} P(r|a) &= \frac{P(r, a)}{P(a)} = \frac{\sum_u P(r|a, u)P(a|u)P(u)}{\sum_u P(a|u)P(u)} \\ &= \sum_u [P(r|a, u) \left(\frac{P(a|u)P(u)}{\sum_u P(a|u)P(u)} \right)]. \end{aligned} \quad (4)$$

Comparing (4) and (3), we observe that $\mathbb{E}[r|a] \neq \mathbb{E}[r|do(a)]$. This is because the term $\frac{P(a|u)P(u)}{\sum_u P(a|u)P(u)}$ in (4) essentially reweights the distribution $P(u)$ in (3) using $P(a|u)$, which is given by the demonstrator's policy function $\pi^*(a|s, u)$. Therefore, the causal effect $\mathbb{E}[r|do(a)]$ is non-identifiable when the context is missing from the observational distribution. And the same holds true for the transition probability $P(s'|s, do(a))$. Nevertheless, as discussed in [14], [19], it is possible to infer lower and upper bounds on both $\mathbb{E}[r|do(a)]$ and $P(s'|s, do(a))$. In the next section, we show that such causal bounds can be computed given the observational data of the demonstrator and they can be used to help the learner learn faster.

III. CAUSAL BOUNDS ON THE VALUE FUNCTIONS

In this section, we first show how to compute upper and lower bounds on the causal effect $\mathbb{E}[r|do(a)]$ at state s given the observational distribution $P(r, a)$. Our approach is a generalization of the method proposed in [14], [19] to compute upper and lower bounds on the reward function using observational data for binary action and reward spaces to the case where these are any discrete and finite spaces. Then, we show how to extend this method to compute causal bounds on the transition probability $P(s'|s, a)$ for any transition (s, a, s') . Finally, we propose a way to obtain upper and lower bounds on the state-based optimal value

function $V^*(s)$ and the action-based optimal value function $Q^*(s, a)$ for any state s and action a , using the causal bounds on the rewards and transition probabilities obtained by the observational data. Similar to [19], throughout this section, we assume that the observational distributions $P(r, a)$ and $P(s', a)$ at state s can be well estimated. This requires the demonstrator to take enough samples so that $P(r, a)$ and $P(s', a)$ are asymptotically correct according to the Law of large numbers.

A. Causal Bounds on the Reward Function

Without loss of generality, let the action and reward spaces be indexed as $\mathcal{A} = \{1, 2, \dots, N_a\}$ and $\mathcal{R} = \{1, 2, \dots, N_r\}$, respectively. For example, in a robot motion planning problem, the actions indexed by 1 and 2 may correspond to 'go up' and 'go right'. The causal effect of taking action a is denoted as $\mathbb{E}[r|do(a)]$, where $\mathbb{E}[r|do(a)] = \sum_{r \in \mathcal{R}} r P(r|do(a))$. At state s , there exist $|\mathcal{R}|^{|\mathcal{A}|}$ mappings from the space \mathcal{A} to \mathcal{R} . To see this, consider an example where $\mathcal{A} = \{1, 2, 3\}$ and $\mathcal{R} = \{1, 2\}$. All the possible mappings are listed in Table I. Mapping 3 in Table I means that actions 1, 2 and 3 receive rewards 1, 2 and 1, respectively.

Next, we build an auxiliary causal graph, similar to the one in Section 8.2 in [14], with two underlying random variables U_a and U_f as shown in Figure 2. Let the supports of U_a and U_f be $\{1, 2, \dots, N_a\}$ and $\{1, 2, \dots, |\mathcal{R}|^{|\mathcal{A}|}\}$, respectively. Here, U_a is the action choice, and U_f is the action-to-reward mapping index. The causal relationship in Figure 2 is defined as

$$a = U_a \quad \text{and} \quad r = f(a, U_f), \quad (5)$$

where $f(a, U_f)$ returns the reward at entry a in the mapping U_f . For instance, in Table I, $f(3, 7) = 1$. Let also $q_{ij} = P(U_a = i, U_f = j)$. As discussed in Section 8.2 in [14], we can construct the joint distribution $\{q_{ij}\}$ in the auxiliary graph in Figure 2 to reflect any possible causal relationship $P(r|do(a))$ in the graph in Figure 1(a), regardless of the value of U_a . To do so, we need to construct $\{q_{ij}\}$ as

$$P(r|do(a)) = \sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_f(r, a)} q_{ij}, \quad (6)$$

where $\mathcal{S}_f(r, a)$ is an index set such that for every $j \in \mathcal{S}_f(r, a)$, $f(a, j) = r$. To see how the distribution $\{q_{ij}\}$ that is constructed according to (6) generates the causal relationship $P(r|do(a))$, consider the example in Table I. Suppose the goal is to construct a distribution $\{q_{ij}\}$, with $i = 1, 2, 3$ and $j = 1, 2, \dots, 8$, so that the auxiliary graph can reflect the causal relationship $P(r = 1|do(a = 2))$. Since the mappings $\{1, 2, 5, 6\}$ assign reward $r = 1$ to action $a = 2$, the probability $P(r = 1|do(a = 2))$ consists of all events $U_f \in \{1, 2, 5, 6\}$. The value of U_a is irrelevant since we fix the action to $do(a = 2)$. Therefore, we have that $P(r = 1|do(a = 2)) = \sum_{i \in \{1, 2, 3\}} \sum_{j \in \{1, 2, 5, 6\}} q_{ij}$.

Index	a = 1	a = 2
1	1	1
2	1	2
3	2	1
4	2	2

Table II

ALL POSSIBLE MAPPINGS FROM THE ACTION SPACE $\mathcal{A} = \{1, 2\}$ TO THE SPACE OF NEXT POSSIBLE STATE $\mathcal{S}' = \{1, 2\}$.

Moreover, we can construct the joint distribution $\{q_{ij}\}$ in the auxiliary graph in Figure 2 to reflect the observational distribution $P(r, a)$ in the graph in Figure 1(a). For this, we need to construct $\{q_{ij}\}$ so that

$$P(r, a) = \sum_{j \in \mathcal{S}_f(r, a)} q_{aj}, \quad \text{for any } (r, a). \quad (7)$$

The difference between (7) and (6) is that (7) requires that $U_a = a$. This is because the observational distribution $P(r, a)$ is generated using the auxiliary graph in Figure 2, where the action a depends on U_a according to the definition in (5). Using (6) and (7), and given the observation distribution $P(r, a)$ of the demonstrator, we can find an upper bound (or lower bound) on $\mathbb{E}[r|do(a)]$ by solving the optimization problem

$$\begin{aligned} \max_{\{q_{ij}\}} \quad & (\text{or } \min_{\{q_{ij}\}}) \sum_{r \in \mathcal{R}} r \sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_f(r, a)} q_{ij} \\ \text{s.t.} \quad & \sum_{j \in \mathcal{S}_f(r, a)} q_{aj} = P(r, a), \quad \text{for all } r \in \mathcal{R}, a \in \mathcal{A}. \end{aligned} \quad (8)$$

Note that problem (8) can be easily modified to incorporate any prior knowledge on the causal relationship of the actions other than the action of interest, $P(r|do(m))$, $m \neq a$. For example, if $P(r|do(m))$ for action m at state s is known and the goal is to infer the causal effect of action $a \neq m$, we can add

$$\sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_f(r, m)} q_{ij} = P(r|do(m))$$

as additional constraints to problem (8) and obtain a tighter bound for action a .

B. Causal Bounds on the Transition Probabilities

We can directly extend the approach described in Section III-A to compute causal bounds on the state transition probabilities. At state s , let $P(a, s')$ be the joint distribution of taking action a to reach state s' , which can be computed from the demonstrator's experience. Moreover, for any state s , let $\mathcal{S}' = \{1, 2, \dots, N_{s'}\}$ be the space of all possible next states. Then, similar to Section III-A, there exist $|\mathcal{S}'|^{|\mathcal{A}|}$ possible mappings from the space \mathcal{A} to \mathcal{S}' . For example, when the spaces of the actions and next states are $\mathcal{A} = \{1, 2\}$ and $\mathcal{S}' = \{1, 2\}$, the mappings are listed in Table II. Let the random variable $U_{s'}$ select a mapping index from the set $\{1, 2, \dots, |\mathcal{S}'|^{|\mathcal{A}|}\}$. We can construct the same auxiliary causal graph as in Figure 2, where the node r is replaced by s' and the causal relationship is defined as

$$a = U_a \quad \text{and} \quad s' = f_s(a, U_{s'}),$$

where $f_s(a, U_{s'})$ returns the state index at entry a in the mapping $U_{s'}$. For example, in Table II, we have that $f_s(1, 3) = 2$. As before, let $q_{ij} = P(U_a = i, U_{s'} = j)$. Similar to (6), we have that

$$P(s'|do(a)) = \sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_{f_s}(s', a)} q_{ij}, \quad (9)$$

where $\mathcal{S}_{f_s}(s', a)$ is an index set such that $f_s(a, j) = s'$ for every $j \in \mathcal{S}_{f_s}(s', a)$. Moreover, similar to (7), we have that

$$P(s', a) = \sum_{j \in \mathcal{S}_{f_s}(s', a)} q_{aj}. \quad (10)$$

Then, given the observational distribution $P(a, s')$, we can find an upper bound (or lower bound) on $P(s'|do(a))$ by solving the optimization problem

$$\begin{aligned} & \max_{\{q_{ij}\}} \text{ (or min)} \sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_{f_s}(s', a)} q_{ij} \\ \text{s.t.} & \sum_{j \in \mathcal{S}_{f_s}(s', a)} q_{aj} = P(a, s'), \text{ for all } s' \in \mathcal{S}', a \in A. \end{aligned} \quad (11)$$

Same as in problem (8), here too any prior knowledge on $P(s'|do(m))$ corresponding to action m at state s can be modeled as additional constraints

$$\sum_{i \in \{1, 2, \dots, N_a\}} \sum_{j \in \mathcal{S}_{f_s}(s', m)} q_{ij} = P(s'|do(m))$$

that can be added to problem (11) so that a tighter bound can be achieved for action $a \neq m$.

C. Causal Bounds on the Value Functions

In this section, we discuss how to use the upper and lower bounds computed in Sections III-A and III-B to find upper and lower bounds on the state-based and action-based functions $V(s)$ and $Q(s, a)$ for the learner agent.

First, we show how to compute such bounds for the function $V(s)$. For this, we rely on the linear programming formulation proposed in [22] to find the optimal value function $V(s)$ for a given discounted MDP problem

$$\begin{aligned} & \min_{\{V(s)\}} \sum_{s \in \mathcal{S}} c(s)V(s) \\ \text{s.t.} & V(s) \geq r(s, a) + \gamma \sum_{s'} P(s'|s, a)V(s'), \forall s, a. \end{aligned} \quad (12)$$

where $c(s) > 0$ for all s are called state-relevance weights. It is shown in [22] that the optimal value function $V(s)$ is the unique solution to problem (12).

Solving problem (12) requires values for $r(s, a)$ and $P(s'|s, a)$, which can not be estimated without bias from the experience data of the demonstrator, as discussed in Section II-C. For this reason, we instead utilize the bounds for $r(s, a)$ and $P(s'|s, a)$, denoted as $r(s, a) \in [\underline{r}(s, a), \bar{r}(s, a)]$ and $P(s'|s, a) \in [\underline{P}(s'|s, a), \bar{P}(s'|s, a)]$, that are computed in Sections III-A and III-B. Using these bounds, we can find

an upper bound on the function $V(s)$ for every state s by solving the optimization problem

$$\begin{aligned} & \max_{\substack{r(s, a) \in \Gamma(s, a) \\ P(s'|s, a) \in \Xi(s, a, s')}} \min_{\{V(s)\}} \sum_{s \in \mathcal{S}} c(s)V(s) \\ \text{s.t.} & V(s) \geq r(s, a) + \gamma \sum_{s'} P(s'|s, a)V(s'), \forall s, a, \end{aligned} \quad (13)$$

where $\Gamma(s, a) := \{r(s, a) \in [\underline{r}(s, a), \bar{r}(s, a)]\}$ and $\Xi(s, a, s') := \{P(s'|s, a) \in [\underline{P}(s'|s, a), \bar{P}(s'|s, a)]\}$ and $\sum_{s'} P(s'|s, a) = 1$ for all s . Similarly, a lower bound on the function $V(s)$ for every s can be found by the solution of the optimization problem

$$\begin{aligned} & \min_{\substack{r(s, a) \in \Gamma(s, a) \\ P(s'|s, a) \in \Xi(s, a, s')}} \min_{\{V(s)\}} \sum_{s \in \mathcal{S}} c(s)V(s) \\ \text{s.t.} & V(s) \geq r(s, a) + \gamma \sum_{s'} P(s'|s, a)V(s'), \forall s, a. \end{aligned} \quad (14)$$

We can replace the inner minimization problem in (13) with its dual form to obtain a maximization problem for (13).¹ Then, problems (13) and (14) become quadratic constrained quadratic programs (QCQP) and can be solved using existing software packages, e.g., CPLEX.

Since problem (13) returns the reward and transition pair that maximizes the summation $\sum_{s \in \mathcal{S}} c(s)V(s)$ rather than the value function $V(s)$ at a specific state s , the solution of (13) cannot be directly used as an upper bound on the value function $V(s)$ for any state s . The same holds true for the solution of problem (14) which can not be used as a lower bound on the function $V(s)$. Next, we discuss how to derive bounds on the value function $V(s)$ for a specific state s given the solutions to problems (13) and (14).

Theorem 3.1: Assume the discrete and finite reward space is bounded, and let \bar{R} and \underline{R} be uniform upper and lower bounds on the reward functions $r(s, a)$. Moreover, let $\{r^*, P^*, V^*(s)\}$ be the solution to problem (13), and let $\bar{V}^*(s)$ be the true upper bound on the value at state s . Then, we have that

$$\bar{V}^*(s) \leq V^*(s) + \frac{1}{c(s)} \left(\sum_{\tilde{s} \neq s} c(\tilde{s})V^*(\tilde{s}) - \frac{\bar{R}}{1 - \gamma} \sum_{\tilde{s} \neq s} c(\tilde{s}) \right). \quad (15)$$

Similarly, let $\{r^*, P^*, V^*(s)\}$ be the solution to problem (14), and let $\underline{V}^*(s)$ be the true lower bound on the

¹The inner linear minimization problem in (13) can be compactly written as $\min_{\vec{V} \in \mathbb{R}^{|\mathcal{S}|}} \vec{c}^T \vec{V}$ s.t. $A\vec{V} \geq \vec{b}$, where the vectors $\vec{c}, \vec{V}, \vec{b}$ and the matrix A contains the variables $c(s), V(s), r(s, a)$ and $P(s'|s, a)$. Then, the corresponding dual problem takes the form $\max_{\vec{\Lambda} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}} \vec{b}^T \vec{\Lambda}$ s.t. $A^T \vec{\Lambda} = \vec{c}, \Lambda \geq 0$, where the vector $\vec{\Lambda}$ contains the Lagrange multipliers for the inequality constraints in the primal problem (13), [23]. By [22], the primal problem has an optimal solution and because of strong duality, the dual problem also has an optimal solution that has the same value with the primal problem, [23].

Algorithm 1: Causal Bound Constrained Q learning

Input: Initial value function estimate $Q(s, a)$ for all (s, a) . Learning rate α_t . Discount factor γ . Number of episodes K . Number of time steps in each episode T . Initial state distribution ρ^0 . Context distribution $\rho(u)$. Causal bounds $[Q(s_t, a_t), \bar{Q}(s_t, a_t)]$ for all (s, a) . Parameter ϵ for the ϵ -greedy policy [1].

```
1 for episode  $1 \leq k \leq K$  do
2   The learner samples the initial state  $s_0$  and context  $u_k$ 
   from distribution  $\rho^0$  and  $\rho(u)$ , respectively ;
3   for time step  $0 \leq t \leq T$  do
4     The learner observes its state  $s_t$  ;
5     The learner takes action  $a_t$  according to the
      $\epsilon$ -greedy policy;
6     The learner receives reward  $r(s_t, a_t)$  and observes
     its next state  $s_{t+1}$  ;
7     The learner updates its value function at  $(s_t, a_t)$ 
     according to (18).
```

value at state s . Then, we have that

$$\underline{V}^*(s) \geq V^*(s) - \frac{1}{c(s)} \left(\frac{\bar{R}}{1-\gamma} \sum_{\tilde{s} \neq s} c(\tilde{s}) - \sum_{\tilde{s} \neq s} c(\tilde{s}) V^*(\tilde{s}) \right). \quad (16)$$

The proof can be found in Appendix A. Theorem 3.1 implies that tight bounds on $V(s)$ can be obtained for any state s by using (15) and (16) and by selecting large enough $c(s)$ while fixing the $c(\tilde{s})$ for all $\tilde{s} \neq s$.

Finally, denote the bounds on the value function $V(s)$ on the right hand side of (15) and (16) as $\bar{V}(s)$ and $\underline{V}(s)$. Since the optimal action-based value function satisfies $Q(s, a) = r(s, a) + \gamma \sum_{s'} P(s'|s, a) V(s')$, we can use the bounds $[\underline{V}(s), \bar{V}(s)]$ together with the bounds $[\underline{r}(s, a), \bar{r}(s, a)]$ and $[\underline{P}(s'|s, a), \bar{P}(s'|s, a)]$ derived in Sections III-A and III-B to obtain bounds $[Q(s, a), \bar{Q}(s, a)]$ on the action-value function $Q(s, a)$, for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$. This process is straightforward and is directly presented in Algorithm 3 in Appendix B.

IV. ALGORITHM DESIGN

In this section, we present two approaches to combine the causal bounds $[Q(s, a), \bar{Q}(s, a)]$ derived in Algorithm 3 in Appendix B with existing RL algorithms to accelerate the learning process of the learner. The first approach uses the causal bounds as the constraints of the value function in the Q learning method [15]. The second approach uses the causal bounds to provide better exploration performance based on existing UCB- Q exploration strategies in [16]–[18].

A. Causal Bound Constrained Q Learning

Let the learner agent interact with the environment at state s_t by taking action a_t , receiving reward $r(s_t, a_t)$, and observing the next state s_{t+1} . In standard Q learning the agent's current action-value function is updated as

$$Q(s_t, a_t) \leftarrow (1 - \alpha_t) Q(s_t, a_t) + \alpha_t (r(s_t, a_t) + \gamma \max_a Q(s_{t+1}, a)), \quad (17)$$

Algorithm 2: Causal Bound Aided UCB- Q learning

Input: Initial upper confidential bound estimate $Q_U(s, a)$ for all (s, a) . Learning rate α_t . Discount factor γ . Number of episodes K . Number of time steps in each episode T . Initial state distribution ρ^0 . Context distribution $\rho(u)$. Causal bounds $[Q(s_t, a_t), \bar{Q}(s_t, a_t)]$ for all (s, a) .

```
1 for episode  $1 \leq k \leq K$  do
2   Sample the initial state  $s_0$  and context  $u_k$  from
   distribution  $\rho^0$  and  $\rho(u)$ , respectively ;
3   for time step  $0 \leq t \leq T$  do
4     Let the agent observe its state  $s_t$  ;
5     Let the agent take action  $a_t$  so that
6        $a_t = \arg \max_a Q_U(s_t, a)$ ;
   The agent receives reward  $r(s_t, a_t)$  and observes
   its next state  $s_{t+1}$  ;
7     The agent updates its upper confidential bound at
      $(s_t, a_t)$  according to (19) ;
8      $Q_U(s_t, a_t) \leftarrow \min(Q_U(s_t, a_t), \bar{Q}(s_t, a_t))$ .
```

where α_t is the learning rate. It is shown in [15] that Q learning almost surely converges to the optimal value with sufficient exploration and an appropriate choice of the learning rate. Using the bounds $[Q(s, a), \bar{Q}(s, a)]$ on the action-value function derived in Section III-C for the case that the context observed by the demonstrator is hidden to the learner, we can modify the above Q learning update by introducing a projection operation as

$$Q(s_t, a_t) \leftarrow \Pi_{[Q(s_t, a_t), \bar{Q}(s_t, a_t)]} \left((1 - \alpha_t) Q(s_t, a_t) + \alpha_t (r(s_t, a_t) + \gamma \max_a Q(s_{t+1}, a)) \right). \quad (18)$$

We formally present the resulting Causal Bound Constrained Q learning (CBC- Q learning) method in Algorithm 1. In the early learning phase of standard Q learning in (17), the value $Q(s, a)$ is incorrect for most state-action pairs. By projecting onto the causal bounds in (18), the values of all state-action pairs are confined to a provably correct region that contains the optimal value, therefore accelerating the learning process. During the later stages of learning, the value iterates $Q(s, a)$ approach the true values, so that the causal bound constraints do not take effect anymore. Therefore, the update (18) is eventually reduced to the standard Q learning update (17), which is unbiased [15].

B. Causal Bound Aided UCB- Q Learning

The convergence of Q learning to the optimal value assumes that all states and actions are sufficiently sampled [15]. This requires the learner to follow a randomized policy function to explore the environment. In [1], an ϵ -greedy policy is discussed, which lets the agent choose the best action according to its current value function estimate with probability $1 - \epsilon$ and choose other actions with probability ϵ . However, this policy can be sample inefficient [24]. The

UCB exploration strategy has been shown to improve on the sampling complexity of the ϵ -greedy policy, first for Multi-Arm Bandit problems [25] and subsequently for general reinforcement learning problems [16]–[18]. The key idea is that instead of estimating the value function, UCB methods keep an estimate of the upper confidential bound of the value function and choose the best action according to this upper confidential bound. Specifically, in [17], [18], the upper confidential bound on the value function $Q_U(s, a)$ is updated as

$$Q_U(s_t, a_t) \leftarrow (1 - \alpha_t)Q_U(s_t, a_t) + \alpha_t(r(s_t, a_t) + \gamma \max_a Q_U(s_{t+1}, a) + b_t), \quad (19)$$

where b_t is an estimate of the confidential interval on the value function at (s_t, a_t) , which is usually of order $O(\frac{1}{\sqrt{k}})$ [17], [18] and k is the number of times (s_t, a_t) is visited at time t . Essentially, the more often every state-action pair (s_t, a_t) is visited, the tighter the corresponding interval b_t becomes [17], [18]. Given the bounds $[Q(s, a), \bar{Q}(s, a)]$ on the action-value function computed using Algorithm 3 in Appendix B, we add an additional step after update (19), i.e.,

$$Q_U(s_t, a_t) \leftarrow \min(Q_U(s_t, a_t), \bar{Q}(s_t, a_t)). \quad (20)$$

We formally present the Causal Bound aided UCB-Q learning method in Algorithm 2. The key idea is that when the estimated upper confidential bound $Q_U(s, a)$ is higher than $\bar{Q}(s, a)$, Algorithm 2 uses the causal bound as the upper confidential bound. This prevents the upper confidential bound at some state-action pairs from being too optimistic, therefore, avoids unnecessary exploration and accelerates the learning process.

V. NUMERICAL EXPERIMENTS

In this section, we provide numerical results for TL in motion planning problems to validate the causal bounds presented in (15) and (16) and the efficacy of Algorithms 1 and 2 proposed in Section IV.

First, we study a scenario where the reward functions are affected by the contextual information. Consider the motion planning problem in Figure 3. Let $s \in \{0, 1, \dots, 4\}^2$ denote the state (position) of the robot and $a \in \{1, 2, 3, 4\}$ denote its actions {go up, go right, go down, go left}. Let the transition be deterministic, that is, when we choose action $a = 1$ at $s = [2, 1]$, the robot transitions to $[2, 2]$, unless it collides with the walls in Figure 3, in which case it remains at its current position. The red and green cells in Figure 3 are goal positions. When the robot takes action a at state s and the next position is not one of the two goal positions, it receives a reward -1 . On the other hand, if the robot is already at the goal position, it remains there and receives a reward 0 regardless of what action it takes. When the robot reaches a goal, it receives a reward according to the random context variable $u = \{0, 1\}$ in the environment. Specifically, when the robot reaches the red cell, if $u = 0$ (or

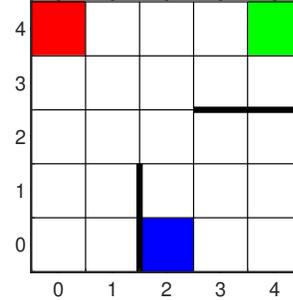


Figure 3. A motion planning problem where the agent starts from the blue cell and moves to the red or green cells. The rewards received when reaching the goals are affected by the context variable u .

(s, a)	$\mathbb{E}[r s, do(a)]$	$\mathbb{E}[r s, a]$	Causal bounds
$([0, 3], 1)$	1.2	1.2	$[0.54, 3.84]$
$([1, 4], 4)$	1.2	3.6	$[0.012, 8.5920]$
$([3, 4], 2)$	3.2	1.8909	$[-0.3640, 4.3160]$
$([4, 3], 1)$	3.2	3.2	$[1.94, 3.74]$

Table III
THE TRUE CAUSAL EFFECTS $\mathbb{E}[r|s, do(a)]$, DIRECT ESTIMATE OF THE EXPECTED REWARDS FROM THE OBSERVATIONAL DISTRIBUTION $P(r, a|s)$ WITHOUT CONSIDERING THE HIDDEN CONTEXTS AND THE CAUSAL BOUNDS OBTAINED BY SOLVING PROBLEM (8).

$u = 1$), it receives a reward 10 with probability 0.6 (or 0.1). Similarly, when the robot reaches the green cell, if $u = 0$ (or $u = 1$), it receives a reward 5 with probability 0.3 (or 0.8). The discount factor is $\gamma = 0.9$. The context variable is sampled from a Bernoulli distribution with parameter 0.8.

In what follows, we focus on computing the causal bounds on the rewards at critical state-action pairs. Specifically, in this example we are interested in $([0, 3], 1)$, $([1, 4], 4)$, $([3, 4], 2)$ and $([4, 3], 1)$. In practice, given the data from the demonstrator, we can identify these critical state-action pairs either by expert knowledge on which state-action pairs are easily affected by potential hidden context, or by using a heuristic, that is, selecting those pairs that receive different rewards. In this example, only the state-action pairs above receive different rewards from the demonstrator’s experience. Note that our approaches to compute the bounds on the value functions in (13) and (14) do not rely on the knowledge of these critical state-action pairs.

Consider the demonstrator agent who can observe the context variable u and also knows the contextual optimal policy

$$\pi^*(s, u) = \begin{cases} 4, & \text{when } s = [1, 4] \text{ and } u = 0, \\ 1, & \text{when } s = [1, 4] \text{ and } u = 1, \\ 4, & \text{when } s = [3, 4] \text{ and } u = 0, \\ 2, & \text{when } s = [3, 4] \text{ and } u = 1, \\ 1, & \text{when otherwise.} \end{cases} \quad (21)$$

The policy in (21) can be obtained using Value Iteration [1] given the true model of the environment. We only present the optimal policy at the critical states because the observational distribution at the critical state-action pairs is only affected by the demonstrator’s policy at these critical

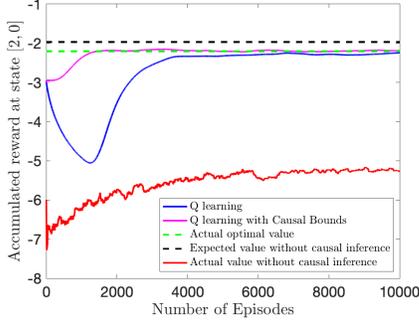


Figure 4. Comparison between the proposed Causal Bound constrained Q learning and the standard Q learning. The blue curve shows the learning progress of the Q learning algorithm [15] without using data from the demonstrator. The magenta curve shows the performance of Algorithm 1. Both curves are obtained by running these algorithms for 10 trials and take the average. Meanwhile, the black (or green) dashed line is the optimal value at state $[2, 0]$ computed using $\mathbb{E}[r|s, a]$ (or $\mathbb{E}[r|s, do(a)]$). The red curve shows the evaluation of the sub-optimal policy computed using $\mathbb{E}[r|s, a]$.

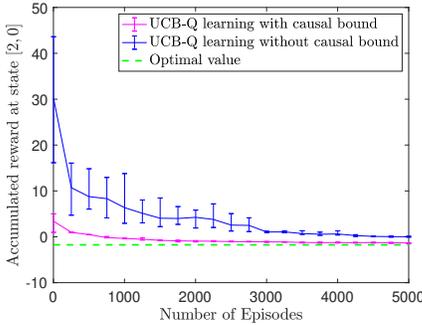


Figure 5. Learning progress of Algorithm 2 (magenta curve) compared to UCB- Q learning algorithm [17] (blue curve) without using causal bounds. The optimal value at state $[2, 0]$ is computed using the true model (green dashed line). Both algorithms are run for 10 trials and the mean-error curves are presented.

states. The demonstrator collects data using an ϵ -greedy policy as described in Section II-B, where $\epsilon = 0.3$. That is, at each state s and observed context u , the demonstrator chooses the optimal action in (21) with probability 0.7 and each of the other three actions with probability 0.1. If the demonstrator collects data for long enough episodes, by the Law of Large Numbers, the observational distribution of reward-action pairs at state s approaches

$$P(r, a|s) = \sum_u P(r|s, a, u)P(a|s, u)P(u). \quad (22)$$

We compute the observational distribution $P(r, a|s)$ at the critical positions using (22) by substituting the definition of the reward function, the policy function (21), and the distribution of u . Next, the learner computes the causal bounds for the rewards at state-action pairs $([0, 3], 1)$, $([1, 4], 4)$, $([3, 4], 2)$ and $([4, 3], 1)$ by solving problem (8) with prior knowledge where, except for the rewards of those critical state-action pairs, all other actions have reward -1 .

The causal bounds on the rewards at these state-action pairs are listed in Table III, together with the true causal

effect $\mathbb{E}[r|s, do(a)]$ and the naive estimation $\mathbb{E}[r|s, a]$ ignoring the hidden context. At $([0, 3], 1)$ and $([4, 3], 1)$, the naive estimation returns the correct results. As we show in Section II-C, this is because the demonstrator chooses action $a = 1$ with same probability regardless of the context variable u . In this case, the naive estimation (4) is equivalent to the causal relationship (3). On the other hand, the demonstrator has higher probability to choose action 4 when $u = 0$ than $u = 1$ at position $[1, 4]$. Then, from Table III, we observe that the naive method over-estimates the expected reward at $([1, 4], 4)$. Therefore, if we directly take the reward estimates $\mathbb{E}[r|s, a]$ in Table III and apply Value Iteration [1] to find the optimal value function and policy, we will obtain a suboptimal policy. This is shown in Figure 4. We observe that the naive method gives an over-optimistic value estimate at state $[2, 0]$ (black dashed line) compared to the true value (green dashed line). Nevertheless, the actual performance of the suboptimal policy (red curve) is far below the true optimal performance (green dashed line). This is because the agent over-estimates the reward at $([1, 4], 4)$ and plans to move to $[1, 4]$, while the actual reward it receives is much lower than the reward it would receive if it moved to $[3, 4]$ or $[4, 3]$.

Next, we validate Algorithms 1 and 2 using the bounds on the value functions $[Q(s_t, a_t), \bar{Q}(s_t, a_t)]$ that are computed based on the causal bounds in Table III. When the causal bounds are only derived at the critical state-action pairs, the bounds on the value functions $[Q(s_t, a_t), \bar{Q}(s_t, a_t)]$ can be found in closed-form, instead of solving the general optimization problems in (13) and (14). The performance of Algorithm 1 is compared to Q learning without demonstrator's data [cf. (17)] and the results are shown in Figure 4. While both Algorithm 1 and plain Q learning converge to the true value, Algorithm 1 requires much fewer data. This is because in plain Q learning (blue curve), the value at state $[2, 0]$ is estimated incorrectly during the early stages due to the incorrect value estimates at its successor states. By projecting onto the causal bounds using (18), these initial errors in value estimation are controlled and the learning process is accelerated. On the other hand, the performance of Algorithm 2 is presented in Figure 5. Specifically, we compare Algorithm 2 with the plain UCB- Q learning algorithm in [17]. UCB- Q learning learns the upper confidential bound on the state values and chooses to explore the actions that maximize these bounds. Since the width of the confidential interval in (19) is of order $O(\frac{1}{\sqrt{k}})$ where k is the total number of times a state-action pair is visited, the upper confidential bounds for most state-action pairs are loose at the beginning while these state-action pairs have not been visited. This explains why the value estimate of the UCB- Q learning algorithm (blue curve) in Figure 5 is initially far from the true value. Furthermore, since the upper confidential bound is loose, UCB- Q explores many low-gain states that have not been visited before and the learning

progress of UCB-Q is slow. The causal bounds tighten the upper confidential bound, therefore, avoiding unnecessary exploration. As a result, in Figure 5, we observe that the Causal Bound aided UCB-Q learning algorithm (magenta curve) converges faster than the standard UCB-Q learning algorithm.

Similar results can be obtained in cases where the transition functions are affected by the contextual information. Similar to the contextual reward case, Algorithms 1 and 2 require fewer samples using the causal bounds, compared to the plain Q learning and UCB-Q learning algorithms. Due to the space limitations, the simulation details and results in this case are presented in Appendix C.

VI. CONCLUSION

In this paper, we studied a transfer reinforcement learning problem, where a context-aware demonstrator agent transfers experience in the form of transition and reward samples to a context-unaware learner agent, excluding the contextual information. The goal of the learner is to find a context-unaware optimal policy. We showed how to obtain provable causal bounds on the reward, transition and value functions, and proposed a Q learning and a UCB-Q learning method that employ these causal bounds to reduce the number of samples required by the learner to find the optimal value function without any bias. We provided numerical examples showing the efficacy of the proposed methods.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research*, vol. 10, no. Jul, pp. 1633–1685, 2009.
- [3] M. E. Taylor, P. Stone, and Y. Liu, "Transfer learning via inter-task mappings for temporal difference learning," *Journal of Machine Learning Research*, vol. 8, no. Sep, pp. 2125–2167, 2007.
- [4] A. Barreto, W. Dabney, R. Munos, J. J. Hunt, T. Schaul, H. P. van Hasselt, and D. Silver, "Successor features for transfer in reinforcement learning," in *Advances in neural information processing systems*, 2017, pp. 4055–4065.
- [5] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1126–1135.
- [6] G. Konidaris, I. Scheidwasser, and A. Barto, "Transfer in reinforcement learning via shared features," *Journal of Machine Learning Research*, vol. 13, no. May, pp. 1333–1371, 2012.
- [7] A. Lazaric, M. Restelli, and A. Bonarini, "Transfer of samples in batch reinforcement learning," in *Proceedings of the 25th international conference on Machine learning*. ACM, 2008, pp. 544–551.
- [8] A. Lazaric and M. Restelli, "Transfer from multiple mdps," in *Advances in Neural Information Processing Systems*, 2011, pp. 1746–1754.
- [9] M. E. Taylor, N. K. Jong, and P. Stone, "Transferring instances for model-based reinforcement learning," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2008, pp. 488–505.
- [10] A. Tirinzoni, A. Sessa, M. Pirotta, and M. Restelli, "Importance weighted transfer of samples in reinforcement learning," in *International Conference on Machine Learning*, 2018, pp. 4943–4952.
- [11] A. Tirinzoni, M. Salvini, and M. Restelli, "Transfer of samples in policy search via multiple importance sampling," in *International Conference on Machine Learning*, 2019, pp. 6264–6274.
- [12] M. T. Rosenstein, Z. Marx, L. P. Kaelbling, and T. G. Dietterich, "To transfer or not to transfer," in *NIPS 2005 workshop on transfer learning*, vol. 898, 2005, p. 3.
- [13] A. Hallak, D. Di Castro, and S. Mannor, "Contextual markov decision processes," *arXiv preprint arXiv:1502.02259*, 2015.
- [14] J. Pearl, *Causality: models, reasoning and inference*, 2nd ed. Cambridge University Press, New York, 2009.
- [15] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [16] M. G. Azar, I. Osband, and R. Munos, "Minimax regret bounds for reinforcement learning," in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 263–272.
- [17] C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan, "Is q-learning provably efficient?" in *Advances in Neural Information Processing Systems*, 2018, pp. 4863–4873.
- [18] K. Dong, Y. Wang, X. Chen, and L. Wang, "Q-learning with ucb exploration is sample efficient for infinite-horizon mdp," *arXiv preprint arXiv:1901.09311*, 2019.
- [19] J. Zhang and E. Bareinboim, "Transfer learning in multi-armed bandits: a causal approach," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. AAAI Press, 2017, pp. 1340–1346.
- [20] M. Hausknecht and P. Stone, "Deep recurrent q-learning for partially observable mdps," in *2015 AAAI Fall Symposium Series*, 2015.
- [21] W. Sun, N. Jiang, A. Krishnamurthy, A. Agarwal, and J. Langford, "Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches," in *Conference on Learning Theory*, 2019, pp. 2898–2933.
- [22] F. d'Epenoux, "A probabilistic production and inventory problem," *Management Science*, vol. 10, no. 1, pp. 98–108, 1963.
- [23] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization*. Athena Scientific Belmont, MA, 1997, vol. 6.
- [24] M. Kearns and S. Singh, "Near-optimal reinforcement learning in polynomial time," *Machine learning*, vol. 49, no. 2-3, pp. 209–232, 2002.
- [25] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.

Algorithm 3: Computing bounds on $Q(s, a)$

Input: The bounds on the reward, transition and state-based value function $[r(s, a), \bar{r}(s, a)]$, $[P(s'|s, a), \bar{P}(s'|s, a)]$ and $[V(s), \bar{V}(s)]$.

- 1 **for** every $(s, a) \in \mathcal{S} \times \mathcal{A}$ **do**
- 2 Set $S' \leftarrow \{\text{all possible next states } s' \text{ at state } s \text{ when action } a \text{ is taken}\}$;
- 3 $p, q \leftarrow 0 \in \mathbb{R}^{|S'|}$;
- 4 $\vec{v} \leftarrow [V(s'_1), V(s'_2), \dots, V(s'_{|S'|})]^T$ so that $V(s'_1) \geq V(s'_2) \geq \dots \geq V(s'_{|S'|})$;
- 5 $p(1) \leftarrow \bar{P}(s'_1|s, a)$, $q(|S'|) \leftarrow \bar{P}(s'_{|S'|}|s, a)$;
- 6 $M \leftarrow \bar{P}(s'_1|s, a)$, $N \leftarrow \bar{P}(s'_{|S'|}|s, a)$;
- 7 **for** $i = 2, 3, \dots, |S'|$ **do**
- 8 **if** $1 - M \leq \bar{P}(s'_i|s, a)$, **then**
- 9 $p(i) \leftarrow 1 - M$;
- 10 $p(j) \leftarrow 0$ for all $i < j \leq |S'|$;
- 11 **break**;
- 12 **else**
- 13 $p(i) \leftarrow \bar{P}(s'_i|s, a)$;
- 14 $M \leftarrow M + P(s'_i|s, a)$;
- 15 **if** $1 - N \leq \bar{P}(s'_{|S'|+1-i}|s, a)$, **then**
- 16 $p(|S'| + 1 - i) \leftarrow 1 - N$;
- 17 $p(j) \leftarrow 0$ for all $1 \leq j < |S'| + 1 - i$;
- 18 **break**;
- 19 **else**
- 20 $p(i) \leftarrow \bar{P}(s'_{|S'|+1-i}|s, a)$;
- 21 $N \leftarrow N + P(s'_{|S'|+1-i}|s, a)$;
- 22 $\bar{Q}(s, a) \leftarrow \bar{r}(s, a) + \gamma \bar{p}^T \vec{v}$;
- 23 $\underline{Q}(s, a) \leftarrow \underline{r}(s, a) + \gamma q^T \vec{v}$.

APPENDIX A

Proof of Theorem 3.1

First, we show the bound in (15). Since $\{V^*(s)\}$ is the solution to problem (13), we obtain that

$$\sum_{s \in \mathcal{S}} c(s) V^*(s) \geq \sum_{s \in \mathcal{S}} c(s) \bar{V}^*(s). \quad (23)$$

Subtracting $\sum_{\tilde{s} \neq s} c(\tilde{s}) \bar{V}^*(\tilde{s})$ from both sides of (23), and dividing by $c(s)$, we get that

$$\bar{V}^*(s) \leq V^*(s) + \frac{1}{c(s)} \left(\sum_{\tilde{s} \neq s} c(\tilde{s}) V^*(\tilde{s}) - \sum_{\tilde{s} \neq s} c(\tilde{s}) \bar{V}^*(\tilde{s}) \right). \quad (24)$$

Since $r(s, a) \geq \underline{R}$ for all (s, a) , we have that any value function satisfies $V(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s] \geq \sum_{t=0}^{\infty} \gamma^t \underline{R} = \frac{\underline{R}}{1-\gamma}$. Replacing $\bar{V}^*(\tilde{s})$ in (24) with $\frac{\underline{R}}{1-\gamma}$, we achieve the bound in (15). Similarly, to show the bound in (16), since $\{V^*(s)\}$ is the solution to problem (14), we obtain that

$$\sum_{s \in \mathcal{S}} c(s) V^*(s) \leq \sum_{s \in \mathcal{S}} c(s) \underline{V}^*(s) \quad (25)$$

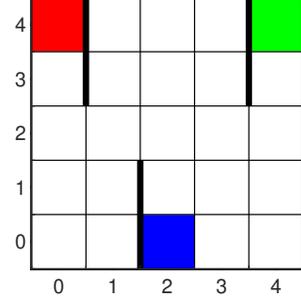


Figure 6. A motion planning problem where the agent starts from the blue cell and moves to the red or green cells. The transition probability at state-action pairs $([0, 2], 1)$ and $([4, 2], 1)$ are affected by the context variable u .

(s, a, s')	$P(s' s, do(a))$	$P(s' s, a)$	Causal bounds
$([0, 2], 1, [0, 3])$	0.18	0.46	$[0.1020, 0.8820]$
$([0, 2], 1, [0, 1])$	0.82	0.54	$[0.1180, 0.8980]$
$([4, 2], 1, [4, 3])$	0.72	0.55	$[0.12, 0.90]$
$([4, 2], 1, [4, 1])$	0.28	0.45	$[0.10, 0.88]$

Table IV

THE TRUE TRANSITION PROBABILITY $P(s'|s, do(a))$, DIRECT ESTIMATE OF THE PROBABILITY FROM THE OBSERVATIONAL DISTRIBUTION $P(r, a|s)$ WITHOUT CONSIDERING THE HIDDEN CONTEXTS AND THE CAUSAL BOUNDS OBTAINED BY SOLVING PROBLEM (8).

Subtracting $\sum_{\tilde{s} \neq s} c(\tilde{s}) \underline{V}^*(\tilde{s})$ from both sides of (25), and dividing by $c(s)$, we get

$$\underline{V}^*(s) \geq V^*(s) - \frac{1}{c(s)} \left(\sum_{\tilde{s} \neq s} c(\tilde{s}) \underline{V}^*(\tilde{s}) - \sum_{\tilde{s} \neq s} c(\tilde{s}) V^*(\tilde{s}) \right). \quad (26)$$

Since $r(s, a) \leq \bar{R}$ for all (s, a) , we have that any value function satisfies $V(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s] \leq \sum_{t=0}^{\infty} \gamma^t \bar{R} = \frac{\bar{R}}{1-\gamma}$. Replacing $\underline{V}^*(\tilde{s})$ in (26) with $\frac{\bar{R}}{1-\gamma}$, we achieve the bound in (16).

APPENDIX B

Computation of Bounds on $Q(s, a)$

In this section, we present an algorithm to compute the upper and lower bounds on the action-based value function $Q(s, a)$ using the causal bounds $[r(s, a), \bar{r}(s, a)]$, $[P(s'|s, a), \bar{P}(s'|s, a)]$ and $[V(s), \bar{V}(s)]$ derived in Section III. This algorithm is illustrated in Algorithm 3.

APPENDIX C

Simulation Results when Transitions are Affected by Contextual Information

Consider the motion planning problem in Figure 6, which is the same as the problem considered in Figure 3, except that the rewards received when reaching the red and green cells are now +10 and +5, respectively, and are not affected by the context variable u . Moreover, under context $u = 0$, when the robot takes action $a = 1$ at position $[0, 2]$ (or $[4, 2]$), it moves upwards with probability 0.7 (or 0.4) and moves downwards otherwise. Under context $u = 1$, when the robot takes action $a = 1$ at position $[0, 2]$ (or $[4, 2]$),

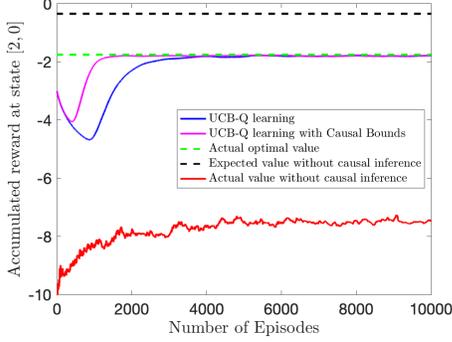


Figure 7. Comparison between the proposed Causal Bound constrained Q learning and the standard Q learning. The blue curve shows the learning progress of the Q learning algorithm [15] without using data from the demonstrator. The magenta curve shows the performance of Algorithm 1. Both curves are obtained by running these algorithms for 10 trials and take the average. The black (or green) dashed line is the optimal value at state $[2, 0]$ computed using $P(s'|s, a)$ (or $P(s'|s, do(a))$). The red curve shows the evaluation of the sub-optimal policy computed using $P(s'|s, a)$.

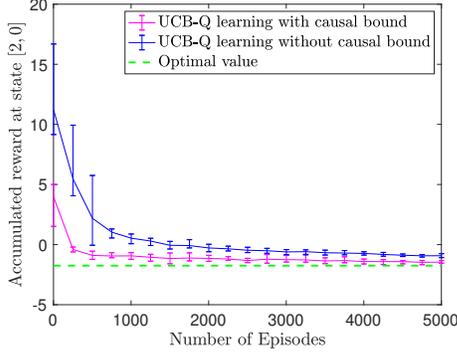


Figure 8. Learning progress of Algorithm 2 (magenta curve) compared to UCB- Q learning algorithm [17] (blue curve) without using causal bounds. The optimal value at state $[2, 0]$ is computed using the true model (green dashed line). Both algorithms are run for 10 trials and the mean-error curves are presented.

it moves upwards with probability 0.05 (or 0.8) and moves downwards otherwise.

Consider a demonstrator agent who knows the optimal contextual policy. Specifically, at the positions $[0, 2]$ and $[4, 2]$, its optimal policy is

$$\pi^*(s, u) = \begin{cases} 1, & \text{when } s = [0, 2] \text{ and } u = 0, \\ 2, & \text{when } s = [0, 2] \text{ and } u = 1, \\ 4, & \text{when } s = [4, 2] \text{ and } u = 0, \\ 1, & \text{when } s = [4, 2] \text{ and } u = 1. \end{cases} \quad (27)$$

The above policy is computed using Value Iteration given the true model of the environment. Same as with the policy in (21), here too we only present the optimal policy at critical states in (27). The demonstrator collects samples using an ϵ -greedy policy at positions $[0, 2]$ and $[4, 2]$, where $\epsilon = 0.3$. The observational distribution of $P(s', a)$ at positions $[0, 2]$ and $[4, 2]$ can be computed as

$$P(s', a|s) = \sum_u P(s'|s, a, u)P(a|s, u)P(u). \quad (28)$$

Given the observational distribution in (28), we can compute the causal bounds on the transition probabilities at $([0, 2], 1)$ and $([4, 2], 1)$ using (11). The causal bounds together with the true transition probabilities and the naive estimation are presented in Table IV. Since the demonstrator's decisions are affected by the context variable at positions $[0, 2]$ and $[4, 2]$, we observe that the naive estimation is biased from the true transition probability. Figure 7 shows the suboptimal policy obtained using the naive estimation $P(s'|s, a)$. We observe that the expected return of the suboptimal policy (black dashed line) is over-optimistic compared to the true optimal return (green dashed line), while the actual return (red curve) we obtain by implementing this suboptimal policy is far lower than the true optimal return. This is because the naive estimation over-estimates the probability that the robot can move upwards to the red cell from position $[0, 2]$ by taking action $a = 1$.

Next, we compute the causal bounds on value functions using the causal bounds in Table IV and use these value function bounds in Algorithms 1 and 2. The performance of Algorithm 1 compared to standard Q learning is shown in Figure 7 and the performance of Algorithm 2 compared to the UCB- Q learning method is shown in Figure 8. The improvement in performance obtained using Algorithms 1 and 2 is similar to that for the contextual reward case.