

Confounding Robust Continuous Control via Automatic Reward Shaping

Mateo Juliani*
Columbia University
New York, NY, United States
msj2164@columbia.edu

Mingxuan Li*
Columbia University
New York, NY, United States
ml@cs.columbia.edu

Elias Bareinboim
Columbia University
New York, NY, United States
eb@cs.columbia.edu

ABSTRACT

Reward shaping has been applied widely to accelerate Reinforcement Learning (RL) agents' training. However, a principled way of designing effective reward shaping functions, especially for complex continuous control problems, remains largely under-explained. In this work, we propose to automatically learn a reward shaping function for continuous control problems from offline datasets, potentially contaminated by unobserved confounding variables. Specifically, our method builds upon the recently proposed causal Bellman equation to learn a tight upper bound on the optimal state values, which is then used as the potentials in the Potential-Based Reward Shaping (PBRS) framework. Our proposed reward shaping algorithm is tested with Soft-Actor-Critic (SAC) on multiple commonly used continuous control benchmarks and exhibits strong performance guarantees under unobserved confounders. More broadly, our work marks a solid first step towards confounding robust continuous control from a causal perspective.

Code: https://github.com/mateojuliani/confounding_robust_control

Full version: <https://arxiv.org/abs/2602.10305>

CCS CONCEPTS

• **Theory of computation** → **Reinforcement learning**.

KEYWORDS

Causal Inference; Reinforcement Learning; Unobserved Confounder

ACM Reference Format:

Mateo Juliani, Mingxuan Li, and Elias Bareinboim. 2026. Confounding Robust Continuous Control via Automatic Reward Shaping. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 20 pages. <https://doi.org/10.65109/UUEW5708>

1 INTRODUCTION

Reinforcement learning (RL) has demonstrated impressive success in continuous control domains such as robotic manipulation, locomotion, and autonomous systems [25, 40, 41]. Despite this progress, learning effective policies in high-dimensional, complex environments remains challenging due to sample inefficiency and high

sensitivity to reward design. When the original task reward is not efficient to learn from, reward shaping can significantly accelerate learning by injecting informative signals that guide exploration and policy improvement. However, designing effective shaping functions remains a persistent challenge, often requiring substantial domain expertise and manual effort to ensure they are helpful without hurting the performance [10, 31, 33, 38].

Potential-Based Reward Shaping (PBRS) [31] offers a principled framework for injecting additional reward signals while preserving the original task's optimal policy. However, the effectiveness of PBRS critically depends on the quality of the potential function used. Recent work has explored learning state potentials automatically from offline data [4, 22, 56], but such approaches typically assume no unobserved confounders (NUC) [28, 29], i.e. access to fully observed (unconfounded) trajectories. Such an assumption can easily break down in many real-world settings. Unobserved confounders can arise from human demonstrations, legacy systems, or sensor capability differences in robotic platforms. When the NUC assumption is violated, the effects of candidate policies become generally unidentifiable. That is, the given model assumptions are insufficient to uniquely recover the value function from offline data, regardless of the sample size [34, 52]. As a result, standard RL methods relying implicitly on NUC can suffer from degraded learning performance in such settings. More recently, Li et al. [19] propose to use partial identification approach to learn confounding robust shaping functions. But their proposed method is limited to discrete and lower dimensional settings. A practical solution for continuous and higher dimensional environments is yet to be discussed.

In this work, we tackle the problem of automatic reward shaping in continuous control settings where offline data may be confounded. Specifically, we utilize offline trajectories collected from unknown, potentially biased behavioral policies to estimate causal upper bounds on the optimal interventional state values. These bounds are then employed as potential functions within the PBRS framework to construct shaping rewards. By incorporating these shaped rewards, we enable model-free RL agents to perform more informed exploration and policy learning, even in the presence of unobserved confounders. We empirically evaluate our framework in confounded MuJoCo and Adroit environments with partial observability. Experiments on a suite of confounded continuous environments show that our method consistently outperforms unshaped and causally unaware shaping baselines (CQL-Shaping, [15], T-Rex-Shaping [3]). These results highlight the robustness and practical effectiveness of our approach in real-world settings.

Our main contributions are as follows:

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

- We derive a Causal Bellman Equation for the stationary infinite-horizon Confounded Markov Decision Process (CMDP), which converges to a tight upper bound of the optimal state values;
- We design a neural learning algorithm that approximates the Causal Bellman Equation in high-dimensional continuous state-action CMDPs;
- We use the learned state value bound as the reward shaping function and empirically demonstrate the superior performance boost when applying to Soft-Actor-Critic [6] in various challenging confounded continuous control environments.

2 BACKGROUND

Confounding Robust Decision-making. We will focus on the sequential decision-making setting in an infinite horizon stationary Markov Decision Process (MDP, Puterman [36]) where the agent intervenes on a sequence of actions X_1, \dots to optimize the cumulative return over reward signals Y_1, \dots given state observations S_1, S_2, \dots at each corresponding time step. Standard MDP formalism focuses on the perspective of the learners who could actively intervene in the environment. Consequently, the data collected from randomized experiments is free from the contamination of unobserved confounding bias and is generally assumed away in the model. However, when considering offline data collected by passive observation [15, 17, 21] where the learner may not necessarily have deliberate control over the behavioral policy generating the data, or when the state attributes are partially observed [11, 13, 46], unobserved confounding arises. Consequently, this could lead to biased estimation and safety/alignment issue in various reinforcement learning tasks, including off-policy learning [14, 19–21, 54], curriculum learning [18], and imitation learning [16, 55].

Continuous Control with Deep Reinforcement Learning. In continuous action space, the sample efficiency problem is exacerbated, rendering commonly used on-policy learning solutions unfavorable, such as TRPO [39], PPO [41] or A3C [24]. At each time step, on-policy algorithms collect new trajectories from the environment only for updating the agents by a single gradient step. As task complexity grows, this procedure becomes increasingly expensive. Off-policy algorithms, on the other hand, reuse past experiences. The direct application of this idea is DQN and its variants [25]. For continuous policy learning, actor-critic based method is preferred for its stability and easy-to-tune hyper-parameters [6, 40]. In this work, we use Soft-Actor-Critic (SAC), a maximum entropy reinforcement learning framework which improves upon the traditional maximum reward framework with substantially better exploration and robustness [5, 57], as our base learning algorithm.

Potential-based reward shaping (PBRS). Reward shaping is a popular line of techniques for incorporating domain knowledge during policy learning. Common approaches such as Potential-Based Reward Shaping (PBRS, Ng et al. [31]) add supplemental signals to the reward function so that it would be easier to learn in future downstream tasks without affecting the optimality of the learned policy. PBRS modifies the reward function in the system by adding the discounted next state potential subtracted by current state potential. This encourages the learning agent to visit states with higher potentials while avoiding visiting states with low potentials. More

importantly, optimal policies remain invariant across this shaping process, i.e., every optimal policy learned in the MDP under PBRS is guaranteed to be optimal in the original MDP, and vice versa.

Notations. We will consistently use capital letters (V) to denote random variables, lowercase letters (v) for their values, and cursive \mathcal{V} to denote their domains. Fix indices $i, j \in \mathbb{N}$. We use bold capital letters (\mathbf{V}) to denote a set of random variables and let $|\mathbf{V}|$ denote its cardinality of the set \mathbf{V} . Finally, $\mathbf{1}_{Z=z}$ is an indicator function that returns 1 if event $Z = z$ holds true; otherwise, it returns 0.

3 THE CHALLENGE OF CONFOUNDED CONTINUOUS CONTROL

In real-world continuous control tasks, agents may have limited sensor information due to cost or environment constraints when deployed in online operation. This leads to poor sample efficiency if the agent is to be tuned fully online. Pre-training with offline data is a common remedy. But people generally expect the behavioral policy generating the offline datasets to have the same sensor capability as the online agent [15, 17, 30]. However, offline data might be collected by agents in a more controlled environment with privileged sensor capability [7, 8, 26]. This mismatch induces unobserved confounding at the decision level: offline decisions reflect privileged information unavailable to the online learning agent. As a result, directly learning from such data without addressing confounding can lead to biased policies and poor online performance. Tackling confounded continuous control is thus critical for reliable and sample-efficient RL under realistic sensory constraints.

In this paper, we consider an extended family of MDPs that explicitly models the presence of unobserved confounders when generating offline data.

Definition 3.1. A Confounded Markov Decision Process (CMDP) \mathcal{M} is a tuple of $\langle \mathcal{S}, \mathcal{X}, \mathcal{Y}, \mathcal{U}, \mathbb{F}, \mathbb{P} \rangle$ where,

- $\mathcal{S}, \mathcal{X}, \mathcal{Y}$ are, respectively, the space of observed states, actions, and rewards;
- \mathcal{U} is the space of unobserved exogenous noise;
- \mathbb{F} is a set consisting of the transition function $\tau : \mathcal{S} \times \mathcal{X} \times \mathcal{U} \mapsto \mathcal{S}$, behavioral policy $\beta : \mathcal{S} \times \mathcal{U} \mapsto \mathcal{X}$, and reward function $r : \mathcal{S} \times \mathcal{X} \times \mathcal{U} \mapsto \mathcal{Y}$;
- \mathbb{P} is a set of distributions P over the unobserved domain \mathcal{U} .

To model general continuous control problems, we assume the space of states, \mathcal{S} , actions, \mathcal{X} , and unobserved exogenous noise, \mathcal{U} , to be multi-dimensional and continuous throughout the paper. Consider a demonstrator agent interacting with a CMDP. For every time step $h = 1, \dots$, the nature draws an exogenous noise U_h from the distribution $P(\mathcal{U})$; the demonstrator performs an action $X_h \leftarrow f_X(S_h, U_h)$, receives a subsequent reward $Y_h \leftarrow f_Y(S_h, X_h, U_h)$, and moves to the next state $S_{h+1} \leftarrow f_S(S_h, X_h, U_h)$. The observed trajectories of the demonstrator (from the learner’s perspective) are thus summarized as the observational distribution $P(\bar{X}, \bar{S}, \bar{Y})$.¹ In the data-generating process described above, for every time step h , the exogenous noise U_h becomes an unobserved confounder affecting the action X_h , reward Y_h , and next state S_{h+1} simultaneously. Therefore, CMDP is also referred to MDP with Unobserved Confounders

¹We will consistently use $\bar{X}, \bar{S}, \bar{Y}$ to represent trajectory sequences.




Figure 1: (a) CMDP causal diagram of the offline data generating process; (b) CMDP causal diagram under policy $\text{do}(\pi)$ during the online learning process. Compared with the standard MDP, the highlighted bi-directed dashed arrows represent confounders affecting both behavioral policy, state transitions and rewards while being unobservable to the online agents.

(MDPUC, [53]) and is a subclass of Confounded Partially Observed MDP [2, 23, 42] where Markov property holds.

The observed distribution of such an offline collected dataset with finite trajectories up to H steps can be written as,

$$P(\bar{X}_{1:H}, \bar{S}_{1:H}, \bar{Y}_{1:H}) = P(s_1) \prod_{h=1}^H \left(\int_{\mathcal{U}} \mathbf{1}_{s_{h+1}=f_S(s_h, x_h, u_h)} \mathbf{1}_{x_h=f_X(s_h, u_h)} \mathbf{1}_{y_h=f_Y(s_h, x_h, u_h)} P(u_h) du_h \right) \quad (1)$$

The causal diagram in Fig. 1a showcases how the confounders are affecting state transitions, reward, and the behavioral agent’s policy while the online learning agent doesn’t have such information when acting in the environment Fig. 1b. By convention [34], we use bi-directed arrows (e.g., $X_h \leftrightarrow S_h$) to indicate the presence of unobserved confounders, U_h , affecting actions, states and rewards.

During the online learning phase, as shown in the causal diagram in Fig. 1b, the agent intervenes on the action variable following a policy $\pi(x_h|s_h)$ that maps from state to a distribution over the action domain \mathcal{X} . This is denoted as policy intervention $\text{do}(\pi)$ replacing the behavioral policy f_X during the offline data collection phase. The online trajectory distribution in CMDP under π , \mathcal{M}_π is,

$$P_\pi(\bar{X}_{1:H}, \bar{S}_{1:H}, \bar{Y}_{1:H}) = P(s_1) \prod_{h=1}^H \left(\pi(x_h|s_h) \mathcal{T}(s_h, x_h, s_{h+1}) \mathcal{R}(s_h, x_h, y_h) \right) \quad (2)$$

where the transition distribution \mathcal{T} and the reward distribution \mathcal{R} are given by,

$$\mathcal{T}(s_h, x_h, s_{h+1}) = \int_{\mathcal{U}} \mathbf{1}_{s_{h+1}=f_S(s_h, x_h, u_h)} P(u_h) du_h \quad (3)$$

$$\mathcal{R}(s_h, x_h, y_h) = \int_{\mathcal{U}} \mathbf{1}_{y_h=f_Y(s_h, x_h, u_h)} P(u_h) du_h \quad (4)$$

In CMDP, the learning goal is still to find the optimal policy π^* that maximizes the cumulative return (often discounted by $\gamma \in (0, 1)$). That is, $\pi^* = \arg \max_{\pi} \mathbb{E}[\sum_{h=1}^{\infty} \gamma^{h-1} y_h]$. This objective function can be solved iteratively using the Bellman Optimality Equation [36],

$$Q^*(s, x) = \mathbb{E}[Y_h + \gamma \max_{x'} Q^*(S_{h+1}, x') | S_h = s, X_h = x] \quad (5)$$

where $Q^*(s, x) = \mathbb{E}[\sum_{t=1}^{\infty} \gamma^{t-1} y_{h+t} | S_h = s, X_h = x]$ by definition is the optimal state action value function denoting the best return after taking action x in state s .

In the offline to online learning setting, one would envision that learning from offline datasets generated by a competitive policy with a good action space coverage should yield near-optimal online policies [15, 30]. This is indeed true under the MDP definition where there are no unobserved confounders. The state transitions and reward functions can be easily identified by the observation distribution of the offline dataset,

$$\mathcal{T}(s_h, x_h, s_{h+1}) = P(s_{h+1} | s_h, x_h) \quad (6)$$

$$\mathcal{R}(s_h, x_h, y_h) = P(y_h | s_h, x_h). \quad (7)$$

When the above identification formula hold, several off-policy algorithms have been proposed to estimate the effect of candidate policies from finite observations [12, 15, 27, 30, 35, 45, 49, 50]. Together with deep learning, these methods could be further extended to complex domains [25, 32, 40, 41, 43]. However, NUC could be fragile in practice and does not necessarily hold due to violations like sensory mismatch in data generation process. In these situations, applying standard off-policy methods may fail to converge to optimal, despite using powerful deep learning models.²

To witness the challenge of confounded continuous control, we instantiate a confounded variant of Hopper [48] where the offline agent has access to the full state observation, a total of 11-dimension vectors while the online agent can observe all of the dimension except the second dimension, the angle of its thigh joint (comparable to the knee in humans). Not observing the position of its thigh joint presents a genuine challenge to the agent given it would be unable to know to how move its thigh joint to generate the power to hop. As shown in Fig. 2 the vanilla SAC [6] agent trained under partial observation cannot recover 1/2 of the performance of the agent trained with complete state observations.

These results expose a critical barrier to reliable learning under restricted sensory capabilities: without access to key hidden factors, even state-of-the-art off-policy RL methods will suffer. And adding good quality offline datasets which supposedly contain valuable information on the optimal state values does not help. As demonstrated in Fig. 2, commonly used algorithms also fall short under the confounded setting due to their fundamental limitation of assuming NUC implicitly followed from the (PO)MDP definition. Then, our central research question is: Can we extract and transfer such contaminated offline knowledge to guide online learning? The answer

²NUC is orthogonal to the issue of action space coverage discussed by the line of research on conservative Q-learning (CQL) [15]. As shown in Fig. 2, CQL learned state values cannot handle the unobserved confounding issue in offline datasets.

is yes. In the next section, we propose the causal bellman equation from which automatic reward shaping with confounded offline data is made possible to facilitate online learning.

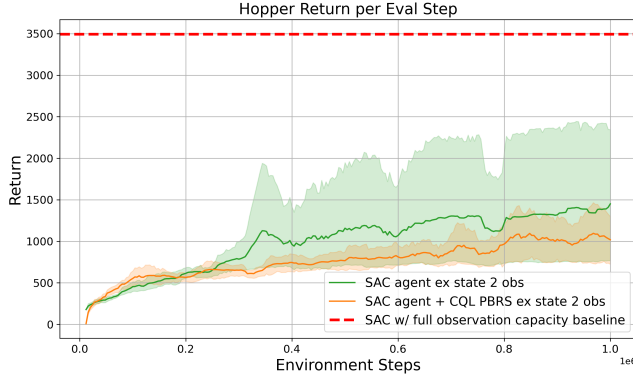


Figure 2: Performance of Hopper SAC Agent with full capacity and SAC agent unable to observe state 2.

4 CONFOUNDED CONTINUOUS CONTROL WITH SHAPED REWARDS

In this section, we will introduce how we learn an optimistic state potential from confounded continuous offline data, which is then used for online fine-tuning. See Sec. 7 in the appendix for proof details of theorems discussed in this section.

4.1 Learn Optimistic State Potentials via Confounding Robust Offline Pretraining

It is well acknowledged that a good state potential function is the optimal state value [31]. But without training the agent, one cannot have easy access to the optimal state values. The premise of using state values estimated from offline dataset as state potentials is that such values are close to optimal state values. However, without deliberate control on the quality of the behavioral policy nor the NUC condition, such offline learned values could be heavily biased and cannot be used for reward shaping as we have seen in Fig. 2. Recently Li et al. [19] proposes to use partial identification methods to upper bound the optimal state values for finite horizon non-stationary CMDPs from such confounded offline datasets robustly. Here we extend the results to stationary infinite-horizon CMDPs.

Theorem 4.1 (Causal Bellman Optimal Equation for Stationary Infinite-Horizon CMDPs). *For a CMDP environment \mathcal{M} with reward $Y_h \leq b, b \in \mathbb{R}$, the optimal value of interventional policies, $V^*(s), \forall s \in \mathcal{S}$, is upper bounded by $V^*(s) \leq \bar{V}(s)$ satisfying the Causal Bellman Optimality Equation,*

$$\bar{V}(s) = \max_x \left[P(x|s) \left(\tilde{\mathcal{R}}(s, x) + \gamma \mathbb{E}_{\tilde{\mathcal{T}}}[\bar{V}(s')] \right) + P(-x|s) \left(b + \gamma \max_{s'} \bar{V}(s') \right) \right] \quad (8)$$

where $\tilde{\mathcal{R}}$ is offline estimated reward distribution and $\tilde{\mathcal{T}}$ is the estimated transition distribution.

Compared with the original Bellman Optimal Equation, the Causal Bellman Equation accounts for the uncertainty brought by confounders in the offline dataset via an extra term,

$$(b + \gamma \max_{s'} \bar{V}_{h+1}(s')) \quad (9)$$

which represents the best return that the agent could have achieved from those “unselected” actions, i.e., $P(-x|s)$. With the Causal Bellman Optimal Equation, we can robustly upper bound the optimal state values from a confounded offline dataset generated by CMDP \mathcal{M} . Next we show that the extended Causal Bellman Optimal Equation converges to a unique fixed point, which is a valid upper bound on the optimal interventional state values for online agents (Fig. 1b).

Theorem 4.2 (Convergence of Causal Bellman Optimal Equation). *The Causal Bellman Optimality Equation converges to a unique fixed point, which is also an upper bound on the optimal interventional state values under the assumption that $P(s, x) > 0, \forall s, x$ in the stationary infinite horizon CMDP \mathcal{M} .*

The proposed Causal Bellman Equation does not translate directly into a practical algorithm for high dimensional continuous control problems with function approximators, though. The propensity score of actions $P(x|s), P(-x|s)$ is ill-defined as any single point in a continuous distribution has zero probability. Furthermore, naively enumerating all the states to calculate $\max_{s'} \bar{V}(s')$ or enumerating actions to get values from Q-values is intractable in continuous state and action space. Thus, we make several practical approximations when implementing the Causal Bellman Equation.

Firstly, we parametrize the causal upper bound state potential to be $V_{\theta_1}(\cdot)$ and its corresponding target network $V_{\theta_1'}(\cdot)$ for smoother learning updates [25]. We also replace the outside \max_x in Eq. (8) with an expectation over the observed action distribution in the offline dataset under the condition that the behavioral policy is competitive. Then, we restrict the observed policy distribution to be within a tractable class of Gaussian policies as P_{θ_2} . To train this policy distribution, we maximize the likelihood of the observed actions given states in the offline dataset \mathcal{D} ,

$$J(\theta_2) = \mathbb{E}_{\mathcal{D}} \left[\log P_{\theta_2}(x|s) \right] \quad (10)$$

When training V_{θ_1} , we apply the reparametrization trick to make sampling from P_{θ_2} differentiable,

$$x = f_{\theta_2}(\epsilon, s) \quad (11)$$

where ϵ is an input noise vector sampled from a fixed standard Gaussian distribution. Lastly, instead of calculating the global best possible next state $\max_{s'} \bar{V}(s')$, we aim at finding the best possible nearby states if an action $x' \neq x$ had been taken. Thus, we choose to model the difference distribution between the current state s and next state s' given state action pair (s, x) . Similarly, we restrict this distribution to be Gaussian, P_{θ_3} , and the training objective is to maximize the log likelihood of observed state differences in \mathcal{D} ,

$$J(\theta_3) = \mathbb{E}_{\mathcal{D}} \left[\log P_{\theta_3}(\Delta_s | s, x) \right] \quad (12)$$

where Δ_s is the state difference between s and s' . And we apply the reparametrization trick again for differentiable sampling,

$$\Delta_s = f_{\theta_3}(\epsilon, s, x). \quad (13)$$

Algorithm 1 Neural Causal Upper Bound State Potential

```

1: Initialize parameters  $\theta_1, \theta'_1, \theta_2, \theta_3$ 
2: while Not Converged do
3:   Sample an offline batch  $\{(s_i, x_i, s'_i, y_i)\}_{i=1}^B$ 
4:   Update observed policy,  $\theta_2 \leftarrow \theta_2 + \lambda_2 \hat{\nabla} J(\theta_2)$  (Eq. (10))
5:   Update state difference,  $\theta_3 \leftarrow \theta_3 + \lambda_3 \hat{\nabla} J(\theta_3)$  (Eq. (12))
6: end while
7: while Not Converged do
8:   Sample an offline batch  $\{(s_i, x_i, s'_i, y_i)\}_{i=1}^B$ 
9:   Update state potential,  $\theta_1 \leftarrow \theta_1 - \lambda_1 \hat{\nabla} J(\theta_1)$  (Eq. (15))
10:  Update target networks,  $\theta'_1 \leftarrow \tau \theta_1 + (1 - \tau) \theta'_1$ 
11: end while
12: return  $V_{\theta_1}(s)$ 

```

Now we can approximate the Causal Bellman Equation backup, \mathcal{B} , with parametrized neural network components as,

$$\hat{\mathcal{B}}V_{\theta_1}(s) = \mathbb{E}_{\mathcal{D}} \left[P_{\theta_2}(x|s) \left(y + \gamma \bar{V}_{\theta_1}(s') \right) + P_{\theta_2}(x'|s) \left(\max_{y \in \mathcal{D}} y + \gamma \bar{V}_{\theta_1}(s + \Delta_s) \right) \right] \quad (14)$$

where $x' = \arg \max_{x'' \sim P_{\theta_2}(\cdot|s), x'' \neq x} V_{\theta_1}(s + \Delta'_s)$, $\Delta'_s \sim P_{\theta_3}(\cdot|s, x')$ is the action that maximize the return of “the road not taken”, and $\Delta_s \sim P_{\theta_3}(\cdot|s, x')$ is the state difference sampled given s, x' . The extra compensation term in Eq. (9) is approximated by the maximum reward observed in the offline dataset plus the discounted best next state value as if the agent had taken x' and transited to a state s' that is Δ_s away from s . In implementation, we also approximate the action propensity score with its corresponding Gaussian density to avoid the zero probability issue. The state potential function \bar{V}_{θ_1} is then trained to minimize the squared residual error,

$$J(\theta_1) = \mathbb{E}_{\mathcal{D}} \left[\frac{1}{2} \left(\bar{V}_{\theta_1}(s) - \hat{\mathcal{B}}V_{\theta_1}(s) \right)^2 \right] \quad (15)$$

See Algo. 1 for the full pseudo-code of learning the state potentials from continuous confounded offline data. In the next section, we will illustrate how to use this learned state potentials as reward shaping functions during online training with SAC.

4.2 Online Fine-tuning with Reward Shaping

We first re-establish the optimal policy invariance property of PBRS [31] in infinite-horizon stationary CMDPs.

Proposition 4.3. *For a CMDP under policy π , \mathcal{M}_{π} , let \mathcal{M}'_{π} be the CMDP obtained from \mathcal{M}_{π} by replacing the reward with the following function, for every time step $h = 1, \dots, H$,*

$$y'_h := y_h + \gamma \phi_h(S_{h+1}) - \phi_h(S_h), \quad (16)$$

where y_h is the original reward returned by \mathcal{M}_{π} ; $\phi_h(\cdot) : \mathcal{S} \mapsto \mathbb{R}$ is a real valued potential function. Then every optimal policy in \mathcal{M}_{π} will also be an optimal policy in \mathcal{M}'_{π} , and vice versa.

This guarantees that the agent could still obtain the same optimal policy as the agent without using reward shaping.

For the learning algorithm, we choose to use a popular value-based maximum entropy learner, SAC [6], for its amenity to PBRS.

Because for policy gradient based methods like PPO [41], multi-step returns are commonly adopted when calculating advantages [40] resulting in the intermediate shaping terms being canceled out without affecting learning process.

The original SAC algorithm learns a soft policy that minimizes the following objective function:

$$J(\phi) = \mathbb{E}_{\mathcal{D}} \left[D_{\text{KL}} \left(\pi_{\phi}(\cdot|s) \left\| \frac{\exp(Q_{\theta}(s, \cdot))}{Z_{\theta}(s)} \right. \right) \right] \quad (17)$$

where the parametrized policy π_{ϕ} is updated towards the exponential of the learned Q function Q_{θ} and $Z_{\theta}(s)$ normalizes the distribution. With the reward shaping function defined in Eq. (16), the Q value function Q_{θ} is now trained to track returns with the shaped reward y' instead of y while other parts of SAC stay unchanged.

5 EXPERIMENTS

In this section, we evaluate the efficacy of the state value upper bounds learned in Algo. 1 by applying them as potentials in Potential-Based Reward Shaping (PBRS, [31]) to augment an SAC agent’s training rewards in a suite of confounded, continuous control environments. We compare our results to 1) a baseline online SAC agent that does not use PBRS, 2) CQL PBRS [15] that uses offline learned conservative Q-values as shaping potentials in PBRS, 3) an online SAC agent using recurrent networks, and 4) T-REX PBRS [3] which uses inverse reinforcement learning to learn a reward function, which we use as the potential. We provide further commentary on which cases the causal reward shaping function can provide the most improvement to online training, and how offline data quality affects performance.

5.1 Experiment Design

Environments and Offline Data: We selected six continuous control tasks from the Gymnasium [48] environment to evaluate our causal PBRS method. We selected four MuJoCo [47] tasks on robotic locomotion including 1) Hopper - control a single legged agent to hop in 2D space without falling, 2) HalfCheetah - control a two legged, Cheetah-like, agent to run forward in 2D space as fast as possible, 3) Walker2D - control a bipedal agent to walk forward in 2D space without falling over, and 4) Ant - control a quadruped robot to walk forward in 3D space. We also selected two tasks from Adroit [37] on controlling a robotic hand to open a door (Door) and to move a ball to the target location (Relocate). We use the offline datasets from Minari [51]. For details, see Sec. 8 in the appendix.

Confounded Environment Setup: To simulate unobserved confounders in our experiments, we remove dimensions from the environment’s observation space³. As a result, we simulate trajectories generated by an agent with access to richer sensory information (the removed dimensions), which is not accessible to the online agent. Consequently, typical off-policy learners cannot identify the behavioral policy or the state / q-value function accurately, given the unobserved confounders in the offline datasets [34, 52]. In practice, removing some of the dimensions will render the task much more challenging than removing others. For example, removing the agent’s velocity observation in the MuJoCo environments - which

³See Sec. 12 in the full version for tests showing this method creates confounding bias between the behavior policy and state transition as depicted in Fig. 1a

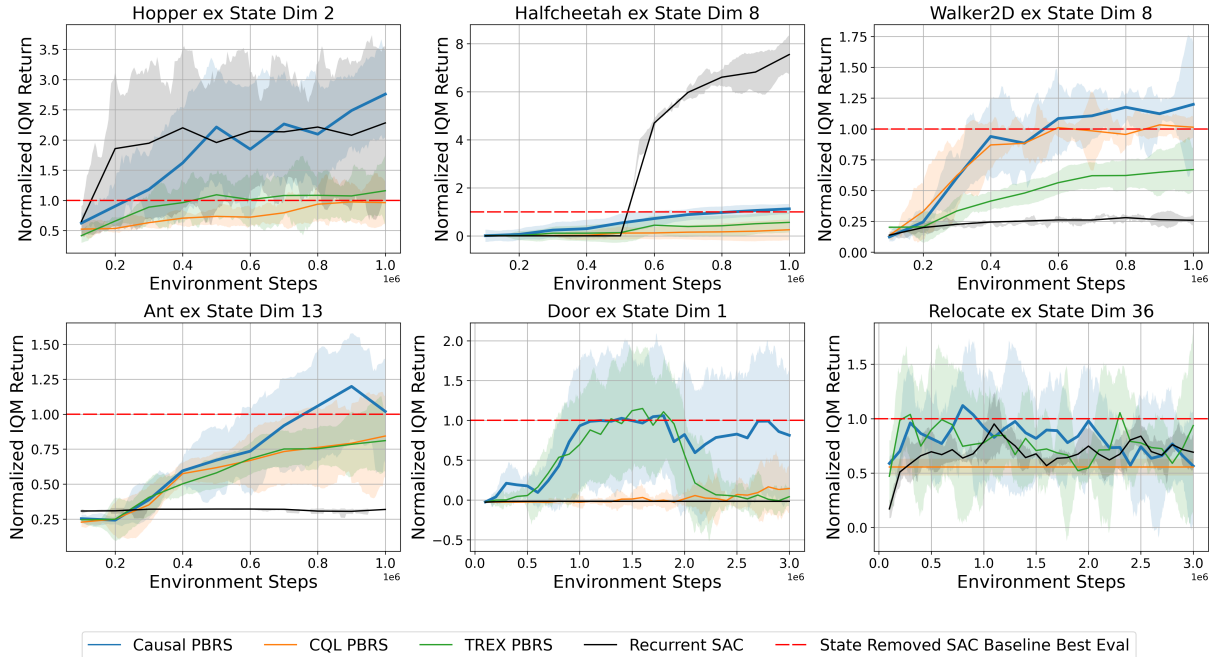


Figure 3: Normalized IQM returns w.r.t State Removed Baseline SAC agent in confounded continuous control benchmarks.

is a key input into the environment’s reward function - would make it difficult to derive the useful state value function from the offline data. On the other hand, removing an observation dimension that is independent of the state / Q-value function would not affect standard off-policy learners significantly. Therefore, to understand which dimensions to remove to create meaningful unobserved confounders, we use the Randomized Conditional Independence Test (RCIT, [44]) to measure the independence degree between an observation dimension and the episode’s reward-to-go, conditioned on the remaining observations and actions. For detailed analysis on the relations between the removal of an observation dimension and the online agent’s performance, see Sec. 5.3.

Offline and Online Training: Once we remove selected state dimensions from the offline dataset, we train our state value upper bounds following Algo. 1. We train the environment models for 50 epochs, and then train the state value upper bounds for 200 epochs. To prevent over-estimation (in particular from the "road not taken" estimate $\gamma \bar{V}_{\theta_1}(s + \Delta_s)$), we clip state value functions using $\frac{\max_{y \in D} y}{1-\gamma}$, the theoretical max value of the value function given the max observed reward in the offline data. As a comparison to our causal shaping functions, we also estimate Q-value functions using CQL [15], and a reward function using T-REX [3].

After training the offline state value functions, we test five online learners: 1) a baseline SAC where we remove the selected state dimension from the observation space (SAC Baseline), 2) a PBRS-based approach where we use the learned Q-value function from CQL (CQL PBRS), 3) a PBRS-based approach using T-REX (T-REX PBRS), 4) an SAC agent using recurrent networks (Recurrent SAC) and 5) a PBRS-based approach where the potentials are the learned state value upper bounds (Causal PBRS, ours). For the Causal, T-REX and CQL PBRS method, we apply a scaled version of the shaping rewards detailed in Prop. 4.3. Specifically, we use $y_{env} +$

$\beta(\gamma\phi(s') - \phi(s))$ as the shaping reward, where y_{env} is the observed environment reward, β is a scaling factor, ϕ is the potential function, and s is the current state and s' is the next state. Empirically, we find the agent benefits the most from the potentials when $\beta \leq 1$.

5.2 Causal Reward Shaping Performance

Fig. 3 shows the interquartile mean (IQM) returns normalized by the state removed SAC baseline’s best evaluation score in each environment. Table 1 provides the returns (including the best average return reached by the agent and the return at the final evaluation step), normalized mean, median, and inter-quartile mean return relative to the SAC baseline under partial state observation.

Overall, the Causal PBRS method outperforms the causally unaware reward shaping baselines (CQL, T-REX Shaping) consistently for an average normalized mean score of 1.29 and normalized IQM of 1.09. In particular, the Causal PBRS method performs the best in cases where the removal of a certain state dimension leads to a larger decrease in the SAC baseline performance. For example, the Causal PBRS method outperforms the SAC baseline in the Hopper environment ex. State 1 and 2, where the SAC baseline does not surpass a score of 2000, versus Hopper ex. State 5, where the SAC baseline performance of around 3300 is comparable to the performance of the Hopper with full observational capabilities [6]. However, we note there is a limit, given that when removing states 1-4 in the Hopper environment (which are the dimensions related to joint positions), both the Causal PBRS and other methods are unable to learn a meaningful policy. Unsurprisingly, the causally unaware CQL/T-REX PBRS methods largely underperforms the SAC baseline (normalized mean of 0.85/0.96), given that the unobserved confounding in the offline reward signal leads to highly biased Q-value estimations. We note the recurrent SAC method is able

Table 1: Best evaluation returns of agents (± 1 standard deviation) on the 18 confounded environments. 1M steps for MuJoCo, 3M for Adroit). All results are averaged over 5 seeds (except Recurrent SAC, which is over 3 seeds), and smoothed over 10 evals. Bold Numbers indicate best best-performing methods. Our Causal PBRS method significantly outperforms the baselines.

Environment	Full State SAC	State Dim Removed	SAC Baseline	CQL	T-REX	Recurrent SAC	Causal PBRS (ours)
Hopper-v5	3500	1	2073 \pm 1424	1778 \pm 1045	1276 \pm 835	562 \pm 72	2252 \pm 1081
		2	1450 \pm 1123	1096 \pm 481	1344 \pm 628	2432 \pm 1067	2762 \pm 836
		3	3283 \pm 127	1478 \pm 444	2780 \pm 859	1601 \pm 1548	3582 \pm 37
		5	3321 \pm 136	1038 \pm 19	2862 \pm 838	2878 \pm 573	3183 \pm 518
		7	2935 \pm 467	1245 \pm 291	2550 \pm 976	2693 \pm 1147	3011 \pm 332
		5,6	655 \pm 155	1143 \pm 728	823 \pm 571	1978 \pm 1347	1037 \pm 210
		1,2,3,4	304 \pm 30	557 \pm 244	310 \pm 67	460 \pm 95	305 \pm 15
		7,8,9,10	372 \pm 19	422 \pm 43	367 \pm 62	3468 \pm 124	1442 \pm 1449
Halfcheetah-v5	12400	1	2013 \pm 32	2039 \pm 31	2059 \pm 30	7957 \pm 645	2055 \pm 24
		3	8931 \pm 1489	9748 \pm 177	9120 \pm 882	10057 \pm 555	8954 \pm 665
		8	896 \pm 494	373 \pm 454	615 \pm 443	7939 \pm 851	1103 \pm 216
Walker2D-v5	4050	6	3981 \pm 154	3243 \pm 283	4045 \pm 306	2210 \pm 1636	3899 \pm 587
		8	3640 \pm 208	3766 \pm 259	2830 \pm 832	1105 \pm 51	4295 \pm 373
		10	3925 \pm 137	3600 \pm 428	4417 \pm 445	1351 \pm 1285	3950 \pm 397
Ant-v5	4000	13	3084 \pm 607	2792 \pm 1013	2830 \pm 832	994 \pm 4	3389 \pm 1137
		16	3750 \pm 896	3190 \pm 464	3662 \pm 1693	1225 \pm 448	2963 \pm 1369
AdroitHandDoor-v1	NA	1	1472 \pm 1385	415 \pm 628	1659 \pm 1189	-27 \pm 1	1725 \pm 1396
AdroitHandRelocate-v1	NA	36	24 \pm 9	13 \pm 2	29 \pm 11	23 \pm 7	30 \pm 13
Normalized Mean (\uparrow)			1.00	0.85	0.96	1.93	1.29
Normalized Median (\uparrow)			1.00	0.85	0.98	0.89	1.09
Normalized IQM (\uparrow)			1.00	0.81	0.96	0.88	1.09

to overcome the confounding bias for some of the simpler environments (Hopper-v5 and Halfcheetah-v5), although results are inconsistent. For more complex and higher dimensional environments (Ant-v5, Relocate-v1, and Walker-v5), the recurrent SAC baseline is unable to learn a useful policy. In each of the following sections, we provide an overview of some of the observation dimensions removed in each environment and the agent’s performance in those environments. For full commentary on all the dimensions removed, see Sec. 14.

5.2.1 Hopper. We removed the following dimensions:

- State Dim 1: The Hopper’s torso angle. One of the termination conditions in Hopper is if the torso angle is bounded between $[-0.2, 0.2]$. Without this dimension, the agent is unable to know how to adjust its torso angle to maintain healthy body positions. The Causal PBRS can improve on the baseline by around 20%, whereas the other methods under perform the baseline by 15% (CQL) to 73% (Recurrent SAC).
- State Dim 2: Angle of the thigh joint. Removing this dimension has a large impact on the Hopper’s performance, reducing its baseline return by 50% vs the Hopper’s performance with full obs capacity using SAC. The Causal PBRS method can drive a large improvement of almost 100% vs the baseline.
- State Dim 1,2,3,4: All of the angular positions of the Hopper’s joints and torso. Without knowing the position of its joints, the Hopper agent is unable to learn a meaningful policy. The Causal PBRS method is similarly unable to learn a meaningful

policy, suggesting a limit to the Causal PBRS’s ability to improve performance in highly confounded environments. The CQL/Recurrent SAC method can generate a slightly better policy, although we note the overall return is still low.

5.2.2 HalfCheetah. We removed the following dimensions:

- State Dim 1: Angle of the front tip. Without the angle of the front tip (used to orient the Cheetah), the Half Cheetah agent’s return drastically decreases from 12,000 to around 2,000. The Recurrent SAC method outperforms all other methods.
- State Dim 8: Velocity of the x-coordinate of the front tip. Unlike the Hopper, removing the forward velocity greatly reduces the HalfCheetah’s return. The Causal PBRS method achieves a higher return vs the baseline policy, however Recurrent SAC method greatly outperforms all other methods.

5.2.3 Walker2D. We removed the following dimensions:

- State Dim 8: Velocity of the x-coordinate of the torso. Removing the forward velocity dimension of the Walker2d slightly decreases performance. The Causal PBRS method achieves a 20% improvement over the second best method (SAC baseline).

5.2.4 Ant. We removed the following dimensions:

- State Dim 13: Velocity of the x-coordinate of the torso. Similar to other environments, the Causal PBRS method helps the Ant recover some of the performance loss from the loss of its forward velocity observation, unlike the CQL PBRS method.

5.2.5 Adroit Door. We removed the following dimensions:




Figure 4: RCIT test statistic v.s. Causal PBRS improvements.

- State Dim 1: Angular position of the horizontal arm joint. The baseline agent and Causal PBRS agent perform comparably. The Causal PBRS method outperforms the baseline by 17%, and avoids catastrophic forgetting (see Fig. 11 for details).

5.2.6 *Adroit Relocate*. We removed the following dimensions:

- State Dim 36: x positional difference from the ball to the target. The overall performance is low, but the Causal PBRS method outperforms both the baseline and CQL PBRS method.

5.3 Relation Between State Return Independence and Performance Gap

To understand when the Causal PBRS method can out perform the SAC baseline without shaping, we plot the RCIT statistic from the conditional independence test of the selected state dimension and the remaining returns-to-go (conditioned on the remaining states and actions) from the Hopper-v5 Environment, and the average gap between the Causal PBRS method and the SAC baseline (Fig. 4). A higher test statistic implies that the observation dimension has a higher conditional dependence on the returns-to-go. In Fig. 4, as the test statistic increases, the gap between the Causal PBRS method and the SAC baseline increases, thus supporting our theory that the Causal PBRS method performs better when there is an increase in confounding bias. We also notice that when we remove dimensions necessary for the task (those with the highest test statistic), the learned causal shaping function may not be informative enough to help recover the performance. See Sec. 10 for a more details.

This independence analysis - in addition to exploring the causal model of the MuJoCo environments - suggests the Causal PBRS method performs best when the confounding variables have a direct causal effect on the reward. In the MuJoCo environments, the agent is rewarded for forward movement (known as the forward reward), which is a major determining factor of the total environment reward. In all four of the MuJoCo experiments, when we remove the x dimension velocity observation, a direct parent of the forward reward, the Causal PBRS method is either able to converge faster to an optimal policy or obtain a better policy than the baseline agent. On the other hand, in cases when the Causal PBRS method under performs (such as Ant dimension 16 - the x-coordinate angular velocity of the torso), the the removed dimension by itself does not have as strong an effect on the total reward / behavioral policy.

5.4 Impact of Offline Data Quality on Online Training with Causal PBRS

For our experiments, the Minari offline dataset contained three levels of expertise - "simple", "medium", and "expert". The results in Table 1 are based on training the optimistic state value functions with all three datasets. To evaluate how the expertise of the agent generating the offline data affects the Causal PBRS method, we trained three different causal reward shaping function, each on one of the Minari datasets (Fig. 5). The Causal PBRS trained solely on the expert data performed the best. However, the Causal PBRS trained with the medium and simple datasets is still able to improve the agent’s policy over the baseline. Notably, the performance of the combined datasets (Combined) is comparable to the Expert performance, albeit slightly slower to converge, suggesting Algo. 1 can converge well in practice, despite the differing data quality.

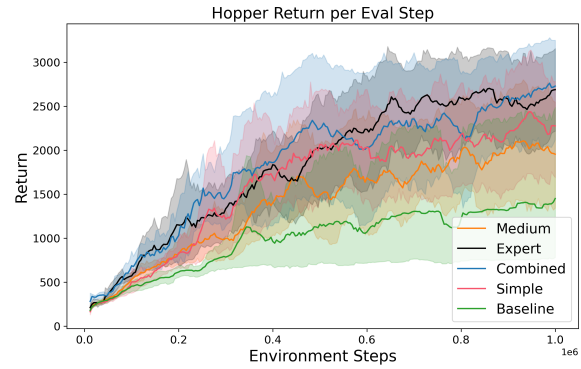


Figure 5: Causal PBRS performance by offline data quality.

6 CONCLUSION

We introduce a causal framework for automatic reward shaping in high-dimensional continuous control with unobserved confounders. By extending the Causal Bellman Equation to continuous settings, our method learns optimistic state potentials that serve as principled shaping functions within the PBRS framework, improving learning efficiency while preserving policy optimality. We evaluate our approach on confounded MuJoCo and Adroit benchmarks, where it consistently outperforms unshaped and causally unaware strong baselines such as CQL, T-REX PBRS from the literature. This work marks an important step toward confounding-robust reinforcement learning and causal reward design for real-world continuous control. It paves the way for several interesting future directions including: 1) tighter Causal Bellman upper bound given more prior knowledge, 2) testing the Causal PBRS method in more diverse confounding settings and under higher-dimensional image observations, and 3) exploring adaptive Causal PBRS for environments with less confounding bias⁴.

ACKNOWLEDGMENTS

This research is supported in part by the NSF, ONR, AFOSR, DoE, Amazon, JP Morgan, and The Alfred P. Sloan Foundation.

⁴see Sec. 11 in the full version for an extended discussion.

REFERENCES

- [1] Stefan Banach. 1922. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae* 3 (1922), 133–181. <https://api.semanticscholar.org/CorpusID:118543265>
- [2] Andrew Bennett and Nathan Kallus. 2024. Proximal Reinforcement Learning: Efficient Off-Policy Evaluation in Partially Observable Markov Decision Processes. *Oper. Res.* 72, 3 (2024), 1071–1086. <https://doi.org/10.1287/OPRE.2021.0781>
- [3] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. 2019. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International conference on machine learning*. PMLR, 783–792.
- [4] Tim Brys, Anna Harutyunyan, Halit Bener Suay, Sonia Chernova, Matthew E. Taylor, and Ann Nowé. 2015. Reinforcement Learning from Demonstration through Shaping. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 3352–3358. <http://ijcai.org/Abstract/15/472>
- [5] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement Learning with Deep Energy-Based Policies. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017 (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 1352–1361. <http://proceedings.mlr.press/v70/haarnoja17a.html>
- [6] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *ICML (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 1856–1865. <http://dblp.uni-trier.de/db/conf/icml/icml2018.html#HaarnojaZAL18>
- [7] Edward S. Hu, James Springer, Oleh Rybkin, and Dinesh Jayaraman. 2024. Privileged Sensing Scaffolds Reinforcement Learning. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=EpVe8jAjdx>
- [8] Dongchi Huang, Jiaqi WANG, Yang Li, Chunhe Xia, Tianle Zhang, and Kaige Zhang. 2025. PIGDreamer: Privileged Information Guided World Models for Safe Partially Observable Reinforcement Learning. In *Forty-second International Conference on Machine Learning*. <https://openreview.net/forum?id=mtk8tTKWs0>
- [9] Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal Mehta, and João G.M. Araújo. 2022. CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms. *Journal of Machine Learning Research* 23, 274 (2022), 1–18. <http://jmlr.org/papers/v23/21-1342.html>
- [10] Sinan Ibrahim, Mostafa Mostafa, Ali Jnadi, Hadi Salloum, and Pavel Osinenko. 2024. Comprehensive Overview of Reward Engineering and Shaping in Advancing Reinforcement Learning Applications. *IEEE Access* 12 (2024), 175473–175500. <https://doi.org/10.1109/ACCESS.2024.3504735>
- [11] Tommi S. Jaakkola, Satinder Singh, and Michael I. Jordan. 1994. Reinforcement Learning Algorithm for Partially Observable Markov Decision Problems. In *Advances in Neural Information Processing Systems 7, [NIPS Conference, Denver, Colorado, USA, 1994]*, Gerald Tesaro, David S. Touretzky, and Todd K. Leen (Eds.). MIT Press, 345–352. https://proceedings.neurips.cc/paper_files/paper/1994/hash/1c1d4df596d01da60385f0bb174a49e0-Abstract.html
- [12] Nan Jiang and Lihong Li. 2016. Doubly Robust Off-policy Value Evaluation for Reinforcement Learning. In *Proceedings of The 33rd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 48)*, Maria Florina Balcan and Kilian Q. Weinberger (Eds.). PMLR, New York, New York, USA, 652–661. <http://proceedings.mlr.press/v48/jjiang16.html>
- [13] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1-2 (1998), 99–134.
- [14] Nathan Kallus and Angela Zhou. 2018. Confounding-Robust Policy Improvement. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (Eds.), 9289–9299. <https://proceedings.neurips.cc/paper/2018/hash/3a09a524440d44d7f19870070a5ad42f-Abstract.html>
- [15] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-Learning for Offline Reinforcement Learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/0d2b2061826a5df322116a5085a6052-Abstract.html>
- [16] Daniel Kumor, Junzhe Zhang, and Elias Bareinboim. 2021. Sequential Causal Imitation Learning with Unobserved Confounders. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.), 14669–14680. <https://proceedings.neurips.cc/paper/2021/hash/7b670d553471ad0fd7491c75bad587ff-Abstract.html>
- [17] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems.
- [18] Mingxuan Li, Junzhe Zhang, and Elias Bareinboim. 2024. Causally Aligned Curriculum Learning. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=hp4yOjhwTs>
- [19] Mingxuan Li, Junzhe Zhang, and Elias Bareinboim. 2025. Automatic Reward Shaping from Confounded Offline Data. In *Forty-second International Conference on Machine Learning*. <https://openreview.net/forum?id=Hu7hUjEMiW>
- [20] Mingxuan Li, Junzhe Zhang, and Elias Bareinboim. 2025. Confounding Robust Reinforcement Learning: A Causal Approach. In *Advances in Neural Information Processing Systems 39: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2 - 7, 2025*. <https://neurips.cc/virtual/2025/poster/119520>
- [21] Miao Lu, Yifei Min, Zhaoran Wang, and Zhuoran Yang. 2023. Pessimism in the Face of Confounders: Provably Efficient Offline Reinforcement Learning in Partially Observable Markov Decision Processes. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. https://openreview.net/forum?id=PbkBDQ5_UbV
- [22] Lina Mezghani, Sainbayar Sukhbaatar, Piotr Bojanowski, Alessandro Lazaric, and Karteek Alahari. 2022. Learning Goal-Conditioned Policies Offline with Self-Supervised Reward Shaping. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand (Proceedings of Machine Learning Research, Vol. 205)*, Karen Liu, Dana Kulic, and Jeffrey Ichnowski (Eds.). PMLR, 1401–1410. <https://proceedings.mlr.press/v205/mezghani23a.html>
- [23] Rui Miao, Zhengling Qi, and Xiaohe Zhang. 2022. Off-Policy Evaluation for Episodic Partially Observable Markov Decision Processes under Non-Parametric Models. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.). http://papers.nips.cc/paper_files/paper/2022/hash/03dfa2a7755635f756b160e9f4c6b789-Abstract-Conference.html
- [24] Volodymyr Mnih, Adria Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016 (JMLR Workshop and Conference Proceedings, Vol. 48)*, Maria-Florina Balcan and Kilian Q. Weinberger (Eds.). JMLR.org, 1928–1937. <http://proceedings.mlr.press/v48/mnih16.html>
- [25] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-Level Control through Deep Reinforcement Learning. *Nature* 518, 7540 (2015), 529–533. <https://doi.org/10.1038/nature14236>
- [26] Gianluca Monaci, Michel Aractingi, and Tomi Silander. 2022. DiPCAN: Distilling Privileged Information for Crowd-Aware Navigation. In *Robotics: Science and Systems XVIII, New York City, NY, USA, June 27 - July 1, 2022*, Kris Hauser, Dylan A. Shell, and Shoudong Huang (Eds.). <https://doi.org/10.15607/RSS.2022.XVIII.045>
- [27] Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. 2016. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*. 1054–1062.
- [28] Susan A Murphy. 2003. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 65, 2 (2003), 331–355.
- [29] Susan A Murphy. 2005. A Generalization Error for Q-Learning. *Journal of Machine Learning Research* 6 (2005), 1073–1097.
- [30] Mitsuhiro Nakamoto, Simon Zhai, Anikait Singh, Max Sobol Mark, Yi Ma, Chelsea Finn, Aviral Kumar, and Sergey Levine. 2023. Cal-QL: Calibrated Offline RL Pre-Training for Efficient Online Fine-Tuning. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/c44a04289beaf0a7d968a94066a1d696-Abstract-Conference.html
- [31] Andrew Y. Ng, Daiishi Harada, and Stuart Russell. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999)*, Bled, Slovenia, June 27 - 30, 1999, Ivan Bratko and Saso Dzeroski (Eds.). Morgan Kaufmann, 278–287.
- [32] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. 2019. Solving Rubik’s Cube with a Robot Hand. (2019). arXiv:1910.07113 <http://arxiv.org/abs/1910.07113> arxiv.

- [33] Alexander Pan, Kush Bhatia, and Jacob Steinhardt. 2022. The Effects of Reward Misspecification: Mapping and Mitigating Misaligned Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25–29, 2022*. OpenReview.net. <https://openreview.net/forum?id=JYtwGwL7ye>
- [34] Judea Pearl. 2009. *Causality: Models, Reasoning, and Inference* (2 ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- [35] Doina Precup, Richard S. Sutton, and Satinder P. Singh. 2000. Eligibility traces for off-policy policy evaluation. In *Proceedings of the Seventeenth International Conference on Machine Learning*. 759–766.
- [36] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (1 ed.). Wiley. <https://doi.org/10.1002/9780470316887>
- [37] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. 2017. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations.
- [38] Jette Randløv and Preben Alstrøm. 1998. Learning to Drive a Bicycle Using Reinforcement Learning and Shaping. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998), Madison, Wisconsin, USA, July 24–27, 1998*, Jude W. Shavlik (Ed.), Morgan Kaufmann, 463–471.
- [39] John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. 2015. Trust region policy optimization. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37 (Lille, France) (ICML '15)*. JMLR.org, 1889–1897.
- [40] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1506.02438>
- [41] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. (2017). arXiv:1707.06347 <http://arxiv.org/abs/1707.06347> arxiv.
- [42] Chengchun Shi, Masatoshi Uehara, Jiawei Huang, and Nan Jiang. 2022. A Minimax Learning Approach to Off-Policy Evaluation in Confounded Partially Observable Markov Decision Processes. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 20057–20094. <https://proceedings.mlr.press/v162/shi22f.html>
- [43] Haochen Shi, Huazhe Xu, Samuel Clarke, Yunzhu Li, and Jiajun Wu. 2023. RoboCook: Long-Horizon Elasto-Plastic Object Manipulation with Diverse Tools. In *Conference on Robot Learning, CoRL 2023, 6–9 November 2023, Atlanta, GA, USA (Proceedings of Machine Learning Research, Vol. 229)*, Jie Tan, Marc Toussaint, and Kourosh Darvish (Eds.). PMLR, 642–660. <https://proceedings.mlr.press/v229/shi23a.html>
- [44] Eric Strobl, Kun Zhang, and Shyam Visweswaran. 2018. Approximate Kernel-Based Conditional Independence Tests for Fast Non-Parametric Causal Discovery. *Journal of Causal Inference* 7 (12 2018). <https://doi.org/10.1515/jci-2018-0017>
- [45] Adith Swaminathan and Thorsten Joachims. 2015. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*. 814–823.
- [46] Guy Tennenholtz, Uri Shalit, and Shie Mannor. 2020. Off-policy evaluation in partially observable environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 10276–10283.
- [47] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033. <https://doi.org/10.1109/IROS.2012.6386109>
- [48] Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U. Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, et al. 2024. Gymnasium: A Standard Interface for Reinforcement Learning Environments.
- [49] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3–4 (1992), 279–292.
- [50] Christopher John Cornish Hellaby Watkins. 1989. *Learning from delayed rewards*. Ph.D. Dissertation. University of Cambridge England.
- [51] Omar G. Younis, Rodrigo Perez-Vicente, John U. Balis, Will Dudley, Alex Davey, and Jordan K Terry. 2024. *Minari*. Farama Foundation. <https://doi.org/10.5281/zenodo.13767625>
- [52] Junzhe Zhang and Elias Bareinboim. 2019. Near-optimal reinforcement learning in dynamic treatment regimes. In *Advances in Neural Information Processing Systems*. 13401–13411.
- [53] Junzhe Zhang and Elias Bareinboim. 2022. Can Humans Be out of the Loop?. In *1st Conference on Causal Learning and Reasoning, CLear 2022, Sequoia Conference Center, Eureka, CA, USA, 11–13 April, 2022 (Proceedings of Machine Learning Research, Vol. 177)*, Bernhard Schölkopf, Caroline Uhler, and Kun Zhang (Eds.). PMLR, 1010–1025. <https://proceedings.mlr.press/v177/zhang22a.html>
- [54] J. Zhang and E. Bareinboim. 2024. *Eligibility Traces for Confounding Robust Off-Policy Evaluation*. Technical Report R-105. Causal Artificial Intelligence Lab, Columbia University.
- [55] Junzhe Zhang, Daniel Kumor, and Elias Bareinboim. 2020. Causal imitation learning with unobserved confounders. *Advances in neural information processing systems* 33 (2020), 12263–12274.
- [56] Yi Zhang, Ruihong Qiu, Jiajun Liu, and Sen Wang. 2024. ROLeR: Effective Reward Shaping in Offline Reinforcement Learning for Recommender Systems. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM 2024, Boise, ID, USA, October 21–25, 2024*, Edoardo Serra and Francesca Spezzano (Eds.). ACM, 3269–3278. <https://doi.org/10.1145/3627673.3679633>
- [57] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13–17, 2008*, Dieter Fox and Carla P. Gomes (Eds.). AAAI Press, 1433–1438. <http://www.aaai.org/Library/AAAI/2008/aaai08-227.php>