

Causal Flow Q-Learning for Robust Offline Reinforcement Learning

Mingxuan Li¹ Junzhe Zhang² Elias Bareinboim¹

Abstract

Expressive policies based on flow-matching have been successfully applied in reinforcement learning (RL) more recently due to their ability to model complex action distributions from offline data. These algorithms build on standard policy gradients, which assume that there is no unmeasured confounding in the data. However, this condition does not necessarily hold for pixel-based demonstrations when a mismatch exists between the demonstrator’s and the learner’s sensory capabilities, leading to implicit confounding biases in offline data. We address the challenge by investigating the problem of confounded observations in offline RL from a causal perspective. We develop a novel causal offline RL objective that optimizes policies’ worst-case performance that may arise due to confounding biases. Based on this new objective, we introduce a practical implementation that learns expressive flow-matching policies from confounded demonstrations, employing a deep discriminator to assess the discrepancy between the target policy and the nominal behavioral policy. Experiments across 25 pixel-based tasks demonstrate that our proposed confounding-robust augmentation procedure achieves a success rate 120% that of confounding-unaware, state-of-the-art offline RL methods.

1. Introduction

Offline reinforcement learning (RL) offers an alternative paradigm to traditional online RL, enabling effective policy learning for decision-making from previously collected observational data when active exploration in the underlying environment is costly, unsafe, or even impractical (Lange et al., 2012; Levine et al., 2020). The standard offline RL problem can be interpreted as a constrained opti-

mization: the agent seeks to maximize long-term rewards while remaining close to the state-action distributions of the observed trajectories (Levine et al., 2020). More recently, there is a growing body of offline RL that attempts to represent more complex, multimodal policy distributions using an expressive policy class that explicitly learns the policy’s velocity field via deep generative modeling (Mandlekar et al., 2022; Lipman et al., 2024). These expressive policy classes include denoising diffusion (Wang et al., 2023) and flow-matching (Park et al., 2025b). They enable modeling more complex policy distributions in the dataset, thereby enforcing accurate behavioral constraints which is critical to offline RL algorithms (Tarasov et al., 2023).

Despite these progresses, all the algorithms described above build on the standard off-policy policy gradient framework (Sutton, 2018), which relies on the key assumption that the demonstrator’s behavioral policy and the learner’s target policy share the same support over the state-action domain (Precup et al., 2000). When such conditions fail to hold, implicit confounding biases can be introduced into the observational data, consequently posing significant challenges for training RL systems from offline data (Bareinboim et al., 2024; Li et al., 2025b).

Example 1 (Confounded Pixel-Based Observations). Fig. 1 shows an example that illustrates challenges of implicit confounding bias in offline RL. Specifically, the expert actions are sampled based on true state observations, whereas the final pixel-based offline dataset pairs those actions with the corresponding image observations. The robot arm is tasked to move the colored cubes into the desired translucent positions. But pixel observations in Fig. 1b, a blurry front-view with occlusion, cannot fully capture the true state information shown in Fig. 1a, which underlies the expert’s decisions. Learning from such datasets under the MDP assumption results in significant performance degradation despite that the same method works well in the same task with a structured state-observation space as the demonstrator does (Park et al., 2025b; Dong et al., 2025). ■

The problem of addressing confounding bias in decision-making has been studied under the rubrics of causal inference (Pearl, 2009) and, more recently, causal reinforcement learning (Bareinboim et al., 2024). Most methods explore additional prior knowledge about the underlying causal mechanisms under which the effect of the target

¹Causal AI Lab, Columbia University, NY, United States

²Department of Electrical Engineering and Computer Science, Syracuse University, NY, United States. Correspondence to: Mingxuan Li <ml@cs.columbia.edu>.

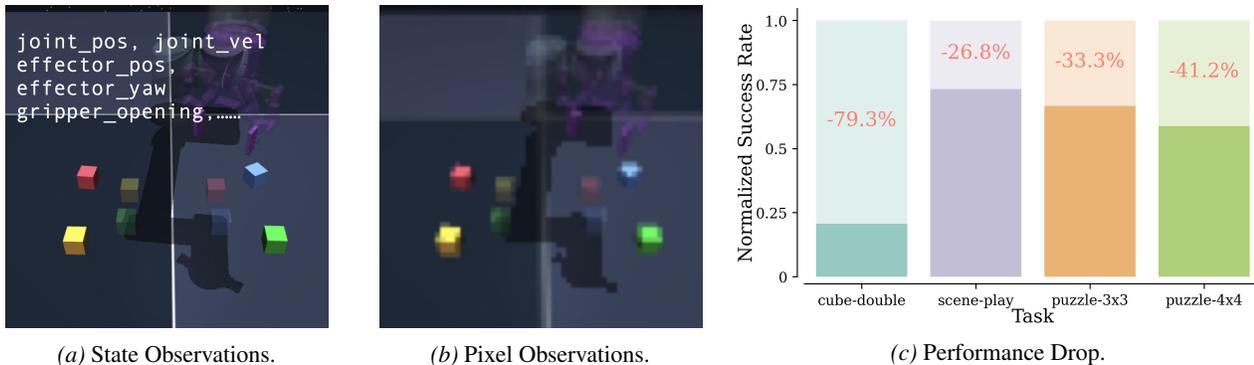


Figure 1. **Confounding biases in offline data causes performance loss.** Left-middle: For pixel based tasks in offline RL datasets, expert actions are sampled based on true state vectors but presented with pixel observations during the offline learning stage. Right: SOTA offline RL algorithm performance drops sharply with pixel observations despite using image augmentations and strong neural encoders.

policy is uniquely determined (i.e., identifiable) from the observational data (Pearl & Robins, 1995; Bareinboim et al., 2024). When the causal knowledge is not available or too weak to ensure identifiability, partial identification methods could be incorporated with standard RL algorithms to reason about the optimal actions in the worst-case environments compatible with observations (Kallus & Zhou, 2018; Zhang & Bareinboim, 2019; Namkoong et al., 2020; Zhang & Bareinboim, 2024; Li et al., 2025a;b; Hess et al., 2025). Functional approximators have been used to evaluate the worst-case policy return from confounded observations for environments with complex states but a simple, discrete action domain. Still, significant challenges remain in learning a robust policy from confounded, offline observations in an expressive policy class with continuous actions.¹

We overcome this challenge by formalizing the confounding issue in offline RL, developing a novel offline RL objective and then employing flow-matching to model the complex state-action distributions in observed trajectories. More specifically, our contributions are as follows: (1) We introduce a novel offline RL objective that is robust against the implicit unobserved confounding in the offline data. Optimizing this objective leads to a safe policy selecting optimal actions in the worst-case environment. (2) We develop an algorithm, called *Causal Flow Q-Learning* (CFQL), to optimize the newly proposed confounding robust objective over an expressive policy class, leveraging flow-matching to model the observed state-action distributions. We evaluate CFQL through extensive experiments. The result demonstrate a 120% success rate improvement on average in 25 pixel-based offline RL tasks with some even surpassing policies using true state observations. Due to the space constraint, all proofs are provided in Sec. C.

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. We use bold

¹For a more detailed discussion on offline RL and causal decision-making, we refer readers to Sec. A.

capital letters (V) to denote a set of random variables and let $|V|$ denote its cardinality of set V . Fix indices $i, j \in \mathbb{N}$. Let $\bar{X}_{i:j}$ stand for a sequence $\{X_i, X_{i+1}, \dots, X_j\}$. We denote by $P(\mathbf{X})$ a probability distribution over variables \mathbf{X} . We consistently use $P(\mathbf{x})$ as abbreviations of probabilities $P(\mathbf{X} = \mathbf{x})$; so does $P(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = P(\mathbf{y} \mid \mathbf{x})$. Finally, $\mathbf{1}_{Z=z}$ is an indicator function that returns 1 if event $Z = z$ holds true; otherwise, it returns 0.

2. Confounding Robust Policy Learning

We investigate a sequential decision-making setting in which the agent chooses a sequence of actions to optimize subsequent rewards. We assume a Confounded Markov Decision Process (CMDP) to explicitly model the challenges of unobserved confounding in offline RL (Zhang & Bareinboim, 2022; Bennett et al., 2021; Kallus & Zhou, 2020; Zhang & Bareinboim, 2025; Li et al., 2025a;b).²

Definition 2.1. A Confounded Markov Decision Process (CMDP) \mathcal{M} is a tuple of $\langle \mathcal{S}, \mathcal{X}, \mathcal{Y}, \mathcal{U}, \mathcal{F}, P \rangle$ where (1) $\mathcal{S}, \mathcal{X}, \mathcal{Y}$ are, respectively, the space of observed states, actions, and rewards; (2) \mathcal{U} is the space of unobserved exogenous noise; (3) \mathcal{F} is a set consisting of the transition function $f_S : \mathcal{S} \times \mathcal{X} \times \mathcal{U} \mapsto \mathcal{S}$, behavioral policy $f_X : \mathcal{S} \times \mathcal{U} \mapsto \mathcal{X}$, and reward function $f_Y : \mathcal{S} \times \mathcal{X} \times \mathcal{U} \mapsto \mathcal{Y}$; (4) P is an exogenous distribution over the domain \mathcal{U} .

Throughout this paper, we assume the action domain \mathcal{X} to be continuous unless specified otherwise, while the state domain \mathcal{S} could be complex and continuous; the reward domain \mathcal{Y} is bounded in a real interval $[a, b] \subset \mathbb{R}$. Consider a demonstrator agent interacting with a CMDP \mathcal{M} , generating the off-policy data. For every time step $t = 1, \dots, T$, the environment first draws an exogenous noise U_t from the distribution $P(\mathcal{U})$; the demonstrator then performs an action $X_t \leftarrow f_X(S_t, U_t)$, receives a subsequent reward $Y_t \leftarrow r_t(S_t, X_t, U_t)$, and moves to the next state $S_{t+1} \leftarrow$

²See Sec. C for a detailed discussion on causal foundations and CMDPs as a relaxation to MDPs.

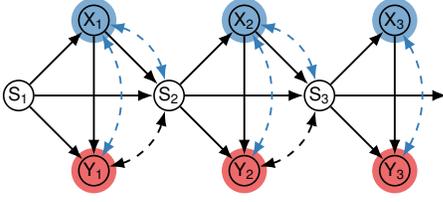


Figure 2. **Causal diagram representing the data-generating mechanisms in a Confounded Markov Decision Process.** Bi-directed arrows represent information used by the expert’s policy but unobservable to the learner in the offline dataset.

$f_S(S_t, X_t, U_t)$. The observed trajectories of the demonstrator (from the learner’s perspective) are summarized as the observational distribution $P(\bar{X}_{1:T}, \bar{S}_{1:T}, \bar{Y}_{1:T})$.

Fig. 2 shows a graphical representation (Bareinboim et al., 2022) illustrating the generative process of the observational data in CMDPs. More specifically, solid nodes represent observed variables X_t, S_t, Y_t , and arrows represent the functional relationships f_X, f_S, f_Y among them. By convention, exogenous variables U_t are often not explicitly shown in the graph; bi-directed arrows $X_t \longleftrightarrow Y_t$ and $X_t \longleftrightarrow S_{t+1}$ indicate the presence of an unobserved confounder (UC) U_t affecting the action, state, and reward simultaneously. These bi-directed arrows (highlighted in blue) represent the unobserved confounders among action X_t , reward Y_t , and state S_{t+1} in the off-policy data, violating the condition of no unobserved confounder (NUC) (Robbins, 1985; Bareinboim et al., 2024). Such violations lead to challenges in offline RL as we have seen above.

A policy π in a CMDP \mathcal{M} is a decision rule $\pi(x_t | s_t)$ mapping from state to a distribution over action domain \mathcal{X} . An intervention $\text{do}(\pi)$ is an operation that replaces the behavioral policy f_X in CMDP \mathcal{M} with the policy π . Let \mathcal{M}_π be the submodel induced by intervention $\text{do}(\pi)$. The interventional distribution $P_\pi(\bar{X}_{1:T}, \bar{S}_{1:T}, \bar{Y}_{1:T})$ is defined as the distribution over observed variables in \mathcal{M}_π ,

$$P_\pi(\bar{x}_{1:T}, \bar{s}_{1:T}, \bar{y}_{1:T}) = P(s_1) \prod_{t=1}^T \left(\pi(x_t | s_t) \mathcal{T}(s_t, x_t, s_{t+1}) \mathcal{R}(s_t, x_t, y_t) \right) \quad (1)$$

where the transition distribution \mathcal{T} and the reward distribution \mathcal{R} are given by, for $t = 1, \dots, T$,

$$\mathcal{T}(s_t, x_t, s_{t+1}) = \int_{\mathcal{U}} \mathbf{1}_{s_{t+1}=f_S(s_t, x_t, u_t)} P(u_t), \quad (2)$$

$$\mathcal{R}(s_t, x_t, y_t) = \int_{\mathcal{U}} \mathbf{1}_{y_t=f_Y(s_t, x_t, u_t)} P(u_t). \quad (3)$$

For convenience, we write the reward function $\mathcal{R}(s, x)$ as the expected value $\sum_y y \mathcal{R}(s, x, y)$. Fix a discounted factor $\gamma \in [0, 1]$. A common objective for an agent is to optimize its cumulative return $R_t = \sum_{i=0}^{\infty} \gamma^i Y_{t+i}$.

Offline RL. In offline RL, when there is no unobserved confounder introducing spurious correlations between actions and subsequent outcomes, one can identify the parameterizations of the transition distribution \mathcal{T} and the reward function \mathcal{R} from offline data. This means that the expected return of candidate policies can be estimated from the sampling process of the offline data. An optimal policy is then obtainable by maximizing the estimated return (under proper behavioral regularization).

For example, let the state-action value function $Q_\phi(s, x) = \mathbb{E}_\pi[R_t | S_t = s, X_t = x]$ denote the expected return of a target policy π conditioning on state s and action x with parameter ϕ . Let θ denote the parameter of π . The basic offline RL objective with behavior regularization is defined as (Wu et al., 2019; Fujimoto & Gu, 2021; Tarasov et al., 2023):

$$\mathcal{L}(\theta) = \mathbb{E}_{x, s \sim \mathcal{D}, x' \sim \pi_\theta} \left[-Q_\phi(s, x') - \alpha \log \pi(x|s) \right], \quad (4)$$

where \mathcal{D} is the offline observations; and α are some coefficients that control the strength of the behavioral cloning (BC) regularizer. The Q-value critic Q_ϕ is trained by empirical Bellman loss on the offline dataset. Target policy is learned by applying gradient-based optimization to the objective in Eq. (4). When the NUC condition holds in the data, the target policy is guaranteed to improve and converge to the optimum under standard convexity assumptions (Zhang et al., 2020).

Causal Reinforcement Learning. However, the NUC condition does not always hold in real-world applications. Confounding bias can arise when the sensory capabilities of the demonstrator and the learner differ, posing challenges for offline learning algorithms.

Example 2 (Confounded Pixel Observations (continued)). Consider the example in Fig. 1 again, which illustrates the presence of implicit confounding bias in offline RL data. Fig. 1c shows the return of flow policies trained on pixels in comparison to the ones trained on the same task but structured state observations (Park et al., 2025b; Dong et al., 2025). We see a consistent trend of performance loss across the board, with some even suffering a nearly 80% drop. All success rate under pixel-based observations are normalized with respect to the performance under structured state observations of the same algorithm and averaged over all 5 tasks in each category from OGBench (Park et al., 2025a). ■

Recently, there has been a growing body of work in causal inference (Kallus & Zhou, 2018; Zhang & Bareinboim, 2019; Joshi et al., 2024b; Namkoong et al., 2020; Li et al., 2025b) and safe reinforcement learning (Kumar et al., 2020; Nakamoto et al., 2023) to address the challenges of data bias and distribution shifts in policy learning. Closest to our setting, Zhang & Bareinboim (2025) addressed the challenges

of confounding bias in off-policy evaluation by deriving a novel *Causal Bellman equation* to bound the value function from confounded observations. Specifically, the value function $Q_\pi(s, x)$ of a policy π is lower bounded by a function $\underline{Q}_\pi(s, x)$ given by,

$$\begin{aligned} \underline{Q}_\pi(s, x) = & \mu(\neg x | s) \left(a + \gamma \min_{s'} \underline{V}_\pi(s') \right) \\ & + \mu(x | s) \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s', x'} \tilde{\mathcal{T}}(s, x, s') \underline{V}_\pi(s') \right) \end{aligned} \quad (5)$$

where state value bound $\underline{V}_\pi(s) = \sum_x \pi(x | s) \underline{Q}_\pi(s, x)$ and $P(\neg x | s) = 1 - P(x | s)$; $\tilde{\mathcal{T}}$ and $\tilde{\mathcal{R}}$ are nominal transition distribution and reward function computed from the observational distribution, i.e.,

$$\tilde{\mathcal{T}}(s, x, s') = P(S_{t+1} = s' | S_t = s, X_t = x), \quad (6)$$

$$\tilde{\mathcal{R}}(s, x) = \mathbb{E}[Y_t | S_t = s, X_t = x]. \quad (7)$$

The learner can then obtain an effective policy from confounded offline data by iteratively optimizing the above value-function lower bound over the action space (Li et al., 2025b). However, the existing causal Bellman equation is limited to discrete actions and, consequently, cannot handle complex, multimodal continuous action distributions. The remainder of this paper will address this challenge.

3. Causal Flow Q-Learning

In this section, we introduce a causal offline RL objective, based on which we present Causal Flow Q-Learning.

We note that for any policy π , the state value function is lower bounded by $\underline{V}_\pi(s) = \sum_x \pi(x | s) \underline{Q}_\pi(s, x)$, where $\underline{Q}_\pi(s, x)$ is given by Eq. (5). After a few simplifications, a closed-form solution can be derived, as shown next.

Theorem 3.1. *For a CMDP environment \mathcal{M} with reward signals $Y_t \in [a, b] \subseteq \mathbb{R}$, fix a policy π . The state value function $V_\pi(s) \geq \underline{V}_\pi(s)$ for any state $s \in \mathcal{S}$, where the lower bound $\underline{V}_\pi(s)$ is given by as follows,*

$$\begin{aligned} \underline{V}_\pi(s) = & \mathbb{E}_{\substack{x \sim p(\cdot | s) \\ x', x^* \sim \pi(\cdot | s)}} \left[\mathbf{1}_{x \neq x'} \left(a + \gamma \min_{s^*} \underline{Q}_\pi(s^*, x^*) \right) \right. \\ & \left. + \mathbf{1}_{x=x'} \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s'} \tilde{\mathcal{T}}(s, x, s') \underline{Q}_\pi(s', x^*) \right) \right] \end{aligned} \quad (8)$$

Fig. 3 shows a backup diagram illustrating this update step. As in the standard Bellman optimality equation (Bellman, 1966), Eq. (8) recursively updates the state value function using the current estimates of the optimal value function.

On the other hand, Eq. (8) explicitly accounts for the off-policy nature of the confounded observations: when the behavior policy takes the same action $x_t = x$ as the target action, the update follows standard Bellman equation and uses the next sampled state s_t ; when the sampled action $x_t \neq x$ differs from the target, our algorithm updates, instead, using the value function associated with the next worst-case or best-case state s^* , corresponding to the estimation of the lower bound and upper bound respectively.

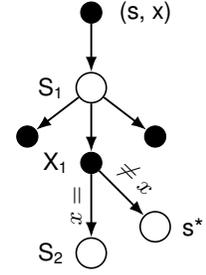


Figure 3. Backup diagram for causal policy gradient.

Thm. 3.1 lower bounds the expected return of a policy π that optimizes a worst-case CMDP instance $\underline{\mathcal{M}}$ compatible with the confounded observations. Maximizing the lower bound in Thm. 3.1 leads to a safe policy with a performance guarantee in the underlying environment. In causal offline RL objective, this bound is used as the Q loss in Eq. (4) and can be calculated with function approximators. Specifically, let θ denote the parameters of the target policy π , and \mathcal{D} be the finite set of offline observations. Then, the causal offline RL objective is estimated from the observational data as,

$$\begin{aligned} \mathcal{L}(\theta) = & \mathbb{E}_{\substack{(s, x, y, s') \sim \mathcal{D} \\ x', x^* \sim \pi_\theta(\cdot | s)}} \left[D(s, x, x') \left(y + \gamma \underline{Q}_{\pi_\theta}(s', x^*) \right) \right. \\ & \left. + (1 - D(s, x, x')) \left(a + \gamma \min_{s^*} \underline{Q}_{\pi_\theta}(s^*, x^*) \right) \right] \end{aligned} \quad (9)$$

Among the above quantities, $D(s, x, x')$ is a discriminator that approximates the indicator function $\mathbf{1}_{x=x'}$ at state s . It takes as input a state action pair, s, x , and returns 1 if the policy distribution $\pi_\theta(x | s)$ agrees with offline data. x^* is an overloaded notation for next step actions in general sampled from π_θ . We will discuss the training details further next.

We incorporate the confounding robust Q-value target in Eq. (9) to offline RL objective, based on Flow Q-Learning (FQL) introduced by Park et al. (2025b). Note that the proposed causal offline RL objective is not limited to FQL, and could be incorporated into other gradient-based policy learning to mitigate the influence of confounding bias. For example, we demonstrate an alternative implementation in Sec. G using value flows (Dong et al., 2025).

Similar to FQL, our proposed augmentation, called Causal FQL (CFQL), focuses on learning two policies during training: (1) a behavioral cloning (BC) policy μ_ω mimicking the nominal behavioral policy distribution $P(X_t | S_t)$ from the observational data; (2) a target policy π_θ deciding the learner's actions. Specifically, the BC policy $\mu_\omega(s, z)$ is a continuous normalizing flow (Lipman et al., 2022) taking a state s and an independent noise z as input. It transforms the

noise z into an action x drawn from the nominal behavioral distribution $P(X_t | S_t)$ following a multi-step denoising procedure. On the other hand, the target policy $\pi_\theta(s, z)$ is a one-step generative model directly mapping a noise z into an instance x in the action domain.

Now we describe the details of our proposed algorithm, Causal FQL (Algo. 1). CFQL contains four main components: an ensemble of critic networks Q_{ϕ_i} , a BC flow policy μ_ω , a one-step target policy π_θ , and a discriminator D_ψ differentiating the actions drawn from the BC and target policies. More specifically, Steps 4-9 train the ensemble of critic Q-networks $Q_{\phi_i}, i = 1, 2, \dots, N$. Each network Q_{ϕ_i} is trained from the confounded observational data following the standard minimization of the one-step Bellman update (Step 8) over the average of Q ensembles $\bar{Q}_\phi(s', x') = \frac{1}{N} \sum_i Q_{\phi_i}(s', x')$.

From Steps 10 - 15, the algorithm trains the velocity field $v_\omega(t, s, x^t)$ of a BC flow policy μ_ω . It models the transformation of an independent noise x^0 to the observed action x^1 as trajectories of a free particle moving at a constant speed. The velocity field $v_\omega(t, s, x^t)$ is a neural network recording the gradient of the trajectory at any time step t . The training follows the standard flow-matching procedure (Lipman et al., 2022; 2024). Steps 16 - 20 train a discriminator network D_ψ in the causal policy gradient Eq. (9) to differentiate the samples drawn from the BC flow μ_ω and target policy π_θ . Specifically, the algorithm obtains samples drawn from $\mu_\omega(s, z)$ and $\pi_\theta(s, z)$ given the observed state s and an independent noise z (Steps 17-19). The samples of the BC flow policy are computed from the trained velocity field following the standard Euler method (Algo. 2) The discriminator D_ψ 's training objective is then given by,

$$\mathcal{L}_{\text{Disc}}(\psi) = \mathbb{E}_{s \sim \mathcal{D}, z \sim \mathcal{N}(0, I_d)} [\log D_\psi(s, \mu_\omega(s, z)) + \log(1 - D_\psi(s, \pi_\theta(s, z)))] \quad (10)$$

Minimizing the above discriminator loss is equivalent to solving a binary classification problem where class 1 represents the action x is from the BC flow policy μ_ω while class 0 represents that the action is sampled from the one-step target policy π_θ given the current state s .

Finally, CFQL trains a one-step target policy π_θ optimizing the worst-case return compatible with the confounded observations in Steps 21 - 24. The key challenge here is to obtain a reliable estimation for the worst-case value function $\min_{s^*} Q_{\pi_\theta}(s^*, x^*)$ in Eq. (9) without being overly pessimistic. To address this challenge, instead of taking the minimum over the whole state space or the whole batch of data (Li et al., 2025b), we propose to learn ensembles of Q-networks and take the minimum over the learned ensembles to simulate the worst case scenario.

Algorithm 1 Causal Flow Q-Learning (Causal FQL)

```

1: Input: Offline dataset  $\mathcal{D}$ .
2: while not converged do
3:   Sample batch  $\{(s, x, y, s')\} \sim \mathcal{D}$ 
    $\nabla$  Train critic  $Q_{\phi_i}$ 
4:   for  $i = 1, \dots, N$  do
5:      $z \sim \mathcal{N}(0, I_d)$ 
6:      $x' \leftarrow \pi_\theta(s', z)$ ,
7:     Update  $\phi_i$  to minimize:
        $\mathbb{E}[(Q_{\phi_i}(s, x) - y - \gamma \bar{Q}_\phi(s', x'))^2]$ 
8:   end for
    $\nabla$  Train BC flow policy  $\mu_\omega$ 
9:    $x^0 \sim \mathcal{N}(0, I_d)$ 
10:   $x^1 \leftarrow x$ 
11:   $t \sim \text{Unif}(0, 1)$ 
12:   $x^t \leftarrow (1 - t)x^0 + x^1$ 
13:  Update  $\omega$  to minimize  $\mathbb{E}[\|v_\omega(t, s, x^t) - (x^1 - x^0)\|^2]$ 
    $\nabla$  Train discriminator  $D_\psi$ 
14:   $z \sim \mathcal{N}(0, I_d)$ 
15:   $x \leftarrow \mu_\omega(s, z)$ 
16:   $x' \leftarrow \pi_\theta(s, z)$ 
17:  Update  $\psi$  to minimize:
      $\mathbb{E}[\log(1 - D_\psi(s, x)) + \log D_\psi(s, x')]$ 
    $\nabla$  Train one-step target policy  $\pi_\theta$ 
18:   $z \sim \mathcal{N}(0, I_d)$ 
19:   $x \leftarrow \pi_\theta(s, z)$ 
20:  Update  $\theta$  to minimize:
      $\mathbb{E}[-\underline{Q}_*(s, x) + \alpha \|x - \mu_\omega(s, z)\|^2]$ 
21: end while
22: Return one-step flow policy  $\pi_\theta$ .
    
```

Algorithm 2 BC Flow Policy $\mu_\omega(s, z)$

```

1: Input: Velocity field  $v_\omega(t, s, z)$ , time step  $M$ .
2: for  $t = 0, 1, \dots, M - 1$  do
3:    $z \leftarrow z + v_\omega(t/M, s, z)/M$ 
4: end for
5: Return sampled action  $z$ 
    
```

With the discriminator D_ψ , the value function evaluation in the causal policy gradient can be approximated as,

$$\underline{Q}_*(s, x) = D_\psi(s, x) \bar{Q}_\phi(s, x) + (1 - D_\psi(s, x)) \min_i Q_{\phi_i}(s, x) \quad (11)$$

And the one-step target policy is trained with optimizing the following objective function:

$$\mathcal{L}_\pi(\theta) = \mathbb{E}_{s \sim \mathcal{D}, x \sim \pi_\theta} [-\underline{Q}_*(s, x)] + \alpha \mathcal{L}_{\text{Distill}}(\theta) \quad (12)$$

where α is a hyperparameter controlling the strength of the target policy regularizer; and $\mathcal{L}_{\text{Distill}}(\theta)$ is the distillation

Table 1. **Aggregated offline evaluation results.** We report aggregated mean success rate and mean deviation across different task groups. Causal-FQL consistently outperforms all baselines across the task groups. When comparing against the original FQL, the confounding robust Causal-FQL achieves a $1.2\times$ success rate improvement on average.

	Gaussian Policies			Flow Policies			
	IQL	ReBRAC	FBRAC	IQN	IFQL	FQL	Causal-FQL
visual-cube-single-play-singletask (5 tasks)	64 ± 17	82 ± 7	64 ± 10	58 ± 7	44 ± 9	65 ± 12	81 ± 6
visual-cube-double-play-singletask (5 tasks)	11 ± 6	1 ± 1	2 ± 1	1 ± 0	2 ± 2	6 ± 1	11 ± 3
visual-scene-play-singletask (5 tasks)	26 ± 5	28 ± 5	11 ± 1	41 ± 6	21 ± 2	41 ± 4	43 ± 1
visual-puzzle-3x3-play-singletask (5 tasks)	2 ± 3	20 ± 1	1 ± 0	19 ± 1	21 ± 0	20 ± 1	26 ± 2
visual-puzzle-4x4-play-singletask (5 tasks)	0 ± 0	5 ± 1	0 ± 0	13 ± 3	4 ± 7	11 ± 3	13 ± 2
Normalized Mean	0.72	0.95	0.55	0.92	0.65	1.0	1.21

loss following the definition of (Park et al., 2025b), i.e.,

$$\mathcal{L}_{\text{Distill}}(\omega) = \mathbb{E}_{\substack{s \sim \mathcal{D}, \\ z \sim \mathcal{N}(0, I_d)}} [\|\pi_{\theta}(s, z) - \mu_{\omega}(s, z)\|_2^2] \quad (13)$$

The CFQL proposed is similar to the original FQL, with an additional discriminator training step and regularization to improve worst-case value function estimation. This means that the gradient-based policy optimization could be consistently performed using the standard automatic differentiation framework (Baydin et al., 2018).

4. Experiments

In this section, we evaluate the newly proposed Causal Flow Q-Learning on 25 OGBench (Park et al., 2025a) visual tasks. We aim to answer the following questions:

- Q1:** How does Causal FQL improve upon the vanilla FQL on pixel-based tasks?
- Q2:** How sample efficient is Causal FQL on offline-to-online pixel-based tasks?
- Q3:** What hyper-parameters matter for Causal FQL?

Experiment Setup. We evaluate our methods on offline RL tasks from the recently proposed OGBench (Park et al., 2025a). As we consider solving the confounding biases in pixel-based offline RL tasks (Sec. 4, previous benchmarks like D4RL (Fu et al., 2020) that only support structured state observations are not within our considerations. OGBench is originally proposed as a goal conditioned RL benchmark, thus we use its `-singletask` variants for compatibility with standard offline RL frameworks. More specifically, we choose manipulation tasks due to their multi-modal nature that is suitable for demonstration of flow based policy. For all tasks, we provide agents $64 \times 64 \times 3$ RGB image observations (e.g., Fig. 1b) with frame stacking and random crop as image augmentations. To maintain a fair comparison, we use similar sized value/policy/actor networks for all methods and use the image encoder from IMPALA (Espeholt et al., 2018) in all tasks. For the full setup details, see Sec. F.



Figure 4. **Pixel-based OGBench Tasks.**

Baselines. We use the following algorithms representing different policy extraction strategies as baselines to the offline learning tasks. Due to high computational requirements of visual tasks, we mainly use baseline algorithms that have demonstrated consistent performance in prior work (Dong et al., 2025; Park et al., 2025b). We choose IQL (Kostrikov et al., 2021) and ReBRAC (Tarasov et al., 2023) as representatives for the class of algorithms learning scalar Q-values with a Gaussian policy. And we choose IQN (Dabney et al., 2018) as distributional Q-values with flow policy baselines, FBRAC, IFQL, and FQL (Park et al., 2025b) as scalar Q-values critic flow policy baselines. FBRAC and IFQL are flow policy based adoption to the original algorithm of Diffusion-QL (DQL) (Wang et al., 2023) and IDQL (Hansen-Estruch et al., 2023), respectively.

For offline-to-online experiments, we mainly use ReBRAC and FQL as baselines because of their relatively strong performance than other baselines in the offline experiments.

Evaluation. For offline experiments, we run 500K steps training for all baselines and Causal-FQL proposed, averaging results over 4 seeds and present the results with mean and standard deviations. In the tables, we highlight entries that are within 95% best performance on each task, following (Park et al., 2025b). For offline-to-online experiments, we use the same setup during the offline stage and extend training to the online stage for an additional 500K steps. See Sec. E and F for implementation and experiment details.

4.1. Offline Performance

From aggregated results in Table 1, Causal-FQL (CFQL) achieves the best or near-best performance in all task groups, surpassing the causal-unaware vanilla FQL by over

Causal Flow Q-Learning

Table 2. Full offline evaluation results. Causal-FQL outperforms the vanilla FQL on all domains and achieves the best or near-best results in 19 out of 25 domains. Due to high computation requirements, we report results averaged over 4 seeds and bold values within 95% of the best performance of each domain and reuse results from prior work when available (Park et al., 2025b; Dong et al., 2025).

	Gaussian Policies			Flow Policies			
	IQL	ReBRAC	FBRAC	IQN	IFQL	FQL	Causal-FQL
visual-cube-single-play-singletask-task1-v0 (*)	70 ± 12	83 ± 6	55 ± 8	5 ± 4	49 ± 7	81 ± 12	85 ± 6
visual-cube-single-play-singletask-task2-v0	59 ± 13	85 ± 6	70 ± 11	58 ± 9	45 ± 11	58 ± 12	81 ± 5
visual-cube-single-play-singletask-task3-v0	64 ± 35	88 ± 4	81 ± 4	83 ± 11	67 ± 6	73 ± 16	84 ± 5
visual-cube-single-play-singletask-task4-v0	68 ± 12	78 ± 9	56 ± 19	78 ± 3	34 ± 13	63 ± 8	74 ± 5
visual-cube-single-play-singletask-task5-v0	59 ± 14	76 ± 9	59 ± 10	64 ± 9	27 ± 10	49 ± 12	79 ± 7
visual-cube-double-play-singletask-task1-v0 (*)	34 ± 23	4 ± 4	6 ± 2	4 ± 1	8 ± 6	23 ± 4	43 ± 11
visual-cube-double-play-singletask-task2-v0	2 ± 1	0 ± 0	2 ± 2	0 ± 0	0 ± 0	0 ± 0	2 ± 1
visual-cube-double-play-singletask-task3-v0	7 ± 4	2 ± 2	2 ± 1	0 ± 0	1 ± 1	4 ± 2	4 ± 2
visual-cube-double-play-singletask-task4-v0	1 ± 1	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0	1 ± 1
visual-cube-double-play-singletask-task5-v0	11 ± 2	0 ± 0	0 ± 0	1 ± 1	2 ± 1	4 ± 1	4 ± 2
visual-scene-play-singletask-task1-v0 (*)	97 ± 2	98 ± 4	46 ± 4	95 ± 2	86 ± 10	98 ± 3	100 ± 0
visual-scene-play-singletask-task2-v0	21 ± 16	30 ± 15	0 ± 0	79 ± 15	0 ± 0	86 ± 8	92 ± 3
visual-scene-play-singletask-task3-v0	12 ± 9	10 ± 7	10 ± 3	31 ± 14	19 ± 2	22 ± 6	22 ± 1
visual-scene-play-singletask-task4-v0	1 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0	1 ± 1	2 ± 1
visual-scene-play-singletask-task5-v0	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0
visual-puzzle-3x3-play-singletask-task1-v0 (*)	7 ± 15	88 ± 4	7 ± 2	84 ± 1	100 ± 0	94 ± 1	98 ± 1
visual-puzzle-3x3-play-singletask-task2-v0	0 ± 0	12 ± 1	0 ± 0	6 ± 2	0 ± 0	0 ± 0	1 ± 0
visual-puzzle-3x3-play-singletask-task3-v0	0 ± 0	1 ± 1	0 ± 0	1 ± 0	2 ± 1	0 ± 0	3 ± 2
visual-puzzle-3x3-play-singletask-task4-v0	1 ± 1	0 ± 1	0 ± 0	3 ± 0	1 ± 0	5 ± 4	8 ± 3
visual-puzzle-3x3-play-singletask-task5-v0	0 ± 0	0 ± 0	0 ± 0	1 ± 0	0 ± 0	1 ± 2	19 ± 4
visual-puzzle-4x4-play-singletask-task1-v0 (*)	0 ± 0	26 ± 6	0 ± 0	21 ± 3	8 ± 15	33 ± 6	37 ± 4
visual-puzzle-4x4-play-singletask-task2-v0	0 ± 0	0 ± 0	1 ± 1	14 ± 2	1 ± 1	1 ± 2	4 ± 1
visual-puzzle-4x4-play-singletask-task3-v0	0 ± 0	0 ± 0	0 ± 0	9 ± 3	9 ± 15	16 ± 6	18 ± 4
visual-puzzle-4x4-play-singletask-task4-v0	0 ± 0	1 ± 1	0 ± 0	12 ± 4	2 ± 2	2 ± 1	6 ± 2
visual-puzzle-4x4-play-singletask-task5-v0	0 ± 0	0 ± 0	1 ± 1	11 ± 4	0 ± 0	1 ± 1	1 ± 1

20%. Notably, CFQL is able to match the prior state-of-the-art Gaussian baseline (ReBRAC) performance in cube-single in which vanilla FQL falls far behind. Sec. 4 shows the full offline result in individual task of each group. CFQL is able to obtain the best or near-best performance in 19 out of 25 tasks. Especially in visual-cube-double-play-singletask-task1, CFQL nearly doubles the performance of vanilla FQL and in visual-puzzle-3x3-play-singletask-task5, our method is the only one that obtains non-near-zero success rate. We find that CFQL performs at least as well as FQL, empirically verifying that the confounding robust objective function is an informative lower bound on optimal Q-values. There are also tasks in which both FQL and CFQL don’t perform well. We hypothesize that the biggest challenges pertain in those tasks may not be the confounding biases but rather on other dimensions of offline RL, like representation learning or simply due to limited state-action space coverage of the offline data.

4.2. Offline-to-Online Performance

We also study the offline to online finetuning performance of Causal-FQL. Overall, Causal-FQL achieves strong fine-tuning performance over vanilla FQL. It brings 8% improvement on average over the prior SOTA, FQL’s

fine-tuning performance, suggesting wide applicability of Causal-FQL in both offline and offline-to-online tasks. On visual-puzzle-4x4, Causal-FQL is also the first algorithm in the literature to fully solve task1. Fig. 5 shows the full evaluation curves during training.

A notable difference between Causal-FQL and prior offline RL work is that we find it necessary to use a different objective during the online phase. Recall that in Fig. 1, confounding biases exist in the offline dataset because the data collection policy (expert) is operating in the structured state observation space while the learner’s policy is operating in the pixel observation space. The actions from the expert depend on information that may not be fully observable in the pixel space, giving rise to the confounding biases. But when it comes to online fine-tuning, the data collection policy is the learner’s policy itself and there is no observation space mismatch. Thus, we should switch back to the vanilla FQL objective function during online fine-tuning in principle.

However, in implementation, the standard practice is to gradually replace the offline trajectories in the replay buffer with online experiences so that the training doesn’t suffer from a sudden distribution shift. As a result, the actual data provided to the agent during online phase is a mixture

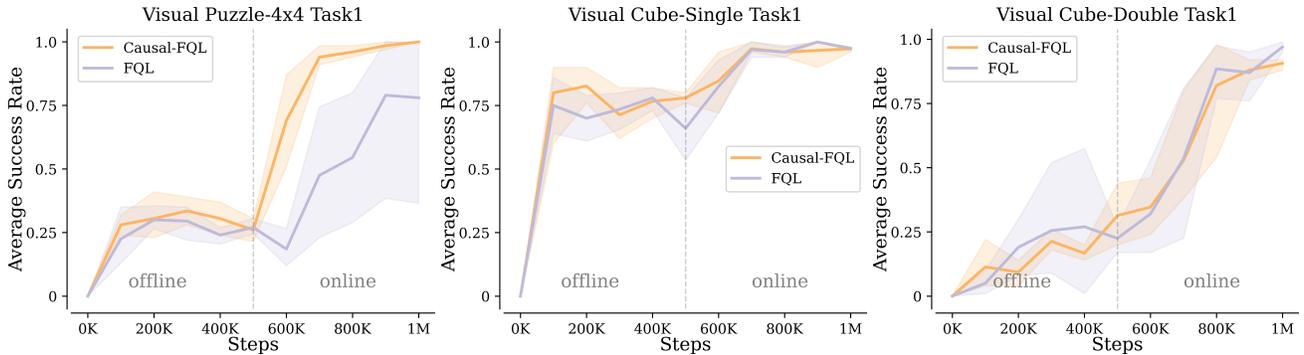


Figure 5. **Offline-to-online evaluation.** Causal-FQL converges to near optimal success rate in three representative tasks selected from task types that are not solved near-optimal in offline evaluation. All results are averaged over 4 seeds.

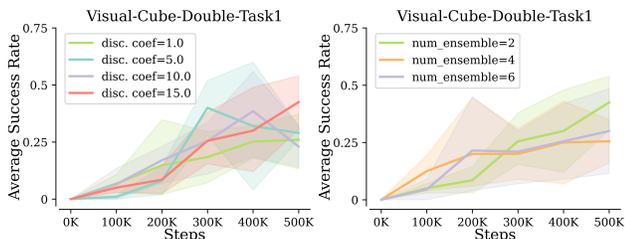


Figure 6. **A suitably tuned discriminator loss coefficient and number of critic ensembles boosts the performance.**

of confounded and confounding free data. The objective function selection is subject to tuning for specific tasks. In `visual-cube-single` and `visual-cube-double`, we find it better to fully switch back to FQL objective during online phase while in `visual-puzzle-4x4`, we find it better to use balanced sampling from both online and offline replay buffer, and use the corresponding Causal-FQL/FQL objective for each half of the batch. We believe a finer control of the objective used during online fine-tuning for Causal-FQL could bring further performance and sample efficiency improvement.

4.3. Hyperparameter Tuning

In this section, we study how to tune the two fundamental components of the confounding robust Q-value objective (Eq. (11)), the discriminator and the Q-value ensembles. We mainly tuned the hyperparameters in task 1 of `visual-cube-double` for its moderate difficulty. For each task group, we further tune the initial discriminator loss coefficient for task 1. Here we present the major findings. Detailed hyperparameters are presented in Sec. F.

Discriminator loss coefficient. We sweep the discriminator loss coefficient over $\{1, 5, 10, 15\}$. Sec. 4.3 shows

the full success rate curve over different coefficients. We notice that for overly small coefficients, the method is close to the original FQL and usually results in unstable convergence towards the end of the offline training session. Similar patterns are observed in other environments as well.

Number of the Q-network Ensembles. We vary the number of Q-value ensembles, using the values $\{2, 4, 6\}$, while keeping all other parameters set to the optimal values identified in previous tuning. Our findings indicate that, in general, fewer ensembles yield better results. This may be due to the tendency to take the minimum from a larger number of ensembles, which can lead to over-pessimism in the results.

5. Conclusion

We study the presence of implicit confounding biases in pixel-based offline reinforcement learning tasks and introduce Causal-FQL, a novel flow-based offline agent that is robust to such confounding factors. Theoretically, we establish a novel objective for the causal policy gradient that optimizes the target policy’s performance in the worst-case environment. In practice, we propose a straightforward yet high-performance approximation to the confounding-robust offline RL objective, utilizing deep ensembles of discriminators and critics. Extensive experiments show that Causal-FQL outperforms prior state-of-the-art methods in both offline and challenging visual control tasks that transition from offline to online settings. Causal-FQL is the first method to fully solve certain tasks in the OGBench visual benchmark. Finally, we emphasize the significance of principled causal analysis in RL problems; although observation-space mismatches may appear benign in prior offline RL literature, we demonstrate that they pose a major obstacle for visual tasks. Future work will focus on exploring alternative implementations of the causal policy gradient, using different critic networks and more expressive policy classes.

Impact Statement

This paper investigates the theoretical and algorithmic framework for robust policy learning from an expressive policy class using confounded offline data obtained by passively observing an expert demonstrator. Implicit confounding bias arises when the input variables that the expert uses to determine the action values are unknown and differ from the observed state used by the learner. Example 1 illustrates a situation where the expert has access to the underlying structured states, while the learner is trying to learn from video recordings of the expert’s natural trajectories.

Our framework can be applied in various fields, including autonomous vehicle development, industrial robotics, and chronic disease management, among others. Especially for safety critical domains, our method mitigates the potential risks associated with offline reinforcement learning (RL) training from demonstrations that involve unobserved confounding. We believe building trusted AI decision making systems are becoming increasingly crucial as the prevalence of black-box AI systems grows while our understanding of their long-term societal implications remains limited.

References

- Agrawalla, B. K., Nauman, M., Agrawal, K., and Kumar, A. floq: Training critics via flow-matching for scaling compute in value-based RL. In *NeurIPS 2025 Workshop: Second Workshop on Aligning Reinforcement Learning Experimentalists and Theorists*, 2025. URL <https://openreview.net/forum?id=WwQoSHGCXg>.
- Balke, A. and Pearl, J. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1172–1176, September 1997.
- Ball, P. J., Smith, L. M., Kostrikov, I., and Levine, S. Efficient online reinforcement learning with offline data. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 1577–1594. PMLR, 2023. URL <https://proceedings.mlr.press/v202/ball123a.html>.
- Bareinboim, E., Correa, J. D., Ibeling, D., and Icard, T. *On Pearl’s Hierarchy and the Foundations of Causal Inference*, pp. 507–556. Association for Computing Machinery, New York, NY, USA, 1 edition, 2022. ISBN 9781450395861. URL <https://doi.org/10.1145/3501714.3501743>.
- Bareinboim, E., Zhang, J., and Lee, S. An introduction to causal reinforcement learning. Technical Report R-65, Causal Artificial Intelligence Lab, Columbia University, December 2024.
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18(1):1–43, 2018. URL <http://jmlr.org/papers/v18/baydin17-468.html>.
- Bellman, R. Dynamic programming. *Science*, 153(3731):34–37, 1966. doi: 10.1126/science.153.3731.34. URL <https://www.science.org/doi/abs/10.1126/science.153.3731.34>.
- Bennett, A. and Kallus, N. Proximal reinforcement learning: Efficient off-policy evaluation in partially observed markov decision processes. *Oper. Res.*, 72(3):1071–1086, 2024. doi: 10.1287/OPRE.2021.0781. URL <https://doi.org/10.1287/opre.2021.0781>.
- Bennett, A., Kallus, N., Li, L., and Mousavi, A. Off-policy evaluation in infinite-horizon reinforcement learning with latent confounders. In *International Conference on Artificial Intelligence and Statistics*, pp. 1999–2007. PMLR, 2021.
- Black, K., Brown, N., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Shi, L. X., Tanner, J., Vuong, Q., Walling, A., Wang, H., and Zhilinsky, U. π_0 : A vision-language-action flow model for general robot control, 2026. URL <https://arxiv.org/abs/2410.24164>.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. Decision transformer: Reinforcement learning via sequence modeling. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 15084–15097, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/7f489f642a0ddb10272b5c31057f0663-Abstract.html>.
- Correa, J. and Bareinboim, E. General transportability of soft interventions: Completeness results. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10902–10912. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/

- 7b497aa1b2a83ec63d1777a88676b0c2-Paper.pdf.
- Dabney, W., Ostrovski, G., Silver, D., and Munos, R. Implicit quantile networks for distributional reinforcement learning. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1096–1105. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/dabney18a.html>.
- Deng, S., Yan, M., Wei, S., Ma, H., Yang, Y., Chen, J., Zhang, Z., Yang, T., Zhang, X., Cui, H., Zhang, Z., and Wang, H. GraspVLA: a grasping foundation model pre-trained on billion-scale synthetic action data. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=zEC8TOXDkH>.
- Dong, P., Zheng, C., Finn, C., Sadigh, D., and Eysenbach, B. Value flows, 2025. URL <https://arxiv.org/abs/2510.07650>.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1406–1415. PMLR, 2018. URL <http://proceedings.mlr.press/v80/espeholt18a.html>.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. D4rl: Datasets for deep data-driven reinforcement learning. arxiv, 2020.
- Fujimoto, S. and Gu, S. A minimalist approach to offline reinforcement learning. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=Q32U7dzWXpc>.
- Garg, D., Hejna, J., Geist, M., and Ermon, S. Extreme q-learning: Maxent rl without entropy, 2023. URL <https://arxiv.org/abs/2301.02328>.
- Ghugare, R. and Eysenbach, B. Normalizing flows are capable models for rl, 2025. URL <https://arxiv.org/abs/2505.23527>.
- Guo, H., Cai, Q., Zhang, Y., Yang, Z., and Wang, Z. Provably efficient offline reinforcement learning for partially observable markov decision processes. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 8016–8038. PMLR, 2022. URL <https://proceedings.mlr.press/v162/guo22a.html>.
- Hansen-Estruch, P., Kostrikov, I., Janner, M., Kuba, J. G., and Levine, S. Idql: Implicit q-learning as an actor-critic method with diffusion policies, 2023. URL <https://arxiv.org/abs/2304.10573>.
- Hess, K., Frauen, D., Melnychuk, V., and Feuerriegel, S. Efficient and sharp off-policy learning under unobserved confounding, 2025. URL <https://arxiv.org/abs/2502.13022>.
- Imbens, G. W. and Angrist, J. D. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.
- Intelligence, P., Black, K., Brown, N., Darpinian, J., Dhabalia, K., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Galliker, M. Y., Ghosh, D., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., LeBlanc, D., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Ren, A. Z., Shi, L. X., Smith, L., Springenberg, J. T., Stachowicz, K., Tanner, J., Vuong, Q., Walke, H., Walling, A., Wang, H., Yu, L., and Zhilinsky, U. $\pi_{0.5}$: a vision-language-action model with open-world generalization, 2025. URL <https://arxiv.org/abs/2504.16054>.
- Janner, M., Li, Q., and Levine, S. Offline reinforcement learning as one big sequence modeling problem. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 1273–1286, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/099fe6b0b444c23836c4a5d07346082b-Abstract.html>.
- Jiang, Y., Cheng, S., Ding, Y., Gao, F., and Qi, B. Asyncvla: Asynchronous flow matching for vision-language-action models, 2025. URL <https://arxiv.org/abs/2511.14148>.
- Joshi, S., Zhang, J., and Bareinboim, E. Towards safe policy learning under partial identifiability: a causal approach. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’24/IAAI’24/EAAI’24*. AAAI Press,

- 2024a. ISBN 978-1-57735-887-9. doi: 10.1609/aaai.v38i12.29198. URL <https://doi.org/10.1609/aaai.v38i12.29198>.
- Joshi, S., Zhang, J., and Bareinboim, E. Towards safe policy learning under partial identifiability: A causal approach. In Wooldridge, M. J., Dy, J. G., and Natarajan, S. (eds.), *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pp. 13004–13012. AAAI Press, 2024b. doi: 10.1609/AAAI.V38I12.29198. URL <https://doi.org/10.1609/aaai.v38i12.29198>.
- Kallus, N. and Zhou, A. Confounding-robust policy improvement. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 9289–9299, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/3a09a524440d44d7f19870070a5ad42f-Abstract.html>.
- Kallus, N. and Zhou, A. Confounding-robust policy evaluation in infinite-horizon reinforcement learning. *Advances in neural information processing systems*, 33: 22293–22304, 2020.
- Kostrikov, I., Nair, A., and Levine, S. Offline reinforcement learning with implicit q-learning, 2021.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/0d2b2061826a5df3221116a5085a6052-Abstract.html>.
- Lange, S., Gabel, T., and Riedmiller, M. Batch reinforcement learning. In *Reinforcement learning: State-of-the-art*, pp. 45–73. Springer, 2012.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Li, M., Zhang, J., and Bareinboim, E. Causally aligned curriculum learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=hp4yOjhwTs>.
- Li, M., Zhang, J., and Bareinboim, E. Automatic reward shaping from confounded offline data. In *Forty-second International Conference on Machine Learning*, 2025a. URL <https://openreview.net/forum?id=Hu7hUjEMiW>.
- Li, M., Zhang, J., and Bareinboim, E. Confounding robust deep reinforcement learning: A causal approach. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b. URL <https://openreview.net/forum?id=9fUr5iFu9j>.
- Lillemark, H. J., Huang, B., Zhan, F., Du, Y., and Keller, T. A. Flow equivariant world models: Memory for partially observed dynamic environments, 2026. URL <https://arxiv.org/abs/2601.01075>.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R. T., Lopez-Paz, D., Ben-Hamu, H., and Gat, I. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*, 2024.
- Liu, T., Zhao, S., and Rhinehart, N. Towards foundational liDAR world models with efficient latent flow matching. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=OyX9cC9WaV>.
- Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Lv, L., Li, Y., Luo, Y., Sun, F., Kong, T., Xu, J., and Ma, X. Flow-based policy for online reinforcement learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=CANUXhPoyn>.
- Mandlekar, A., Xu, D., Wong, J., Nasiriany, S., Wang, C., Kulkarni, R., Fei-Fei, L., Savarese, S., Zhu, Y., and Martín-Martín, R. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning*, pp. 1678–1690. PMLR, 2022.
- Manski, C. F. Nonparametric bounds on treatment effects. *The American Economic Review*, 80:319–323, 1989. URL <https://api.semanticscholar.org/CorpusID:153656329>.

- Miao, R., Qi, Z., and Zhang, X. Off-policy evaluation for episodic partially observable markov decision processes under non-parametric models. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *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*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/03dfa2a7755635f756b160e9f4c6b789-Abstract-Conference.html.
- Nair, A., Gupta, A., Dalal, M., and Levine, S. Awac: Accelerating online reinforcement learning with offline datasets, 2021. URL <https://arxiv.org/abs/2006.09359>.
- Nakamoto, M., Zhai, S., Singh, A., Mark, M. S., Ma, Y., Finn, C., Kumar, A., and Levine, S. Cal-ql: Calibrated offline RL pre-training for efficient online fine-tuning. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *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*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/c44a04289beaf0a7d968a94066a1d696-Abstract-Conference.html.
- Namkoong, H., Keramati, R., Yadlowsky, S., and Brunskill, E. Off-policy policy evaluation for sequential decisions under unobserved confounding. *Advances in Neural Information Processing Systems*, 33:18819–18831, 2020.
- Nandiraju, G., Ju, Y., Xu, C., and Wang, H. HDFlow: Hierarchical diffusion-flow planning for long-horizon robotic assembly. In *NeurIPS 2025 Workshop on Embodied World Models for Decision Making*, 2025. URL <https://openreview.net/forum?id=NDGP58PuQX>.
- Narvekar, S., Peng, B., Leonetti, M., Sinapov, J., Taylor, M. E., and Stone, P. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research*, 21:181:1–181:50, 2020. URL <http://jmlr.org/papers/v21/20-212.html>.
- Nguyen, K., Le, A. T., Pham, T., Huber, M., Peters, J., and Vu, M. N. Flowmp: Learning motion fields for robot planning with conditional flow matching, 2025. URL <https://arxiv.org/abs/2503.06135>.
- Park, S., Frans, K., Eysenbach, B., and Levine, S. OGBench: Benchmarking offline goal-conditioned RL. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=M992mjgKzI>.
- Park, S., Li, Q., and Levine, S. Flow q-learning. In *Forty-second International Conference on Machine Learning*, 2025b. URL <https://openreview.net/forum?id=KVf2SFLlpi>.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009. doi: 10.1017/CBO9780511803161.
- Pearl, J. and Robins, J. Probabilistic evaluation of sequential plans from causal models with hidden variables. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, UAI’95, pp. 444–453, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1558603859.
- Precup, D., Sutton, R. S., and Singh, S. P. Eligibility traces for off-policy policy evaluation. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 759–766, 2000.
- Robbins, H. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pp. 169–177. Springer, 1985.
- Rohbeck, M., Bunne, C., Brouwer, E. D., Huetter, J.-C., Bittan, A., Chen, K. Y., Regev, A., and Lopez, R. Modeling complex system dynamics with flow matching across time and conditions. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=hwnObmOTrV>.
- Shi, C., Uehara, M., Huang, J., and Jiang, N. A minimax learning approach to off-policy evaluation in confounded partially observable markov decision processes. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 20057–20094. PMLR, 2022. URL <https://proceedings.mlr.press/v162/shi22f.html>.
- Song, Y., Zhou, Y., Sekhari, A., Bagnell, D., Krishnamurthy, A., and Sun, W. Hybrid RL: using both offline and online data can make RL efficient. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=yyBis80iUuU>.
- Sutton, R. S. Reinforcement learning: An introduction. A Bradford Book, 2018.

- Tarasov, D., Kurenkov, V., Nikulin, A., and Kolesnikov, S. Revisiting the minimalist approach to offline reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=vqGWslLeEw>.
- Tchetgen, E. J. T., Ying, A., Cui, Y., Shi, X., and Miao, W. An introduction to proximal causal learning, 2020.
- Tian, J. and Pearl, J. Probabilities of causation: Bounds and identification. *Ann. Math. Artif. Intell.*, 28(1-4):287–313, 2000. doi: 10.1023/A:1018912507879. URL <https://doi.org/10.1023/A:1018912507879>.
- Tiofack, F. N., Hellard, T. L., Schramm, F., Perrin-Gilbert, N., and Carpentier, J. Guided flow policy: Learning from high-value actions in offline reinforcement learning, 2025. URL <https://arxiv.org/abs/2512.03973>.
- Tong, A., FATRAS, K., Malkin, N., Huguët, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. Improving and generalizing flow-based generative models with mini-batch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=CD9Snc73AW>. Expert Certification.
- Wang, Z., Hunt, J. J., and Zhou, M. Diffusion policies as an expressive policy class for offline reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=AHvFDPi-FA>.
- Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning, 2019. URL <https://arxiv.org/abs/1911.11361>.
- Yang, H., Yu, C., Sun, P., and Chen, S. Hybrid policy optimization from imperfect demonstrations. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023. Curran Associates Inc.
- Zhang, J. and Bareinboim, E. Near-optimal reinforcement learning in dynamic treatment regimes. In *Advances in Neural Information Processing Systems*, pp. 13401–13411, 2019.
- Zhang, J. and Bareinboim, E. Can humans be out of the loop? In *Conference on Causal Learning and Reasoning*, pp. 1010–1025. PMLR, 2022.
- Zhang, J. and Bareinboim, E. Eligibility traces for confounding robust off-policy evaluation. Technical Report R-105, Causal Artificial Intelligence Lab, Columbia University, May 2024.
- Zhang, J. and Bareinboim, E. Eligibility traces for confounding robust off-policy evaluation: A causal approach. In *41st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2025.
- Zhang, K., Koppel, A., Zhu, H., and Basar, T. Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612, 2020.
- Zhang, Y., Yu, S., Zhang, T., Guang, M., Hui, H., Long, K., Wang, Y., Yu, C., and Ding, W. Sac flow: Sample-efficient reinforcement learning of flow-based policies via velocity-reparameterized sequential modeling, 2026. URL <https://arxiv.org/abs/2509.25756>.

A. Related Work

Off-policy Causal Reinforcement Learning. When the no unobserved confounding assumption does not hold, one would need to either identify the reward and transition distributions before evaluating policy values or bound the possible policy values. There is a rich line of literature in identifying policy values directly from confounded data (Shi et al., 2022; Miao et al., 2022; Guo et al., 2022; Bennett & Kallus, 2024). But they usually invoke other critical learning assumptions such as the existence of bridge functions in the line of proximal causal inference literature (Tchetgen et al., 2020). On the other hand, without further assumptions, one can utilize the bounding method to account for the whole range of possible policy values. Seminal work of Manski (Manski, 1989) developed the first bounds on causal effects in non-identifiable settings using observational data in the single-stage treatment model with contextual information (i.e., a contextual bandit model). These bounds were then expanded to the instrumental variable setting (Balke & Pearl, 1997; Imbens & Angrist, 1994), to partially identify counterfactual probabilities of causation (Tian & Pearl, 2000), to construct reward shaping functions automatically (Li et al., 2025a). This work is inspired by a recent work in partially identifying the policy values via bounding (Zhang & Bareinboim, 2025) and confounding robust off-policy learning in discrete action space (Li et al., 2025b).

Offline Reinforcement Learning. Offline reinforcement learning (Levine et al., 2020) concerns the problem of learning optimal policies from a static dataset collected by other behavioral policies while maintaining close affinity to the behavioral state-action distributions in the dataset. Several different strategies have been studied in the literature to enforce this constraint, including but not limited to conservatism (Kumar et al., 2020; Nakamoto et al., 2023), generative modeling (Chen et al., 2021; Janner et al., 2021), OOD detection (Garg et al., 2023), and behavioral regularization (Fujimoto & Gu, 2021; Tarasov et al., 2023; Park et al., 2025b; Dong et al., 2025; Agrawalla et al., 2025). After offline training, an online finetuning phase can be added to further improve the agent’s performance. To avoid sudden distribution shift caused unlearning (Nakamoto et al., 2023), one could use balanced sampling (Song et al., 2023; Yang et al., 2023), maintaining offline data in the replay buffer (Park et al., 2025b; Nair et al., 2021; Ball et al., 2023) or penalize Q-values (Nakamoto et al., 2023).

Flow Matching in RL. Due to the powerful multi-modal distribution modeling capability of flow matching and relatively easy implementation compared against solving PDEs as in diffusion based methods, flow matching (Lipman et al., 2022; Liu et al., 2022; Tong et al., 2024) has gained great attention in the RL community recently. Its applications in reinforcement learning spans from planning (Nandiraju et al., 2025; Nguyen et al., 2025), learning world models (Liu et al., 2025; Rohbeck et al., 2025; Lillemark et al., 2026), online learning (Lv et al., 2025), offline learning (Park et al., 2025b; Dong et al., 2025; Agrawalla et al., 2025; Tiofack et al., 2025; Zhang et al., 2026; Ghugare & Eysenbach, 2025) to multi-modal vision-language-action models (VLAs) (Black et al., 2026; Intelligence et al., 2025; Deng et al., 2025; Jiang et al., 2025).

B. Limitations

There are mainly two types of limitations of the proposed Causal-FQL. The first type of limitations pertain to FQL (Park et al., 2025b) that 1) it requires numerically solving ODEs during training resulting in potentially slow training; 2) FQL only deploys the simplest form of flow matching. Combining with recent advancements like rectified flows (Liu et al., 2022) and optimal transport conditional flow matching (Tong et al., 2024) should bring faster convergence, stabler training and better performance; 3) FQL does not have a built in exploration mechanism, which could limit its online fine-tuning performance. The second type of limitations are from approximating the causal offline RL objective. As we introduce more components to the pipeline of FQL, the combinatorial space of tunable hyper-parameters is even harder to tune than before. Given the computational heavy nature of pixel-based offline RL tasks, a simpler confounding robust offline RL method is our next step before larger scale applications.

C. Extended Preliminaries and Proof

Causal Foundations. We briefly recap on the foundations of causal inference to facilitate our discussion. Our definition of CMDP is closely tied to the definition of Structural Causal Models (Pearl, 2009).

Definition C.1. An SCM \mathcal{M} is a tuple $\langle U, V, \mathcal{F}, P \rangle$ where

- U is a set of exogenous variables;
- V is a set of endogenous variables;

- \mathcal{F} is a set of functions s.t. each $f_V \in \mathcal{F}$;
- the exogenous distribution $P(U)$.

f_V decides values of an endogenous variable $V \in \mathbf{V}$ taking as argument a combination of other variables in the system. That is, $V \leftarrow f_V(\mathbf{PA}_V, \mathbf{U}_V)$, $\mathbf{PA}_V \subseteq \mathbf{V}$, $\mathbf{U}_V \subseteq \mathbf{U}$. Values of exogenous variables \mathbf{U} are drawn from $P(U)$.

A policy π over a subset of variables $\mathbf{X} \subseteq \mathbf{V}$ is a sequence of decision rules $\{\pi(X|\mathbf{S}_X)\}_{X \in \mathbf{X}}$, where every $\pi(X|\mathbf{S}_X)$ is a probability distribution mapping from domains of a set of covariates $\mathbf{S}_X \subseteq \mathbf{V}$ to the domain of action X . An intervention following a policy π over variables \mathbf{X} , denoted by $\text{do}(\pi)$, is an operation which sets values of every $X \in \mathbf{X}$ to be decided by policy $X \sim \pi(X|\mathbf{S}_X)$ (Correa & Bareinboim, 2020), replacing the functions $f_{\mathbf{X}} = \{f_X : \forall X \in \mathbf{X}\}$ that would normally determine their values. For an SCM \mathcal{M} , let \mathcal{M}_π be a submodel of \mathcal{M} induced by intervention $\text{do}(\pi)$. For a set $\mathbf{Y} \subseteq \mathbf{V}$, the interventional distribution $P_\pi(\mathbf{Y})$ is defined as the distribution over \mathbf{Y} in the submodel \mathcal{M}_π , i.e., $P_{\mathcal{M}}(\mathbf{Y}; \pi) \triangleq P_{\mathcal{M}_\pi}(\mathbf{Y})$; restriction \mathcal{M} is left implicit when it is obvious.

Each SCM \mathcal{M} is also associated with a causal diagram \mathcal{G} (e.g., Fig. 2), which is a directed acyclic graph (DAG) where nodes represent endogenous variables \mathbf{V} and arrows represent the arguments $\mathbf{PA}_V, \mathbf{U}_V$ of each structural function $f_V \in \mathcal{F}$. Exogenous variables \mathbf{U} are often not explicitly shown by convention. However, a bi-directed arrow $V_i \leftrightarrow V_j$ indicates the presence of an unobserved confounder (UC), $U_{i,j} \in \mathbf{U}$ affecting V_i, V_j , simultaneously (Bareinboim et al., 2022). We will use standard graph-theoretic family abbreviations to represent graphical relationships, such as parents (*pa*), children (*ch*), descendants (*de*), and ancestors (*an*). For example, the set of parent nodes of \mathbf{X} in \mathcal{G} is denoted by $\text{pa}(\mathbf{X})_{\mathcal{G}} = \cup_{X \in \mathbf{X}} \text{pa}(X)_{\mathcal{G}}$. Capitalized versions *Pa*, *Ch*, *De*, *An* include the argument as well, e.g., $\text{Pa}(\mathbf{X})_{\mathcal{G}} = \text{pa}(\mathbf{X})_{\mathcal{G}} \cup \mathbf{X}$. A path from a node X to a node Y in \mathcal{G} is a sequence of edges that does not include a particular node more than once. Two sets of nodes \mathbf{X}, \mathbf{Y} are said to be d-separated by a third set \mathbf{Z} in a DAG \mathcal{G} , denoted by $(\mathbf{X} \perp\!\!\!\perp \mathbf{Y} | \mathbf{Z})_{\mathcal{G}}$, if every edge path from nodes in \mathbf{X} to nodes in \mathbf{Y} is “blocked” by nodes in \mathbf{Z} . The criterion of blockage follows Pearl (2009, Def. 1.2.3). For more details on SCMs, we refer readers to Pearl (2009); Bareinboim et al. (2022).

From MDP to CMDP. In MDP, all the state variables are observable to the agent. When transitioning from the demonstrator’s perspective to the learner’s perspective, this assumption is implicitly assumed a priori. However, as we have demonstrated in the main text, this may not be usually the case. And we argue that the same problem also lurks in other aspects of reinforcement learning (Bareinboim et al., 2024), like and off-policy learning (Joshi et al., 2024a; Zhang & Bareinboim, 2019; Li et al., 2025b), curriculum learning (Narvekar et al., 2020; Li et al., 2024), and reward shaping (Li et al., 2025a). This motivates our proposal of CMDP as a proper modeling of the confounding biases while maintaining the Markov property for efficient policy learning.

Compared to MDP graphically, CMDP only adds three bi-directed arrows within each time step. From the definition of causal diagram we discussed above, this indicates that the three variables, current time step action X_t , reward Y_t and next state S_{t+1} , may be affected by the same latent confounding variables. In the pixel-based offline RL tasks, those variables could be occluded joint angles, exact angular velocities of the end effector or the precise location of an object, all kind of information that the demonstrator’s policy relies on but not fully observable to the learner via pixel observations.

Proof of Theorem 3.1 We are now ready to provide the detailed proof for the causal policy gradient objective in Thm. 3.1.

Theorem C.2 (Restatement of Thm. 3.1). *For a CMDP environment \mathcal{M} with reward signals $Y_t \in [a, b] \subseteq \mathbb{R}$, fix a policy π . The state value function $V_\pi(s) \geq \underline{V}_\pi(s)$ for any state $s \in \mathcal{S}$, where the lower bound $\underline{V}_\pi(s)$ is given by as follows,*

$$\underline{V}_\pi(s) = \mathbb{E}_{\substack{x \sim p(\cdot|s) \\ x', x^* \sim \pi(\cdot|s)}} \left[\mathbf{1}_{x \neq x'} \left(a + \gamma \min_{s^*} \underline{Q}_\pi(s^*, x^*) \right) + \mathbf{1}_{x=x'} \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s'} \tilde{\mathcal{T}}(s, x, s') \underline{Q}_\pi(s', x^*) \right) \right] \quad (14)$$

Proof. Since $\underline{V}_\pi(s) = \sum_x \pi(x|s) \underline{Q}_\pi(s, x)$, it follows from Eq. (5) that the state value function is lower bounded by

$$\underline{V}_\pi(s) = \sum_x \pi(x|s) \mu(\neg x|s) \left(a + \gamma \min_{s'} \underline{V}_\pi(s') \right) \quad (15)$$

$$+ \sum_x \pi(x|s) \mu(x|s) \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s', x'} \tilde{\mathcal{T}}(s, x, s') \underline{V}_\pi(s') \right) \quad (16)$$

Following the simplification step used in (Zhang & Bareinboim, 2024, Theorem 1),

$$\underline{V}_\pi(s) = \sum_x \mu(x | s) \pi(\neg x | s) \left(a + \gamma \min_{s'} \underline{V}_\pi(s') \right) \quad (17)$$

$$+ \sum_x \mu(x | s) \pi(x | s) \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s', x'} \tilde{\mathcal{T}}(s, x, s') \underline{V}_\pi(s') \right) \quad (18)$$

The above equation could be further written as

$$\underline{V}_\pi(s) = \sum_{x, x'} \mu(x | s) \pi(x' | s) \mathbf{1}_{x \neq x'} \left(a + \gamma \min_{s'} \underline{V}_\pi(s') \right) \quad (19)$$

$$+ \sum_{x, x'} \mu(x | s) \pi(x' | s) \mathbf{1}_{x=x'} \left(\tilde{\mathcal{R}}(s, x) + \gamma \sum_{s', x'} \tilde{\mathcal{T}}(s, x, s') \underline{V}_\pi(s') \right) \quad (20)$$

Finally, replacing the summation over the policy product $\mu(x | s) \pi(x' | s)$ with the expectation proves the statement. \square

D. Additional Results

Other than manipulation tasks, we also test Causal-FQL in visual locomotion tasks like `visual-antmaze` series. The results are not as amazing as those on manipulation tasks, which are as expected. We hypothesize that this is because FQL and Causal-FQL sample actions from the policy directly without considering the noisy distribution of Q-values while Gaussian policy extraction with rejection sampling handles uni-modal reward distributions better. Thus, prior Gaussian SOTA like BRAC performs well in locomotion tasks. As an indirect proof to our hypothesis, in Sec. D, ReBRAC (Tarasov et al., 2023) also outperforms its own flow-based counterpart, FBRAC (Park et al., 2025b), and IFQL (FQL with rejection sampling) also outperforms FQL.

Table 3. **Locomotion offline evaluation results.** Both Causal-FQL and vanilla FQL are not on par with Gaussian based methods. We report results averaged over 4 seeds and bold values within 95% of the best performance of each domain following prior work (Park et al., 2025b).

	Gaussian Policies		Flow Policies				
	IQL	ReBRAC	FBRAC	IQN	IFQL	FQL	Causal-FQL
<code>visual-antmaze-medium-navigate-singletask-task1-v0</code> (*)	78 ± 9	54 ± 15	27 ± 3	62 ± 7	81 ± 3	32 ± 3	40 ± 15
<code>visual-antmaze-medium-navigate-singletask-task2-v0</code>	90 ± 3	96 ± 1	42 ± 4	88 ± 2	87 ± 1	60 ± 2	61 ± 2
<code>visual-antmaze-medium-navigate-singletask-task3-v0</code>	80 ± 6	97 ± 1	32 ± 4	64 ± 5	92 ± 1	35 ± 8	45 ± 3
<code>visual-antmaze-medium-navigate-singletask-task4-v0</code>	89 ± 4	93 ± 2	23 ± 2	71 ± 6	84 ± 3	35 ± 2	39 ± 2
<code>visual-antmaze-medium-navigate-singletask-task5-v0</code>	84 ± 2	97 ± 1	25 ± 4	84 ± 1	89 ± 2	29 ± 7	29 ± 6
<code>visual-antmaze-teleport-navigate-singletask-task1-v0</code> (*)	5 ± 2	2 ± 0	1 ± 1	2 ± 1	7 ± 4	2 ± 1	4 ± 2
<code>visual-antmaze-teleport-navigate-singletask-task2-v0</code>	10 ± 2	10 ± 3	6 ± 5	7 ± 3	13 ± 3	6 ± 1	7 ± 4
<code>visual-antmaze-teleport-navigate-singletask-task3-v0</code>	7 ± 7	4 ± 1	10 ± 4	6 ± 4	8 ± 9	9 ± 4	14 ± 5
<code>visual-antmaze-teleport-navigate-singletask-task4-v0</code>	4 ± 6	4 ± 0	10 ± 2	4 ± 2	18 ± 2	9 ± 1	11 ± 3
<code>visual-antmaze-teleport-navigate-singletask-task5-v0</code>	2 ± 1	2 ± 1	2 ± 1	2 ± 1	4 ± 2	1 ± 1	4 ± 2

E. Implementation Details

Causal-FQL is built upon the codebase of FQL (Park et al., 2025b)³. Causal-FQL inherits most of the architectural design of FQL with minimum changes implemented for the causal offline RL objective. First, we add an action discriminator and a confounding-robust Q weighting scheme: it builds an extra action discriminator network (the same size as the Q-critic), uses it in actor loss to classify flow actions vs one-step target flow actions, computes factual weights from discriminator logits, and mixes ensemble Q-values as Eq. (11); it also adds a discriminator loss term with optional exponential decay and a switch during online learning phase that disables discriminator weighting and trains with plain mean Q values over the ensembles. Due to limited computation resources, we don't have the capacity to fully tune the full hyperparameter space of Causal-FQL on visual tasks. But we do observe consistent performance improvement in all visual tasks over FQL using the default FQL hyperparameters.

³<https://github.com/seohongpark/fql>

F. Experiment Details

Tasks and Environments. We use pixel based tasks from OGBench (Park et al., 2025a) for evaluation. We use a total of 25 visual tasks across 5 environments for evaluation. Below is a list of the datasets we used, each of which defines 5 reward maximizing tasks indexed from `-task1` to `-task5`. In the main text, we evaluated manipulation tasks involving multi-step of subtasks with rewards bounded by the negative of number of subtasks and 0. In Sec. D, we evaluate locomotion tasks with rewards -1 or 0 indicating whether the goal is reached.

- `visual-cube-single-play-singletask`
- `visual-cube-double-play-singletask`
- `visual-scene-play-singletask`
- `visual-puzzle-3x3-play-singletask`
- `visual-puzzle-4x4-play-singletask`

Methods and Hyperparameters. We adopt the baseline implementation from FQL (Park et al., 2025b) and Value Flows (Dong et al., 2025).⁴ Below we list the baseline methods we used. We take default hyper-parameters recommended in FQL and Value Flows code base. See the attached link in footnote for detailed command to replicate the results.

- **IQL** (Kostrikov et al., 2021). Implicit Q-Learning use expectile regression to represent Q-values which is then used to select optimal actions. We use the default inverse temperature α recommended in the codebase.
- **ReBRAC** (Tarasov et al., 2023). ReBRAC is a Gaussian policy based offline RL method that uses TD3 to learn values with behavioral regularizations.
- **FBRAC** (Park et al., 2025b). A variant of BRAC with flow policies instead. It doesn’t use the one-step flow policy extraction thus requiring backpropagation through time.
- **IQN** (Dabney et al., 2018). IQN is a distributional RL baseline that represents the Q-value distribution with quantile values. Action selection is based on taking argmax over sampled actions’ Q-values.
- **IFQL** (Park et al., 2025b). A variant of Implicit Diffusion Q (IDQL, (Hansen-Estruch et al., 2023)) with flow policy and rejection sampling action selection, which is first proposed in (Park et al., 2025b).
- **FQL** (Park et al., 2025b). FQL uses one-step flow policy to maximize the Q-values learned by an ensemble of critics. It uses a BC flow policy as the behavioral regularization.

G. Causal Value Flows

Here we describe the implementation details of Causal Value Flows and preliminary offline evaluation results for applying the newly proposed causal offline RL objective to Value Flows (Dong et al., 2025) as a demonstration on the versatility of our method. We adopt the discriminator coefficient and number of ensembles setup from Causal-FQL and hold other parameters as default in value flows. As in Causal-FQL, we add an extra action discriminator network, the same size as the Q-critic network. When calculating Eq. (11), we use the values sampled by Euler method with the trained return vector field ensembles. Then we reweight the mean Q-values and add the minimum Q-value as the worst case value. Other parts of value flows stay unchanged.

Without further tuning, on task 1 of `visual-cube-double-play`, the augmented causal value flow is able to achieve over 0.6 success rate within 400K offline training steps, **almost doubles the vanilla value flow’s performance**, which only achieves 0.32 with 1M offline training steps.

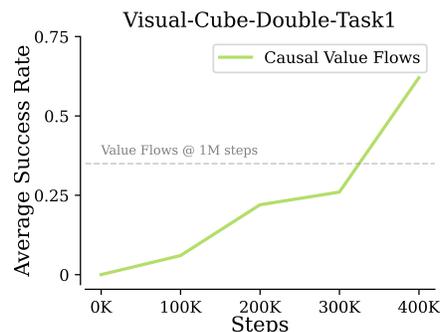


Figure 7. **Causal Value Flows outperform vanilla Value Flows.**

⁴<https://github.com/chongyi-zheng/value-flows>