

# Scaling Marginalized Importance Sampling to High-Dimensional State-Spaces via State Abstraction

Brahma S. Pavse, Josiah P. Hanna

University of Wisconsin – Madison  
pavse@wisc.edu, jphanna@cs.wisc.edu

## Abstract

We consider the problem of off-policy evaluation (OPE) in reinforcement learning (RL), where the goal is to estimate the performance of an evaluation policy,  $\pi_e$ , using a fixed dataset,  $\mathcal{D}$ , collected by one or more policies that may be different from  $\pi_e$ . Current OPE algorithms may produce poor OPE estimates under policy distribution shift i.e., when the probability of a particular state-action pair occurring under  $\pi_e$  is very different from the probability of that same pair occurring in  $\mathcal{D}$ . In this work, we propose to improve the accuracy of OPE estimators by projecting the high-dimensional state-space into a low-dimensional state-space using concepts from the state abstraction literature. Specifically, we consider marginalized importance sampling (MIS) OPE algorithms which compute state-action distribution correction ratios to produce their OPE estimate. In the original ground state-space, these ratios may have high variance which may lead to high variance OPE. However, we prove that in the lower-dimensional abstract state-space the ratios can have lower variance resulting in lower variance OPE. We then highlight the challenges that arise when estimating the abstract ratios from data, identify sufficient conditions to overcome these issues, and present a minimax optimization problem whose solution yields these abstract ratios. Finally, our empirical evaluation on difficult, high-dimensional state-space OPE tasks shows that the abstract ratios can make MIS OPE estimators achieve lower mean-squared error and more robust to hyperparameter tuning than the ground ratios.

## 1 Introduction

One of the key challenges when applying reinforcement learning (RL) (Sutton and Barto 2018) to real-world tasks is the problem of off-policy evaluation (OPE) (Fu et al. 2021; Voloshin et al. 2021). The goal of OPE is to evaluate a policy of interest by leveraging offline data generated by possibly different policies. Solving the OPE problem would enable us to estimate the performance of a potentially risky policy without having to actually deploy it. This capability is especially important for sensitive real-world tasks such as healthcare and autonomous driving.




Figure 1: Left side: AntUMaze domain. Right side: Projecting high-dimensional ant into lower-dimensional point-mass.

The core OPE problem is to produce accurate policy value estimates under policy distribution shift. This problem is particularly difficult on tasks with high-dimensional state-spaces (Voloshin et al. 2021; Fu et al. 2021). For example, consider the AntUMaze problem illustrated on the left side of Figure 1. In this task, an ant-like robot with a high-dimensional state representation moves in a U-shaped maze and receives a reward only for reaching a specific 2D coordinate goal location. The state-space of this task includes information such as 2D location, ant limb angles, torso orientation etc., resulting in a 29-dimensional state-space. The OPE task is to evaluate the performance of a particular policy’s ability to take the ant to the 2D goal location using data that may be collected by different policies. Policy distribution shift is common in this type of high-dimensional task since the chances of different policies inducing similar limb angles, torso orientations, paths traversed etc. are incredibly slim. Notice, however, that while different policies may induce different body configurations, they may traverse similar 2D paths since all (successful) policies must move the ant through roughly the same path to reach the goal. Moreover, the only critical information needed from the state-space to determine the ant’s per-step reward are its 2D coordinates. Motivated by this observation, we propose to improve the accuracy of OPE algorithms on high-dimensional state-space tasks by projecting the high-dimensional state-space into a lower-dimensional space. This idea is illustrated

on the right side of Figure 1 where the ant is reduced to a 2D point-mass.

With this general motivation in mind, in this paper, we leverage concepts from the state abstraction literature (Li, Walsh, and Littman 2006) to improve the accuracy of marginalized importance sampling (MIS) OPE algorithms which estimate state-action density correction ratios to compute a policy value estimate (Liu et al. 2018a; Xie, Ma, and Wang 2019; Yin and Wang 2020). Due to the low chances of similarity between states of policies in high-dimensional state-spaces, current MIS algorithms can produce high variance state-action density ratios, resulting in high variance OPE estimates. However, if we are given a suitable state abstraction function, we can project the high-dimensional *ground* state-space into a lower-dimensional *abstract* state-space. The projection step increases the chances of similarity between these lower-dimensional states, resulting in low variance density ratios and OPE estimates. To the best of our knowledge, this work is the first to leverage concepts from state abstraction to improve OPE. We make the following contributions:

1. Theoretical analyses showing: (a) the variance of abstract state-action ratios is at most that of ground state-action ratios; and (b) that our abstract MIS OPE estimator is unbiased, strongly consistent, and can have lower variance than its ground equivalent.
2. An algorithm, based on a popular class of MIS algorithms, to estimate the abstract state-action ratios.
3. An empirical analysis of our estimator on a variety of high-dimensional state-space tasks.

## 2 Preliminaries

In this section, we discuss relevant background information.

### 2.1 Notation and Problem Setup

We consider an infinite-horizon Markov decision process (MDP),  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, P, \gamma, d_0 \rangle$ , where  $\mathcal{S}$  is the state-space,  $\mathcal{A}$  is the action-space,  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta([0, \infty))$  is the reward function,  $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  is the transition dynamics function,  $\gamma \in [0, 1)$  is the discount factor, and  $d_0 \in \Delta(\mathcal{S})$  is the initial state distribution. The agent acting, according to policy  $\pi$ , in the MDP generates a trajectory:  $s_0, a_0, r_0, s_1, a_1, r_1, \dots$ , where  $s_0 \sim d_0$ ,  $a_t \sim \pi(\cdot|s_t)$ ,  $r_t \sim \mathcal{R}(s_t, a_t)$ , and  $s_{t+1} \sim P(\cdot|s_t, a_t)$  for  $t \geq 0$ . We define  $r(s, a) := \mathbb{E}_{r \sim \mathcal{R}(s, a)}[r]$  and the agent’s discounted state-action occupancy measure under policy  $\pi$  as  $d_\pi$ :

$$d_\pi(s, a) := \lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} \gamma^t d_\pi(s_t, a_t)}{\sum_{t=0}^{T-1} \gamma^t}$$

where  $d_\pi(s_t, a_t)$  is the probability the agent will be in state  $s$  and take action  $a$  at time-step  $t$  under policy  $\pi$ .

In the infinite-horizon setting, we define the performance of policy  $\pi$  to be its average reward,  $\rho(\pi) := \mathbb{E}_{(s, a) \sim d_\pi, r \sim \mathcal{R}(s, a)}[r]$ .

Note that  $\mathbb{E}_{(s, a) \sim d_\pi, r \sim \mathcal{R}(s, a)}[r] = (1 - \gamma) \mathbb{E}_{s_0 \sim d_0, a_0 \sim \pi} [q^\pi(s_0, a_0)]$  where  $q^\pi(s, a) :=$

$\mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r_{k+t} | s_t = s, a_t = a]$  is the action-value function which satisfies  $q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(s, a), a' \sim \pi(s')} [q^\pi(s', a')]$ .

### 2.2 Off-Policy Evaluation (OPE)

In behavior-agnostic off-policy evaluation (OPE), the goal is to estimate the performance of an evaluation policy  $\pi_e$  given only a fixed offline data set of transition tuples,  $\mathcal{D} := \{(s_i, a_i, s'_i, r_i)\}_{i=1}^{mT}$ , where  $(s_i, a_i) \sim d_{\mathcal{D}}$ ,  $m$  is the batch size (number of trajectories), and  $T$  is the fixed length of each trajectory, generated by *unknown* and possibly *multiple* behavior policies. The difficulty in OPE is to estimate  $\rho(\pi_e)$  under  $d_{\pi_e}$  given samples only from  $d_{\mathcal{D}}$ .

We define the average-reward in dataset  $\mathcal{D}$  to be  $\bar{r}_{\mathcal{D}} := \mathbb{E}_{(s, a) \sim d_{\mathcal{D}}, r \sim \mathcal{R}(s, a)}[r]$ . As in prior OPE work, we assume that if  $d_{\pi_e}(s, a) > 0$  then  $d_{\mathcal{D}}(s, a) > 0$ . Empirically, we measure the accuracy of an estimate  $\hat{\rho}(\pi_e)$  by generating  $M$  datasets and then computing the *relative* mean-squared error (MSE):  $\text{MSE}(\hat{\rho}(\pi_e)) := \frac{1}{M} \sum_{i=1}^M \frac{(\rho(\pi_e) - \hat{\rho}_i(\pi_e))^2}{(\rho(\pi_e) - \bar{r}_{\mathcal{D}_i})^2}$ , where  $\hat{\rho}_i(\pi_e)$  is computed using dataset  $\mathcal{D}_i$  and  $\bar{r}_{\mathcal{D}_i}$  is the average reward in  $\mathcal{D}_i$ .

**Marginalized Importance Sampling (MIS)** In this work, we focus on MIS methods, which evaluate  $\pi_e$  by using the ratio between  $d_{\pi_e}$  and  $d_{\mathcal{D}}$ . That is, MIS methods evaluate  $\pi_e$  by estimating  $\rho(\pi_e) := \mathbb{E}_{(s, a) \sim d_{\mathcal{D}}, r \sim \mathcal{R}(s, a)}[\zeta(s, a)r]$ , where

$$\zeta(s, a) := \frac{d_{\pi_e}(s, a)}{d_{\mathcal{D}}(s, a)}$$

is the state-action density ratio for state-action pair  $(s, a)$  and  $d_\pi(s, a) = d_\pi(s)\pi(a|s)$ . When the true  $\zeta$  is known, the empirical estimate of  $\rho(\pi_e)$  is:

$$\hat{\rho}(\pi_e) := \frac{1}{N} \sum_{i=1}^N \zeta(s_i, a_i) r(s_i, a_i) \quad (1)$$

where  $N$  is the number of samples. In practice, however,  $\zeta$  is unknown and must be estimated.

One set of  $\zeta$ -estimation algorithms, which have also shown potential for good OPE performance (Voloshin et al. 2021), is the DICE family (Yang et al. 2020). While there are many variations, the general DICE optimization problem is:

$$\begin{aligned} \max_{\zeta: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}} \min_{\nu: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}} J(\zeta, \nu) := \\ \mathbb{E}_{(s, a, s') \sim d_{\mathcal{D}}, a' \sim \pi_e} [\zeta(s, a)(\nu(s, a) - \gamma \nu(s', a'))] \\ - (1 - \gamma) \mathbb{E}_{s_0 \sim d_0, a_0 \sim \pi_e} [\nu(s_0, a_0)] \end{aligned} \quad (2)$$

where the solution to the optimization,  $\zeta^*(s, a)$ , are the true ratios. The estimator we present in Section 3 builds upon the DICE framework.

### 2.3 State Abstractions

We define a state abstraction function as a mapping  $\phi : \mathcal{S} \rightarrow \mathcal{S}^\phi$ , where  $\mathcal{S}$  is called the *ground* state-space and  $\mathcal{S}^\phi$  is called the *abstract* state-space. We consider state abstraction functions that partition the ground state-space into disjoint sets.

We can use  $\phi$  to project the original MDP into a new abstract MDP with the same action-space  $\mathcal{A}$  and reward and transition dynamics functions defined as:

$$\begin{aligned}\mathcal{R}^\phi(s^\phi, a) &= \sum_{s \in \phi^{-1}(s^\phi)} w(s) \mathcal{R}(s, a) \\ P^\phi(s'^\phi | s^\phi, a) &= \sum_{s \in \phi^{-1}(s^\phi), s' \in \phi^{-1}(s'^\phi)} w(s) P(s' | s, a)\end{aligned}$$

where  $w : \mathcal{S} \rightarrow [0, 1]$  is a ground state weighting function where for each  $s^\phi \in \mathcal{S}^\phi$ ,  $\sum_{s \in \phi^{-1}(s^\phi)} w(s) = 1$  (Li, Walsh, and Littman 2006). Similarly a policy can be transformed into its abstract equivalent as:

$$\pi^\phi(a | s^\phi) = \sum_{s \in \phi^{-1}(s^\phi)} w(s) \pi(a | s).$$

We define the following state-weighting function for an arbitrary policy  $\pi$ :  $w_\pi(s) = \frac{d_\pi(s)}{\sum_{s' \in \phi^{-1}(s^\phi)} d_\pi(s')}$  and only consider abstractions that satisfy:

**Assumption 1** (Reward distribution equality).  $\forall s_1, s_2 \in \mathcal{S}$  such that  $\phi(s_1) = \phi(s_2)$ ,  $\forall a, \mathcal{R}(s_1, a) = \mathcal{R}(s_2, a)$ .

Assumption 1 implies that, regardless of the choice of the state-weighting function, for given action  $a$ ,  $\forall s \in s^\phi$ ,  $\mathcal{R}^\phi(s^\phi, a) = \mathcal{R}(s, a)$  i.e. the reward distribution of an abstract state equals that of the ground states within that abstract state.

### 3 Abstract Marginalized Importance Sampling

Marginalized IS methods may suffer from high variance in high-dimensional state-spaces. To potentially reduce this high variance, we propose to first use  $\phi$  to project  $\mathcal{D}$  into the abstract state-space to obtain:  $\mathcal{D}^\phi := \{(s^\phi, a, r^\phi, s'^\phi)\}$  where  $s^\phi = \phi(s)$  and  $r^\phi(s^\phi, a) = r(s, a) \forall s \in s^\phi$ , and then use the following estimator on  $\mathcal{D}^\phi$  to estimate  $\rho(\pi_e^\phi)$ :

**Definition 1** (Abstract estimator). We define our estimator of  $\rho(\pi_e^\phi)$  as follows:

$$\hat{\rho}(\pi_e^\phi) := \frac{1}{N} \sum_{i=1}^N \frac{d_{\pi_e^\phi}(s_i^\phi, a_i)}{d_{\mathcal{D}^\phi}(s_i^\phi, a_i)} r^\phi(s_i^\phi, a_i) \quad (3)$$

where  $N$  is the number of samples,  $d_{\pi_e^\phi}(s^\phi, a) = d_{\pi_e}(s^\phi) \pi^\phi(a | s^\phi)$  with  $d_{\pi_e}(s^\phi) = \sum_{s \in \phi^{-1}(s^\phi)} d_\pi(s)$  and  $\pi^\phi$  constructed using  $w_\pi$ .

In the remainder of this section, we first give an example to build intuition for why the abstract ratios  $\frac{d_{\pi_e^\phi}(s^\phi, a)}{d_{\mathcal{D}^\phi}(s^\phi, a)}$  can have lower variance than the ground ratios and then show theoretically that the OPE estimator given in Equation (3) is strongly consistent and can produce lower variance OPE estimates of  $\rho(\pi_e)$  than the ground equivalent (Equation (1)) when the true ratios are known. Finally, when the true abstract ratios are unknown, we identify sufficient conditions under which unbiased estimation of the ratios is possible and adapt an existing DICE algorithm to estimate them.

### 3.1 A Hard Example for Ground MIS Ratios

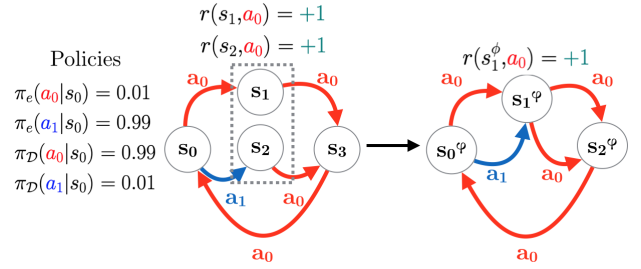


Figure 2: TwoPath MDP where ground density ratios for  $(s_1, a_0)$  and  $(s_2, a_0)$  are high variance. However, upon aggregation of equivalent states (grey dotted lines), the abstract density ratio of  $(s_1^\phi, a_0)$  is low variance.

We present a hard example for ground MIS ratios in Figure 2 that provides intuition for why the abstract MIS ratios can have lower variance ratios than the ground ratios. Consider two symmetric policies,  $\pi_e$  and  $\pi_D$ , executed in the ground MDP (left side). In this example, the high variance of the true state-action density ratios  $\frac{d_{\pi_e}(s_1, a_0)}{d_{\pi_D}(s_1, a_0)} \approx 0$  and  $\frac{d_{\pi_e}(s_2, a_0)}{d_{\pi_D}(s_2, a_0)} \approx 100$  can lead to high variance estimates of  $\rho(\pi_e)$ . Notice, however, that states  $s_1$  and  $s_2$  are essentially equivalent i.e.  $r(s_1, a) = r(s_2, a) \forall a \in \mathcal{A}$  and can be aggregated together into a single state,  $s_1^\phi$  (Assumption 1). We find that the state-action density ratio in this abstract MDP (right side)  $\frac{d_{\pi_e^\phi}(s_1^\phi, a_0)}{d_{\pi_D^\phi}(s_1^\phi, a_0)} = 1$  is of low variance, which can lead to low variance estimates of  $\rho(\pi_e)$ .

In general, we prove that the abstract ratios are guaranteed to have variance at most that of the ground ratios (proof in Appendix):

$$\text{Theorem 1. } \text{Var} \left( \frac{d_{\pi_e^\phi}(s^\phi, a)}{d_{\mathcal{D}^\phi}(s^\phi, a)} \right) \leq \text{Var} \left( \frac{d_{\pi_e}(s, a)}{d_{\mathcal{D}}(s, a)} \right)$$

where equality holds if either or both of the following are true: 1)  $\phi$  is the identity function i.e.  $\phi(s) = s, \forall s \in \mathcal{S}$  and 2) if  $\forall s_1, s_2 \in \mathcal{S}$  such that  $\phi(s_1) = \phi(s_2) = s^\phi$  and for a given action  $a$ ,  $\frac{d_{\pi_e}(s_1, a)}{d_{\mathcal{D}}(s_1, a)} = \frac{d_{\pi_e}(s_2, a)}{d_{\mathcal{D}}(s_2, a)}, \forall s^\phi \in \mathcal{S}^\phi, a \in \mathcal{A}$ . Thus, Theorem 1 implies that projecting  $\mathcal{S} \rightarrow \mathcal{S}^\phi$  can lower the variance of density ratios.

### 3.2 MIS OPE with True Abstract Ratios

We now present the statistical properties of our estimator assuming it has access to the true abstract state-action ratios. Due to space constraints, we defer proofs to Appendix.

We prove our abstract estimator is unbiased (Theorem 4 in Appendix) and strongly consistent (Theorem 2 and Corollary 1):

**Theorem 2.** Our estimator,  $\hat{\rho}(\pi_e^\phi)$ , given in Equation 3 is an asymptotically consistent estimator of  $\rho(\pi_e)$  in terms of MSE:  $\lim_{N \rightarrow \infty} \mathbb{E}[(\hat{\rho}(\pi_e^\phi) - \rho(\pi_e))^2] = 0$ .

We also compare the variances between our abstract estimator and the ground equivalent. If we assume that  $\mathcal{D}$  is i.i.d as done in previous work (Sutton et al. 2008; Uehara, Huang, and Jiang 2020; Nachum et al. 2019; Zhang\* et al. 2020), then  $\text{Var}(\hat{\rho}(\pi_e^\phi)) \leq \text{Var}(\hat{\rho}(\pi_e))$ , where equality holds only under the same conditions as described below Theorem 1. In general, however, variance reduction depends on the covariances between the weighted per-step rewards (Liu, Bacon, and Brunskill 2020):

**Theorem 3.** *If Assumption 1 and if for any fixed  $1 \leq t < k \leq T$ ,*

$$\text{Cov}\left(\frac{d_{\pi_e^\phi}(s_t^\phi, a_t)}{d_{\mathcal{D}^\phi}(s_t^\phi, a_t)} r^\phi(s_t^\phi, a_t), \frac{d_{\pi_e^\phi}(s_k^\phi, a_k)}{d_{\mathcal{D}^\phi}(s_k^\phi, a_k)} r^\phi(s_k^\phi, a_k)\right) \leq$$

$$\text{Cov}\left(\frac{d_{\pi_e}(s_t, a_t)}{d_{\mathcal{D}}(s_t, a_t)} r(s_t, a_t), \frac{d_{\pi_e}(s_k, a_k)}{d_{\mathcal{D}}(s_k, a_k)} r(s_k, a_k)\right) \text{ hold then}$$

$$\text{Var}(\hat{\rho}(\pi_e^\phi)) \leq \text{Var}(\hat{\rho}(\pi_e)).$$

### 3.3 Estimating the Abstract Ratios

Thus far, we have assumed access to the true abstract ratios. However, in practice, these ratios are unknown and must be estimated from  $\mathcal{D}^\phi$ . In this section, we highlight the challenges in estimating the abstract ratios and identify sufficient conditions on  $\phi$  that allow for accurate estimation. Once  $\phi$  satisfies these conditions, any off-the-shelf state-action density estimation method can be used to estimate the abstract ratios. In this work, we focus on showing that DICE estimates the true abstract ratios. We note that the following new conditions on  $\phi$  are only needed for unbiased estimation of the abstract ratios; accurate OPE on the abstracted MDP can still be realized under only Assumption 1.

We first observe that evaluating  $\pi_e$  using  $\mathcal{D}$  is equivalent to evaluating  $\pi_e^\phi$  using  $\mathcal{D}^\phi$  if  $\mathcal{D}^\phi$  is generated from an abstract MDP with transition dynamics constructed according to  $w_{\pi_e}$ . This equivalence is given by the following proposition:

**Proposition 1.** *If Assumption 1 holds, the average reward of ground policy  $\pi$  executed in ground MDP  $\mathcal{M}$ ,  $\rho(\pi)$ , is equal to the average reward of abstract policy  $\pi^\phi$  executed in abstract MDP  $\mathcal{M}^\phi$  constructed with  $w_\pi$ ,  $\rho(\pi^\phi)$ . That is,  $\rho(\pi) = \rho(\pi^\phi)$ .*

Proposition 1 suggests that we can estimate  $\rho(\pi_e)$  by applying any OPE algorithm to evaluate  $\pi_e^\phi$  using  $\mathcal{D}^\phi$  if the abstract transition dynamics of  $\mathcal{D}^\phi$  are distributed according to  $P_{w_{\pi_e}}^\phi$ . Unfortunately, since  $w_{\pi_e}$  is unknown, two challenges arise. Fortunately, there are special cases where unbiased estimation of the abstract ratios is still possible using only existing MIS algorithms.

#### Challenge 1: Transition dynamics distribution shift.

The off-policy data  $\mathcal{D}^\phi := \{(s^\phi, a, r^\phi, s'^\phi)\}$  is distributed in the following way:  $(s^\phi, a) \sim d_{\mathcal{D}^\phi}, r^\phi \sim \mathcal{R}^\phi, s'^\phi \sim P_{w_{\mathcal{D}^\phi}}^\phi(s^\phi, a)$  where  $P_{w_{\mathcal{D}^\phi}}^\phi$  are the transition dynamics of the abstract MDP constructed with  $w_{\mathcal{D}^\phi}$  as the state-weighting function. Thus, in addition to the original policy distribution shift problem, we also encounter a *transition dynamics distribution shift* problem due to the projection where we want to evaluate  $\pi_e^\phi$  in an MDP with  $P_{w_{\pi_e}}^\phi$ , but we only have samples of data generated in an MDP with  $P_{w_{\mathcal{D}^\phi}}^\phi$ . Moreover, since

$w_{\pi_e}$  is unknown, we cannot compute  $P_{w_{\pi_e}}^\phi$  and correct the distribution shift, say through importance sampling (Precup, Sutton, and Singh 2000), as we would correct for policy distribution shift. However, one condition on  $\phi$  that will avoid the transition distribution shift is:

**Assumption 2** (Transition dynamics similarity).  $\forall s_1, s_2 \in \mathcal{S}$  such that  $\phi(s_1) = \phi(s_2), \forall a \in \mathcal{A}, x \in \mathcal{S}^\phi$ , we have  $\sum_{s' \in \phi^{-1}(x)} P(s'|s_1, a) = \sum_{s' \in \phi^{-1}(x)} P(s'|s_2, a)$ .

Together with Assumption 1,  $\phi$  is now the so-called *bisimulation* abstraction (Ferns, Panangaden, and Precup 2011; Castro 2020). This property of  $\phi$  eliminates the transition dynamics distribution shift since now  $P_{w_{\pi_e}}^\phi(s'^\phi|s^\phi, a) = P_{w_{\mathcal{D}^\phi}}^\phi(s'^\phi|s^\phi, a) = P^\phi(s'^\phi|s^\phi, a) = \sum_{s' \in s'^\phi} P(s'|s, a) \forall s \in s^\phi, \forall a \in \mathcal{A}$  (applying Assumption 2 and definition of  $P^\phi$  from Section 2.3). Thus, we have  $\mathcal{D}^\phi$  distributed with  $P_{w_{\pi_e}}^\phi = P^\phi$ , which then allows us to apply any MIS algorithm to compute the abstract state-action ratios using  $\mathcal{D}^\phi$ .

**Challenge 2: Inaccessible  $\pi_e^\phi$ .** To the best of our knowledge, all existing MIS algorithms require access to the evaluation policy to estimate the density ratios. However, in the abstract MDP,  $\pi_e^\phi$  is inaccessible since  $w_{\pi_e}$  is unknown. To overcome this issue, we identify the following condition on  $\phi$ :

**Assumption 3** ( $\pi_e$  action-distribution equality).  $\forall s_1, s_2 \in \mathcal{S}$  such that  $\phi(s_1) = \phi(s_2), \pi_e(s_1) = \pi_e(s_2)$

Assumption 3 then gives us  $\pi_e^\phi(s^\phi) = \pi_e(s), \forall s \in s^\phi$  (applying Assumption 3 and definition of  $\pi^\phi$  from Section 2.3), which allows us to simulate sampling from  $\pi_e^\phi(\cdot|\phi(s))$  by just sampling from  $\pi_e(\cdot|s)$ .

Given a  $\phi$  with these properties, we can compute the abstract ratios needed to estimate  $\rho(\pi_e)$  by applying a suitable MIS algorithm to  $\mathcal{D}^\phi$ . We use BestDICE (Yang et al. 2020), and call our algorithm AbstractBestDICE, which solves the following optimization problem:

$$\begin{aligned} \min_{\nu, \lambda} \max_{\zeta} J(\nu, \zeta, \lambda) := & -\mathbb{E}_{\mathcal{D}^\phi} \left[ \frac{1}{2} \zeta(s^\phi, a)^2 \right] \\ & + \mathbb{E}_{\mathcal{D}^\phi} \left[ \zeta(s^\phi, a) \left( \gamma \mathbb{E}_{a' \sim \pi_e^\phi} [\nu(s'^\phi, a')] - \nu(s^\phi, a) - \lambda \right) \right] \\ & + (1 - \gamma) \mathbb{E}_{s_0^\phi \sim d_{0^\phi}, a_0 \sim \pi_e^\phi} [\nu(s_0^\phi, a_0)] + \lambda \end{aligned} \quad (4)$$

where,  $\nu : \mathcal{S}^\phi \times \mathcal{A} \rightarrow \mathbb{R}, \lambda \in \mathbb{R}$ , and  $\zeta : \mathcal{S}^\phi \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ . The solution to the optimization,  $\zeta^*(s^\phi, a) = d_{\pi_e^\phi}(s^\phi, a)/d_{\mathcal{D}^\phi}(s^\phi, a)$  is the true abstract ratios. Since the derivation of AbstractBestDICE follows the same steps as BestDICE, we defer it to Appendix.

We note that it may be difficult to validate Assumptions 2 and 3 in practice, which may result in loss of the consistency guarantee of Theorem 2. Nevertheless, in Section 4 we show that AbstractBestDICE leads to accurate OPE in high-dimensional state-spaces even when assumptions may not hold.

## 4 Empirical Study

We will now show how projecting  $\mathcal{S} \rightarrow \mathcal{S}^\phi$  can produce more accurate OPE estimates in practice. We answer the fol-

lowing questions:

1. Do the true abstract ratios produce lower variance OPE estimates than the true ground ratios?
2. Does AbstractBestDICE: (a) compute the true ratios when Assumptions 1, 2, and 3 are satisfied and (b) produce data-efficient and stable estimates of  $\rho(\pi_e)$  even when Assumptions 2 and 3 fail to hold?

#### 4.1 Empirical Setup

In this section, we describe the algorithms and domains of our empirical study. Due to space constraints, we defer specific details to the Appendix.

**Algorithms** We compare AbstractBestDICE to ground BestDICE (Yang et al. 2020). As also reported by Yang et al. (2020); Fu et al. (2021), we found in preliminary experiments (see Appendix) that BestDICE performed much better than other MIS methods such as DualDICE (Nachum et al. 2019), Minimax-Weight Learning (Uehara, Huang, and Jiang 2020), etc.

**Domains** We focus on high-dimensional state-space tasks, which have been known to be particularly challenging for DICE methods (Fu et al. 2021). For each environment below we specify a fixed  $\phi$ .

- **TwoPath MDP.** This MDP is pictured in Figure 2. In this domain, Assumptions 1, 2, and 3 are satisfied. We also run the same experiments for when these assumptions are violated (see Appendix).
- **Reacher** (Brockman et al. 2016). A robotic arm tries to move to a goal location. Here,  $s \in \mathbb{R}^{11}$  and  $a \in \mathbb{R}^2$ . Since the reward function is the Euclidean distance between the arm and goal,  $\phi$  extracts only the arm-to-goal vector, and angular velocities from the ground state, resulting in  $s^\phi \in \mathbb{R}^4$ .
- **Walker2D** (Brockman et al. 2016). A bi-pedal robot tries to move as fast as possible. Here,  $s \in \mathbb{R}^{18}$  and  $a \in \mathbb{R}^6$ . We use the Euclidean distance from the start location as the reward function and use a  $\phi$  that extracts  $x$  and  $z$  coordinates and top angle of the walker’s body, resulting in  $s^\phi \in \mathbb{R}^3$ .
- **Pusher** (Brockman et al. 2016). A robotic arm tries to push an object to a goal location. Here,  $s \in \mathbb{R}^{23}$  and  $a \in \mathbb{R}^7$ . Since the reward function is the Euclidean distance between object and arm and object and goal,  $\phi$  extracts only object-to-arm and object-to-goal vectors, resulting in  $s^\phi \in \mathbb{R}^6$ .
- **AntUMaze** (Fu et al. 2020). This sparse-reward task requires an ant to move from one end of the U-shaped maze to the other end. Here,  $s \in \mathbb{R}^{29}$  and  $a \in \mathbb{R}^8$ . We use the “play” version where the goal location is fixed. Since the reward function is +1 only if the 2D location of the ant is at a certain Euclidean distance from the 2D goal location,  $\phi$  extracts only the 2D coordinates of the ant, resulting in  $s^\phi \in \mathbb{R}^2$ .

For Reacher, Walker2D, Pusher, and AntUMaze,  $\phi$  satisfies only Assumption 1.


#### 4.2 Empirical Results

In this section, we describe our main empirical results; additional experiments can be found in Appendix.

**True Ratios for OPE** We conduct an experiment on the TwoPath MDP to estimate  $\rho(\pi_e)$  where we apply the ground estimator given in Equation (1) and our abstract estimator given in Equation (3), assuming *both have access to their respective true ratios*. The results of this experiment are illustrated in Figure 3(a). We can observe that the abstract estimator with the true abstract ratios produces substantially more data-efficient and lower variance OPE estimates for different batch sizes compared to the ground equivalent.



(a) MSE vs. Batch size (# of trajectories). The vertical axis axis is log-scaled. Errors are computed over 15 trials with 95% confidence intervals. Lower is better. Since  $\rho(\pi_e) = \rho(\pi_D)$  in this domain, we use regular MSE instead of relative.



(b) Estimated  $\hat{d}_{\pi_e^\phi}$  (vertical axis) vs. true  $d_{\pi_e^\phi}$  (horizontal axis). Values are averaged over 15 trials. We expect the dots to be as close to the black line as possible. Each dot is for each  $(s^\phi, a)$ . Dots are only at extreme ends due to choice of  $\pi_e, \pi_D$ , and ToyMDP design. Only 4 out of 6 dots are visible due to overlap between the dots for  $(s_1^\phi, a_1)$  and  $(s_2^\phi, a_1)$ , and  $(s_1^\phi, a_0)$  and  $(s_2^\phi, a_0)$ .

Figure 3: True ratio experiments.

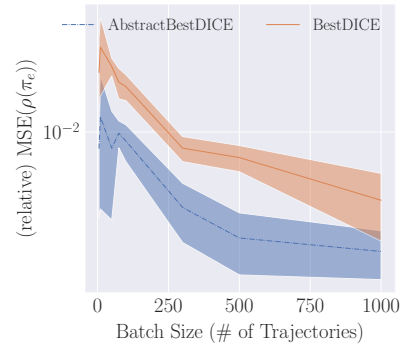
**True Ratio Estimation** To verify if AbstractBestDICE accurately estimates the true ratios we conduct the following experiment on the TwoPath MDP. We give AbstractBestDICE data of batch size 300 to estimate the abstract ratios  $\hat{\zeta}^\phi$  and then use  $\hat{\zeta}^\phi$  to estimate the abstract state-action densities of  $\pi_e^\phi$ ,  $\hat{d}_{\pi_e^\phi}(s^\phi, a) = \hat{\zeta}^\phi(s^\phi, a)d_{\pi_D^\phi}(s^\phi, a)$ , where we have access to the true  $d_{\pi_D^\phi}(s^\phi, a)$ . We then compare  $\hat{d}_{\pi_e^\phi}$  to the true  $d_{\pi_e^\phi}$ , which we compute using a batch size of 300 trajectories collected from  $\pi_e^\phi$  roll-outs, using a correlation plot shown in Figure 3(b). From the figure we can see AbstractBestDICE accurately estimates the abstract state-action density ratios. When assumptions are violated, however, ratio estimation accuracy can reduce (see Appendix).

**Data-Efficiency** Figure 4 shows the results of our (relative) MSE vs. batch size experiment for the function approximation case. For a given batch size, we train each algorithm for 100k epochs with different hyperparameters sets, record the (relative) MSE on the last epoch by each hyperparameter set, and plot the lowest MSE achieved by these hyperparameter sets. We find that AbstractBestDICE is able to achieve lower MSE than BestDICE for a given batch size. We note that while hyperparameter tuning is difficult in OPE, in this experiment, we aim to evaluate the performance of each algorithm assuming both had favorable hyperparameters.

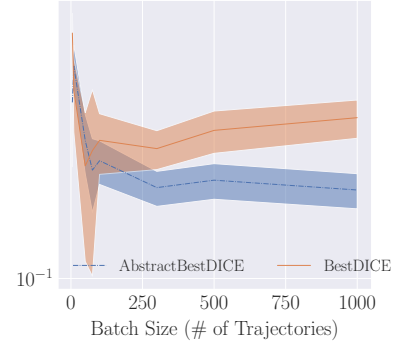
**Hyperparameter Robustness** Finally, we study the robustness of these algorithms to hyperparameters tuning. In practical OPE, hyperparameter tuning with respect to MSE is impractical since the true  $\rho(\pi_e)$  is unknown (Fu et al. 2021; Paine et al. 2020). Thus, we want OPE algorithms to be as robust as possible to hyperparameter tuning. The main hyperparameters for DICE are the learning rates of  $\zeta$  and  $\nu$ ,  $\alpha_\zeta$  and  $\alpha_\nu$ . For these experiments, we focus on very small batch sizes, where we would expect high sensitivity. The results of this study are in Figure 5. We find that our algorithm has a less volatile MSE than BestDICE (also see Appendix for more similar results). In a related experiment, we also find AbstractBestDICE can be more stable than BestDICE during training (see Appendix).

## 5 Related Work

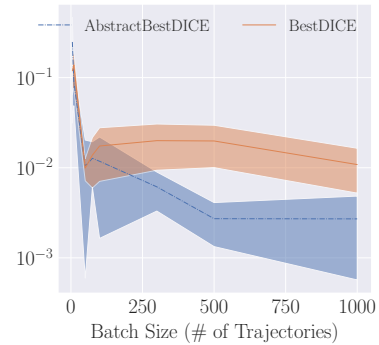
**MIS and Off-Policy Evaluation.** There have been broadly three families of MIS algorithms in the OPE literature to estimate state-action density ratios. One is the DICE family, which includes: minimax-weight learning (Uehara, Huang, and Jiang 2020), DualDice (Nachum et al. 2019), GenDICE (Zhang\* et al. 2020), GradientDICE (Zhang, Liu, and Whiteson 2020), and BestDICE (Yang et al. 2020). In our work, we adapt BestDICE to estimate the abstract ratios. The second family of MIS algorithms is the COP-TD algorithm (Hallak and Mannor 2017; Gelada and Bellemare 2019), which learns the state density ratios with an on-line TD-styled update. The third family is the variational power method (Wen et al. 2020) algorithm which generalizes the power iteration method to estimate density ratios. While our focus has been on MIS algorithms, there are many other OPE algorithms such as model-based methods (Zhang



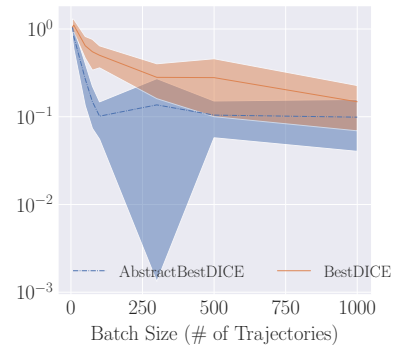
(a) Reacher



(b) Walker2D



(c) Pusher



(d) AntUMaze

Figure 4: Relative MSE vs. Batch Size (# of trajectories). Vertical axis is log-scaled. Errors are computed over 15 trials with 95% confidence intervals. Lower is better.

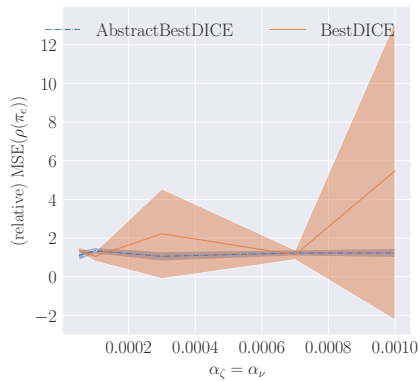


Figure 5: Robustness of BestDICE and AbstractBestDICE to hyperparameters on the AntUMaze domain for batch size (# of trajectories) of 5. Errors are computed over 15 trials with 95% confidence intervals. Lower is better.

et al. 2021b; Hanna, Stone, and Niekum 2017; Liu et al. 2018b), fitted-Q evaluation (Le, Voloshin, and Yue 2019), doubly-robust methods (Jiang and Li 2016; Thomas and Brunskill 2016), and IS (Precup, Sutton, and Singh 2000; Thomas 2015; Hanna, Niekum, and Stone 2019; Thomas, Theocharous, and Ghavamzadeh 2015).

**State Abstraction and Representation Learning.** The literature on state abstraction is extensive (Singh, Jaakkola, and Jordan 1994; Dietterich 2000; Ferns, Panangaden, and Precup 2011; Li, Walsh, and Littman 2006; Abel 2020). However, much of this work has been exclusively focused on building a theory of abstraction and on learning optimal policies. A related topic to state abstraction is representation learning. Recently, there has been much work showing the importance of good representations for offline RL (Wang et al. 2021; Yin et al. 2022; Geng et al. 2022; Zhan et al. 2022; Chen and Jiang 2022). To the best of our knowledge, no work has leveraged state abstraction techniques to improve the accuracy of OPE algorithms.

## 6 Summary and Future Work

In this work, we showed that we can improve the accuracy of OPE estimates by projecting the original ground state-space into a lower-dimensional abstract state-space using state abstraction and performing OPE in the resulting abstract Markov decision process. Our theoretical results proved that: 1) abstract state-action ratios have variance at most that of the ground ratios; and 2) the abstract MIS OPE estimator is unbiased, strongly consistent, and can have lower variance than the ground equivalent. We then highlighted the challenges that arise when estimating the abstract ratios from data, identified sufficient conditions to overcome these issues, and adapted BestDICE into AbstractBestDICE to estimate the abstract ratios. In our empirical results, we obtained more accurate OPE estimates with added hyperparameter robustness on difficult, high-dimensional state-space tasks.

There are several directions for future work. First, As-

sumptions 2 and 3 are strict. Further investigation is needed to see if these assumptions can be relaxed. Second, we assumed the abstraction function was given. It would be interesting to leverage existing ideas (Gelada et al. 2019; Zhang et al. 2021a) to learn  $\phi$ . Finally, we want to emphasize that this work instantiates the general abstraction + OPE direction. While this work focused exclusively on MIS algorithms, a promising direction will be to apply abstraction techniques to model-based, trajectory IS, and value-function based OPE.

## Acknowledgements

Support for this research was provided by the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin — Madison with funding from the Wisconsin Alumni Research Foundation. The authors thank the anonymous reviewers, Nicholas Corrado, Ishan Durugkar, and Subhojyoti Mukherjee for their helpful comments in improving this work.

## References

- Abel, D. 2020. *A Theory of Abstraction in Reinforcement Learning*. Ph.D. thesis, Brown University.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *CoRR*, abs/1606.01540.
- Castro, P. S. 2020. Scalable Methods for Computing State Similarity in Deterministic Markov Decision Processes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(06): 10069–10076.
- Chen, J.; and Jiang, N. 2022. Offline Reinforcement Learning Under Value and Density-Ratio Realizability: The Power of Gaps. In *The 38th Conference on Uncertainty in Artificial Intelligence*.
- Dietterich, T. G. 2000. Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition. *J. Artif. Int. Res.*, 13(1): 227–303.
- Ferns, N.; Panangaden, P.; and Precup, D. 2011. Bisimulation Metrics for Continuous Markov Decision Processes. *SIAM Journal on Computing*, 40(6): 1662–1714.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4RL: Datasets for Deep Data-Driven Reinforcement Learning. arXiv:2004.07219.
- Fu, J.; Norouzi, M.; Nachum, O.; Tucker, G.; Wang, Z.; Novikov, A.; Yang, M.; Zhang, M. R.; Chen, Y.; Kumar, A.; Paduraru, C.; Levine, S.; and Paine, T. 2021. Benchmarks for Deep Off-Policy Evaluation. In *ICLR*.
- Gelada, C.; and Bellemare, M. G. 2019. Off-Policy Deep Reinforcement Learning by Bootstrapping the Covariate Shift. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 3647–3655.
- Gelada, C.; Kumar, S.; Buckman, J.; Nachum, O.; and Bellemare, M. G. 2019. DeepMDP: Learning Continuous Latent Space Models for Representation Learning. *CoRR*, abs/1906.02736.

- Geng, X.; Li, K.; Gupta, A.; Kumar, A.; and Levine, S. 2022. Effective Offline RL Needs Going Beyond Pessimism: Representations and Distributional Shift. In *Decision Awareness in Reinforcement Learning Workshop at ICML 2022*.
- Hallak, A.; and Mannor, S. 2017. Consistent On-Line off-Policy Evaluation. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, 1372–1383. JMLR.org.
- Hanna, J.; Niekum, S.; and Stone, P. 2019. Importance Sampling Policy Evaluation with an Estimated Behavior Policy. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*.
- Hanna, J.; Stone, P.; and Niekum, S. 2017. Bootstrapping with Models: Confidence Intervals for Off-Policy Evaluation. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Jiang, N.; and Li, L. 2016. Doubly Robust Off-Policy Value Evaluation for Reinforcement Learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, 652–661. JMLR.org.
- Le, H.; Voloshin, C.; and Yue, Y. 2019. Batch Policy Learning under Constraints. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3703–3712. PMLR.
- Li, L.; Walsh, T. J.; and Littman, M. L. 2006. Towards a Unified Theory of State Abstraction for MDPs. In *Proceedings of the Ninth International Symposium on Artificial Intelligence and Mathematics*, 531–539.
- Liu, Q.; Li, L.; Tang, Z.; and Zhou, D. 2018a. Breaking the Curse of Horizon: Infinite-Horizon Off-Policy Estimation. *CoRR*, abs/1810.12429.
- Liu, Y.; Bacon, P.-L.; and Brunskill, E. 2020. Understanding the Curse of Horizon in Off-Policy Evaluation via Conditional Importance Sampling. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Liu, Y.; Gottesman, O.; Raghu, A.; Komorowski, M.; Faisal, A. A.; Doshi-Velez, F.; and Brunskill, E. 2018b. Representation Balancing MDPs for Off-policy Policy Evaluation. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Nachum, O.; Chow, Y.; Dai, B.; and Li, L. 2019. DualDICE: Behavior-Agnostic Estimation of Discounted Stationary Distribution Corrections. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Paine, T. L.; Paduraru, C.; Michi, A.; Gülçehre, Ç.; Zolna, K.; Novikov, A.; Wang, Z.; and de Freitas, N. 2020. Hyperparameter Selection for Offline Reinforcement Learning. *CoRR*, abs/2007.09055.
- Precup, D.; Sutton, R. S.; and Singh, S. P. 2000. Eligibility Traces for Off-Policy Policy Evaluation. In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML '00*, 759–766. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1558607072.
- Singh, S.; Jaakkola, T.; and Jordan, M. 1994. Reinforcement Learning with Soft State Aggregation. In Tesauro, G.; Touretzky, D.; and Leen, T., eds., *Advances in Neural Information Processing Systems*, volume 7. MIT Press.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. The MIT Press, second edition.
- Sutton, R. S.; Szepesvári, C.; Geramifard, A.; and Bowling, M. 2008. Dyna-Style Planning with Linear Function Approximation and Prioritized Sweeping. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI'08*, 528–536. Arlington, Virginia, USA: AUAI Press. ISBN 0974903949.
- Thomas, P.; Theodorou, G.; and Ghavamzadeh, M. 2015. High-Confidence Off-Policy Evaluation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
- Thomas, P. S. 2015. *Safe Reinforcement Learning*. Ph.D. thesis, University of Massachusetts Amherst.
- Thomas, P. S.; and Brunskill, E. 2016. Data-Efficient Off-Policy Policy Evaluation for Reinforcement Learning. arXiv:1604.00923.
- Uehara, M.; Huang, J.; and Jiang, N. 2020. Minimax Weight and Q-Function Learning for off-Policy Evaluation. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Voloshin, C.; Le, H. M.; Jiang, N.; and Yue, Y. 2021. Empirical Study of Off-Policy Policy Evaluation for Reinforcement Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Wang, R.; Wu, Y.; Salakhutdinov, R.; and Kakade, S. 2021. Instabilities of Offline RL with Pre-Trained Neural Representation. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 10948–10960. PMLR.
- Wen, J.; Dai, B.; Li, L.; and Schuurmans, D. 2020. Batch Stationary Distribution Estimation. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Xie, T.; Ma, Y.; and Wang, Y.-X. 2019. Towards Optimal Off-Policy Evaluation for Reinforcement Learning with Marginalized Importance Sampling. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Yang, M.; Nachum, O.; Dai, B.; Li, L.; and Schuurmans, D. 2020. Off-Policy Evaluation via the Regularized Lagrangian. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6551–6561. Curran Associates, Inc.
- Yin, M.; Duan, Y.; Wang, M.; and Wang, Y.-X. 2022. Near-optimal Offline Reinforcement Learning with Linear Representation: Leveraging Variance Information with Pessimism. In *International Conference on Learning Representations*.



- Yin, M.; and Wang, Y. 2020. Asymptotically Efficient Off-Policy Evaluation for Tabular Reinforcement Learning. *CoRR*, abs/2001.10742.
- Zhan, W.; Huang, B.; Huang, A.; Jiang, N.; and Lee, J. 2022. Offline Reinforcement Learning with Realizability and Single-policy Concentrability. In Loh, P.-L.; and Raginsky, M., eds., *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, 2730–2775. PMLR.
- Zhang, A.; McAllister, R. T.; Calandra, R.; Gal, Y.; and Levine, S. 2021a. Learning Invariant Representations for Reinforcement Learning without Reconstruction. In *International Conference on Learning Representations*.
- Zhang, M. R.; Paine, T.; Nachum, O.; Paduraru, C.; Tucker, G.; ziyu wang; and Norouzi, M. 2021b. Autoregressive Dynamics Models for Offline Policy Evaluation and Optimization. In *International Conference on Learning Representations*.
- Zhang\*, R.; Dai\*, B.; Li, L.; and Schuurmans, D. 2020. GenDICE: Generalized Offline Estimation of Stationary Values. In *International Conference on Learning Representations*.
- Zhang, S.; Liu, B.; and Whiteson, S. 2020. GradientDICE: Rethinking Generalized Offline Estimation of Stationary Values. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 11194–11203. PMLR.