

# OPTIDICE: OFFLINE POLICY OPTIMIZATION VIA STATIONARY DISTRIBUTION CORRECTION ESTIMATION

Jongmin Lee<sup>\*,1</sup>, Weonseok Jeon<sup>\*,3,4</sup>, Byung-Jun Lee<sup>1,5</sup>, Joelle Pineau<sup>3,4,6</sup>, Kee-Eung Kim<sup>1,2</sup>

<sup>1</sup>School of Computing, KAIST, Daejeon, Republic of Korea

<sup>2</sup>Graduate School of AI, KAIST, Daejeon, Republic of Korea

<sup>3</sup>Mila, Quebec AI Institute

<sup>4</sup>School of Computer Science, McGill University

<sup>5</sup>Gauss Labs Inc., Seoul, Republic of Korea

<sup>6</sup>Facebook AI Research

jmlee@ai.kaist.ac.kr, jeonwons@mila.quebec

## ABSTRACT

We consider the offline reinforcement learning (RL) setting where the agent aims to optimize the policy solely from the data without further environment interactions. In offline RL, the distributional shift becomes the primary source of difficulty, which arises from the deviation of the target policy being optimized from the behavior policy used for data collection. This typically causes overestimation of action values, which poses severe problems for model-free algorithms that use bootstrapping. To mitigate the problem, prior offline RL algorithms often used sophisticated techniques that encourage underestimation of action values, which introduces an additional set of hyperparameters that need to be tuned properly. In this paper, we present an offline RL algorithm that prevents overestimation in a more principled way. Our algorithm, OptiDICE, tightly integrates the optimization of the target policy and the stationary distribution ratio estimation of the target policy and the behavior policy. Using an extensive set of benchmark datasets for offline RL, we show that OptiDICE performs competitively with the state-of-the-art methods.

## 1 INTRODUCTION

The availability of large-scale datasets has been one of the important factors contributing to the recent success in machine learning, for real-world tasks such as computer vision (Deng et al., 2009; Krizhevsky et al., 2012) and natural language processing (Devlin et al., 2019). The standard workflow in developing systems for typical machine learning tasks is to train and validate the model on the dataset, and then to deploy the model with its parameter fixed when we are satisfied with training. This offline training supports various operational requirements of the deployed system, such as guaranteeing the minimum prediction accuracy once the system goes online.

However, this workflow is not straightforwardly applicable to the standard setting of reinforcement learning (RL) (Sutton & Barto, 1998) because of the online learning assumption: the RL agent needs to continuously explore the environment and learn from its trial-and-error experiences to be properly trained. This aspect has been one of the fundamental bottlenecks for the practical adoption of RL in many real-world domains, where the exploratory behaviors are costly or even dangerous, e.g. autonomous driving (Yu et al., 2020b) and clinical treatment (Yu et al., 2020a).

Offline RL (also referred to as batch RL) (Ernst et al., 2005; Lange et al., 2012; Fujimoto et al., 2019; Levine et al., 2020) casts the RL problem in the offline training setting. Due to its capability of learning from offline dataset, offline RL has been often adopted to continual and skill learning problems; Cabi et al. (2020) utilize offline RL with dataset gathered in continual learning process;

---

\*Equal contribution

Pertsch et al. (2020) learn skill priors via offline RL and uses them to improve exploration; Lu et al. (2021) train model and policy via offline learning and fine-tune the policy with interactions, showing their methods are robust to non-stationarity.

One of the most relevant areas of research in this regard is the off-policy RL (Lillicrap et al., 2016; Haarnoja et al., 2018; Fujimoto et al., 2018), since we need to deal with the distributional shift resulting from the trained policy being deviated from the policy used to collect the data. However, without the data continuously collected online, this distributional shift cannot be reliably corrected and poses a significant challenge to RL algorithms that employ bootstrapping together with function approximation: it causes compounding overestimation of the action values for model-free algorithms (Fujimoto et al., 2019; Kumar et al., 2019), which arises from computing the bootstrapped target using the predicted values of *out-of-distribution* actions. To mitigate the problem, most of the current offline RL algorithms have proposed sophisticated techniques to encourage underestimation of action values, introducing an additional set of hyperparameters that needs to be tuned properly (Fujimoto et al., 2019; Kumar et al., 2019; Jaques et al., 2019; Lee et al., 2020; Kumar et al., 2020).

In this paper, we present an offline RL algorithm that essentially eliminates the need to evaluate out-of-distribution actions, thus avoiding the problematic overestimation of values. Our algorithm, *offline policy Optimization via stationary DIstribution Correction Estimation* (OptiDICE), estimates stationary distribution ratios that correct the discrepancy between the data distribution and the *optimal* policy’s stationary distribution. We first show that such optimal stationary distribution corrections can be estimated via a minimax optimization that does not involve sampling from the target policy. Then, we derive and exploit the closed-form solution to the sub-problem of the aforementioned minimax optimization, which reduces the overall problem into an unconstrained convex optimization, and thus greatly stabilizing our method. To the best of our knowledge, OptiDICE is the first deep offline RL algorithm that optimizes policy purely in the space of *stationary distributions*, rather than in the space of either Q-functions or policies (Nachum et al., 2019b). In the experiments, we demonstrate that OptiDICE performs competitively with the state-of-the-art methods using the D4RL offline RL benchmarks (Fu et al., 2021).

## 2 BACKGROUND

We consider the reinforcement learning problem with the environment modeled as a Markov Decision Process (MDP)  $M = \langle S, A, T, R, p_0, \gamma \rangle$  (Sutton & Barto, 1998), where  $S$  is a set of states  $s$ ,  $A$  is a set of actions  $a$ ,  $R : S \times A \rightarrow \mathbb{R}$  is a reward function,  $T : S \times A \rightarrow \Delta(S)$  is a transition probability,  $p_0 \in \Delta(S)$  is an initial state distribution, and  $\gamma \in [0, 1]$  is a discount factor. The policy  $\pi : S \rightarrow \Delta(A)$  is a mapping from state to distribution over actions. While  $T(s, a)$  and  $\pi(s)$  indicate distributions by definition, we let  $T(s'|s, a)$  and  $\pi(a|s)$  denote their evaluations for brevity. For the given policy  $\pi$ , the stationary distribution  $d^\pi$  is defined as

$$d^\pi(s, a) = \begin{cases} (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s, a_t = a) & \text{if } \gamma < 1, \\ \lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T \Pr(s_t = s, a_t = a) & \text{if } \gamma = 1, \end{cases}$$

where  $s_0 \sim p_0$  and  $a_t \sim \pi(s_t)$ ,  $s_{t+1} \sim T(s_t, a_t)$  for all non-negative integer  $t$ . The goal of RL is to learn an optimal policy that maximizes rewards through interactions with the environment:  $\max_{\pi} \mathbb{E}_{(s,a) \sim d^\pi} [R(s, a)]$ . The value functions of policy  $\pi$  is defined as  $Q^\pi(s, a) := \mathbb{E}_{\pi, M} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) | s_0 = s, a_0 = a]$  and  $V^\pi(s) := \mathbb{E}_{a \sim \pi(s)} [Q^\pi(s, a)]$ , where the action-value function  $Q^\pi$  is a unique solution of the Bellman equation:

$$Q^\pi(s, a) = R(s, a) + \gamma \mathbb{E}_{s' \sim T(s, a), a' \sim \pi(s')} [Q^\pi(s', a')].$$

In offline RL, the agent optimizes the policy from static dataset  $D = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^N$  collected before the training phase. We denote the empirical distribution of the dataset by  $d^D$  and will abuse the notation  $d^D$  to represent  $s \sim d^D$ ,  $(s, a) \sim d^D$ , and  $(s, a, s') \sim d^D$ .

Prior offline RL algorithms (Fujimoto et al., 2019; Kumar et al., 2019; Wu et al., 2019; Lee et al., 2020; Kumar et al., 2020; Nachum et al., 2019b) rely on estimating Q-values for optimizing the target policy. This procedure often yields unreasonably high Q-values due to the compounding error from bootstrapped estimation with out-of-distribution actions sampled from the target policy (Kumar et al., 2019).

### 3 OPTIDICE

In this section, we present *offline policy Optimization via stationary DIstribution Correction Estimation* (OptiDICE). Instead of the *optimism in the face of uncertainty* principle (Szita & Lörincz, 2008) in online RL, we discourage the uncertainty as in most offline RL algorithms (Kidambi et al., 2020; Yu et al., 2020c); otherwise, the resulting policy may fail to improve on the data-collection policy, or even suffer from severe performance degradation (Petrik et al., 2016; Laroche et al., 2019). Specifically, we consider the regularized policy optimization framework (Nachum et al., 2019b)

$$\pi^* := \arg \max_{\pi} \mathbb{E}_{(s,a) \sim d^\pi} [R(s,a)] - \alpha D_f(d^\pi || d^D), \quad (1)$$

where  $D_f(d^\pi || d^D) := \mathbb{E}_{(s,a) \sim d^D} \left[ f\left(\frac{d^\pi(s,a)}{d^D(s,a)}\right) \right]$  is the  $f$ -divergence between the stationary distribution  $d^\pi$  and the dataset distribution  $d^D$ , and  $\alpha > 0$  is a hyperparameter that balances between pursuing the reward-maximization and penalizing the deviation from the distribution of the offline dataset (i.e. penalizing distributional shift). We assume  $d^D > 0$  and that  $f$  being strictly convex and continuously differentiable. Note that we impose regularization in the space of stationary distributions rather than in the space of policies (Wu et al., 2019). However, (1) in its form involves evaluation of  $d^\pi$ , which is not directly accessible in the offline RL setting.

To make the optimization tractable, we reformulate (1) in terms of optimizing a stationary distribution  $d : S \times A \rightarrow \mathbb{R}$ . For brevity, we consider discounted MDPs ( $\gamma < 1$ ) and then generalize the result to undiscounted MDPs ( $\gamma = 1$ ). Using  $d$ , we rewrite (1) as

$$\max_{d \geq 0} \mathbb{E}_{(s,a) \sim d} [R(s,a)] - \alpha D_f(d || d^D) \quad (2)$$

$$\text{s.t. } (\mathcal{B}_* d)(s) = (1 - \gamma)p_0(s) + \gamma(\mathcal{T}_* d)(s) \quad \forall s, \quad (3)$$

where  $(\mathcal{B}_* d)(s) := \sum_{\bar{a}} d(s, \bar{a})$  is a marginalization operator, and  $(\mathcal{T}_* d)(s) := \sum_{\bar{s}, \bar{a}} T(s|\bar{s}, \bar{a})d(\bar{s}, \bar{a})$  is a transposed Bellman operator. Note that when  $\alpha = 0$ , the optimization (2-3) is exactly the dual formulation of the linear program (LP) for finding an optimal policy of the MDP (Puterman, 1994), where the constraints (3-3) are often called the Bellman flow constraints. Once the optimal stationary distribution  $d^*$  is obtained, we can recover the optimal policy  $\pi^*$  in (1) from  $d^*$  by  $\pi^*(a|s) = \frac{d^*(s,a)}{\sum_{\bar{a}} d^*(s,\bar{a})}$ .

We can then obtain the following Lagrangian for the constrained optimization problem in (2-3):

$$\max_{d \geq 0} \min_{\nu} \mathbb{E}_{(s,a) \sim d} [R(s,a)] - \alpha D_f(d || d^D) + \sum_s \nu(s) \left( (1 - \gamma)p_0(s) + \gamma(\mathcal{T}_* d)(s) - (\mathcal{B}_* d)(s) \right), \quad (4)$$

where  $\nu(s)$  are the Lagrange multipliers. Lastly, we eliminate the direct dependence on  $d$  and  $\mathcal{T}_*$  by rearranging the terms in (4) and optimizing the distribution ratio  $w$  instead of  $d$  (Nachum et al., 2019a):

$$\begin{aligned} & \mathbb{E}_{(s,a) \sim d} [R(s,a)] - \alpha D_f(d || d^D) + \sum_s \nu(s) \left( (1 - \gamma)p_0(s) + \gamma(\mathcal{T}_* d)(s) - (\mathcal{B}_* d)(s) \right) \\ &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ -\alpha f\left(\frac{d(s,a)}{d^D(s,a)}\right) \right] + \sum_{s,a} d(s,a) \left( \underbrace{R(s,a) + \gamma(\mathcal{T}\nu)(s,a) - (\mathcal{B}\nu)(s,a)}_{=: e_\nu(s,a) \text{ ('advantage' using } \nu)} \right) \end{aligned} \quad (5)$$

$$\begin{aligned} &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ -\alpha f\left(\frac{d(s,a)}{d^D(s,a)}\right) \right] + \mathbb{E}_{(s,a) \sim d^D} \left[ \underbrace{\frac{d(s,a)}{d^D(s,a)}}_{=: w(s,a)} (e_\nu(s,a)) \right] \\ &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} [-\alpha f(w(s,a))] + \mathbb{E}_{(s,a) \sim d^D} [w(s,a) (e_\nu(s,a))] =: L(w, \nu). \end{aligned} \quad (6)$$

The first equality holds due to the property of the adjoint (transpose) operators  $\mathcal{B}_*$  and  $\mathcal{T}_*$ , i.e. for any  $\nu$ ,

$$\sum_s \nu(s) (\mathcal{B}_* d)(s) = \sum_{s,a} d(s,a) (\mathcal{B}\nu)(s,a), \quad \sum_s \nu(s) (\mathcal{T}_* d)(s) = \sum_{s,a} d(s,a) (\mathcal{T}\nu)(s,a),$$

where  $(\mathcal{T}\nu)(s,a) = \sum_{s'} T(s'|s,a)\nu(s')$  and  $(\mathcal{B}\nu)(s,a) = \nu(s)$ . Note that  $L(w, \nu)$  in (6) does not involve expectation over  $d$ , but only expectation over  $p_0$  and  $d^D$ , which allows us to perform optimization only with the offline data.

**Remark.** The terms  $w(s, a)$  and  $\nu(s)$  in (6) are estimated only by using the samples from the dataset distribution  $d^D$ :

$$\hat{L}(w, \nu) := (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s, a, s') \sim d^D}[-\alpha f(w(s, a)) + w(s, a)(\hat{e}_\nu(s, a, s'))]. \quad (7)$$

Here,  $\hat{e}_\nu(s, a, s') := R(s, a) + \gamma\nu(s') - \nu(s)$  is sample-based estimation for advantage  $e_\nu(s, a)$ . On the other hand, prior offline RL algorithms often involve estimations using out-of-distribution actions sampled from the target policy, e.g. employing a critic to compute bootstrapped targets for the value function. Thus, our method is free from the compounding error in the bootstrapped estimation due to using out-of-distribution actions.

In short, OptiDICE solves the problem

$$\max_{w \geq 0} \min_{\nu} L(w, \nu), \quad (8)$$

where the optimal solution  $w^*$  of the optimization (8) represents the stationary distribution corrections between the *optimal* policy's stationary distribution and data distribution:  $w^*(s, a) = \frac{d^{\pi^*}(s, a)}{d^D(s, a)}$ .

### 3.1 A CLOSED-FORM SOLUTION

When the state and/or action spaces are large or continuous, it is a standard practice to use function approximators to represent terms like  $w$  and  $\nu$ , and perform gradient-based optimization of  $L$ . However, this could break nice properties for optimizing  $L$ , such as concavity in  $w$  and convexity in  $\nu$ , which causes numerical instability and poor convergence for the maximin optimization (Goodfellow et al., 2014). We mitigate this issue by obtaining the closed-form solution of the inner optimization, which reduces the overall problem into a unconstrained convex optimization.

Since the optimization problem (2-3) is an instance of convex optimization, one can easily show that the strong duality holds by the Slater's condition (Boyd et al., 2004). Hence we can reorder the optimization from maximin to minimax:

$$\min_{\nu} \max_{w \geq 0} L(w, \nu). \quad (9)$$

Then, for any  $\nu$ , a closed-form solution to the inner maximization of (9) can be derived as follows:

**Proposition 1.** The closed-form solution of the inner maximization of (9),  $w_\nu^* := \arg \max_{w \geq 0} L(w, \nu)$ , is

$$w_\nu^*(s, a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s, a)}{\alpha} \right) \right) \quad \forall s, a, \quad (10)$$

where  $(f')^{-1}$  is the inverse function of the derivative  $f'$  of  $f$  and is strictly increasing by strict convexity of  $f$ . (Proof in Appendix D.)

**Proposition 1** implies that for a fixed  $\nu$ , the optimal stationary distribution correction  $w_\nu^*(s, a)$  is larger for a state-action pair with larger advantage  $e_\nu(s, a)$ . This solution property has a natural interpretation as follows. As  $\alpha \rightarrow 0$ , the term in (5) becomes the Lagrangian of the primal LP for solving the MDP, where  $d(s, a)$  serve as Lagrange multipliers to impose constraints  $R(s, a) + \gamma(\mathcal{T}\nu)(s, a) \leq \nu(s) \forall s, a$ . Also, each  $\nu(s)$  serves as the optimization variable representing the optimal state value function (Puterman, 1994). Thus,  $e_{\nu^*}(s, a) = Q^*(s, a) - V^*(s)$ , i.e. the advantage function of the *optimal policy*, should be zero for the optimal action while it should be lower for sub-optimal actions. For  $\alpha > 0$ , the convex regularizer  $f\left(\frac{d(s, a)}{d^D(s, a)}\right)$  in (5) relaxes those constraints into soft ones, but it still prefers the actions with higher  $e_{\nu^*}(s, a)$  over those with lower  $e_{\nu^*}(s, a)$ . From this perspective,  $\alpha$  adjusts the softness of the constraints  $e_\nu(s, a) \leq 0 \forall s, a$ , and  $f$  determines the relation between advantages and stationary distribution corrections.

Finally, we reduce the nested optimization in (9) to the following single minimization problem by plugging  $w_\nu^*$  into  $L(w, \nu)$ :

$$\begin{aligned} \min_{\nu} L(w_\nu^*, \nu) &= (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s, a) \sim d^D} \left[ -\alpha f \left( \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} e_\nu(s, a) \right) \right) \right) \right] \\ &\quad + \mathbb{E}_{(s, a) \sim d^D} \left[ \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} e_\nu(s, a) \right) \right) (e_\nu(s, a)) \right]. \end{aligned} \quad (11)$$

**Proposition 2.**  $L(w_\nu^*, \nu)$  is convex with respect to  $\nu$ . (Proof in Appendix E.)

The minimization of this convex objective can be performed much more reliably than the nested minimax optimization problem. To solve minimization in (11), we devise the objective that can be easily optimized via sampling from  $D$ :

$$\begin{aligned}\tilde{L}(\nu) := & (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s, a, s') \sim d^D} \left[ -\alpha f \left( \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} \hat{e}_\nu(s, a, s') \right) \right) \right) \right. \\ & \left. + \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} \hat{e}_\nu(s, a, s') \right) \right) (\hat{e}_\nu(s, a, s')) \right].\end{aligned}\quad (12)$$

However, careful readers may notice that  $\tilde{L}(\nu)$  can be a biased estimate of our target objective  $L(w_\nu^*, \nu)$  in (11) due to non-linearity of  $(f')^{-1}$  and suffer from double-sample problem (Baird, 1995). We justify  $\tilde{L}(\nu)$  by proving that  $\tilde{L}(\nu)$  is the upper bound of  $L(w_\nu^*, \nu)$ :

**Corollary 3.**  $\tilde{L}(\nu)$  in (12) is an upper bound of  $L(w_\nu^*, \nu)$  in (11), i.e.  $L(w_\nu^*, \nu) \leq \tilde{L}(\nu)$  always holds, where equality holds when the MDP is deterministic. (Proof in Appendix E.)

### 3.2 STATIONARY DISTRIBUTION CORRECTION ESTIMATION W/ FUNCTION APPROXIMATION

Based on the results from the previous section, we assume that  $\nu$  and  $w$  are parameterized by  $\theta$  and  $\phi$ , respectively, and that both models are sufficiently expressive, e.g. using deep neural networks. Using these models, we optimize  $\theta$  by

$$\min_{\theta} J_\nu(\theta) := \min_{\theta} \tilde{L}(\nu_\theta).\quad (13)$$

After obtaining a solution  $\theta^*$  to the optimization problem, we need a way to evaluate  $w^*(s, a) = \frac{d^{\pi^*}(s, a)}{d^D(s, a)}$  for any  $(s, a)$  to finally obtain the optimal policy  $\pi^*$ . However, the closed-form solution  $w_{\nu_\theta}^*(s, a)$  in **Proposition 1** can be evaluated only on  $(s, a)$  in  $D$  since it requires both  $R(s, a)$  and  $\mathbb{E}_{s' \sim T(s, a)}[\nu_\theta(s')]$  to evaluate the advantage  $e_\nu$ . Therefore, we use a parametric model  $e_\phi$  that approximates the advantage inside the analytic formula presented in (10), so that

$$w_\phi(s, a) := \max \left( 0, (f')^{-1} \left( \frac{e_\phi(s, a)}{\alpha} \right) \right).\quad (14)$$

We consider two options to optimize  $\phi$  once we obtain  $\nu_\theta$  from (13). First,  $\phi$  can be optimized via

$$\min_{\phi} J_w(\phi; \theta) := \min_{\phi} -\hat{L}(w_\phi, \nu_\theta)\quad (15)$$

which corresponds to solving the original minimax problem (9). We also consider

$$\min_{\phi} J_w^{\text{MSE}}(\phi; \theta) := \min_{\phi} \mathbb{E}_{(s, a, s') \sim d^D} [(e_\phi(s, a) - \hat{e}_{\nu_\theta}(s, a, s'))^2],\quad (16)$$

which minimizes the mean squared error (MSE) between the advantage  $e_\phi(s, a)$  and the target induced by  $\nu_\theta$ . We observe that using either  $J_w$  or  $J_w^{\text{MSE}}$  works effectively, which will be detailed in our experiments. In our implementation, we perform joint training of  $\theta$  and  $\phi$ , rather than optimizing  $\phi$  after convergence of  $\theta$ . Additional considerations and details for the complete OptiDICE algorithm (policy extraction and how to deal with  $\gamma = 1$ ) are described in Appendix B.

## 4 EXPERIMENTS

In this section, we evaluate OptiDICE in continuous MDPs (See Appendix F.1 for tabular MDPs). For the  $f$ -divergence, we use  $f(x) = \frac{1}{2}(x - 1)^2$  which corresponds to the  $\chi^2$ -divergence for the tabular-MDP experiment, while we use its softened version for continuous MDPs (See Appendix H).

### 4.1 D4RL BENCHMARK (CONTINUOUS CONTROL TASKS)

We evaluate OptiDICE in continuous MDPs using D4RL offline RL benchmarks (Fu et al., 2021). We use Maze2D (3 tasks) and Gym-MuJoCo (12 tasks) domains from the D4RL dataset (See Appendix I.1 for task description). We interpret terminal states as absorbing states and use the absorbing-state implementation proposed by Kostrikov et al. (2019). For obtaining  $\pi_\beta$  discussed in Section 3.3, we use the tanh-squashed *mixture* of Gaussians policy  $\pi_\beta$  to embrace the multi-modality of data collected from heterogeneous policies. For the target policy  $\pi_\psi$ , we use a tanh-squashed

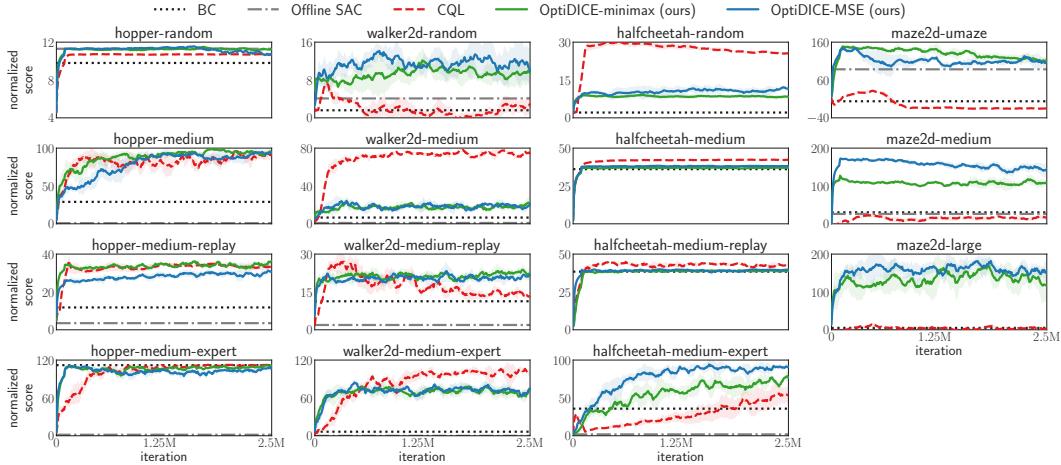


Figure 1: Performance of BC, offline SAC, CQL, OptiDICE-minimax (15) and OptiDICE-MSE (16) on D4RL benchmark (Fu et al., 2021) for  $\gamma = 0.99$ . OptiDICE strictly outperforms CQL on 6 tasks, while performing on par with CQL on 4 tasks. We report mean scores and their 95% confidence intervals obtained from 5 runs for each task.

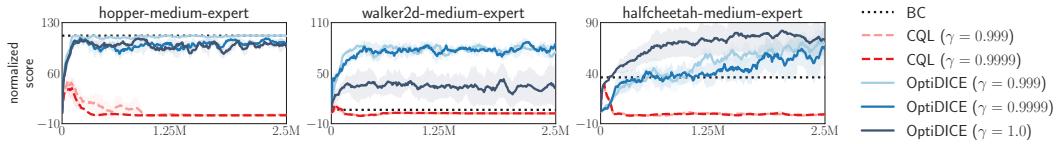


Figure 2: Performance of BC, CQL and OptiDICE (with MSE objective in (16)) in D4RL benchmark (Fu et al., 2021) for  $\gamma = 0.999, 0.9999$  and  $1.0$ .

Gaussian policy, following conservative Q Learning (CQL) (Kumar et al., 2020)—the state-of-the-art model-free offline RL algorithm. We provide detailed information of the experimental setup in Appendix I.1.

The normalized performance of OptiDICE and the best model-free algorithm for each domain is presented in *Table 2*, and learning curves for CQL and OptiDICE are shown in *Figure 1*, where  $\gamma = 0.99$  used for all algorithms. Most notably, OptiDICE achieves state-of-the-art performance for all tasks in the Maze2D domain, by a large margin. In Gym-MuJoCo domain, OptiDICE achieves the best mean performance for 4 tasks (hopper-medium, hopper-medium-expert, walker2d-random, and halfcheetah-medium-expert). Another noteworthy observation is that OptiDICE overwhelmingly outperforms AlgaeDICE (Nachum et al., 2019b) in all domains (*Table 2* and *Table 3* in Appendix for detailed performance of AlgaeDICE), although both AlgaeDICE and OptiDICE stem from the same objective in (1). This is because AlgaeDICE optimizes a nested max-min-max problem, which can suffer from severe overestimation by using out-of-distribution actions and numerical instability. In contrast, OptiDICE solves a simpler minimization problem and does not rely on out-of-distribution actions, exhibiting stable optimization.

As discussed in Section B.2, OptiDICE can naturally be generalized to undiscounted problems ( $\gamma = 1$ ). In *Figure 2*, we vary  $\gamma \in \{0.999, 0.9999, 1.0\}$  to validate OptiDICE’s robustness in  $\gamma$  by comparing with CQL in {hopper-medium-expert, walker2d-medium-expert, halfcheetah-medium-expert} (See Appendix J for other tasks). The performance of OptiDICE stays stable, while CQL easily becomes unstable as  $\gamma$  increases, due to the divergence of Q-function. This is because OptiDICE relies on *normalized* stationary distribution corrections, whereas CQL learns the action-value function whose values becomes unbounded as  $\gamma$  gets close to 1, resulting in numerical instability.

## 5 CONCLUSION

We presented OptiDICE, an offline RL algorithm that aims to estimate optimal stationary distribution corrections. We formulated the estimation problem as a minimax optimization that does not involve sampling from the target policy, which essentially circumvents the overestimation issue incurred by bootstrapped target with out-of-distribution actions. In the experiments, we demonstrated that OptiDICE performs competitively with the state-of-the-art offline RL baselines.

## ACKNOWLEDGMENTS

This work was supported by the National Research Foundation (NRF) of Korea (NRF-2019M3F2A1072238 and NRF-2019R1A2C1087634), and the Ministry of Science and Information communication Technology (MSIT) of Korea (IITP No. 2019-0-00075, IITP No. 2020-0-00940 and IITP No. 2017-0-01779 XAI).

## REFERENCES

- Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Machine Learning (ICML)*, 1995.
- Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Serkan Cabi, Sergio Gómez Colmenarejo, Alexander Novikov, Ksenia Konyushkova, Scott Reed, Rae Jeong, Konrad Zolna, Yusuf Aytar, David Budden, Mel Vecerik, et al. Scaling data-driven robotics with reward sketching and batch reinforcement learning. In *Robotics: Science and Systems (RSS) Proceedings*, 2020.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research (JMLR)*, 2005.
- Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: Datasets for deep data-driven reinforcement learning, 2021. URL [https://openreview.net/forum?id=px0-N3\\_KjA](https://openreview.net/forum?id=px0-N3_KjA).
- Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 2672–2680, 2014.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- Garud N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 2005.

- Natasha Jaques, Asma Ghandeharioun, Judy Hanwen Shen, Craig Ferguson, Agata Lapedriza, Noah Jones, Shixiang Gu, and Rosalind Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog, 2019.
- Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. MOrEL : Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. Discriminator-Actor-Critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.
- Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy Q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative Q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. *Reinforcement learning: State-of-the-art*. Springer Berlin Heidelberg, 2012.
- Romain Laroche, Paul Trichelair, and Remi Tachet Des Combes. Safe policy improvement with baseline bootstrapping. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- Byung-Jun Lee, Jongmin Lee, Peter Vrancx, Dongho Kim, and Kee-Eung Kim. Batch reinforcement learning with hyperparameter gradients. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems, 2020.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *Proceedings of the 4th International Conference on Learning Representations (ICLR)*, 2016.
- Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. Reset-free lifelong learning with skill-space planning. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2021.
- Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. DualDICE: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019a.
- Ofir Nachum, Bo Dai, Ilya Kostrikov, Yinlam Chow, Lihong Li, and Dale Schuurmans. AlgaeDICE: Policy gradient from arbitrary experience. *arXiv preprint arXiv:1912.02074*, 2019b.
- Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 2005.
- Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning, 2019.
- Karl Pertsch, Youngwoon Lee, and Joseph J. Lim. Accelerating reinforcement learning with learned skill priors. In *Conference on Robot Learning (CoRL)*, 2020.
- Marek Petrik, Mohammad Ghavamzadeh, and Yinlam Chow. Safe policy improvement by minimizing robust baseline regret. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

- Martin L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 1st edition, 1994.
- John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft Q-learning, 2017.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: An introduction*. MIT Press, 1998.
- Richard S Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 1999.
- István Szita and András Lörincz. The many faces of optimism: A unifying approach. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, 2008.
- Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning, 2019.
- Chao Yu, Jiming Liu, and Shamim Nemati. Reinforcement learning in healthcare: A survey, 2020a.
- F. Yu, H. Chen, X. Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, V. Madhavan, and Trevor Darrell. BDD100K: A diverse driving dataset for heterogeneous multitask learning. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020b.
- Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. MOPO: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020c.
- Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. GenDICE: Generalized offline estimation of stationary values. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020.

## A ILLUSTRATIVE EXAMPLE

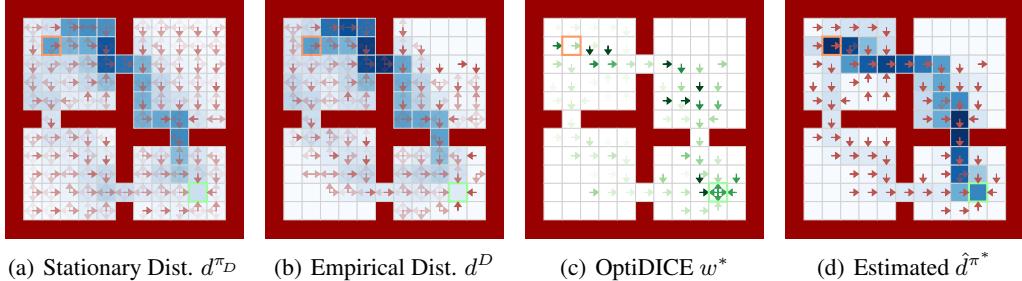


Figure 3: Illustrative example of how OptiDICE estimates the optimal policy’s stationary distribution in the Four Rooms domain (Sutton et al., 1999; Nachum et al., 2019b). The initial state and the goal state are denoted by orange and green squares, respectively. Based on a sub-optimal data-collection policy  $\pi_D$ , which induces  $d^{\pi_D}$  shown in (a), a static dataset is sampled and its empirical distribution  $d^D$  ( $\neq d^{\pi_D}$ ) shown in (b). The opacity of each square is determined by the state marginals of each stationary distribution, where the opacity of the arrow shows the policy induced by each stationary distribution. By multiplying the OptiDICE  $w^*$  obtained by solving (8) (shown in (c)), a near-optimal policy  $\pi^*$  is obtained from  $\hat{d}^{\pi^*}(s, a) = d^D(s, a)w^*(s, a)$  shown in (d).

*Figure 3* outlines how our approach works in the Four Rooms domain (Sutton et al., 1999) where the agent aims to navigate to a goal location in a maze composed of four rooms. We collected static dataset  $D$  consisting of 50 episodes with maximum time step 50 using the data-collection policy  $\pi_D = 0.5\pi_{\text{true}}^* + 0.5\pi_{\text{rand}}$ , where  $\pi_{\text{true}}^*$  is the optimal policy of the underlying true MDP and  $\pi_{\text{rand}}$  is the random policy sampled from the Dirichlet distribution, i.e.  $\pi_{\text{rand}}(s) \sim \text{Dir}(1, 1, 1, 1) \forall s$ .

In this example, we explicitly construct Maximum Likelihood Estimate (MLE) MDP  $\hat{M}$  based on the static dataset  $D$ . OptiDICE performs the minimization (11) to obtain  $\nu^*$  while using  $\hat{M}$  to compute  $e_\nu(s, a)$ . Then,  $w_{\nu^*}^*$  is estimated directly via Eq. (10) (*Figure 3(c)*). Finally,  $w^*$  is multiplied by  $d^D$  to correct  $d^D$  towards an optimal policy, resulting in  $\hat{d}^{\pi^*}$ , which is the stationary distribution of the estimated optimal policy (*Figure 3(d)*). For tabular MDPs, the global optima  $(\nu^*, w^*)$  can always be obtained. We describe OptiDICE for finite MDPs in Appendix F.

## B ALGORITHMIC CONSIDERATIONS

### B.1 POLICY EXTRACTION

As the last step, we need to extract the optimal policy  $\pi^*$  from the optimal stationary distribution corrections  $w_\phi(s, a) = \frac{d^{\pi^*}(s, a)}{d^D(s, a)}$ . While the optimal policy can be easily obtained by  $\pi^*(a|s) = \frac{d^D(s, a)w_\phi(s, a)}{\sum_{\bar{a}} d^D(s, \bar{a})w_\phi(s, \bar{a})}$  for tabular domains, this procedure is not straightforwardly applicable to continuous domains.

One of the ways to address continuous domains is to use importance-weighted behavioral cloning: we optimize the parameterized policy  $\pi_\psi$  by maximizing the log-likelihood on  $(s, a)$  that would be sampled from the optimal policy  $\pi^*$ :

$$\max_{\psi} \mathbb{E}_{(s, a) \sim d^{\pi^*}} [\log \pi_\psi(a|s)] = \max_{\psi} \mathbb{E}_{(s, a) \sim d^D} [w_\phi(s, a) \log \pi_\psi(a|s)].$$

Despite its simplicity, this approach does not work well in practice, since  $\pi_\psi$  will be trained only on samples from the intersection of the supports of  $d^{\pi^*}$  and  $d^D$ , which becomes very scarce when  $\pi^*$  deviates significantly from the data collection policy  $\pi_D$ .

We thus use the information projection (I-projection) for training the policy:

$$\min_{\psi} \mathbb{KL}(d^D(s)\pi_\psi(a|s)||d^D(s)\pi^*(a|s)), \quad (17)$$

where we replace  $d^{\pi^*}(s)$  by  $d^D(s)$  for  $d^{\pi^*}(s, a)$ . This results in minimizing the discrepancy between  $\pi_\psi(a|s)$  and  $\pi^*(a|s)$  on the stationary distribution over states from  $\pi_D$ . This approach is motivated by the desideratum that the policy  $\pi_\psi$  should be trained at least on the states observed in  $D$  to be robust upon deployment. Now, rearranging the terms in (17), we obtain

$$\begin{aligned} \mathbb{KL}(d^D(s)\pi_\psi(a|s)||d^D(s)\pi^*(a|s)) &= -\mathbb{E}_{\substack{s \sim d^D \\ a \sim \pi_\psi(s)}} \left[ \underbrace{\log \frac{d^*(s, a)}{d^D(s, a)}}_{= w_\phi(s, a)} - \underbrace{\log \frac{\pi_\psi(a|s)}{\pi_D(a|s)}}_{\text{constant for } \pi} - \underbrace{\log \frac{d^*(s)}{d^D(s)}}_{\text{constant for } \pi} \right] \\ &= -\mathbb{E}_{\substack{s \sim d^D \\ a \sim \pi_\psi(s)}} [\log w_\phi(s, a) - D_{\text{KL}}(\pi_\psi(a|s) || \pi_D(a|s))] + C \\ &=: J_\pi(\psi; \phi, \pi_D) \end{aligned} \quad (18)$$

We can interpret this I-projection objective as a KL-regularized actor-critic architecture (Fox et al., 2016; Schulman et al., 2017), where  $\log w_\phi(s, a)$  taking the role of the critic and  $\pi_\psi$  being the actor. Note that I-projection requires us to evaluate  $\pi_D$  for the KL regularization term. We therefore employ another parameterized policy  $\pi_\beta$  to approximate  $\pi_D$ , trained via simple behavioral cloning (BC).

### B.2 GENERALIZATION TO $\gamma = 1$

For  $\gamma = 1$ , our original problem (2-3) for the stationary distribution  $d$  is an ill-posed problem: for any  $d$  that satisfies the Bellman flow constraints (3-3) and a constant  $c \geq 0$ ,  $cd$  also satisfies the Bellman flow constraints (3-3) (Zhang et al., 2020). We address this issue by adding additional normalization constraint  $\sum_{s, a} d(s, a) = 1$  to (2-3). By using analogous derivation from (2) to (9) with the normalization constraint—introducing a Lagrange multiplier  $\lambda \in \mathbb{R}$  and changing the variable  $d$  to  $w$ —we derive the following minimax objective for  $w, \nu$  and  $\lambda$ :

$$\begin{aligned} \min_{\nu, \lambda} \max_{w \geq 0} L(w, \nu, \lambda) &:= L(w, \nu) + \lambda(1 - \mathbb{E}_{(s, a) \sim d^D} [w(s, a)]) \\ &= (1 - \gamma)\mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s, a) \sim d^D} [-\alpha f(w(s, a))] - \mathbb{E}_{(s, a) \sim d^D} [w(s, a)(e_\nu(s, a) - \lambda)] + \lambda. \end{aligned} \quad (19)$$

Similar to (7), we define  $\hat{L}(w, \nu, \lambda)$ , an unbiased estimator for  $L(w, \nu, \lambda)$  such that

$$\hat{L}(w, \nu, \lambda) := (1 - \gamma)\mathbb{E}_{s \sim p_0} [\nu(s)] + \lambda + \mathbb{E}_{(s, a, s') \sim d^D} [-\alpha f(w(s, a)) + w(s, a)(\hat{e}_{\nu, \lambda}(s, a, s'))], \quad (20)$$

where  $\hat{e}_{\nu, \lambda}(s, a, s') := \hat{e}_\nu(s, a, s') - \lambda$ . We then derive a closed-form solution for the inner maximization in (19):

**Proposition 4.** *The maximizer  $w_{\nu,\lambda}^* : S \times A \rightarrow \mathbb{R}$  of the inner optimization of (19), which is defined by  $w_{\nu,\lambda}^* := \arg \max_{w \geq 0} L(w, \nu, \lambda)$ , is*

$$w_{\nu,\lambda}^*(s, a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s, a) - \lambda}{\alpha} \right) \right).$$

(Proof in Appendix G.)

Similar to (12), we minimize the biased estimate  $\tilde{L}(\nu, \lambda)$ , which is an upper bound of  $L(w_{\nu,\lambda}^*, \nu, \lambda)$ , by applying the closed-form solution from **Proposition 4**:

$$\begin{aligned} \tilde{L}(\nu, \lambda) := & (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s, a, s') \sim d^D} \left[ -\alpha f \left( \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} \hat{e}_{\nu,\lambda}(s, a, s') \right) \right) \right) \right. \\ & \left. + \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} \hat{e}_{\nu,\lambda}(s, a, s') \right) \right) (\hat{e}_{\nu,\lambda}(s, a, s')) \right] + \lambda. \end{aligned} \quad (21)$$

By using the above estimators, we correspondingly update our previous objectives for  $\theta$  and  $\phi$  as follows. First, the objective for  $\theta$  is modified to

$$\min_{\theta} J_{\nu}(\theta, \lambda) := \min_{\theta} \tilde{L}(\nu_{\theta}, \lambda). \quad (22)$$

For  $\phi$ , we modify our approximator in (14) by including the Lagrangian  $\lambda' \in \mathbb{R}$ :

$$w_{\phi, \lambda'}(s, a) := \max \left( 0, (f')^{-1} \left( \frac{e_{\phi}(s, a) - \lambda'}{\alpha} \right) \right).$$

Note that  $\lambda' \neq \lambda$  is used to stabilize the learning process. For optimizing over  $\phi$ , the minimax objective (15) is modified as

$$\min_{\phi} J_w(\phi, \lambda'; \theta) := \min_{\phi} -\hat{L}(w_{\phi, \lambda'}, \nu_{\theta}, \lambda'), \quad (23)$$

while the same objective  $J_w^{\text{MSE}}(\phi; \theta)$  in (16) is used for the MSE objective. We additionally introduce learning objectives for  $\lambda$  and  $\lambda'$ , which is required for the normalization constraint discussed in this subsection:

$$\min_{\lambda} J_{\nu}(\theta, \lambda) \quad \text{and} \quad \min_{\lambda'} J_w(\phi, \lambda'; \theta). \quad (24)$$

Finally, by using the above objectives in addition to BC objective and policy extraction objective in (18), we describe our algorithm, OptiDICE, in **Algorithm 1**, where we train neural network parameters via stochastic gradient descent. In our algorithm, we use a warm-up iteration—optimizing all networks except  $\pi_{\psi}$ —to prevent  $\pi_{\psi}$  from converging to sub-optimal policies during its initial training. Even when  $\gamma < 1$ , we empirically observe that using the normalization constraint stabilizes OptiDICE’s learning process, and we thus use the normalization constraint in all experiments (Zhang et al., 2020).

## C ALGORITHM

---

**Algorithm 1** OptiDICE
 

---

**Input:** A dataset  $D := \{(s_i, a_i, r_i, s'_i)\}_{i=1}^N$ , a set of initial states  $D_0 := \{s_{0,i}\}_{i=1}^{N_0}$ , neural networks  $\nu_\theta$  and  $e_\phi$  with parameters  $\theta$  and  $\phi$ , learnable parameters  $\lambda$  and  $\lambda'$ , policy networks  $\pi_\beta$  and  $\pi_\psi$  with parameter  $\beta$  and  $\psi$ , a learning rate  $\eta$

- 1: **for** each iteration **do**
- 2:   Sample mini-batches from  $D$  and  $D_0$ , respectively.
- 3:   Compute  $\theta$ -gradient to optimize (22):  

$$g_\theta \approx \nabla_\theta J_\nu(\theta, \lambda) \quad (minimax\ obj.\ (23))$$
- 4:   Compute  $\phi$ -gradient for either one of objectives:  

$$g_\phi \approx \nabla_\phi J_w(\phi, \lambda'; \theta) \quad (MSE\ obj.\ (16))$$
- 5:   Compute  $\lambda$  and  $\lambda'$  gradients to optimize (24):  

$$g_\lambda \approx \nabla_\lambda J_\nu(\theta, \lambda), g_{\lambda'} \approx \nabla_{\lambda'} J_w(\phi, \lambda'; \theta)$$
- 6:   Compute  $\beta$ -gradient  $g_\beta$  for BC.
- 7:   Compute  $\psi$ -gradient via (18) (*policy extraction*):  

$$g_\psi \approx \nabla_\psi J_\pi(\psi; \phi, \pi_\beta)$$
- 8:   Perform SGD updates:  

$$\begin{aligned} \theta &\leftarrow \theta - \eta g_\theta, & \lambda &\leftarrow \lambda - \eta g_\lambda, & \beta &\leftarrow \beta - \eta g_\beta, \\ \phi &\leftarrow \phi - \eta g_\phi, & \lambda' &\leftarrow \lambda' - \eta g_{\lambda'}, & \psi &\leftarrow \psi - \eta g_\psi. \end{aligned}$$
- 9: **end for**

**Output:**  $\nu_\theta \approx \nu^*, w_{\phi, \lambda'} \approx w^*, \pi_\psi \approx \pi^*$ ,

---

## D PROOF OF PROPOSITION 1

We first show that our original problem (2-3) is an instance of convex programming due to the convexity of  $f$ .

**Lemma 5.** *The constraint optimization (2-3) is a convex optimization.*

*Proof.*

$$\max_{d \geq 0} \mathbb{E}_{(s,a) \sim d}[R(s,a)] - \alpha D_f(d||d^D) \quad (2)$$

$$\text{s.t. } (\mathcal{B}_*d)(s) = (1-\gamma)p_0(s) + \gamma(\mathcal{T}_*d)(s) \quad \forall s, \quad (3)$$

The objective function  $\mathbb{E}_{(s,a) \sim d}[R(s,a)] - \alpha D_f(d||d^D)$  is concave for  $d : S \times A \rightarrow \mathbb{R}$  (not only for probability distribution  $d \in \Delta(S \times A)$ ) since  $D_f(d||d^D)$  is convex in  $d$ : for  $t \in [0, 1]$  and any  $d_1 : S \times A \rightarrow \mathbb{R}, d_2 : S \times A \rightarrow \mathbb{R}$ ,

$$\begin{aligned} D_f((1-t)d_1 + td_2 || d^D) &= \sum_{s,a} d^D(s,a) f\left((1-t)\frac{d_1(s,a)}{d^D(s,a)} + t\frac{d_2(s,a)}{d^D(s,a)}\right) \\ &< \sum_{s,a} d^D(s,a) \left\{ (1-t)f\left(\frac{d_1(s,a)}{d^D(s,a)}\right) + t f\left(\frac{d_2(s,a)}{d^D(s,a)}\right) \right\} \\ &= (1-t)D_f(d_1 || d^D) + t D_f(d_2 || d^D), \end{aligned}$$

where the strict inequality follows from assuming  $f$  is strictly convex. In addition, the equality constraints (3) are affine in  $d$ , and the inequality constraints  $d \geq 0$  are linear and thus convex in  $d$ . Therefore, our problem is an instance of a convex programming, as we mentioned in Section 3.1.  $\square$

In addition, by using the strong duality and the change-of-variable from  $d$  to  $w$ , we can rearrange the original maximin optimization to the minimax optimization.

**Lemma 6.** *We assume that all states  $s \in S$  are reachable for a given MDP. Then,*

$$\max_{w \geq 0} \min_\nu L(w, \nu) = \min_\nu \max_{w \geq 0} L(w, \nu).$$

*Proof.* Let us define the Lagrangian of the constraint optimization (2-3)

$$\begin{aligned}\mathcal{L}(d, \nu, \mu) := & \mathbb{E}_{(s,a) \sim d}[R(s, a)] - \alpha D_f(d || d^D) \\ & + \sum_s \nu(s) \left( (1 - \gamma)p_0(s) + \gamma \sum_{\bar{s}, \bar{a}} T(s|\bar{s}, \bar{a})d(\bar{s}, \bar{a}) - \sum_{\bar{a}} d(s, \bar{a}) \right) \\ & + \sum_{s,a} \mu(s, a)d(s, a)\end{aligned}$$

with Lagrange multipliers  $\nu(s) \forall s$  and  $\mu(s, a) \forall s, a$ . With the Lagrangian  $\mathcal{L}(d, \nu, \mu)$ , the original problem (2-3) can be represented by

$$\max_{d \geq 0} \min_{\nu} \mathcal{L}(d, \nu, 0) = \max_d \min_{\nu, \mu \geq 0} \mathcal{L}(d, \nu, \mu).$$

For an MDP where every  $s \in S$  is reachable, there always exists  $d$  such that  $d(s, a) > 0 \forall s, a$ . From Slater's condition for convex problems (the condition that there exists a strictly feasible  $d$  (Boyd et al., 2004)), the strong duality holds, i.e., we can change the order of optimizations:

$$\max_d \min_{\nu, \mu \geq 0} \mathcal{L}(d, \nu, \mu) = \min_{\nu, \mu \geq 0} \max_d \mathcal{L}(d, \nu, \mu) = \min_{\nu} \max_{d \geq 0} \mathcal{L}(d, \nu, 0).$$

Here, the last equality holds since  $\max_{d \geq 0} \mathcal{L}(d, \nu, 0) = \max_d \min_{\mu \geq 0} \mathcal{L}(d, \nu, \mu) = \min_{\mu \geq 0} \max_d \mathcal{L}(d, \nu, \mu)$  for fixed  $\nu$  due to the strong duality. Finally, by applying the change of variable  $w = d/d^D$ , we have

$$\max_{w \geq 0} \min_{\nu} L(w, \nu) = \min_{\nu} \max_{w \geq 0} L(w, \nu).$$

□

Finally, the solution of the inner maximization  $\max_{w \geq 0} L(w, \nu)$  can be derived as follows:

**Proposition 1.** *The closed-form solution of the inner maximization of (9), i.e.*

$$w_{\nu}^* := \arg \max_{w \geq 0} (1 - \gamma) \mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} [-\alpha f'(w(s, a))] + \mathbb{E}_{(s,a) \sim d^D} [w(s, a)(e_{\nu}(s, a))]$$

is given as

$$w_{\nu}^*(s, a) = \max \left( 0, (f')^{-1} \left( \frac{e_{\nu}(s, a)}{\alpha} \right) \right) \quad \forall s, a, \tag{25}$$

where  $(f')^{-1}$  is the inverse function of the derivative  $f'$  of  $f$  and is strictly increasing by strict convexity of  $f$ .

*Proof.* For a fixed  $\nu$ , let the maximization  $\max_{w \geq 0} L(w, \nu)$  be the primal problem. Then, its corresponding dual problem is

$$\max_w \min_{\mu \geq 0} L(w, \nu) + \sum_{s,a} \mu(s, a)w(s, a).$$

Since the strong duality holds, satisfying KKT condition is both necessary and sufficient conditions for the solutions  $w^*$  and  $\mu^*$  of primal and dual problems (we will use  $w^*$  and  $\mu^*$  instead of  $w_{\nu}^*$  and  $\mu_{\nu}^*$  for notational brevity).

**Condition 1 (Primal feasibility).**  $w^* \geq 0 \forall s, a$ .

**Condition 2 (Dual feasibility).**  $\mu^* \geq 0 \forall s, a$ .

**Condition 3 (Stationarity).**  $d^D(s, a)(-\alpha f'(w^*(s, a)) + e_{\nu}(s, a) + \mu^*(s, a)) = 0 \forall s, a$ .

**Condition 4 (Complementary slackness).**  $w^*(s, a)\mu^*(s, a) = 0 \forall s, a$ .

From **Stationarity** and  $d^D > 0$ , we have

$$f'(w^*(s, a)) = \frac{e_{\nu}(s, a) + \mu^*(s, a)}{\alpha} \quad \forall s, a$$

and since  $f'$  is invertible due to the strict convexity of  $f$ ,

$$w^*(s, a) = (f')^{-1} \left( \frac{e_\nu(s, a) + \mu^*(s, a)}{\alpha} \right) \quad \forall s, a.$$

Now for fixed  $(s, a) \in S \times A$ , let us consider two cases: either  $w^*(s, a) > 0$  or  $w^*(s, a) = 0$ , where **Primal feasibility** is always satisfied in either way:

**Case 1** ( $w^*(s, a) > 0$ ).  $\mu^*(s, a) = 0$  due to **Complementary slackness**, and thus,

$$w^*(s, a) = (f')^{-1} \left( \frac{e_\nu(s, a)}{\alpha} \right) > 0.$$

Note that **Dual feasibility** holds. Since  $f'$  is a strictly increasing function,  $e_\nu(s, a) > \alpha f'(0)$  should be satisfied if  $f'(0)$  is well-defined.

**Case 2** ( $w^*(s, a) = 0$ ).  $\mu^*(s, a) = \alpha f'(0) - e_\nu(s, a) \geq 0$  due to **Stationarity** and **Dual feasibility**, and thus,  $e_\nu(s, a) \leq \alpha f'(0)$  should be satisfied if  $f'(0)$  is well-defined.

In summary, we have

$$w_\nu^*(s, a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s, a)}{\alpha} \right) \right).$$

□

## E PROOFS OF PROPOSITION 2 AND COROLLARY 3

**Proposition 2.**  $L(w_\nu^*, \nu)$  is convex with respect to  $\nu$ .

*Proof by Lagrangian duality.* Let us consider Lagrange dual function

$$g(\nu, \mu) := \max_d \mathcal{L}(d, \nu, \mu),$$

which is always convex in Lagrange multipliers  $\nu, \mu$  since  $\mathcal{L}(d, \nu, \mu)$  is affine in  $\nu, \mu$ . Also, for any  $\mu_1, \mu_2 \geq 0$  and its convex combination  $(1-t)\mu_1 + t\mu_2$  for  $0 \leq t \leq 1$ , we have

$$\min_{\mu \geq 0} g((1-t)\nu_1 + t\nu_2, \mu) \leq g((1-t)\nu_1 + t\nu_2, (1-t)\mu_1 + t\mu_2) \leq (1-t)g(\nu_1, \mu_1) + tg(\nu_2, \mu_2)$$

by using the convexity of  $g(\nu, \mu)$ . Since the above statement holds for any  $\mu_1, \mu_2 \geq 0$ , we have

$$\min_{\mu \geq 0} g((1-t)\nu_1 + t\nu_2, \mu) \leq (1-t) \min_{\mu_1 \geq 0} g(\nu_1, \mu_1) + t \min_{\mu_2 \geq 0} g(\nu_2, \mu_2).$$

Therefore, a function

$$G(\nu) := \min_{\mu \geq 0} g(\nu, \mu) = \min_{\mu \geq 0} \max_d \mathcal{L}(d, \nu, \mu) = \max_{d \geq 0} \mathcal{L}(d, \nu, 0)$$

is convex in  $\nu$ . By following the change-of-variable, we have

$$\max_{d \geq 0} \mathcal{L}(d, \nu, 0) = \max_{w \geq 0} L(w, \nu) = L(\arg \max_{w \geq 0} L(w, \nu), \nu) = L(w_\nu^*, \nu)$$

is convex in  $\nu$ .

□

*Proof by exploiting second-order derivative.* Suppose  $((f')^{-1})'$  is well-defined, where  $f$  we consider in this work satisfies the condition. Let us define

$$h(x) := -f \left( \max \left( 0, (f')^{-1}(x) \right) \right) + \max \left( 0, (f')^{-1}(x) \right) \cdot x. \quad (26)$$

Then,  $L(w_\nu^*, \nu)$  can be represented by using  $h$ :

$$\begin{aligned} L(w_\nu^*, \nu) &= (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ -\alpha f \left( \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} e_\nu(s, a) \right) \right) \right) \right. \\ &\quad \left. + \max \left( 0, (f')^{-1} \left( \frac{1}{\alpha} e_\nu(s, a) \right) \right) e_\nu(s, a) \right] \end{aligned} \quad (27)$$

$$= (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ \alpha h \left( \frac{1}{\alpha} e_\nu(s, a) \right) \right] \quad (28)$$

We prove that  $h(x)$  is convex in  $x$  by showing  $h''(x) \geq 0 \forall x$ . Recall that  $f'$  is a strictly increasing function by the strict convexity of  $f$ , which implies that  $(f')^{-1}$  is also a strictly increasing function.

**Case 1.** If  $(f')^{-1}(x) > 0 \forall x$ ,

$$\begin{aligned} h(x) &= -f((f')^{-1}(x)) + (f')^{-1}(x) \cdot x, \\ h'(x) &= -\underbrace{f'((f')^{-1}(x))}_{\text{(identity function)}} ((f')^{-1})'(x) + ((f')^{-1})'(x) \cdot x + (f')^{-1}(x) \\ &= -x \cdot ((f')^{-1})'(x) + ((f')^{-1})'(x) \cdot x + (f')^{-1}(x) \\ &= (f')^{-1}(x), \\ h''(x) &= ((f')^{-1})'(x) > 0, \end{aligned}$$

where  $((f')^{-1})'(x) > 0$  since it is the derivative of the strictly increasing function  $(f')^{-1}$ .

**Case 2.** If  $(f')^{-1}(x) \leq 0 \forall x$ ,

$$h(x) = -f(0) \Rightarrow h'(x) = 0 \Rightarrow h''(x) = 0.$$

Therefore,  $h''(x) \geq 0$  holds for all  $x$ , which implies that  $h(x)$  is convex in  $x$ . Finally, for  $t \in [0, 1]$  and any  $\nu_1 : S \rightarrow \mathbb{R}$ ,  $\nu_2 : S \rightarrow \mathbb{R}$ ,

$$\begin{aligned} L(w_{t\nu_1+(1-t)\nu_2}^*, t\nu_1 + (1-t)\nu_2) &= (1 - \gamma)\mathbb{E}_{s \sim p_0}[t\nu_1(s) + (1-t)\nu_2(s)] \\ &\quad + \mathbb{E}_{(s,a) \sim d^D} \left[ \alpha h \left( \frac{1}{\alpha} \left( t\{R(s, a) + \gamma(\mathcal{T}\nu_1)(s, a) - (\mathcal{B}\nu_1)(s, a)\} \right. \right. \right. \\ &\quad \left. \left. \left. + (1-t)\{R(s, a) + \gamma(\mathcal{T}\nu_2)(s, a) - (\mathcal{B}\nu_2)(s, a)\} \right) \right) \right] \\ &\leq t \left\{ (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu_1(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ \alpha h \left( \frac{1}{\alpha} (R(s, a) + \gamma(\mathcal{T}\nu_1)(s, a) - (\mathcal{B}\nu_1)(s, a)) \right) \right] \right\} \\ &\quad + (1-t) \left\{ (1 - \gamma)\mathbb{E}_{s \sim p_0}[\nu_2(s)] + \mathbb{E}_{(s,a) \sim d^D} \left[ \alpha h \left( \frac{1}{\alpha} (R(s, a) + \gamma(\mathcal{T}\nu_2)(s, a) - (\mathcal{B}\nu_2)(s, a)) \right) \right] \right\} \\ &\quad \quad \quad \text{(by convexity of } h\text{)} \\ &= tL(w_{\nu_1}^*, \nu_1) + (1-t)L(w_{\nu_2}^*, \nu_2) \end{aligned}$$

which concludes the proof.  $\square$

**Corollary 3.**  $\tilde{L}(\nu)$  in (12) is an upper bound of  $L(w_\nu^*, \nu)$  in (11), i.e.  $L(w_\nu^*, \nu) \leq \tilde{L}(\nu)$  always holds, where equality holds when the MDP is deterministic.

*Proof by Lagrangian duality.* Let us consider a function  $h$  in (26). From **Proposition 2**, we have  $\mathbb{E}_{(s,a) \sim d^D}[h(\frac{1}{\alpha}e_\nu(s, a))]$  is convex in  $\nu$ , i.e., for  $t \in [0, 1]$ ,  $\nu_1 : S \rightarrow \mathbb{R}$  and  $\nu_2 : S \rightarrow \mathbb{R}$ ,

$$\begin{aligned} &\mathbb{E}_{(s,a) \sim d^D} [h(\frac{1}{\alpha}e_{(1-t)\nu_1+t\nu_2}(s, a))] \\ &= \mathbb{E}_{(s,a) \sim d^D} [h((1-t) \cdot \frac{1}{\alpha}e_{\nu_1}(s, a) + t \cdot \frac{1}{\alpha}e_{\nu_2}(s, a))] \\ &\leq (1-t)\mathbb{E}_{(s,a) \sim d^D} [h(\frac{1}{\alpha}e_{\nu_1}(s, a))] + t\mathbb{E}_{(s,a) \sim d^D} [h(\frac{1}{\alpha}e_{\nu_2}(s, a))] \\ &= \mathbb{E}_{(s,a) \sim d^D} [(1-t) \cdot h(\frac{1}{\alpha}e_{\nu_1}(s, a)) + t \cdot h(\frac{1}{\alpha}e_{\nu_2}(s, a))]. \end{aligned}$$

Since **Proposition 2** should be satisfied for any MDP and  $d^D > 0$ , we have

$$h((1-t) \cdot \frac{1}{\alpha} e_{\nu_1}(s, a) + t \cdot \frac{1}{\alpha} e_{\nu_2}(s, a)) \leq (1-t) \cdot h(\frac{1}{\alpha} e_{\nu_1}(s, a)) + t \cdot h(\frac{1}{\alpha} e_{\nu_2}(s, a)) \quad \forall s, a.$$

To prove this, if

$$h((1-t) \cdot \frac{1}{\alpha} e_{\nu_1}(s, a) + t \cdot \frac{1}{\alpha} e_{\nu_2}(s, a)) > (1-t) \cdot h(\frac{1}{\alpha} e_{\nu_1}(s, a)) + t \cdot h(\frac{1}{\alpha} e_{\nu_2}(s, a)) \quad \exists s, a,$$

we can always find out  $d^D > 0$  that contradicts **Proposition 2**. Also, since  $\frac{1}{\alpha} e_{\nu}(s, a)$  can have an arbitrary real value,  $h$  should be a convex function. Therefore, it can be shown that

$$h(\mathbb{E}_{s' \sim T(s, a)} [\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s')]) \leq \mathbb{E}_{s' \sim T(s, a)} [h(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s'))] \quad \forall s, a,$$

due to Jensen's inequality, and thus,

$$\begin{aligned} \mathbb{E}_{(s, a) \sim d^D} [h(\frac{1}{\alpha} e_{\nu}(s, a))] &= \mathbb{E}_{(s, a) \sim d^D} [h(\mathbb{E}_{s' \sim T(s, a)} [\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s')])] \\ &\leq \mathbb{E}_{(s, a, s') \sim d^D} [h(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s'))]. \end{aligned}$$

Also, the inequality becomes tight when the transition model is deterministic since  $h(\frac{1}{\alpha} \mathbb{E}_{s' \sim T(s, a)} [\hat{e}_{\nu}(s, a, s')]) = \mathbb{E}_{s' \sim T(s, a)} [h(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s'))]$  should always hold for the deterministic transition  $T$ .  $\square$

*Proof by exploiting second-order derivative.* We start from (28) in the proof of **Proposition 2**.

$$\begin{aligned} L(w_{\nu}^*, \nu) &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s, a) \sim d^D} \left[ \alpha h\left(\frac{1}{\alpha} e_{\nu}(s, a)\right) \right] \\ &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s, a) \sim d^D} \left[ \alpha h\left(\frac{1}{\alpha} \mathbb{E}_{s' \sim T(s, a)} [\hat{e}_{\nu}(s, a, s')]\right) \right] \\ &\leq (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{\substack{(s, a) \sim d^D \\ s' \sim T(s, a)}} \left[ \alpha h\left(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s')\right) \right] \end{aligned} \tag{28}$$

(by Jensen's inequality with the convexity of  $h$ )

$$\begin{aligned} &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s, a, s') \sim d^D} \left[ -\alpha f\left(\max\left(0, (f')^{-1}\left(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s')\right)\right)\right) \right. \\ &\quad \left. + \max\left(0, (f')^{-1}\left(\frac{1}{\alpha} \hat{e}_{\nu}(s, a, s')\right)\right)(\hat{e}_{\nu}(s, a, s')) \right] = \tilde{L}(\nu) \end{aligned} \tag{12}$$

(by definition of  $h$ )

(29)

Also, Jensen's inequality becomes tight when the transition model is deterministic for the same reason we describe in *Proof by Lagrangian duality*.  $\square$

## F OPTIDICE FOR FINITE MDPs

For tabular MDP experiments, we assume that the data-collection policy is given to OptiDICE for a fair comparison with SPIBB (Laroche et al., 2019) and BOPAH (Lee et al., 2020), which directly exploit the data-collection policy  $\pi_D$ . However, the extension of tabular OptiDICE to not assuming  $\pi_D$  is straightforward.

As a first step, we construct an MLE MDP  $\hat{M} = \langle S, A, T, R, p_0, \gamma \rangle$  using the given offline dataset. Then, we compute a stationary distribution of the data-collection policy  $\pi_D$  on the MLE MDP, denoted as  $d^{\pi_D}$ . Finally, we aim to solve the following policy optimization problem on the MLE MDP:

$$\pi^* := \arg \max_{\pi} \mathbb{E}_{(s, a) \sim d^{\pi}} [R(s, a)] - \alpha D_f(d^{\pi} || d^{\pi_D}),$$

which can be reformulated in terms of optimizing the stationary distribution corrections  $w$  with Lagrange multipliers  $\nu$ :

$$\begin{aligned} &\min_{\nu} \max_{w \geq 0} L(w, \nu) \\ &= (1 - \gamma) \mathbb{E}_{s \sim p_0(s)} [\nu(s)] \\ &\quad + \mathbb{E}_{(s, a) \sim \mathcal{D}} \left[ -\alpha f(w(s, a)) + w(s, a) \left( R(s, a) + \gamma(\mathcal{T}\nu)(s, a) - (\mathcal{B}\nu)(s, a) \right) \right]. \end{aligned} \tag{30}$$

For tabular MDPs, we can describe the problem using vector-matrix notation. Specifically,  $\nu \in \mathbb{R}^{|S|}$  is represented as a  $|S|$ -dimensional vector,  $w \in \mathbb{R}^{|S||A|}$  by  $|S||A|$ -dimensional vector, and  $R \in \mathbb{R}^{|S||A|}$  by  $|S||A|$ -dimensional reward vector. Then, we denote  $D = \text{diag}(d^{\pi_D}) \in \mathbb{R}^{|S||A| \times |S||A|}$  as a diagonal matrix,  $\mathcal{T} \in \mathbb{R}^{|S||A| \times |S|}$  as a matrix, and  $\mathcal{B} \in \mathbb{R}^{|S||A| \times |S|}$  as a matrix that satisfies

$$\begin{aligned}\mathcal{T}\nu &\in \mathbb{R}^{|S||A|} \quad \text{s.t.} \quad (\mathcal{T}\nu)((s, a)) = \sum_{s'} T(s'|s, a)\nu(s') \\ \mathcal{B}\nu &\in \mathbb{R}^{|S||A|} \quad \text{s.t.} \quad (\mathcal{B}\nu)((s, a)) = \nu(s)\end{aligned}$$

For brevity, we only consider the case where  $f(x) = \frac{1}{2}(x - 1)^2$  that corresponds to  $\chi^2$ -divergence-regularized policy optimization, and the problem (30) becomes

$$\min_{\nu} \max_{w \geq 0} L(w, \nu) = (1 - \gamma)p_0^\top v - \frac{\alpha}{2}(w - 1)^\top D(w - 1) + w^\top D(R + \gamma\mathcal{T}\nu - \mathcal{B}v) \quad (31)$$

From **Proposition 1**, we have the closed-form solution of the inner maximization as  $w_\nu^* = \max(0, \frac{1}{\alpha}(R + \gamma\mathcal{T}\nu - \mathcal{B}v) + 1)$  since  $(f')^{-1}(x) = x + 1$ . By plugging  $w_\nu^*$  into  $L(w, \nu)$ , we obtain

$$\min_{\nu} L(w_\nu^*, \nu) = L(\nu) := (1 - \gamma)p_0^\top \nu - \frac{\alpha}{2}(w_\nu^* - 1)D(w_\nu^* - 1) + w_\nu^{*\top} D(R + \gamma\mathcal{T}\nu - \mathcal{B}v) \quad (32)$$

Since  $L(\nu)$  is convex in  $\nu$  by **Proposition 2**, we perform a second-order optimization, i.e., Newton's method, to compute an optimal  $\nu^*$  efficiently. For almost every  $\nu$ , we can compute the first and second derivatives as follows:

$$\begin{aligned}e_\nu &:= R + \gamma\mathcal{T}\nu - \mathcal{B}\nu && \text{(advantage using } \nu\text{)} \\ m &:= \mathbf{1} \left( \frac{1}{\alpha}e_\nu + 1 \geq 0 \right) && \text{(binary masking vector)} \\ w_\nu^* &:= \left( \frac{1}{\alpha}e_\nu + 1 \right) \odot m && \text{(where } \odot m \text{ denotes element-wise masking)} \quad \text{(closed-form solution)} \\ J &:= \frac{\partial w_\nu^*}{\partial \nu} = \frac{1}{\alpha}(\gamma\mathcal{T} - \mathcal{B}) \odot m && \text{(where } \odot m \text{ denotes row-wise masking)} \quad \text{(Jacobian matrix)} \\ g &:= \frac{\partial L(\nu)}{\partial \nu} = (1 - \gamma)p_0 - \alpha J^\top D(w_\nu^* - 1) + J^\top D e_\nu + (\gamma\mathcal{T} - \mathcal{B})^\top D w_\nu^* && \text{(first-order derivative)} \\ H &:= \frac{\partial^2 L(\nu)}{\partial \nu^2} = -\alpha J^\top D J + J^\top D(\gamma\mathcal{T} - \mathcal{B}) + (\gamma\mathcal{T} - \mathcal{B})^\top D J && \text{(second-order derivative).}\end{aligned}$$

We iteratively update  $\nu$  in the direction of  $-H^{-1}g$  until convergence. Finally,  $w_{\nu^*}^*$  and the corresponding optimal policy  $\pi^*(a|s) \propto w_{\nu^*}^*(s, a) \cdot d^{\pi_D}(s, a)$  are computed. The pseudo-code of these procedures is presented in **Algorithm 2**.

## F.1 RANDOM MDPs

We validate tabular OptiDICE's efficiency and robustness using randomly generated MDPs with varying numbers of trajectories and the degree of optimality of the data-collection policy, where we follow the experimental protocol of (Laroche et al., 2019; Lee et al., 2020). We conduct repeated experiments for 10,000 runs. For each run, an MDP is generated randomly, and a data-collection policy is constructed according to the given degree of optimality  $\zeta \in \{0.9, 0.5\}$ . Then,  $N$  trajectories for  $N \in \{10, 20, 50, 100, 200, 500, 1000, 2000\}$  are collected using the generated MDP and the data-collection policy  $\pi_D$ . Finally, the constructed data-collection policy and the collected trajectories are given to each offline RL algorithm, and we measure the mean performance and the CVaR 5% performance.

### F.1.1 RANDOM MDP GENERATION

We generate random MDPs with  $|S| = 50$ ,  $|A| = 4$ ,  $\gamma = 0.95$ , and a deterministic initial state distribution, i.e.  $p_0(s) = 1$  for a fixed  $s = s_0$ . The transition model has connectivity 4: for each  $(s, a)$ , non-zero probabilities of transition to next states are given to four different states  $(s'_1, s'_2, s'_3, s'_4)$ , where the random transition probabilities are sampled from a Dirichlet distribution  $[p(s'_1|s, a), p(s'_2|s, a), p(s'_3|s, a), p(s'_4|s, a)] \sim \text{Dir}(1, 1, 1, 1)$ . The reward of 1 is given to one state that minimizes the optimal state value at the initial state; other states have zero rewards. This design of the reward function can be understood as we choose a goal state that is the most difficult to reach from the initial state. Once the agent reaches the rewarding goal state, the episode terminates.

---

**Algorithm 2** Tabular OptiDICE ( $f(x) = \frac{1}{2}(x - 1)^2$ )

---

**Input:** MLE MDP  $\hat{M} = \langle S, A, \mathcal{T}, r, \gamma, p_0 \rangle$ , data-collection policy  $\pi_D$ , regularization hyperparameter  $\alpha > 0$ .

$$\begin{aligned} d^{\pi_D} &\leftarrow \text{COMPUTESTATIONARYDISTRIBUTION}(\hat{M}, \pi_D) \\ D &\leftarrow \text{diag}(d^{\pi_D}) \\ v &\leftarrow (\text{random initialization}) \\ \text{while } v \text{ is not converged do} \\ &\quad e \leftarrow r + \gamma \mathcal{T}v - \mathcal{B}v \\ &\quad m \leftarrow \mathbf{1}(\frac{1}{\alpha}e_\nu + 1 \geq 0) \\ &\quad w \leftarrow (\frac{1}{\alpha}e + 1) \odot m \\ &\quad J \leftarrow \frac{1}{\alpha}(\gamma \mathcal{T} - \mathcal{B}) \odot m \\ &\quad g \leftarrow (1 - \gamma)p_0 - \alpha J^\top D(w - 1) + J^\top De + (\gamma \mathcal{T} - \mathcal{B})^\top Dw \\ &\quad H \leftarrow -\alpha J^\top D J + J^\top D(\gamma \mathcal{T} - \mathcal{B}) + (\gamma \mathcal{T} - \mathcal{B})^\top D J \\ &\quad v \leftarrow v - \eta H^{-1}g \quad (\text{where } \eta \text{ is a step-size}) \\ \text{end while} \\ \pi^*(a|s) &\leftarrow \frac{w(s, a)d^{\pi_D}(s, a)}{\sum_{a'} w(s, a')d^{\pi_D}(s, a')} \quad \forall s, a \end{aligned}$$

**Output:**  $\pi^*, w$

---

### F.1.2 DATA-COLLECTION POLICY CONSTRUCTION

The notion of  $\zeta$ -optimality of a policy is defined as a relative performance with respect to a uniform random policy  $\pi_{\text{unif}}$  and an optimal policy  $\pi^*$ :

$$(\zeta\text{-optimal policy } \pi\text{'s performance } V^\pi(s_0)) = \zeta V^*(s_0) + (1 - \zeta)V^{\pi_{\text{unif}}}(s_0)$$

However, there are infinitely many ways to construct a  $\zeta$ -optimal policy. In this work, we follow the way introduced in Laroche et al. (2019) to construct a  $\zeta$ -optimal data-collection policy, and the process proceeds as follows. First, an optimal policy  $\pi^*$  and the optimal value function  $Q^*$  are computed. Then, starting from  $\pi_{\text{soft}} := \pi^*$ , the policy  $\pi_{\text{soft}}$  is softened via  $\pi_{\text{soft}} \propto \exp(Q^*(s, a)/\tau)$  by increasing the temperature  $\tau$  until the performance reaches  $\frac{\zeta+1}{2}$ -optimality. Finally, the softened policy  $\pi_{\text{soft}}$  is perturbed by discounting action selection probability of an optimal action at randomly selected state. This perturbation continues until the performance of the perturbed policy reaches  $\zeta$ -optimality. The pseudo-code for the process of the data-collection policy construction is presented in **Algorithm 3**.

---

**Algorithm 3** Data-collection policy construction

---

**Input:** MDP  $M$ , Degree of optimality of the data-collection policy  $\zeta$   
 Compute the optimal policy  $\pi^*$  and its value function  $Q^*(s, a)$  on the given MDP  $M$ .  
 Initialize  $\pi_{\text{soft}} \leftarrow \pi^*$   
 Initialize a temperature parameter  $\tau \leftarrow 10^{-7}$

**while**  $V^{\pi_{\text{soft}}}(s_0) > \frac{1}{2}V^*(s_0) + \frac{1}{2}(\zeta V^*(s_0) + (1 - \zeta)V^{\pi_{\text{unif}}}(s_0))$  **do**

- Set  $\pi_{\text{soft}}$  to  $\pi_{\text{soft}}(a|s) \propto \exp\left(\frac{Q^*(s, a)}{\tau}\right) \quad \forall s, a$
- $\tau \leftarrow \tau/0.9$

**end while**

Initialize  $\pi_D \leftarrow \pi_{\text{soft}}$

**while**  $V^{\pi_D}(s_0) > \zeta V^*(s_0) + (1 - \zeta)V^{\pi_{\text{unif}}}(s_0)$  **do**

- Sample  $s \in S$  uniformly at random.
- $\pi_D(a^*|s) \leftarrow 0.9\pi_D(a^*|s)$  where  $a^* = \arg \max_a Q^*(s, a)$ .
- Normalize  $\pi_D(\cdot|s)$  to ensure  $\sum_a \pi_D(a|s) = 1$ .

**end while**

**Output:** The data-collection policy  $\pi_D$

---

### F.1.3 HYPERPARAMETERS

We compare our tabular OptiDICE with BasicRL, RaMDP (Petrik et al., 2016), RobustMDP (Nilim & El Ghaoui, 2005; Iyengar, 2005), SPIBB (Laroche et al., 2019), and BOPAH (Lee et al., 2020). For the hyperparameters, we follow the setting in the public code of SPIBB and BOPAH, which are listed as follows:

**RaMDP.**  $\kappa = 0.003$  is used for the reward-adjusting hyperparameter.

**RobustMDP.**  $\delta = 0.001$  is used for the confidence interval hyperparameter to construct an uncertainty set.

**SPIBB.**  $N_\wedge = 5$  is used for the data-collection policy bootstrapping threshold.

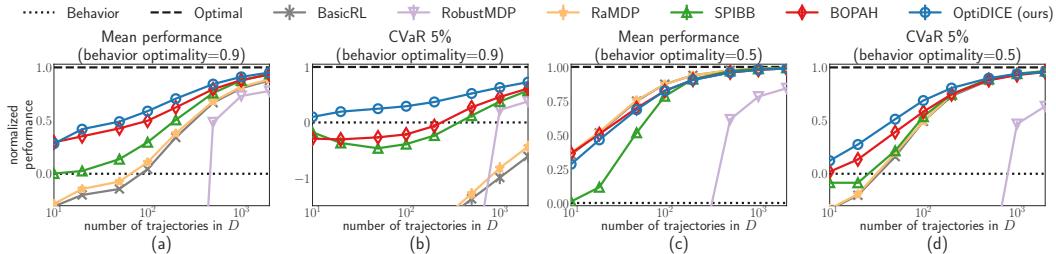
**BOPAH.** The 2-fold cross validation criteria and fully state-dependent KL-regularization is used.

**OptiDICE.**  $\alpha = N^{-1}$  for the number  $N$  of trajectories is used for the reward-regularization balancing hyperparameter. We also use  $f(x) = \frac{1}{2}(x - 1)^2$  which corresponds to  $\chi^2$ -divergence.

### F.1.4 EXPERIMENTAL RESULTS

We validate tabular OptiDICE’s efficiency and robustness using randomly generated MDPs by following the experimental protocol from Laroche et al. (2019) and Lee et al. (2020), which is described above. We consider a data-collection policy  $\pi_D$  characterized by the behavior optimality parameter  $\zeta$  that relates to  $\pi_D$ ’s performance  $\zeta V^*(s_0) + (1 - \zeta)V^{\pi_{\text{unif}}}(s_0)$  where  $\pi_{\text{unif}}$  denotes the uniformly random policy. We evaluate each algorithm in terms of the normalized performance of the policy  $\pi$ , given by  $(V^*(s_0) - V^{\pi_D}(s_0))/(V^*(s_0) - V^{\pi_{\text{unif}}}(s_0))$ , which intuitively measures the performance enhancement of  $\pi$  over  $\pi_D$ . Each algorithm is tested for 10,000 runs, and their mean and 5% conditional value at risk (5%-CVaR) are reported, where the mean of the worst 500 runs is considered for 5%-CVaR. Note that CVaR implicitly stands for the robustness of each algorithm.

We describe the performance of tabular OptiDICE and baselines in *Figure 4*. For  $\zeta = 0.9$ , where  $\pi_D$  is near-deterministic and thus  $d^D$ ’s support is relatively small, OptiDICE outperforms the baselines in both mean and CVaR (*Figure 4(a),(b)*). For  $\zeta = 0.5$ , where  $\pi_D$  is highly stochastic and thus  $d^D$ ’s support is relatively large, OptiDICE outperforms the baselines in CVaR, while performing competitively in mean. In summary, OptiDICE was more sample-efficient and stable than the baselines.



**Figure 4:** Performance of tabular OptiDICE and baseline algorithms in random MDPs. For baselines, we use *BasicRL* (a model-based RL computing an optimal policy via MLE MDP), *RobustMDP* (Nilim & El Ghaoui, 2005; Iyengar, 2005), *Reward-adjusted MDP (RaMDP)* (Petrik et al., 2016), *SPIBB* (Laroche et al., 2019), *BOPAH* (Lee et al., 2020). For varying numbers of trajectories and two types of data-collection policies ( $\zeta = 0.9, 0.5$ ), the mean and the 5%-CVaR of normalized performances for 10,000 runs are reported with 95% confidence intervals. OptiDICE performs better than (for  $\zeta = 0.9$ ) or on par with (for  $\zeta = 0.5$ ) the baselines in the mean performance measure, while always outperforming the baselines in the CVaR performance measure.

## G PROOF OF PROPOSITION 4

**Proposition 4.** *The closed-form solution of the inner maximization with normalization constraint, i.e.,*

$$w_{\nu,\lambda}^* := \arg \max_{w \geq 0} (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \mathbb{E}_{(s,a) \sim d^D} [-\alpha f(w(s,a)) + w(s,a)(e_\nu(s,a) - \lambda)] + \lambda$$

is given as

$$w_{\nu,\lambda}^*(s,a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s,a) - \lambda}{\alpha} \right) \right).$$

*Proof.* Similar to the proof for **Proposition 1**, we consider the maximization problem

$$\max_{w \geq 0} L(w, \nu, \lambda)$$

for fixed  $\nu$  and  $\lambda$ , where we consider this maximization as a primal problem. Then, its dual problem is

$$\max_w \min_{\mu \geq 0} L(w, \nu, \lambda) + \sum_{s,a} \mu(s,a) w(s,a).$$

Since the strong duality holds, KKT condition is both necessary and sufficient conditions for primal and dual solutions  $w^*$  and  $\mu^*$ , where dependencies on  $\nu, \lambda$  are ignored for brevity. While the KKT conditions on **Primal feasibility**, **Dual feasibility** and **Complementary slackness** are the same as those in the proof of **Proposition 1**, the condition on **Stationarity** is slightly different due to the normalization constraint:

**Condition 1 (Primal feasibility).**  $w^* \geq 0 \forall s, a$ .

**Condition 2 (Dual feasibility).**  $\mu^* \geq 0 \forall s, a$ .

**Condition 3 (Stationarity).**  $d^D(s,a)(-\alpha f'(w^*(s,a)) + e_\nu(s,a) + \mu^*(s,a) - \lambda) = 0 \forall s, a$ .

**Condition 4 (Complementary slackness).**  $w^*(s,a)\mu^*(s,a) = 0 \forall s, a$ .

The remainder of the proof is similar to the proof of **Proposition 1**. From **Stationarity** and  $d^D > 0$ , we have

$$f'(w^*(s,a)) = \frac{e_\nu(s,a) + \mu^*(s,a) - \lambda}{\alpha} \forall s, a,$$

and since  $f'$  is invertible due to the strict convexity of  $f$ ,

$$w^*(s,a) = (f')^{-1} \left( \frac{e_\nu(s,a) + \mu^*(s,a) - \lambda}{\alpha} \right) \forall s, a.$$

Given  $s, a$ , assume either  $w^*(s,a) > 0$  or  $w^*(s,a) = 0$ , satisfying **Primal feasibility**.

**Case 1** ( $w^*(s,a) > 0$ ).  $\mu^*(s,a) = 0$  due to **Complementary slackness**, and thus,

$$w^*(s,a) = (f')^{-1} \left( \frac{e_\nu(s,a) - \lambda}{\alpha} \right) > 0.$$

Note that **Dual feasibility** holds. Since  $f'$  is a strictly increasing function due to the strict convexity of  $f$ ,  $e_\nu(s,a) - \lambda > \alpha f'(0)$  should be satisfied if  $f'(0)$  is well-defined.

**Case 2** ( $w^*(s,a) = 0$ ).  $\mu^*(s,a) = \alpha f'(0) - e_\nu(s,a) + \lambda \geq 0$  due to **Stationarity** and **Dual feasibility**, and thus,  $e_\nu(s,a) - \lambda \leq \alpha f'(0)$  should be satisfied if  $f'(0)$  is well-defined.

In summary, we have

$$w_{\nu,\lambda}^*(s,a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s,a) - \lambda}{\alpha} \right) \right).$$

□

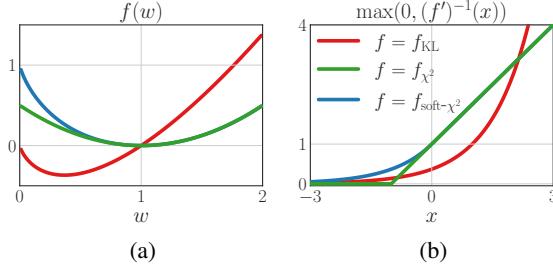


Figure 5: We depict (a) generator functions  $f$  of  $f$ -divergences and (b) corresponding functions  $\max(0, (f')^{-1}(\cdot))$  used to define the closed-form solution in **Proposition 4**. While  $f_{\text{KL}}(x)$  has a numerical instability for large  $x$  and  $f_{\chi^2}(x)$  provides zero gradients for negative  $x$ ,  $f_{\text{soft-}\chi^2}$  does not suffer from both issues.

## H $f$ -DIVERGENCE

Pertinent to the result of **Proposition 4**, one can observe that the choice of the function  $f$  of  $f$ -divergence can affect the numerical stability of optimization especially when using the closed-form solution of  $w_{\nu,\lambda}^*$ :

$$w_{\nu,\lambda}^*(s, a) = \max \left( 0, (f')^{-1} \left( \frac{e_\nu(s, a) - \lambda}{\alpha} \right) \right).$$

For example, for the choice of  $f(x) = f_{\text{KL}}(x) := x \log x$  that corresponds to KL-divergence, we have  $(f'_{\text{KL}})^{-1}(x) = \exp(x - 1)$ . This yields the following closed-form solution of  $w_{\nu,\lambda}^*$ :

$$w_{\nu,\lambda}^*(s, a) = \exp \left( \frac{e_\nu(s, a) - \lambda}{\alpha} - 1 \right).$$

However, the choice of  $f_{\text{KL}}$  can incur numerical instability due to its inclusion of an  $\exp(\cdot)$ , i.e. for values of  $\frac{1}{\alpha}(e_\nu(s, a) - \lambda)$  in order of tens, the value of  $w_{\nu,\lambda}^*(s, a)$  easily explodes and so does the gradient  $\nabla_\nu w_{\nu,\lambda}^*(s, a)$ .

Alternatively, for the choice of  $f(x) = f_{\chi^2}(x) := \frac{1}{2}(x - 1)^2$  that corresponds to  $\chi^2$ -divergence, we have  $(f'_{\chi^2})^{-1}(x) = x + 1$ . This yields the following closed-form solution of  $w_{\nu,\lambda}^*$ :

$$w_{\nu,\lambda}^*(s, a) = \text{ReLU} \left( \frac{e_\nu(s, a) - \lambda}{\alpha} + 1 \right),$$

where  $\text{ReLU}(x) := \max(0, x)$ . Still, this choice may suffer from dying gradient problem: for values of negative  $\frac{1}{\alpha}(e_\nu(s, a) - \lambda) + 1$ , the gradient  $\nabla_\nu w_{\nu,\lambda}^*(s, a)$  becomes zero, which can make training  $\nu$  slow or even fail.

Consequently, we adopt the function  $f = f_{\text{soft-}\chi^2}$  that combines the form of  $f_{\text{KL}}$  and  $f_{\chi^2}$ , which can prevent both of the aforementioned issues:

$$f_{\text{soft-}\chi^2}(x) := \begin{cases} x \log x - x + 1 & \text{if } 0 < x < 1 \\ \frac{1}{2}(x - 1)^2 & \text{if } x \geq 1. \end{cases} \Rightarrow (f_{\text{soft-}\chi^2}(x)')^{-1}(x) = \begin{cases} \exp(x) & \text{if } x < 0 \\ x + 1 & \text{if } x \geq 0 \end{cases}$$

This particular choice of  $f$  yields the following closed-form solution of  $w_{\nu,\lambda}^*$ :

$$w_{\nu,\lambda}^*(s, a) = \text{ELU} \left( \frac{e_\nu(s, a) - \lambda}{\alpha} \right) + 1.$$

Here,  $\text{ELU}(x) := \exp(x) - 1$  if  $x < 0$  and  $x$  for  $x \geq 0$ . Note that the solution for  $f = f_{\text{soft-}\chi^2}$  is numerically stable for large  $\frac{1}{\alpha}(e_\nu(s, a) - \lambda)$  and always gives non-zero gradients. We use  $f = f_{\text{soft-}\chi^2}$  for the D4RL experiments.

## I EXPERIMENTAL SETTINGS

### I.1 D4RL BENCHMARK

#### I.1.1 TASK DESCRIPTIONS

We use Maze2D and Gym-MuJoCo environments of D4RL benchmark (Fu et al., 2021) to evaluate OptiDICE and CQL (Kumar et al., 2020) in continuous control tasks. We summarize the descriptions of tasks in D4RL paper (Fu et al., 2021) as follows:

**Maze2D.** This is a navigation task in 2D state space, while the agent tries to reach a fixed goal location. By using priorly gathered trajectories, the goal of the agent is to find out a shortest path to reach the goal location. The complexity of the maze increases with the order of "maze2d-umaze", "maze2d-medium" and "maze2d-large".

**Gym-MuJoCo.** For each task in {hopper, walker2d, halfcheetah} of MuJoCo continuous controls, the dataset is gathered in the following ways.

*random.* The dataset is generated by a randomly initialized policy in each task.

*medium.* The dataset is generated by using the policy trained by SAC (Haarnoja et al., 2018) with early stopping.

*medium-replay.* The "replay" dataset consists of the samples gathered during training the policy for "medium" dataset. The "medium-replay" dataset includes both "medium" and "replay" datasets.

*medium-expert.* The dataset is given by using the same amount of expert trajectories and suboptimal trajectories, where those suboptimal ones are gathered by using either a randomly uniform policy or a medium-performance policy.

#### I.1.2 HYPERPARAMETER SETTINGS FOR CQL

We follow the hyperparameters specified by Kumar et al. (2020). For learning both the Q-funtions and the policy, fully-connected multi-layer perceptrons (MLPs) with three hidden layers and ReLU activations are used, where the number of hidden units on each layer is equal to 256. A Q-function learning rate of 0.0003 and a policy learning rate of 0.0001 are used with Adam optimizer for these networks. CQL( $\mathcal{H}$ ) is evaluated, with an approximate max-backup (see Appendix F of (Kumar et al., 2020) for more details) and a static  $\alpha = 5.0$ , which controls the conservativeness of CQL. The policy of CQL is updated for 2,500,000 iterations, while we use 40,000 warm-up iterations where we update Q-functions as usual, but the policy is updated according to the behavior cloning objective.

### I.1.3 HYPERPARAMETER SETTINGS FOR OPTIDICE

For neural networks  $\nu_\theta, e_\phi, \pi_\psi$  and  $\pi_\beta$  in **Algorithm 1**, we use fully-connected MLPs with two hidden layers and ReLU activations, where the number of hidden units on each layer is equal to 256. For  $\pi_\psi$ , we use tanh-squashed normal distribution. We regularize the entropy of  $\pi_\psi$  with learnable entropy regularization coefficients, where target entropies are set to be the same as those in SAC (Haarnoja et al., 2018) ( $-\dim(A)$  for each task). For  $\pi_\beta$ , we use tanh-squashed mixture of normal distributions, where we build means and standard deviations of each mixture component upon shared hidden outputs. No entropy regularization is applied to  $\pi_\beta$ . For both  $\pi_\psi$  and  $\pi_\beta$ , means are clipped within  $(-7.24, 7.24)$ , while log of standard deviations are clipped within  $(-5, 2)$ . For the optimization of each network, we use stochastic gradient descent with Adam optimizer and its learning rate 0.0003. The batch size is set to be 512. Before training neural networks, we preprocess the dataset  $D$  by standardizing observations and rewards. We additionally scale the rewards by multiplying 0.1. We update the policy  $\pi_\psi$  for 2,500,000 iterations, while we use 500,000 warm-up iterations for other networks, i.e., those networks other than  $\pi_\psi$  are updated for 3,000,000 iterations.

For each task and OptiDICE methods (OptiDICE-minimax, OptiDICE-MSE), we search the number  $K$  of mixtures (for  $\pi_\beta$ ) within  $\{1, 5, 9\}$  and the coefficient  $\alpha$  within  $\{0.0001, 0.001, 0.01, 0.1, 1\}$ , while we additionally search  $\alpha$  over  $\{2, 5, 10\}$  for hopper-medium-replay. The hyperparameters  $K$  and  $\alpha$  showing the best mean performance were chosen, which are described as follows:

Table 1: Hyperparameters

Task	OptiDICE-MSE		OptiDICE-minimax	
	$K$	$\alpha$	$K$	$\alpha$
maze2d-umaze	5	0.001	1	0.01
maze2d-medium	5	0.0001	1	0.01
maze2d-large	1	0.01	1	0.01
hopper-random	5	1	5	1
hopper-medium	9	0.1	9	0.1
hopper-medium-replay	9	10	1	2
hopper-medium-expert	9	1	5	1
walker2d-random	9	0.0001	1	0.0001
walker2d-medium	9	0.01	5	0.01
walker2d-medium-replay	9	0.1	9	0.1
walker2d-medium-expert	5	0.01	5	0.01
halfcheetah-random	5	0.0001	9	0.001
halfcheetah-medium	1	0.01	1	0.1
halfcheetah-medium-replay	9	0.01	1	0.1
halfcheetah-medium-expert	9	0.01	9	0.01

## J EXPERIMENTAL RESULTS

### J.1 EXPERIMENTAL RESULTS FOR $\gamma = 0.99$

Table 2: Normalized performance of OptiDICE compared with the best model-free baseline in the D4RL benchmark tasks (Fu et al., 2021). In the *Best baseline* column, the algorithm with the best performance among 8 algorithms (offline SAC (Haarnoja et al., 2018), BEAR (Kumar et al., 2019), BRAC (Wu et al., 2019), AWR (Peng et al., 2019), cREM (Agarwal et al., 2020), BCQ (Fujimoto et al., 2019), AlgaeDICE (Nachum et al., 2019b), CQL (Kumar et al., 2020)) is presented, taken from (Fu et al., 2021). OptiDICE achieved highest scores in 7 tasks.

D4RL Task	Best baseline	OptiDICE
maze2d-umaze	88.2 Offline SAC	<b>111.0</b>
maze2d-medium	33.8 BRAC-v	<b>145.2</b>
maze2d-large	40.6 BRAC-v	<b>155.7</b>
hopper-random	<b>12.2</b> BRAC-v	11.2
hopper-medium	58.0 CQL	<b>94.1</b>
hopper-medium-replay	<b>48.6</b> CQL	36.4
hopper-medium-expert	110.9 BCQ	<b>111.5</b>
walker2d-random	7.3 BEAR	<b>9.9</b>
walker2d-medium	<b>81.1</b> BRAC-v	21.8
walker2d-medium-replay	<b>26.7</b> CQL	21.6
walker2d-medium-expert	<b>111.0</b> CQL	74.8
halfcheetah-random	<b>35.4</b> CQL	11.6
halfcheetah-medium	<b>46.3</b> BRAC-v	38.2
halfcheetah-medium-replay	<b>47.7</b> BRAC-v	39.8
halfcheetah-medium-expert	64.7 BCQ	<b>91.1</b>

Table 3: Normalized performance of OptiDICE compared with baselines. We report the performance of CQL (Kumar et al., 2020) by using the code released by authors, while mean scores for other baselines—BEAR (Kumar et al., 2019), BRAC (Wu et al., 2019) and AlgaeDICE (Nachum et al., 2019b)—come from D4RL benchmark (Fu et al., 2021). OptiDICE achieves the best performance on 8 tasks compared to our baselines. Note that 3-run mean scores without confidence intervals were reported on each task by Fu et al. (2021). For CQL and OptiDICE, we use 5 runs and report means and 95% confidence intervals.

Task	BC	Offline SAC	BEAR	BRAC-v	Algae DICE	CQL	OptiDICE minimax	MSE
maze2d-umaze	3.8	88.2	3.4	-16.0	-15.7	-14.9 $\pm$ 0.7	<b>111.0 <math>\pm</math> 8.3</b>	105.8 $\pm$ 17.5
maze2d-medium	30.3	26.1	29.0	33.8	10.0	17.2 $\pm$ 8.7	109.9 $\pm$ 7.7	<b>145.2 <math>\pm</math> 17.5</b>
maze2d-large	5.0	-1.9	4.6	40.6	-0.1	1.6 $\pm$ 3.8	116.1 $\pm$ 43.1	<b>155.7 <math>\pm</math> 33.4</b>
hopper-random	9.8	11.3	11.4	<b>12.2</b>	0.9	10.7 $\pm$ 0.0	11.2 $\pm$ 0.1	10.7 $\pm$ 0.2
hopper-medium	29.0	0.8	52.1	31.1	1.2	89.8 $\pm$ 7.6	92.9 $\pm$ 2.6	<b>94.1 <math>\pm</math> 3.7</b>
hopper-medium-replay	11.8	3.5	33.7	0.6	1.1	33.3 $\pm$ 2.2	<b>36.4 <math>\pm</math> 1.1</b>	30.7 $\pm$ 1.2
hopper-medium-expert	111.9	1.6	96.3	0.8	1.1	<b>112.3 <math>\pm</math> 0.2</b>	111.5 $\pm$ 0.6	106.7 $\pm$ 1.8
walker2d-random	1.6	4.1	7.3	1.9	0.5	3.0 $\pm$ 1.8	9.4 $\pm$ 2.2	<b>9.9 <math>\pm</math> 4.3</b>
walker2d-medium	6.6	0.9	59.1	<b>81.1</b>	0.3	73.7 $\pm$ 2.7	21.8 $\pm$ 7.1	20.8 $\pm$ 3.1
walker2d-medium-replay	11.3	1.9	19.2	0.9	0.6	13.4 $\pm$ 0.8	21.5 $\pm$ 2.9	<b>21.6 <math>\pm</math> 2.1</b>
walker2d-medium-expert	6.4	-0.1	40.1	81.6	0.4	<b>99.7 <math>\pm</math> 7.2</b>	74.7 $\pm$ 7.5	74.8 $\pm$ 9.2
halfcheetah-random	2.1	30.5	25.1	<b>31.2</b>	-0.3	25.5 $\pm$ 0.5	8.3 $\pm$ 0.8	11.6 $\pm$ 1.2
halfcheetah-medium	36.1	-4.3	41.7	<b>46.3</b>	-2.2	42.3 $\pm$ 0.1	37.1 $\pm$ 0.1	38.2 $\pm$ 0.1
halfcheetah-medium-replay	38.4	-2.4	38.6	<b>47.7</b>	-2.1	43.1 $\pm$ 0.7	38.9 $\pm$ 0.5	39.8 $\pm$ 0.3
halfcheetah-medium-expert	35.8	1.8	53.4	41.9	-0.8	53.5 $\pm$ 13.3	76.2 $\pm$ 7.0	<b>91.1 <math>\pm</math> 3.7</b>

## J.2 EXPERIMENTAL RESULTS WITH IMPORTANCE-WEIGHTED BC

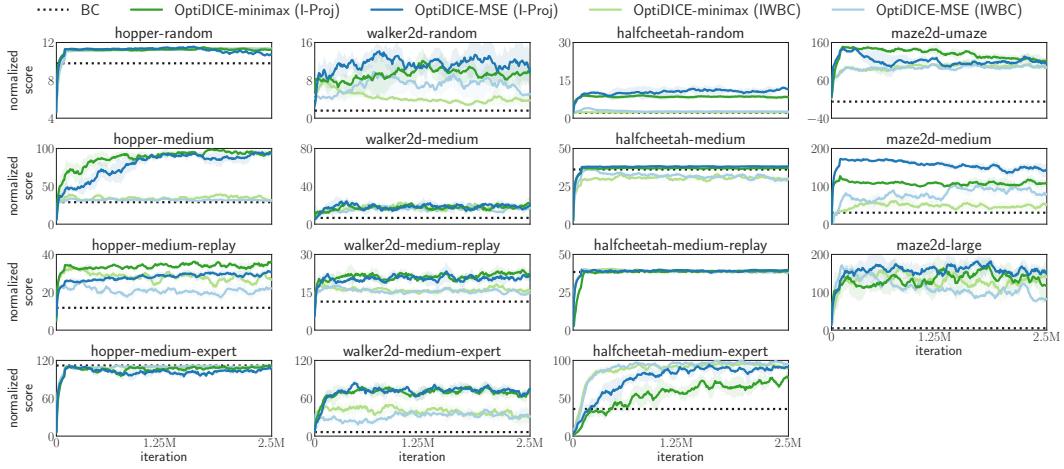


Figure 6: Performance of BC, OptiDICE with importance-weighted BC (IWBC) and information projection (I-Proj) methods on D4RL benchmark.  $\gamma = 0.99$  is used.

The empirical results of OptiDICE for different policy extraction methods (information-weighted BC, I-projection methods) are depicted in *Figure 6*. For those results with IWBC, we search  $\alpha$  within  $\{0.0001, 0.001, 0.01, 0.1, 1\}$  and choose one with the best mean performance, which are summarized in *Table 4*. The hyperparameters other than  $\alpha$  are the same as those used for information-projection methods, which is described in Section I.1.3. We empirically observe that policy extraction with information projection method performs better than the extraction with importance-weighted BC, as discussed in Section B.1.

Table 4: Hyperparameters for importance-weighted BC

Task	OptiDICE-MSE		OptiDICE-minimax	
	$\alpha$	$\alpha$	$\alpha$	$\alpha$
maze2d-umaze	0.001		0.001	
maze2d-medium	0.0001		0.001	
maze2d-large	0.001		0.001	
hopper-random		1		1
hopper-medium		0.01		0.1
hopper-medium-replay		0.1		0.1
hopper-medium-expert		0.1		0.1
walker2d-random	0.0001		0.0001	
walker2d-medium	0.1		0.1	
walker2d-medium-replay	0.1		0.1	
walker2d-medium-expert	0.01		0.01	
halfcheetah-random	0.001		0.01	
halfcheetah-medium	0.0001		0.1	
halfcheetah-medium-replay	0.01		0.01	
halfcheetah-medium-expert	0.01		0.01	

### J.3 EXPERIMENTAL RESULTS FOR $\gamma \in \{0.99, 0.999, 0.9999, 1.0\}$

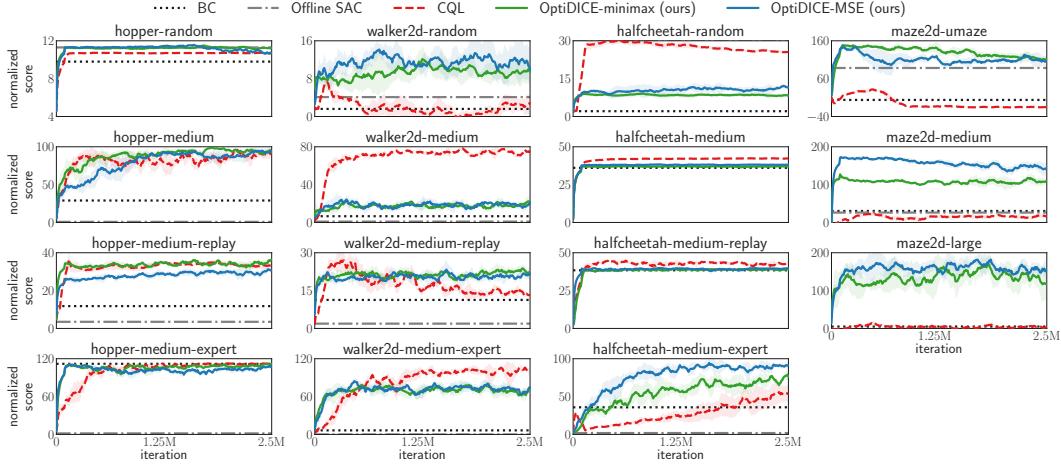


Figure 7: Performance of BC, CQL, OptiDICE-minimax and OptiDICE-MSE on D4RL benchmark for  $\gamma = 0.99$ .

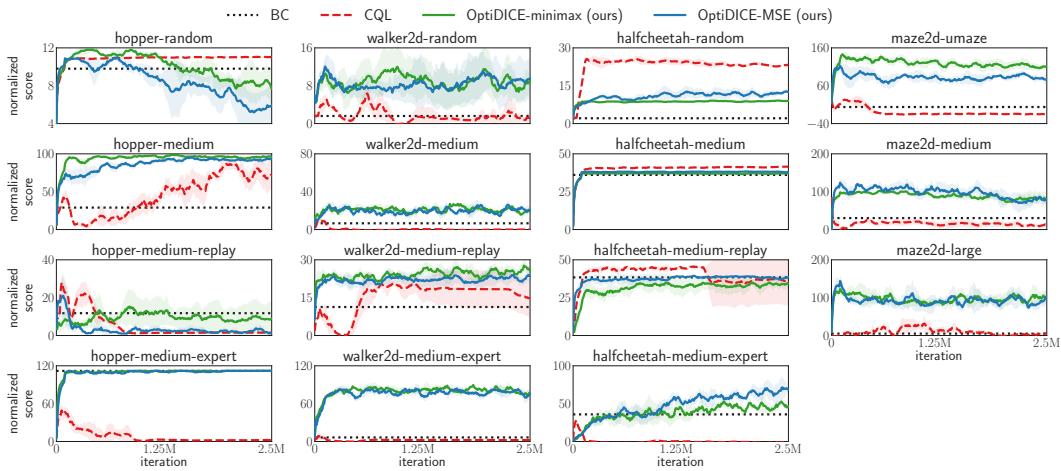


Figure 8: Performance of BC, CQL, OptiDICE-minimax and OptiDICE-MSE on D4RL benchmark for  $\gamma = 0.999$ .

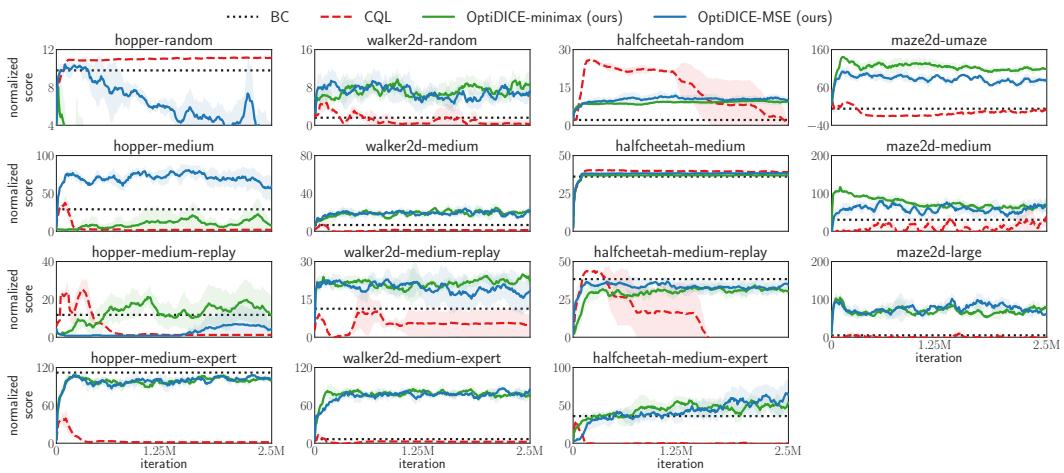


Figure 9: Performance of BC, CQL, OptiDICE-minimax and OptiDICE-MSE on D4RL benchmark for  $\gamma = 0.9999$ .

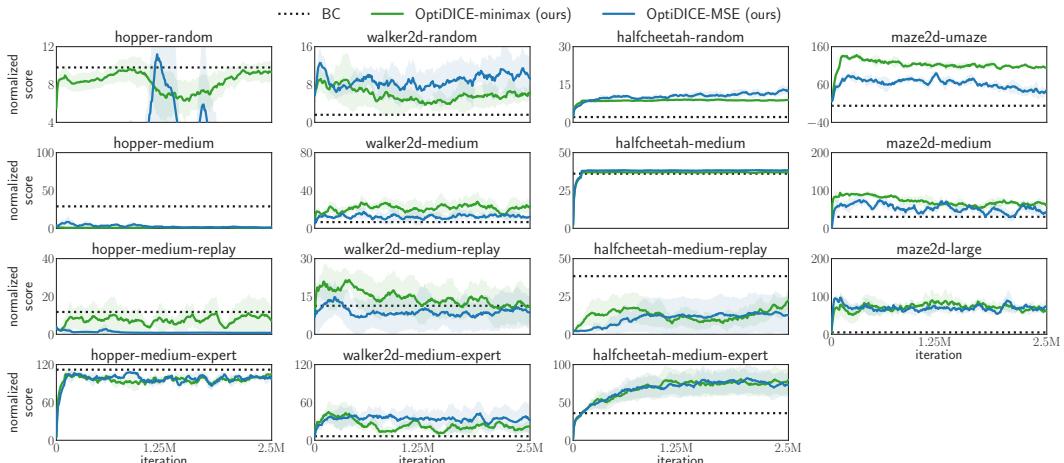


Figure 10: Performance of BC, OptiDICE-minimax and OptiDICE-MSE on D4RL benchmark for  $\gamma = 1$ . Note that CQL cannot deal with  $\gamma = 1$  case. Thus, we only provide the results of OptiDICE and BC.