

General Flexible f -divergence for Challenging Offline RL Datasets with Low Stochasticity and Diverse Behavior Policies

Jianxun Wang

North Carolina State University
Raleigh, United States
jwang75@ncsu.edu

Leonardo Villalobos-Arias

North Carolina State University
Raleigh, United States
lvillal@ncsu.edu

Grant C. Forbes

North Carolina State University
Raleigh, United States
gforbes@alumni.ncsu.edu

David L. Roberts

North Carolina State University
Raleigh, United States
drober4@ncsu.edu

ABSTRACT

Offline RL algorithms aim to improve upon the behavior policy that produces the collected data while constraining the learned policy to be within the support of the dataset. However, practical offline datasets often contain examples with little diversity or limited exploration of the environment, and from multiple behavior policies with diverse expertise levels. Limited exploration can impair the offline RL algorithm’s ability to estimate Q or V values, while constraining towards diverse behavior policies can be overly conservative. Such datasets call for a balance between the RL objective and behavior policy constraints. We first identify the connection between f -divergence and optimization constraint on the Bellman residual through a more general Linear Programming form for RL and the convex conjugate. Following this, we introduce the general flexible function formulation for the f -divergence to incorporate an adaptive constraint on algorithms’ learning objectives based on the offline training dataset. Results from experiments on the MuJoCo, Fetch, and AdroitHand environments show the correctness of the proposed LP form and the potential of the flexible f -divergence in improving performance for learning from a challenging dataset when applied to a compatible constrained optimization algorithm.

KEYWORDS

Offline Reinforcement Learning, f divergence, LEARN

ACM Reference Format:

Jianxun Wang, Grant C. Forbes, Leonardo Villalobos-Arias, and David L. Roberts. 2026. General Flexible f -divergence for Challenging Offline RL Datasets with Low Stochasticity and Diverse Behavior Policies. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/CKXI1923>

1 INTRODUCTION

Recent progress in offline Reinforcement Learning (RL) significantly improved RL algorithms’ practical value by alleviating the bottleneck problem of RL model training, which is the need for online

environment interaction. The capability of training an RL agent on a static dataset opens up application opportunities such as AI-conversation and robot manipulation [12, 29]. A key idea behind the general design principle of offline RL algorithms is to optimize the learning policy with the given reward while constraining it to stay within the support of the static dataset [18, 25]. Algorithms designed following this pessimism principle [7, 15, 16] have demonstrated considerable performance in benchmarks such as D4RL [6].

However, the implicit exploration setting and the limited combination of behavior policies suggest a gap between the characteristics of datasets from the benchmark and the real-world application. Practical considerations [14, 20] like safety and costs of experiments can restrict the practitioner who collects the data to follow an almost deterministic or even rule-based behavior. The limited exploration problem in the offline RL setting can be more severe. In addition, applications of RL may call for large-scale data collection from different practitioners or crowd-sourced data collection [34]. Such practices can lead to datasets with a high mixture of expertise levels, even for learning a single skill.

While previous works [26, 31] have explored characterizing offline RL datasets in terms of trajectories’ performance and the amount of explorations present in the dataset, these measurements may only reflect dataset properties on the distribution level at best, and cannot identify any individual dataset. As the top row in Figure 3 shows, both measurements on dataset returns (PSV, NER) and dataset exploration (SACo) exhibit significant overlap among different settings, in terms of the number of behavior policies and behavior policy stochasticity. This makes it impossible to reliably tell what a single dataset’s setting is from those measurements. Furthermore, as the bottom row in Figure 3 shows, each algorithm can suffer significant performance loss under different dataset mixtures and with reduced stochasticity. These observations suggest that two offline RL datasets can appear empirically similar but differ drastically in the learning algorithms’ performance.

One of the potential causes of this performance degradation can be the algorithm’s universal pessimism towards the dataset. Different expertise levels and stochasticity levels of behavior policies call for different constraints. Previous work on heteroscedastic datasets [28] and mixtures of two behavior policies [11] share our opinion. The former study applied a relaxed constraint in the policy evaluation, while the latter re-weighted their learning data samples by trajectory returns. Yu et al. [37] focused on imitation learning



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

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CKXI1923>

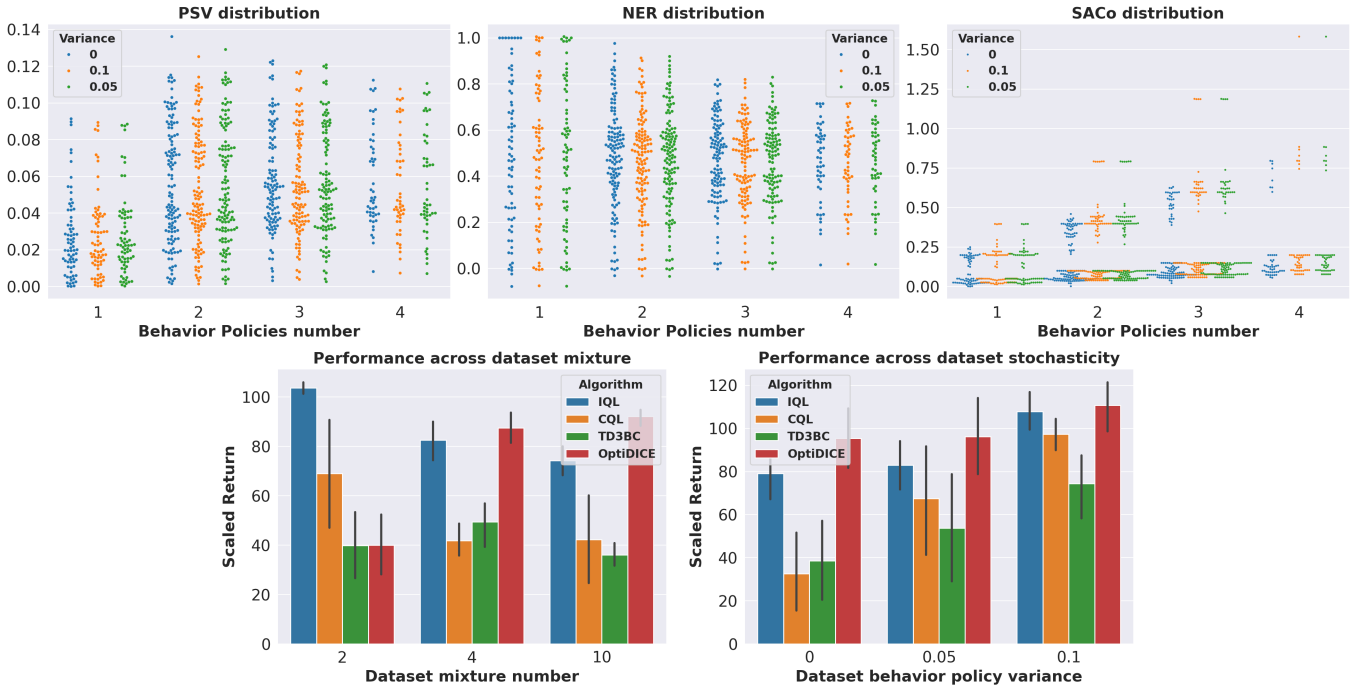


Figure 1: Dataset measurements across different compositions and algorithms’ performance in them. (Top) Positive Scaled Variance (PSV) of the return, Normalized Expected Return (NER), and SACo as a measure of exploration; (Bottom) Algorithm performances across different dataset mixture and behavior policy stochasticity.

by applying a relaxed f -divergence using a small expert dataset and a larger unlabeled dataset.

In this paper, we focus on using the f -divergence to control constraint level. We first highlight the implicit role of f -divergence for any RL algorithms that minimize the Bellman error by generalizing the unconstrained optimization derived from the Linear Programming (LP) formulation of RL and connecting it to traditional RL algorithms. We then show a unified form of LP for RL as optimizing the sum of an initial loss term and an f -divergence penalty on either the density ratio estimation or the estimated value, under the Bellman constraint. These theoretical results show that RL algorithms naturally incorporate the f -divergence against the training data, either explicitly or implicitly. Inspired by these results, we propose the flexible f -divergence, a general function form for f -divergence that is applicable to both traditional Bellman error minimization algorithms and dual optimization algorithms based on density ratio estimation, to enable flexible constraint control. We show that the Flexible f -divergence can achieve improved performance over its base algorithms in challenging settings where base algorithms fail.

The extended version of this paper with appendix is available at <https://arxiv.org/abs/2602.11087>.

2 BACKGROUND

The traditional Reinforcement Learning problem formulates the agent-environment interaction as a Markov Decision Process (MDP) [30]. The MDP can be defined as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, p_0, r, \gamma)$, with a state space \mathcal{S} , an action space \mathcal{A} , a state transition dynamics

$\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})^1$, a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, an initial state distribution p_0 , and a discount factor $\gamma \in [0, 1]$. The agent’s policy $\pi(a_t | s_t)$ is the conditional probability of action to take with the given state. With $s_0 \sim p_0, a_t \sim \pi(a_t | s_t), s_{t+1} \sim \mathcal{T}(s_{t+1} | s_t, a_t)$, the expected discounted return is defined as $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$. Practical evaluations of the expected return for a given policy π are the action-value function $Q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{a' \sim \pi, s' \sim \mathcal{T}} [Q^\pi(s', a')]$ in the Bellman equation format and the state value function $V^\pi(s) = \mathbb{E}_{a \sim \pi} [Q^\pi(s, a)]$. The goal of an RL algorithm is to find the optimal policy π^* that maximizes the expected return, $\pi^* = \operatorname{argmax}_\pi \mathbb{E}_\pi [\mathcal{R}]$, whose value function and action-value function are V^* and Q^* .

In the offline RL setting, the RL algorithm seeks to optimize the policy using a static dataset $\mathcal{D} = \{(s, a, s', r)\}_{i=1}^N$ without access to the actual environment. To provide a measure of data sample distribution, we can express the visitation of state-action pairs as the state-action occupancy measure $d^\pi(s, a) := (1-\gamma) \mathbb{E}_\pi [\sum_t \gamma^t P(s_t, a_t)]$ and the state occupancy measure $d^\pi(s) = \sum_a d^\pi(s, a)$. $d^\mathcal{D}(s, a)$ and $d^\mathcal{D}(s)$ are occupancy measure of the dataset.

The optimization of the RL problem can be formulated as a Linear Programming (LP) problem of the V function (V -LP) with the following as the primal problem [24]:

$$\min_v (1-\gamma) p_0(s) v(s) \tag{1}$$

$$\text{s.t. } (\mathcal{B}v)(s, a) \geq r(s, a) + \gamma (\mathcal{T}v)(s, a), \forall s, a \tag{2}$$

where $(\mathcal{T}v)(s, a) = \sum_{s'} \mathcal{T}(s' | s, a) v(s')$ and $(\mathcal{B}v)(s, a) = v(s)$. The intuition behind the primal problem is to minimize the upper bound

¹ $\Delta(\mathcal{S})$ is the space of probability distribution over \mathcal{S} [1].

of V since it is the fixed point solution of the Bellman equation upon convergence. Therefore the solution of $v(s)$ would be $V^*(s)$. The following optimization problem is the more common dual optimization form of the primal problem.

$$\begin{aligned} \max_{d \geq 0} d(s, a)r(s, a) \quad (3) \\ \text{s.t. } (\mathcal{B}_*d)(s) = (1 - \gamma)p_0(s) + \gamma(\mathcal{T}_*d)(s), \forall s \quad (4) \end{aligned}$$

where $(\mathcal{T}_*d)(s) = \sum_{\bar{s}, \bar{a}} \mathcal{T}(s|\bar{s}, \bar{a})d(\bar{s}, \bar{a})$ and $(\mathcal{B}_*d)(s) = \sum_a d(s, a)$. The dual problem has a more natural interpretation as maximizing the expected return while ensuring the occupancy measure satisfies the Bellman flow constraint 4.

To simplify notation, we define the TD error (or equivalently, the advantage estimation) as

$$e_v(s, a) = r(s, a) + \underbrace{\gamma(\mathcal{T}v)(s, a)}_{Q(s, a)} - v(s). \quad (5)$$

We also define the density ratio estimation of occupancy measures between $d(s, a)$ and $d^{\mathcal{D}}(s, a)$ as $\zeta(s, a) = \frac{d(s, a)}{d^{\mathcal{D}}(s, a)}$. The LP problem results in the Lagrangian below and its dual optimization problem:

$$\begin{aligned} \max_{\zeta \geq 0} \min_v \sum_{s, a} [\underbrace{(1 - \gamma)p_0(s)v(s)}_{L_P} + \underbrace{\zeta(s, a)r(s, a)}_{L_D} \\ + \gamma\zeta(s, a)(\mathcal{T}v)(s, a) - \zeta(s, a)v(s)]. \quad (6) \end{aligned}$$

The summation $\sum_{s, a} [\cdot]$ is only possible in theory, and it implicitly assumes a universal data occupancy measure, i.e., $d^{\mathcal{D}}(s, a) = 1$. In practical offline RL settings, we substitute it with $\mathbb{E}_{(s, a) \sim d^{\mathcal{D}}}[\cdot]$. To accommodate the offline RL setting, one common approach is to augment the dual problem's maximization problem (Eq. 3) and correspondingly the Lagrange equation with a negative f -divergence between d and $d^{\mathcal{D}}$ as $D_f(\frac{d(s, a)}{d^{\mathcal{D}}(s, a)}) = \mathbb{E}_{d^{\mathcal{D}}}[f(\frac{d(s, a)}{d^{\mathcal{D}}(s, a)})]$. However, this added regularization creates a disconnection with the primal problem. We will show that the constraint term can be naturally derived from the primal problem under a relaxed condition and is important in the optimization when using challenging datasets.

Our theoretical results also use the conjugate of the function [4]

$$g(y) = \sup_{x \in \text{dom } g^*} y * x - g^*(x). \quad (7)$$

If $g(x)$ is a convex function, the conjugate of its conjugate is itself. We can evaluate the optimal x value by taking the RHS of Eq. 7 and setting its partial derivative over x to 0.

$$x^* = g^{*'}^{-1}(y) \quad (8)$$

3 THE GENERAL LP FORMULATION FOR RL

Here, we provide a general view of RL in the format of Linear Programming to highlight the role of f -divergence in RL, either explicitly or implicitly. We first connect the traditional RL algorithms, whose optimization aims to minimize the Bellman residual through some loss function, to the LP formulation. Then, we draw the connection from the dual problem (Eq. 3) with the f -divergence constraint to a modified version of its primal problem (Eq. 1).

3.1 Alternative form of L_P

The original L_P in RL's LP form is defined as $\alpha(s)v(s)$ [24]. In Eq. 1, $\alpha(s) = (1 - \gamma)p_0(s)$. In a given MDP, p_0 is stationary and therefore $(1 - \gamma)p_0(s)$ remains constant. This suggests that additional options such as $(1 - \gamma)v(s)$ may also be valid for L_P .

A less obvious form for L_P is $\mathbb{E}_{(s, a) \sim \mathcal{D}}[-e_v(s, a)]$, or simplified as $-e_v(s, a)$ since the offline algorithm will sample from the dataset by definition. By comparing the value function of the dataset $V^{\mathcal{D}}(s)$ and $v(s)$, we can derive $e_v(s, a)$ and transform it to a valid L_P option using the performance difference lemma [1]:

$$V^{\mathcal{D}}(s) - v(s) = \frac{1}{1 - \gamma} \mathbb{E}_{(s, a) \sim \mathcal{D}}[e_v(s, a)] \quad (9)$$

$$(1 - \gamma)v(s) = (1 - \gamma)V^{\mathcal{D}}(s) - \mathbb{E}_{(s, a) \sim \mathcal{D}}[e_v(s, a)]. \quad (10)$$

Since the LHS of Eq. 10 is valid for L_P , its RHS is equivalently valid. While $V^{\mathcal{D}}(s)$ may be unknown, it is a fixed constant for a given static dataset. This provides us with a strong analysis tool for later Sections. From a vector perspective of $\alpha(s)$, this is also suggested in [24] by choosing $\alpha = (I - \gamma\mathcal{T})^T$.

3.2 Connecting Bellman minimization to LP

We start by rewriting the Lagrangian in Eq. 6 into a more compact form with the removal of $\zeta \geq 0$ and the inclusion of f -divergence using a valid function $g^*(\cdot)$:

$$\mathcal{L}(v, \zeta) = \sum_{s, a} \alpha(s)v(s) + \zeta(s, a)e_v(s, a) - g^*(\zeta(s, a)) \quad (11)$$

Previous studies [17, 27] show that the dual optimization of Eq. 11 as $\max_{\zeta} \min_v \mathcal{L}(v, \zeta)$ can be converted into an unconstrained optimization problem either using the convex conjugate or the solution to its Lagrange dual function as follows:

$$\max_{\zeta} \min_v \mathcal{L}(v, \zeta) = \min_v \alpha(s)v(s) + g(e_v(s, a)). \quad (12)$$

By substituting $\alpha(s)v(s)$ with $-e_v(s, a)$, we can derive the following optimization objective:

$$\min_v -e_v(s, a) + g(e_v(s, a)). \quad (13)$$

THEOREM 1. For a convex $g(\cdot)$ with $g^*(1) = 0$ and $g^{*'}(1) = 0$, $-x + g(x)$ is convex and $\min_x g(x) - x = 0$ when $x = 0$

PROOF. Because $g(\cdot)$ is convex, $g'(\cdot)$ is monotonic and therefore $g'(\cdot) - 1$ is monotonic. The minimum value of $-x + g(x)$ will be 0 at $g'(0) = 1$, from $g^{*'}(1) = 0$, and $x^* = 0$. \square

Theorem 1 suggests that the optimization objective in Eq. 13 is equivalent to minimizing v such that the Bellman error, i.e. e_v , is 0, with a careful selection of $g(\cdot)$ and its corresponding conjugate function $g^*(\cdot)$. This also suggests that any existing RL algorithm, as long as it uses a compatible loss function $g(\cdot)$ to minimizing the Bellman error, has an implicit f -divergence implication between the value function in learning and observed data. For example, the traditional MSE loss for minimizing the Bellman error in Q-learning is equivalent to setting $g(e_v) = \frac{1}{2}e_v^2 + e_v$, which is the conjugate of the χ^2 -divergence function.

A key distinction here compared against prior works is the removal of $\zeta \geq 0$. This removal corresponds to changing the inequality constraint in the primal problem constraint (Eq. 2) into the

equality constraint. Note that the constraint change will still yield a valid optimization problem because the optimal value function must satisfy the Bellman equation, which leads to the Bellman error being zero. Such a relaxation provides theoretical justification for prior works with relaxed positivity [13] and opens up additional possible function forms for f -divergence.

REMARK. We follow the definition of f -divergence with domain over \mathbb{R} [3, 32, 35], instead of $[0, +\infty)$. This allows negative probability measures and thus $\zeta < 0$. However, $\zeta < 0$ makes the importance sampling interpretation of the dual problem invalid in the traditional probability. Recent work about the likelihood ratio estimation in quasiprobability may provide some insight about its interpretation [5].

3.3 Unified form of LP for RL

Here, we show a unified form of LP for RL by connecting the dual problem with an f -divergence augmentation to its primal problem with a similar augmentation. We start by rewriting the dual problem in Eq. 3 and 4 using density ratio estimation $\zeta(s, a)$ ² and the standard L_P term with the inclusion of f -divergence:

$$\max_{\zeta} \zeta(s, a)r(s, a) - g^*(\zeta(s, a)) \quad (14)$$

$$\text{s.t. } (\mathcal{B}_* \zeta)(s) - \alpha(s) - \gamma(\mathcal{T}_* \zeta)(s) = 0, \forall s, a \quad (15)$$

The constraint is still the Bellman flow constraint, but rewritten in the format of a residual equal to zero. We define $e_{\zeta}(s, a) = (\mathcal{B}_* \zeta)(s) - \alpha(s) - \gamma(\mathcal{T}_* \zeta)(s)$ as the Bellman occupancy residual. Following Section 3.2, its Lagrangian is Eq. 11, and it can be transformed into the following unconstrained optimization problem with Eq. 13 being its special case:

$$\min_{\nu} \alpha(s)v(s) + g(e_{\nu}(s, a)). \quad (16)$$

We can also rewrite it into a singular unconstrained optimization problem³:

$$\max_{\zeta} \sum_{s,a} \zeta(s, a)r(s, a) - g^*(\zeta(s, a)) - \delta_0(e_{\zeta}(s, a)) \quad (17)$$

Now consider a modified primal problem with a similar augmentation of a convex function $\varphi(\nu)$ and conversion of the inequality constraint to the equality constraint:

$$\min_{\nu} \alpha(s)v(s) + \varphi(\nu(s)) \quad (18)$$

$$\text{s.t. } e_{\nu}(s, a) = 0, \forall s, a \quad (19)$$

²The following conversion is straightforward in the theoretical setting, i.e. $d^{\mathcal{D}}(s, a) = 1$. In the practical sampling setting, $(\mathcal{T}_* \zeta)(s) = \sum_{\bar{s}, \bar{a}} \mathcal{T}(s|\bar{s}, \bar{a})d^{\mathcal{D}}(\bar{s}, \bar{a})\zeta(\bar{s}, \bar{a}) = \mathbb{E}_{(s, \bar{a}) \sim \mathcal{D}}[\mathcal{T}(s|\bar{s}, \bar{a})\zeta(\bar{s}, \bar{a})]$. $(\mathcal{B}_* \zeta)(s) = \sum_a d^{\mathcal{D}}(s, a)\zeta(s, a) = \mathbb{E}_{a \sim \pi_{\beta}(\cdot|s)}[d^{\mathcal{D}}(s)\zeta(s, a)]$, which would transform to full sampling when aggregate over all state s , i.e. $\sum_s (\mathcal{B}_* \zeta)(s) = \mathbb{E}_{(s, a) \sim \mathcal{D}}[\zeta(s, a)]$.

³We define $\delta_0(x) = \begin{cases} 0 & \text{if } x = 0, \\ \infty & \text{otherwise} \end{cases}$.

The min-max optimization over its Lagrangian can be written as

$$\begin{aligned} \max_{\zeta} \min_{\nu} \mathcal{L}(\nu, \zeta) &= \max_{\zeta} \min_{\nu} \sum_{s,a} \alpha(s)v(s) + \zeta(s, a)r(s, a) \\ &\quad + \zeta(s, a) \cdot \gamma(\mathcal{T}\nu)(s, a) - \zeta(s, a) \cdot (\mathcal{B}\nu)(s, a) \\ &\quad + \varphi(\nu(s)) \end{aligned} \quad (20)$$

$$\begin{aligned} &= \min_{\zeta} \max_{\nu} \sum_{s,a} -\zeta(s, a)r(s, a) - \alpha(s)v(s) \\ &\quad - \nu(s) \cdot \gamma(\mathcal{T}_* \zeta)(s) + \nu(s) \cdot (\mathcal{B}_* \zeta)(s) \\ &\quad - \varphi(\nu(s)) \end{aligned} \quad (21)$$

$$\begin{aligned} &= \min_{\zeta} \sum_{s,a} -\zeta(s, a)r(s, a) \\ &\quad + \max_{\nu} [\nu(s) \cdot \underbrace{((\mathcal{B}_* \zeta)(s) - \alpha(s) - \gamma(\mathcal{T}_* \zeta)(s))}_{e_{\zeta}(s, a)} - \varphi(\nu(s))] \end{aligned} \quad (22)$$

$$= \min_{\zeta} \sum_{s,a} -\zeta(s, a)r(s, a) + \max_{\nu} [\nu(s) \cdot e_{\zeta}(s, a) - \varphi(\nu(s))] \quad (23)$$

$$= \min_{\zeta} \sum_{s,a} -\zeta(s, a)r(s, a) + \varphi^*(e_{\zeta}(s, a)) \quad (24)$$

$$= \max_{\zeta} \sum_{s,a} \zeta(s, a)r(s, a) - \varphi^*(e_{\zeta}(s, a)) \quad (25)$$

Eq. 21 changes the min-max optimization direction by negating the Lagrangian. Eq. 24 is the result of applying convex conjugate.

Eq. 25 shares a close similarity with Eq. 16. We can also rewrite the constrained optimization problem in Eq. 18 and 19 into a singular unconstrained optimization problem, similar to Eq. 17:

$$\min_{\nu} \sum_{s,a} \alpha(s)v(s) + \varphi(\nu(s)) + \delta_0(e_{\nu}(s, a)) \quad (26)$$

It is straightforward to spot the similarity between Eq. 16 and 26, as well as the similarity between Eq. 17 and 25. In fact, with either $g(\cdot) = \delta_0(\cdot)$ or $\varphi^*(\cdot) = \delta_0(\cdot)$, i.e., applying the equality constraint to either e_{ν} or e_{ζ} , we can derive the following equivalence:

$$\begin{aligned} &\min_{\nu} \sum_{s,a} \alpha(s)v(s) + \varphi(\nu(s)) + g(e_{\nu}(s, a)) \\ &\equiv \max_{\zeta} \sum_{s,a} \zeta(s, a)r(s, a) - \varphi^*(e_{\zeta}(s, a)) - g^*(\zeta(s, a)) \end{aligned} \quad (27)$$

From the convex conjugate perspective, the equivalence holds because the conjugate function of $\delta_0(\cdot)$ is $\delta_0^*(\cdot) = 0$. From the constrained optimization using Lagrangian, the equivalence holds because any valid solutions need to have $e_{\nu}(s, a) = 0$ or $e_{\zeta}(s, a) = 0$. By applying the equality constraint on $e_{\nu}(s, a)$, the primal problem is equivalent to an unconstrained version of the dual problem, with its equality constraint on $e_{\zeta}(s, a)$ changed into a penalty function characterized by $\varphi^*(\cdot)$. *Vice versa* for the dual problem. As a result, the selection of the penalty function, either $g(\cdot)$ or $\varphi^*(\cdot)$, and its conjugate function directly impact the optimization process.

4 FLEXIBLE f -DIVERGENCE

Theoretical results from the previous section show that the LP optimization of RL, either implicitly or explicitly, contains the f -divergence regularization by choosing the penalty function $g(\cdot)$ and its conjugate $g^*(\cdot)$. As a result, the selection of $g^*(\cdot)$ is crucial.

Previous theoretical analyses of offline RL algorithms' performance guarantees are based on assumptions about the "concentrability" of a dataset's coverage [1, 25, 38]. Formally, a concentrability coefficient C^π is defined as the smallest constant such that $\frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)} \leq C^\pi, \forall \pi \in \mathcal{F}^\pi, s \in S, a \in A$. The concrete family of the policy \mathcal{F}^π determines the concrete type of the concentrability coefficient. For example, setting it to be the family of π^* yields the π^* -concentrability. The concentrability coefficient characterizes the lower bound of $d^{\mathcal{D}}$ against the policy family. Under such an assumption, the traditional application of f -divergence in the optimization can be over-conservative since it places the same constraint over $d^{\mathcal{D}}(s,a)$ from a policy outside of the policy family of our interest, usually those producing higher returns. This is more likely with more fundamentally different behavior policies.

Here, we introduce the function form of the Flexible f -divergence that is generalized and flexible to apply different levels of penalties or constraints by selecting corresponding parameters and base functions. We express the function form in the context of the divergence function $g^*(\cdot)$, but we also provide concrete steps to compute the corresponding $g(\cdot)$ such that the Flexible f -divergence is applicable to the general form of RL algorithms with the Bellman error minimization.

4.1 Functional Form for Flexible f -divergence

We propose a general function format $g^*(\zeta; \alpha_-, \alpha_+, \beta)$, shorten for $g_{\alpha_\pm, \beta}^*(\zeta)$, as the function for the Flexible f -divergence by joining two scaled closed convex functions g_-^* and g_+^* at the threshold value β to constrain ζ with different levels and functions. We denote its conjugate function as $g_{\alpha_\pm, \beta}(\cdot)$. We use α_- and α_+ to scale g_-^* and g_+^* respectively. For general optimization, we would want the function to be continuous at β , i.e., $\alpha_- g_-^*(\beta) = \alpha_+ g_+^*(\beta)$. To ensure $g^*(\zeta; \alpha_-, \alpha_+, \beta)$'s conjugate is also a meaningful optimization objective, both g_-^* and g_+^* must be invertible. Lastly, its Fenchel conjugate needs to be continuous at β , which requires $\alpha_- g_-^*(\beta) = \alpha_+ g_+^*(\beta)$.

We start by picking two closed, strictly convex and differentiable functions \bar{g}_-^* and \bar{g}_+^* from the family of valid functions for f -divergence, i.e., $\mathcal{F} = \{f|f: \mathbb{R}^+ \rightarrow \mathbb{R}, f(1) = 0, f'(1) = 0\}$. The condition follows Theorem 1 to ensure the corresponding $-x + g_{\alpha_\pm, \beta}(x)$ is a full fledged loss function. The choice of functions guarantees their derivative functions to be monotonic and thus invertible. Based on these two base functions, we apply scaling and a linear difference:

$$g_{\alpha_\pm, \beta}^*(\zeta) = \begin{cases} \alpha_- \bar{g}_-^*(\zeta) - \mathbb{1}_{\beta < 1}(k_g \zeta + C_g) & \zeta < \beta \\ \alpha_+ \bar{g}_+^*(\zeta) + \mathbb{1}_{\beta \geq 1}(k_g \zeta + C_g) & \zeta \geq \beta. \end{cases} \quad (28)$$

Depending on whether $\beta < 1$ or not, we apply a linear difference to either \bar{g}_-^* or \bar{g}_+^* . The function without added linear difference ($g_\pm^*(\zeta) = \bar{g}_\pm^*(\zeta)$) will retain its function form scaled by its coefficient and still satisfy the requirement of being a function for f -divergence,

especially $g_\pm^*(1) = 0$.⁴ The function with added linear difference ($g_\pm^*(\zeta) = \bar{g}_\pm^*(\zeta) \pm (k_g \zeta + C_g)$) is not guaranteed to behave the same as the original function or even be a valid function by itself for f -divergence. Evaluating ζ over the support of $d^{\mathcal{D}}$, we have

$$\mathbb{E}_{d^{\mathcal{D}}} [g_\pm^* \left(\frac{d(s,a)}{d^{\mathcal{D}}(s,a)} \right) \pm (k_g \frac{d(s,a)}{d^{\mathcal{D}}(s,a)} + C_g)] \quad (29)$$

$$= \mathbb{E}_{d^{\mathcal{D}}} [g_\pm^* \left(\frac{d(s,a)}{d^{\mathcal{D}}(s,a)} \right)] \pm \mathbb{E}_{d^{\mathcal{D}}} [k_g \left(\frac{d(s,a)}{d^{\mathcal{D}}(s,a)} - 1 \right)] \pm \mathbb{E}_{d^{\mathcal{D}}} [C_g + k] \quad (30)$$

$$= D_{\bar{g}_\pm^*}(\zeta) \pm D_{f(x)=k_g(x-1)}(\zeta) \pm \mathbb{E}_{d^{\mathcal{D}}} [C_g + k] \quad (31)$$

$$= D_{g_\pm^*}(\zeta) \pm (C_g + k). \quad (32)$$

We can derive Eq. 32 from Eq. 31 by using f -divergence's property of $D_f(\frac{d}{d^{\mathcal{D}}}) = 0$ iff $f(x) = c(x-1)$. From Eq. 32, we argue that the function component with added difference can work equivalently as its original f -divergence plus a constant bias and does not impact its usage as a learning objective.

We pick the linear coefficient $k_g = \alpha_- \bar{g}_-^{\prime}(\beta) - \alpha_+ \bar{g}_+^{\prime}(\beta)$ to satisfy the condition of $\alpha_- \bar{g}_-^{\prime}(\beta) = \alpha_+ \bar{g}_+^{\prime}(\beta)$, as illustrated by evaluating the derivative Function 33 with $\zeta = \beta$. This introduces additional difference between $g_-^*(\beta)$ and $g_+^*(\beta)$ terms in Eq. 28. So we choose $C_g = \alpha_- \bar{g}_-^*(\beta) - \alpha_+ \bar{g}_+^*(\beta) - \beta k_g$ to remove the value difference.

For the complete dual optimization algorithms, for example, using dual gradient descent [4], $g^*(\zeta)$ in Eq. 28 would be enough. However, for algorithms such as OptiDICE that use intermediate e_v to estimate its current optimal ζ^* or algorithms directly optimizing Eq. 13, we derive the estimation of $\zeta^* = g_{\alpha_\pm, \beta}^{\prime -1}(e_v)$ by following Eq. 8 to get

$$g_{\alpha_\pm, \beta}^{\prime}(\zeta) = \begin{cases} \alpha_- \bar{g}_-^{\prime}(\zeta) - \mathbb{1}_{\beta < 1}(k_g) & \zeta < \beta \\ \alpha_+ \bar{g}_+^{\prime}(\zeta) + \mathbb{1}_{\beta \geq 1}(k_g) & \zeta \geq \beta \end{cases} \quad (33)$$

$$g_{\alpha_\pm, \beta}^{\prime -1}(e_v) = \begin{cases} \bar{g}_-^{\prime -1} \left(\frac{e_v + \mathbb{1}_{\beta < 1}(k_g)}{\alpha_-} \right) & e_v < \beta_{e_v} \\ \bar{g}_+^{\prime -1} \left(\frac{e_v - \mathbb{1}_{\beta \geq 1}(k_g)}{\alpha_+} \right) & e_v \geq \beta_{e_v} \end{cases}. \quad (34)$$

Since we use β as the threshold to separate the evaluation of ζ , correspondingly, there will be a threshold for e_v as the result of $g(e_v)$ and $g^*(\zeta)$ being the conjugate of each other. We compute the threshold value for e_v as β_{e_v} by evaluating $g^{\prime}(\zeta = \beta)$ as $\beta_{e_v} = \mathbb{1}_{\beta \geq 1}(\alpha_- \bar{g}_-^{\prime}(\beta)) + \mathbb{1}_{\beta < 1}(\alpha_+ \bar{g}_+^{\prime}(\beta))$. For algorithms directly optimizing Eq. 13, We can then acquire the optimization objective of e_v by substituting the variable ζ in the convex conjugate with ζ^* as $L(v) = -e_v + \zeta^* e_v - g_{\alpha_\pm, \beta}^*(\zeta^*)$, with the possible addition of other terms depending on the concrete algorithm.

4.2 Heuristic estimation of α_\pm and β

The native form of $g_{\alpha_\pm, \beta}^*(\zeta)$ relies on the selection of α_\pm, β , and $g_\pm^*(\cdot)$. While the fully automatic selection algorithm for these components is not the main contribution of our paper, we provide heuristic methods to estimate values for α_\pm and β using the learning agent's predicted values from the offline dataset.

Estimating α_\pm . Inspired by Han et al. [10], we train an additional behavior cloning policy π_b from the offline dataset and compare its predicted action probability against learned $\exp(e_v)$. Specifically, we predict the action probability $\pi_b(a|s)$ and $\exp(e_v(s,a))$ for each

⁴We use \bar{g}_\pm^* to indicate either \bar{g}_-^* or \bar{g}_+^* .

state-action pair in the sampled batch. We treat the list of computed probability and the list of $\exp(e_v)$ as two vectors, i.e. $\vec{\pi}_b$ and $\vec{\exp}(e_v)$, and compute α_{\pm} as follows:

$$\alpha_+ = \frac{1}{\cos\langle \vec{\pi}_b, \vec{\exp}(e_v) \rangle} \quad (35)$$

$$\alpha_- = \frac{1}{1 - \cos\langle \vec{\pi}_b, \vec{\exp}(e_v) \rangle}, \quad (36)$$

where $\cos\langle \cdot, \cdot \rangle$ computes the cosine similarity. We use an exponential moving average (EMA) to smooth the estimation throughout training.

Estimating β . Because $g(e_v)$ and $g^*(\zeta)$ are conjugate functions to each other, e_v and ζ has a direct correspondence. We compute the average e_v as \bar{e}_v from sampled batches and smooth it with EMA. We then compute the beta value by $\beta = g_{\alpha_{\pm}, 0}^{*-1}(\bar{e}_v)$.

4.3 Effect of base function, α_{\pm} and β

Table 1 provides some examples of existing f -divergence approaches, with the adjustment to ensure $g^*(1) = 0$, and their corresponding functions. Figure 2 shows graphs of selected flexible f -divergence, in terms of the divergence function $g^*(\cdot)$, the prime inverse function $g^{*-1}(\cdot)$, and conjugate function $g(\cdot)$, as well as the $-e_v + g(e_v)$ as the loss function in the traditional RL. We also illustrate their estimation effect using a simple two-dimensional dataset with Gaussian noise injected into the y -dimension. The top row of Figure 2 compares the effect of different base functions. While all of $g^*(\cdot)$ are convex with the minimum value at $g^*(1) = 0$ and all of $-e_v + g(e_v)$ are convex with the minimum value at $-0 + g(0) = 0$, their derivatives are vastly different and thus have different impact in the optimization effect, as illustrated in the estimation effect. For example, using the Hellinger function as g^* will approximate the maximum value. The bottom row of Figure 2 shows the effect of α_{\pm} and β . A high-level interpretation of the impact from the perspective of the resulting $-e_v + g(e_v)$ is that the larger the magnitude of the derivative is at the positive or the negative side, the more penalty the positive or negative Bellman error e_v will receive. However, such a difference cannot be captured by the coefficient α_{\pm} alone since the derivative is often non-linear.

5 RELATED WORKS

LP and DRE for offline RL LP algorithm for RL relies on the concept of density ratio estimation (DRE), which is ζ , to provide estimation or correction for the density ratio between d and d^D . DualDICE [21] applied the Fenchel Duality to derive the LP format of RL for the MDP with 0 reward values for policy evaluation. AlgaeDICE [22] and GenDICE [39] derived the LP between Q and π for policy improvement. OptiDICE [17] aimed to resolve the instability from the dual optimization by solving the closed-form solution for ζ . ODICE [19] inspected the gradient of LP. f -DVL [27] extended the framework of LP and proved that several offline RL algorithms, such as CQL [16], SQL [36], and XQL [8] are special cases of LP for RL. Here, we further extend the LP framework by identifying $-e_v$ as L_P and relaxing the constraint of $\zeta \geq 0$.

f -divergence and relaxed regularization Some functions for the f -divergence in prior works are special cases of Ada- f . The $f_{\text{soft-}\chi^2}$ function in OptiDICE [17] uses the α -divergence function

as \bar{g}_- and the χ^2 -divergence function as \bar{g}_+ , with $\alpha_{\pm} = \beta = 1$. Similarly, RelaxDICE [37] adds a relaxed f divergence upon the base KL-divergence, resulting in the special case with both base functions being KL-divergence and $\alpha_- = 1, \alpha_+ = 2$. Indirectly, the expectile regression for updating $V(s)$ in IQL [15] can be considered as optimizing Eq. 13 directly with g and g^* being χ^2 . For a selected expectile value τ , we can use $\alpha_- = \frac{1}{1-\tau}$ and $\alpha_+ = \frac{1}{\tau}$ to achieve the same effect. PORelDICE [13] relaxes the regularization by setting $\zeta \geq \epsilon$, which shares a similar effect by using Le-Cam divergence for \bar{g}_- and set $\alpha_- = \epsilon$.

6 EXPERIMENTS AND ANALYSIS

Our experiments serve two main objectives. The first one is to verify whether the proposed general constrained LP format for RL is a valid RL algorithm. Our second objective is to examine the effect of flexible f -divergence in challenging learning settings to see whether it can improve the performance over its base algorithm. Note that the main focus of our experiments is to inspect flexible f -divergence's effect by substituting the base algorithm's default f -divergence function, either in the explicit form of $g^*(\zeta(s, a))$ or the implicit form of $g(e_v(s, a))$. It is not to compare against algorithms with a completely different optimization procedure, though the flexible f -divergence could substitute those algorithms' f -divergence function as well.

We devised two algorithms following the theoretical framework as *Flex-f-Q* and *Flex-f-DICE*. *Flex-f-Q* approximates e_v as $Q_{\phi} - v_{\theta}$, optimizes Eq. 13 directly to update v_{θ} through semi-gradient descent, and use $-e_v$ as L_P . It is conceptually similar to IQL [15] but with *Flex-f* as its divergence. *Flex-f-DICE* is the OptiDICE algorithm with its divergence substituted with *Flex-f*. Close performance between *Flex-f-Q* and IQL can empirically show that the proposed general constrained LP format is a valid RL learning algorithm, with the swapped L_P and removal of $\zeta \geq 0$. Comparison between both pairs of algorithms can show the impact of *Flex-f*.

We focus on environments with continuous action space settings, including continuous control and goal-oriented tasks. For continuous control tasks, we conducted experiments on MuJoCo [33] environments: Hopper-v4, Walker2d-v4, Ant-v4, and HalfCheetah-v4. The main objective of these environments is to control various types of simulated robots to move forward as far as possible without falling into an unhealthy state. For goal-oriented tasks, we adopted Fetch environments [23], including Push-v2 and PickAndPlace-v2. These environments require the RL agent to control a robot arm and move the target object to the specified goal location. We also included Pen-v1 and Hammer-v1 from ArdoitHand with *cloned* and *human* datasets from D4RL [6].

We collected new datasets for the included MuJoCo and Fetch environments. Using the Soft Actor-Critic (SAC) algorithm [9], we trained the behavior policy following D4RL's data collection procedure except for the policy's variance. For goal-oriented environments in Fetch, we used Hindsight Experience Replay (HER) [2] with the future resampling strategy and resampling size $k = 4$. We limited the explorative behavior of the policy during data collection by fixing the variance of the policy distribution to be 0.0 as fully deterministic policies. We constructed datasets by mixing

| f -divergence | $g^*(\zeta)$ | $g(e_\nu)$ | $g'(\zeta)$ | $g'^{-1}(e_\nu)$ |
|-----------------|--|----------------------------------|---|-----------------------------------|
| χ^2 | $\frac{1}{2}(\zeta - 1)^2$ | $\frac{1}{2}e_\nu^2 + e_\nu$ | $\zeta - 1$ | $e_\nu + 1$ |
| KL-divergence | $\zeta \ln \zeta - \zeta + 1$ | $\exp(e_\nu) - 1$ | $\ln \zeta$ | $\exp(e_\nu)$ |
| Reverse-KL | $-\ln \zeta + \zeta - 1$ | $-\ln(1 - e_\nu)$ | $-\frac{1}{\zeta} + 1$ | $\frac{1}{1 - e_\nu}$ |
| Hellinger | $\frac{1}{2}(\sqrt{\zeta} - 1)^2$ | $\frac{e_\nu}{1 - 2e_\nu}$ | $\frac{\sqrt{\zeta} - 1}{\sqrt{\zeta}} * \frac{1}{2}$ | $\frac{1}{1 - 2e_\nu}^2$ |
| Le-Cam | $\frac{1 - \zeta}{2(\zeta + 1)} + \frac{\zeta - 1}{4}$ | $-\sqrt{1 - 4e_\nu} - e_\nu + 1$ | $-\frac{1}{(\zeta + 1)^2} + \frac{1}{4}$ | $\sqrt{\frac{4}{1 - 4e_\nu}} - 1$ |

Table 1: Example divergence functions g^* , conjugate g , derivative, and inverse of derivative, adjusted to ensure $g'^{-1}(1) = 0$.

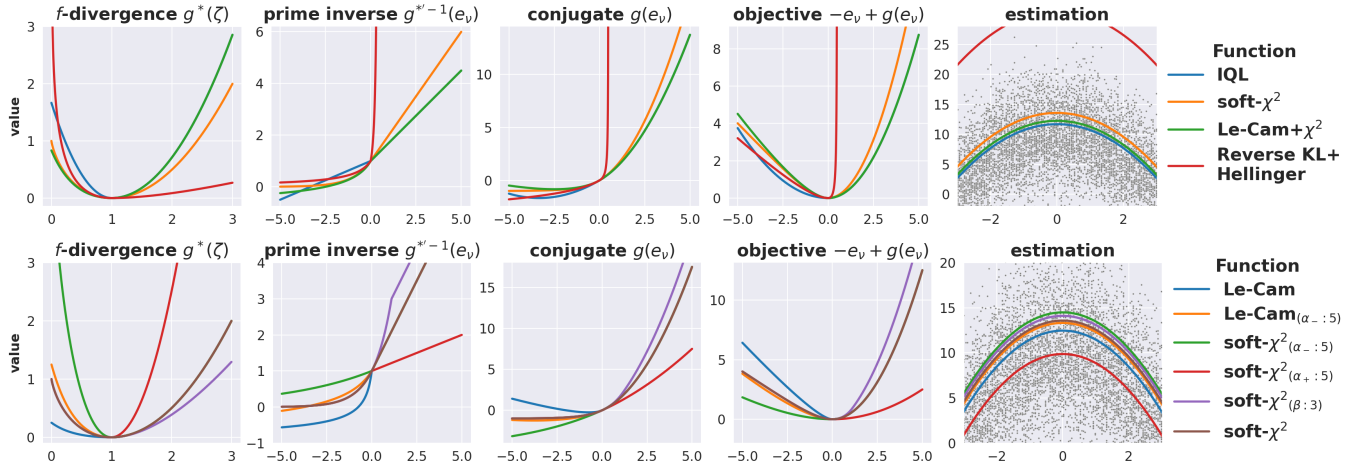


Figure 2: (Top) Example f divergence function. The function for IQL is χ^2 with $\alpha_- = \frac{10}{3}$ and $\alpha_+ = \frac{10}{7}$, corresponding to the 70%-expectile regression. Le-Cam+ χ^2 shares the same coefficient. (Bottom) Illustration of α_- , α_+ , and β 's effect. Default values are 1.0. All functions use χ^2 for as $g^*_+(\cdot)$.

data collected from 2, 4, and 10 behavior policies at different expertise levels, referred to as 2- p , 4- p , and 10- p , respectively. Please see the supplementary material for details about the data collection procedure.

Flex-f-DICE uses χ^2 as g^*_+ and KL as g^*_- in MuJoCo environments. *Flex-f-Q* uses KL as g^*_+ and χ^2 as g^*_- in MuJoCo environments. g^*_+ is χ^2 and g^*_- is Le-Cam for all other environments. Other hyperparameters between the two pairs of compared algorithms, IQL vs. *Flex-f-Q* and OptiDICE vs. *Flex-f-DICE*, are the same, respectively. We trained each algorithm in each setting for five different seeds and report the average normalized returns.

Table 2 shows each algorithm’s average return for each environment and dataset. We highlight the entry if it is higher than the compared counterpart, and underline the entry if the compared algorithms achieved similar performance ($\Delta \leq 5$). *Flex-f-Q*’s performance is overall close to IQL, empirically showing the correctness of the proposed general LP form as the Bellman minimization algorithm. In some cases, the performance of *Flex-f-Q* is higher than IQL. *Flex-f-DICE* can almost always achieve improved performance over OptiDICE. This suggests the potential of *Flex-f* in increasing its base algorithm’s performance. *Flex-f-DICE*’s higher performance in Mujoco environments, where it shares the same soft- χ^2 base function as OptiDICE, verifies the effect of heuristically estimated

α_\pm and β . We also include the result from TD3BC [7] and CQL [16] in the supplementary material for comparison.

Figure 3 shows an example of how α_\pm and β change throughout the training process. α_\pm can change drastically over the course of training. The decrease of α_+ suggests the estimated e_ν is more aligning with the behavior in the dataset and therefore leading to lower regularization. α_+ ’s increase means the estimated e_ν diverges away from dataset’s behavior, leading to higher regularization. Due to the numeric stability, we limited the e_ν value for estimation β to be in $[-0.2, 0.15]$. But different function approximations and gradient update format still lead to different estimated β value.

We conducted an ablation study to investigate the effect of base functions using Hopper and Walker2D environments with 4- p datasets. For $g^*_+(\cdot)$, we considered χ^2 and KL; for $g^*_-(\cdot)$, we considered χ^2 , KL, Le-Cam, and Hellinger. Table 3 shows both *Flex-f-Q* and *Flex-f-DICE*’s performance. From the high level, Table 3 shows that there is unlikely a universal function form that works well for every environments because the combination achieving the highest performance is different for each environment and algorithm, though Hellinger- χ^2 appears to be relatively consistent for *Flex-f-DICE*. This supports our hypothesis that different environments and algorithms call for a different level of constraint and correspondingly different f -divergence as optimization objective.

Table 2: Min-Max Normalized Returns

| env | dataset | IQL | Flex- f -Q | Opti-DICE | Flex- f -DICE |
|--------------|---------|--------------|--------------|--------------|-----------------|
| hopper | 4-p | 62.8 | 76.4 | 77.9 | 99.0 |
| | 2-p | 100.7 | 106.0 | 45.2 | 40.8 |
| | 10-p | 63.5 | 68.5 | 89.9 | 94.5 |
| walker | 4-p | <u>50.3</u> | <u>50.7</u> | 75.7 | 94.2 |
| | 2-p | <u>85.6</u> | <u>88.3</u> | <u>130.6</u> | <u>130.3</u> |
| | 10-p | <u>101.2</u> | <u>102.0</u> | 79.8 | 86.3 |
| ant | 4-p | 113.8 | 123.1 | 126.0 | 136.0 |
| | 2-p | 120.3 | 146.1 | 126.6 | 129.3 |
| | 10-p | <u>92.7</u> | <u>91.1</u> | 74.9 | 81.5 |
| half-cheetah | 4-p | <u>52.3</u> | <u>54.7</u> | 15.9 | 47.0 |
| | 2-p | <u>44.9</u> | <u>42.6</u> | 13.4 | 40.1 |
| | 10-p | <u>78.2</u> | <u>80.3</u> | 80.5 | 79.4 |
| push | 4-p | <u>91.6</u> | <u>92.4</u> | 81.9 | 88.3 |
| | 2-p | <u>95.2</u> | <u>95.2</u> | 87.4 | 93.0 |
| | 10-p | <u>84.1</u> | <u>81.3</u> | <u>69.5</u> | <u>68.4</u> |
| pick & place | 4-p | <u>95.7</u> | <u>97.0</u> | 87.2 | 90.3 |
| | 2-p | <u>92.0</u> | <u>92.7</u> | <u>91.6</u> | 93.2 |
| | 10-p | <u>45.2</u> | <u>45.3</u> | 26.6 | 37.9 |
| pen | cloned | 56.2 | 58.9 | 22.9 | 40.4 |
| | human | 52.9 | 53.4 | 10.2 | 39.6 |
| hammer | cloned | 2.2 | 0.9 | <u>0.2</u> | 1.3 |
| | human | <u>1.0</u> | <u>1.9</u> | 3.4 | 0.9 |

Table 3: Ablation study of base functions using 4-p

| env | Algorithm | | Flex- f -Q | Flex- f -DICE | |
|--------|-----------|-----------|--------------|-----------------|-------|
| | g_+ | g_- | | | |
| hopper | χ^2 | χ^2 | 54.7 | 75.4 | |
| | | KL | 71.9 | 99.0 | |
| | | Le-Cam | 49.2 | 94.4 | |
| | | Hellinger | 55.2 | 100.6 | |
| | | KL | χ^2 | 76.4 | 81.9 |
| walker | χ^2 | KL | 80.8 | 72.3 | |
| | | Le-Cam | 58.1 | 74.6 | |
| | | Hellinger | 54.8 | 88.6 | |
| | | KL | χ^2 | 45.4 | 97.6 |
| | | KL | χ^2 | 55.0 | 94.2 |
| walker | χ^2 | Le-Cam | 46.0 | 97.8 | |
| | | Hellinger | 48.7 | 110.0 | |
| | | KL | χ^2 | 50.7 | 101.8 |
| | | KL | - | - | 95.1 |
| | | Le-Cam | 60.2 | 96.9 | |
| | | Hellinger | 45.0 | 91.7 | |

alleviate the risk of Q value overestimation in the semi-gradient style optimization like IQL or Flex- f -Q.

7 CONCLUSION

In this paper, we identify the challenge of offline RL using datasets with limited stochasticity and a high mixture of expertise levels. We propose a general form of LP with expanded options for L_P and the removed constraint of $\zeta \geq 0$. We show that such an LP form for RL yields an unconstrained optimization problem that is equivalent to any RL algorithms relying on the Bellman error minimization. We also show a unified form of LP for RL, under which the existing dual optimization problem can be viewed as minimizing the sum of an initial loss term and a penalty function, under the Bellman constraint on either learned value function or density ratio estimation. We further derived the general function form of the flexible f -divergence to provide alternative flexible constraints on the dataset support. Such a flexible f -divergence function form opens up future research opportunities of constructing a fully automated adaption algorithm to apply an adaptive constraint based on received data. Experiment results empirically showed the correctness of the proposed LP formulation and the potential of flexible f -divergence in improving algorithm performance with identified challenging datasets.

We note that using flexible f -divergence can introduce new hyperparameters, and deployed heuristic methods for adaptively estimating them during training. Our future work will aim to incorporate the optimization of α_{\pm} and β directly into optimization algorithms, thereby converting the flexible f -divergence into a fully optimizable format.

REFERENCES

[1] Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. 2019. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep 32* (2019), 96.

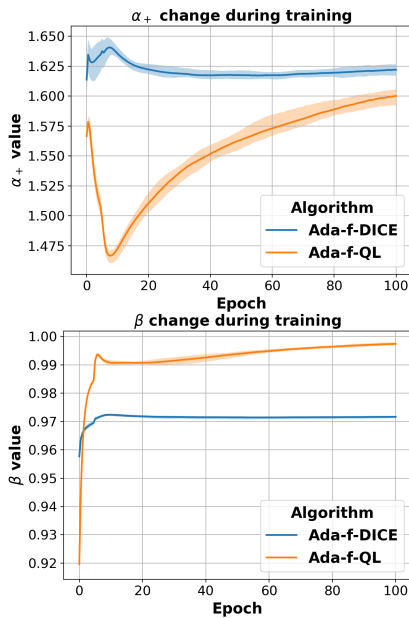


Figure 3: α_{\pm} and β changes throughout training.

An interesting observation from Table 3 is that using KL as $g_+^*(\cdot)$ for Flex- f -Q only exploded when using KL as $g_-^*(\cdot)$, which is why there is no entry, but was still able to function using other functions as $g_-^*(\cdot)$. This suggests that a different $g_-^*(\cdot)$ can effectively

- [2] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *Advances in neural information processing systems* 30 (2017).
- [3] Jonathan M Borwein and Adrian S Lewis. 1993. Partially-finite programming in L_1 and the existence of maximum entropy estimates. *SIAM Journal on Optimization* 3, 2 (1993), 248–267.
- [4] Stephen P Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge university press.
- [5] Matthew Drnevich, Stephen Jiggins, Judith Katzy, and Kyle Cranmer. 2024. Neural quasiprobabilistic likelihood ratio estimation with negatively weighted data. *arXiv preprint arXiv:2410.10216* (2024).
- [6] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219* (2020).
- [7] Scott Fujimoto and Shixiang Shane Gu. 2021. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems* 34 (2021), 20132–20145.
- [8] Divyansh Garg, Joey Hejna, Matthieu Geist, and Stefano Ermon. 2023. Extreme q-learning: Maxent rl without entropy. *arXiv preprint arXiv:2301.02328* (2023).
- [9] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [10] Kinchen Han, Hossam Affi, and Michel Marot. 2025. Cosine Similarity Based Adaptive Implicit Q-Learning for Offline Reinforcement Learning. In *2025 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
- [11] Zhang-Wei Hong, Pulkit Agrawal, Rémi Tachet des Combes, and Romain Laroche. 2023. Harnessing mixed offline reinforcement learning datasets via trajectory weighting. *arXiv preprint arXiv:2306.13085* (2023).
- [12] Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine. 2021. How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research* 40, 4-5 (2021), 698–721.
- [13] Woosung Kim, Donghyeon Ki, and Byung-Jun Lee. 2024. Relaxed stationary distribution correction estimation for improved offline policy optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 13185–13192.
- [14] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE transactions on intelligent transportation systems* 23, 6 (2021), 4909–4926.
- [15] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169* (2021).
- [16] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative q-learning for offline reinforcement learning. *Advances in neural information processing systems* 33 (2020), 1179–1191.
- [17] Jongmin Lee, Wonseok Jeon, Byungjun Lee, Joelle Pineau, and Kee-Eung Kim. 2021. Optidice: Offline policy optimization via stationary distribution correction estimation. In *International Conference on Machine Learning*. PMLR, 6120–6130.
- [18] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* (2020).
- [19] Liyuan Mao, Haoran Xu, Weinan Zhang, and Xianyuan Zhan. 2024. Odice: Revealing the mystery of distribution correction estimation via orthogonal-gradient update. *arXiv preprint arXiv:2402.00348* (2024).
- [20] Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. 2001. Marginal mean models for dynamic regimes. *J. Amer. Statist. Assoc.* 96, 456 (2001), 1410–1423.
- [21] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. 2019. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. *Advances in neural information processing systems* 32 (2019).
- [22] Ofir Nachum, Bo Dai, Ilya Kostrikov, Yinlam Chow, Lihong Li, and Dale Schuurmans. 2019. Algaedice: Policy gradient from arbitrary experience. *arXiv preprint arXiv:1912.02074* (2019).
- [23] Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464* (2018).
- [24] Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- [25] Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. 2021. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Advances in Neural Information Processing Systems* 34 (2021), 11702–11716.
- [26] Kajetan Schweighofer, Andreas Radler, Marius-Constantin Dinu, Markus Hofmarcher, Vihang Patil, Angela Bitto-Nemling, Hamid Eghbal-zadeh, and Sepp Hochreiter. 2021. A Dataset Perspective on Offline Reinforcement Learning. *arXiv preprint arXiv:2111.04714* (2021).
- [27] Harshit Sikchi, Qingqing Zheng, Amy Zhang, and Scott Niekum. 2023. Dual rl: Unification and new methods for reinforcement and imitation learning. *arXiv preprint arXiv:2302.08560* (2023).
- [28] Anikait Singh, Aviral Kumar, Quan Vuong, Yevgen Chebotar, and Sergey Levine. 2022. Offline rl with realistic datasets: Heteroskedasticity and support constraints. *arXiv preprint arXiv:2211.01052* (2022).
- [29] Shashi Kant Singh, Shubham Kumar, and Pawan Singh Mehra. 2023. Chat gpt & google bard ai: A review. In *2023 International Conference on IoT, Communication and Automation Technology (ICICAT)*. IEEE, 1–6.
- [30] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [31] Phillip Swazinna, Steffen Udluft, and Thomas Runkler. 2021. Measuring data quality for dataset selection in offline reinforcement learning. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 1–8.
- [32] Dávid Terjék. 2021. Moreau-Yosida f -divergences. In *International Conference on Machine Learning*. PMLR, 10214–10224.
- [33] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033. <https://doi.org/10.1109/IROS.2012.6386109>
- [34] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. 2023. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*. PMLR, 1723–1736.
- [35] Haoming Wang and Lek-Heng Lim. 2025. Glivenko-Cantelli for f -divergence. *arXiv preprint arXiv:2503.17355* (2025).
- [36] Haoran Xu, Li Jiang, Jianxiong Li, Zhuoran Yang, Zhaoran Wang, Victor Wai Kin Chan, and Xianyuan Zhan. 2023. Offline rl with no ood actions: In-sample learning via implicit value regularization. *arXiv preprint arXiv:2303.15810* (2023).
- [37] Lantao Yu, Tianhe Yu, Jiaming Song, Willie Neiswanger, and Stefano Ermon. 2023. Offline imitation learning with suboptimal demonstrations via relaxed distribution matching. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 11016–11024.
- [38] Wenhao Zhan, Baihe Huang, Audrey Huang, Nan Jiang, and Jason Lee. 2022. Offline reinforcement learning with realizability and single-policy concentrability. In *Conference on Learning Theory*. PMLR, 2730–2775.
- [39] Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. 2020. Gendice: Generalized offline estimation of stationary values. *arXiv preprint arXiv:2002.09072* (2020).