
Offline Minimax Soft-Q-learning Under Realizability and Partial Coverage

Masatoshi Uehara*

Genentech

uehara.masatoshi@gene.com

Nathan Kallus

Cornell University

kallus@cornell.edu

Jason D. Lee

Princeton University

jasonlee@princeton.edu

Wen Sun

Cornell University

ws455@cornell.edu

Abstract

In offline RL, we have no opportunity to explore so we must make assumptions that the data is sufficient to guide picking a good policy, and we want to make these assumptions as harmless as possible. In this work, we propose value-based algorithms for offline RL with PAC guarantees under just partial coverage, specifically, coverage of just a single comparator policy, and realizability of the soft (entropy-regularized) Q-function of the single policy and a related function defined as a saddle point of certain minimax optimization problem. This offers refined and generally more lax conditions for offline RL. We further show an analogous result for vanilla Q-functions under a soft margin condition. To attain these guarantees, we leverage novel minimax learning algorithms and analyses to accurately estimate either soft or vanilla Q-functions with strong L^2 -convergence guarantees. Our algorithms' loss functions arise from casting the estimation problems as nonlinear convex optimization problems and Lagrangifying. Surprisingly we handle partial coverage even without explicitly enforcing pessimism.

1 Introduction

In offline Reinforcement Learning (RL), we must learn exclusively from offline data and are unable to actively interact with the environment (Levine et al., 2020). Offline RL has garnered considerable interest in a range of applications where experimentation may be prohibitively costly or risky.

Offline RL is generally based on two types of assumptions: sufficient coverage in the offline data and sufficient function approximation. For instance, classical Fitted-Q-iteration (Antos et al., 2008; Chen and Jiang, 2019) requires (a) full coverage in the offline data, $\max_{(s,a)} d_{\pi,\mu_0}(s,a)/P_{\pi_b}(s,a) < \infty$ for any policy π where $P_{\pi_b}(s,a)$ is the offline data's distribution on the states and actions and $d_{\pi,\mu_0}(s,a)$ is the state-action occupancy distribution under a policy π and initial-state distribution $\mu_0(s)$; (b) realizability of the Q^* -function in a hypothesis class; and (c) Bellman completeness, *i.e.*, the Bellman operator applied to any function in the hypothesis class remains in the class. Full coverage (a) and Bellman completeness (c) can be particularly stringent because offline data is often insufficiently exploratory and Bellman completeness significantly restricts transition dynamics.

To overcome these challenges, we here propose algorithms with guarantees under realizability of single functions and refined partial coverage of single policies, and without Bellman completeness. We tackle this by introducing two novel value-based algorithms. The first algorithm, MSQP (mimimax soft-Q-learning with penalization), comprises of two steps: learning soft Q-functions (a.k.a., entropy-

*This work is done when he was at Cornell University

Table 1: Summary of partial-coverage-type guarantees with model-free general function approximation. Here, $w^* := d_{\pi^*, \mu_0} / P_b$ where d_{π^*, μ_0} is the occupancy distribution under the optimal policy π^* starting from μ_0 and P_b is the distribution over the offline data. A function \tilde{w}_α^* is a regularized marginal density ratio that satisfies $\tilde{w}_0^* = w^*$. Functions q^*, q_α^*, q_π are the optimal Q^* -function, the soft Q-function, and the Q-function under a policy π , respectively. Functions v_α^*, l_α^* are Lagrange multipliers of specific minimax optimization problems. The operator \mathcal{B}^π is a Bellman operator under a policy π . Function classes $\mathcal{W}, \mathcal{Q}, \mathcal{L}, \mathcal{V}$ consist of functions that map states (and actions) to real numbers. Note the guarantees provided by Jiang and Huang (2020); Xie et al. (2021) are more general than the below in that the output policy can compete with any policy in the policy class Π . For simplicity, we set the comparator policy to be the optimal policy π^* in this table. Note that other studies (Ozdaglar et al., 2023; Rashidinejad et al., 2022; Zhu et al., 2023) proposing model-free general function approximation under partial coverage rely on the completeness-type assumption as in (Xie et al., 2021) or realizability for any π as in Jiang and Huang (2020).

Primary Assumptions	
Jiang and Huang (2020)	$w^* \in \mathcal{W}, q_\pi \in \mathcal{Q} \forall \pi \in \Pi$
Xie et al. (2021)	$q_\pi \in \mathcal{Q}, \mathcal{B}^\pi \mathcal{Q} \subset \mathcal{Q} \forall \pi \in \Pi$
Zhan et al. (2022)	$\tilde{w}_\alpha^* \in \mathcal{W}, v_\alpha^* \in \mathcal{V}$
MSQP	$q_\alpha^* \in \mathcal{Q}, l_\alpha^* \in \mathcal{L}$
Chen and Jiang (2022)	Hard margin, $w^* \in \mathcal{W}, q^* \in \mathcal{Q}$
MQP	Soft margin, $q^* \in \mathcal{Q}, l^* \in \mathcal{L}$

regularized Q-functions, as defined in Fox et al., 2015; Schulman et al., 2017) from offline data, and using the softmax policies of the learned soft Q-functions. The second algorithm, MQP (mimimax Q-learning with penalization), consists of two steps: learning standard Q-functions from offline data and employing the greedy policy of the learned Q function on the offline data.

Using the above-mentioned two algorithms, we attain PAC guarantees under partial coverage and realizability, yet without Bellman completeness. In particular, in MSQP using soft Q-functions, we ensure strong performance under the realizability of q_α^*, l_α^* and the (density-ratio-based) partial coverage $\max_{(s,a)} d_{\pi_\alpha^*, \mu_0}(s, a) / P_b(s, a) < \infty$. Here q_α^* is a soft Q-function, l_α^* is a function that possesses a certain dual relation to q_α^* , π_α^* is the soft-max optimal policy, and α is the temperature parameter for the entropy-regularization. Notably, $\max_{(s,a)} d_{\pi_\alpha^*, \mu_0}(s, a) / P_b(s, a) < \infty$ is significantly less stringent than the uniform coverage in that the coverage is only imposed against a policy π_α^* . In MQP using Q-functions, we similarly ensure strong performance under a soft margin, the realizability of q^*, l^* , and the partial coverage $\max_{(s,a)} d_{\pi^*, \mu_0}(s, a) / P_b(s, a) < \infty$. Here q^* is the vanilla Q-function and l^* is a function that possesses a certain dual relation to q^* , and π^* is the usual optimal policy. Note the soft margin is introduced to allow realizability on standard Q-functions rather than soft Q-functions. However, the conditions $\max_{(s,a)} d_{\pi_\alpha^*, \mu_0}(s, a) / P_b(s, a) < \infty$ or $\max_{(s,a)} d_{\pi^*, \mu_0}(s, a) / P_b(s, a) < \infty$ may still be strong as these marginal density ratios may not exist in large-scale MDPs. For example, this condition is easily violated when the initial distribution μ_0 is not covered by P_b (i.e., $\max_s \mu_0(s) / P_b(s) = \infty$ where $P_b(s) := \sum_a P_b(s, a)$). Therefore, as an additional innovation, in our algorithms we can further relax these density-ratio-based partial coverage conditions. Specifically, we can demonstrate results under a *refined partial coverage*, which is adaptive to Q-function classes, even when the initial distribution μ_0 is not covered by P_b .²

The primary challenge lies in the design of loss functions for effectively learning soft Q-functions and vanilla Q-functions from offline data *without Bellman completeness*. To tackle this, we devise new minimax loss functions with certain regularization terms to achieve favorable L^2 -convergence rates on the offline data (i.e., in terms of $\mathbb{E}_{(s,a) \sim P_b} [\{\hat{q}_\alpha - q\}^2(s, a)]$ given an estimator \hat{q}). This result serves as the key building block for obtaining refined partial coverage under realizability and is of independent interest in its own right. Existing results are often constrained to specific models, such as linear models (Shi et al., 2022a), or they require Bellman completeness (Antos et al., 2008; Chen and Qi, 2022; Chen and Jiang, 2019). In contrast, our guarantee is applicable to any function

²Note μ_0 and P_b could be generally different even in the contextual bandit setting. This important setting is often considered in the literature on external validity/transportability in causal inference, as results of randomized clinical trials cannot be directly transported because covariate distributions in offline data and target data are different (Cole and Stuart, 2010; Pearl and Bareinboim, 2014; Dahabreh et al., 2019).

approximation method, without the need for Bellman completeness. To the best of our knowledge, this is the first guarantee of its kind.

Our work exhibits marked improvements over two closely related studies (Zhan et al., 2022; Chen and Jiang, 2022). Similar to our work, they propose algorithms that operate under the realizability of specific functions and partial coverage, yet without Bellman completeness. Zhan et al. (2022) ensures a PAC guarantee under (a’) partial coverage in the offline data $\max_{(s,a)} d_{\tilde{\pi}_\alpha^*, \mu_0}(s, a) / P_b(s, a) < \infty$ where $\tilde{\pi}_\alpha^*$ is a specific near-optimal policy under the regularization, which differs from the soft optimal policy, and (b’) realizability of $d_{\tilde{\pi}_\alpha^*, \mu_0} / P_b$ and the regularized value function. However, unlike MSQP, it is unclear how to refine the abovementioned coverage, *i.e.*, the guarantee could be vacuous when the initial distribution is not covered by offline data. A similar guarantee, but without regularization, is obtained under the additional hard margin (a.k.a., gap) condition in Chen and Jiang (2022). Our soft margin is a strict relaxation of the hard margin, which is important because, unlike the soft margin, the hard margin generally does not hold in continuous state spaces and involves very large constants in discrete state spaces. Lastly, although Chen and Jiang (2022); Zhan et al. (2022) use completely different algorithms and attain guarantees for regularized value-functions and non-regularized value functions, respectively, our guarantee can afford guarantees for regularized and non-regularized value-functions in a *unified* manner since MQP can be seen as a limit of MSQP when α goes to 0.

Our contributions are summarized below and in Table 1.

1. We establish that the optimal policy can be learned under partial coverage and realizability of the optimal soft Q-function and its dual. Notably, we abstain from the use of possibly stronger conditions in offline RL, such as full coverage, Bellman completeness, and uniform realizability over the policy class (such as $q_\pi \in \mathcal{Q}$ for any π as in Jiang and Huang, 2020). In particular, while a similar guarantee is provided in Zhan et al. (2022), our partial coverage guarantee has an advantage in that we are able to potentially accommodate scenarios where the initial distribution is not covered by P_b . This is feasible because our algorithm is value-based in nature, which allows us to leverage the structure of the Q-function classes and refine the coverage condition.
2. We demonstrate that the optimal policy can be learned under partial coverage, realizability of the Q-function and its dual, and a soft margin. While a similar guarantee is obtained in Chen and Jiang (2022), our guarantee has the advantage that the soft margin is significantly less stringent than the hard margin required therein.

1.1 Related Works

We summarize related works as follows. Further related works is discussed in Section A.

Offline RL under partial coverage. There is a growing number of results under partial coverage following the principle of pessimism in offline RL (Yu et al., 2020; Kidambi et al., 2020). In comparison to works that focus on tabular (Rashidinejad et al., 2021; Li et al., 2022; Shi et al., 2022b; Yin and Wang, 2021) or linear models (Jin et al., 2020; Chang et al., 2021; Zhang et al., 2022; Nguyen-Tang et al., 2022; Bai et al., 2022), our emphasis is on general function approximation (Jiang and Huang, 2020; Uehara and Sun, 2021; Xie et al., 2021; Zhan et al., 2022; Zhu et al., 2023; Rashidinejad et al., 2022; Zanette and Wainwright, 2022; Ozdaglar et al., 2023). Among them, we specifically focus on model-free methods. The representative work is summarized in Table 1.

Soft (entropy-regularized) Q-functions. Soft Q-functions are utilized in various contexts in RL (Geist et al., 2019; Neu et al., 2017). They have been shown to improve performance in online RL settings, as demonstrated in Soft Q-Learning (Fox et al., 2015; Schulman et al., 2017) and Soft Actor Critic (Haarnoja et al., 2018). In the field of imitation learning, they play a crucial role in Maximum Entropy IRL (Ziebart et al., 2008, 2010). Furthermore, within the realm of offline RL, these soft Q-functions are utilized to make the learned policy and behavior policy sufficiently similar (Wu et al., 2019; Fakoor et al., 2021). However, to the best of the authors’ knowledge, none of these proposals in the context of offline RL have provided sample complexity results under partial coverage.

Lagrangian view of offline RL. In the realm of offline policy evaluation (OPE), Nachum and Dai (2020); Yang et al. (2020); Huang and Jiang (2022) have formulated the problem as a constrained

linear optimization problem. Notably, within the context of policy optimization, [Zhan et al. \(2022\)](#) have proposed estimators for regularized density ratios with L^2 -convergence guarantees, which is a crucial step in obtaining a near-optimal policy. Our work is similarly motivated, but with a key distinction: our target functions are the soft Q-function and Q-function, rather than the regularized density ratio, which presents additional analytical challenges due to the nonlinear constraint.

2 Preliminaries

We consider an infinite-horizon discounted MDP $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, r, \gamma, \mu_0 \rangle$ where \mathcal{S} is the state space, \mathcal{A} is the finite action space, $\gamma \in [0, 1)$ is the discount factor, reward r is a random variable following $P_r(\cdot | s, a)$ on $[R_{\min}, R_{\max}]$ ($R_{\min} \geq 0$), μ_0 is the initial distribution. A policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ is a map from the state to the distribution over actions. We denote the discounted state-action occupancy distribution under a policy π starting from an initial distribution μ_0 by $d_{\pi, \mu_0}(s, a)$. With slight abuse of notation, we denote $d_{\pi, \mu_0}(s) = \sum_a d_{\pi, \mu_0}(s, a)$. We define the value under π as $J(\pi) := \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t \tilde{r}(s_t, a_t)]$ where the expectation is taken under π . We denote the optimal policy $\arg \max_\pi J(\pi)$ by π^* , and its Q-function $\mathbb{E}_{\pi^*}[\sum_t \gamma^t \tilde{r}(s_t, a_t) | s_0 = s, a_0 = a]$ by $q^*(s, a)$.

In offline RL, using offline data $\mathcal{D} = \{(s_i, a_i, r_i, s'_i) : i = 1, \dots, n\}$, we search for the policy π^* that maximizes the policy value. We suppose each (s_i, a_i, r_i, s'_i) is sampled i.i.d. from $s_i \sim P_b, a_i \sim \pi_b(\cdot | s), r_i \sim P_r(\cdot | s_i, a_i), s'_i \sim P(\cdot | s_i, a_i)$. We denote the sample average of f by $\mathbb{E}_n[f(s, a, r, s')] = \frac{1}{n} \sum_{i=1}^n f(s_i, a_i, r_i, s'_i)$, and the expectation of f with respect to the offline data distribution by $\mathbb{E}[f(s, a, r, s')]$ (without any scripts). The policy π_b used to collect data is typically referred to as a behavior policy. With slight abuse of notation, we denote $P_b(s, a) = P_b(s)\pi_b(a | s)$.

Notation. We denote the support of $P_b(\cdot)$ by $(\mathcal{S} \times \mathcal{A})_b$, and the L^∞ -norm on $(\mathcal{S} \times \mathcal{A})_b$ by $\|\cdot\|_{\infty, b}$. The L^∞ -norm on $(\mathcal{S} \times \mathcal{A})$ is denoted by $\|\cdot\|_\infty$. We define $w_\pi(s, a) = d_{\pi, \mu_0}(s, a)/P_b(s, a)$ (if it exists). We define $\text{softmax}(h) = \frac{\exp(h(s, a))}{\sum_a \exp(h(s, a))}$ and $\|h\|_2 = \mathbb{E}_{(s, a) \sim P_b}[h^2(s, a)]^{1/2}$ for $h : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. We denote universal constants by c_1, c_2, \dots . We use the convention $a/0 = \infty$ when $a \neq 0$ and $0/0 = 0$.

3 Algorithms

In this section, we present two algorithms. The first algorithm aims to estimate the soft optimal policy by first estimating a soft Q-function. The second algorithm estimates the optimal policy after estimating the Q-function.

3.1 Minimax Soft-Q-learning with Penalization

Our ultimate aim is to mimic the optimal policy π^* . As a first step, we begin by finding a policy that maximizes the following regularized objective: $\arg \max_\pi J_\alpha(\pi)$ where for $\alpha > 0$ we define

$$J_\alpha(\pi) = (1 - \gamma)^{-1} \mathbb{E}_{(s, a) \sim d_{\pi, \mu_0}, r \sim P_r(\cdot | s, a)} [r - \underbrace{\alpha \log\{\pi(a | s)/\pi_b(a | s)\}}_{\text{KL penalty (between } \pi \text{ and } \pi_b)}]$$

This objective function is used in a variety of contexts in RL as mentioned in Section 1.1. The optimal policy that maximizes $J_\alpha(\pi)$ with respect to π is

$$\pi_\alpha^* = \text{softmax}(q_\alpha^*/\alpha + \log \pi_b), \quad (1)$$

where $q_\alpha^* : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the soft Q-function uniquely characterized by the soft Bellman equation:

$$\forall (s, a); \mathbb{E}_{s' \sim P(\cdot | s, a)} [\gamma \Omega_{\alpha, \pi_b}(q_\alpha^*)(s') + r - q_\alpha^*(s, a) | s, a] = 0,$$

where $\Omega_{\alpha, \pi_b} : [\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}] \rightarrow [\mathcal{S} \rightarrow \mathbb{R}]$ has $\Omega_{\alpha, \pi_b}(q)(s) = \alpha \log \sum_a \{\exp(q(s, a')/\alpha) \pi_b(a' | s)\}$. As opposed to the standard objective function with $\alpha = 0$, the KL penalty term serves as a regularization term that renders π_α^* sufficiently proximate to π_b . As α approaches ∞ , the optimal policy π_α^* approaches π_b . On the other hand, when $\alpha = 0$, π_α^* is π^* . Thus, in order to compete with π^* , it is necessary to keep α sufficiently small. We elaborate on this selection procedure in Section 5.

The natural method for offline RL using this formulation involves learning q_α^* from the offline data and plugging it into (1). The question that remains is how to accurately learn q_α^* from the offline data. We consider the following optimization problem:

$$\arg \min_{q \in \mathcal{Q}} 0.5 \mathbb{E}_{(s, a) \sim P_b}[q^2(s, a)] \quad (2)$$

Algorithm 1 MSQP (Minimax Soft-Q-learning with Penalization)

- 1: **Require:** Parameter $\alpha \in \mathbb{R}^+$, Models $\mathcal{Q}, \mathcal{L} \subset [\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+]$.
- 2: Estimate q_α^* as follows:

$$\hat{q}_\alpha \in \arg \min_{q \in \mathcal{Q}} \max_{l \in \mathcal{L}} \mathbb{E}_n [q^2(s, a)/2 + \{\gamma \Omega_{\alpha, \pi_b}(q)(s') + r - q(s, a)\}l(s, a)]. \quad (5)$$

- 3: Estimate the soft optimal policy: $\hat{\pi}_\alpha = \text{softmax}(\hat{q}_\alpha/\alpha + \log \pi_b)$.
-

Algorithm 2 MQP (Minimax Q^* -learning with Penalization)

- 1: **Require:** Models $\mathcal{Q}, \mathcal{L} \subset [\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+]$.
- 2: Estimate q^* as follows:

$$\hat{q}_0 \in \arg \min_{q \in \mathcal{Q}} \max_{l \in \mathcal{L}} \mathbb{E}_n [q^2(s, a)/2 + \{\gamma \max_{a'} q(s', a') + r - q(s, a)\}l(s, a)]. \quad (8)$$

- 3: Estimate the optimal policy: $\hat{\pi}_0(a | s) = \arg \max_{a: \pi_b(a|s) > 0} \hat{q}_0(s, a)$.
-

where \mathcal{Q}' consists of all functions $q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ satisfying

$$\mathbb{E}_{s' \sim P(\cdot | s, a)} [\gamma \Omega_{\alpha, \pi_b}(q)(s') + r - q(s, a) | s, a] = 0 \quad \forall (s, a) \in (\mathcal{S} \times \mathcal{A})_b. \quad (3)$$

Here, because of the constraint (3), the solution is q_α^* . Furthermore, we use $q^2(s, a)$ in (2) because this choice relaxes the equality in (3) to an inequality ≤ 0 as we will demonstrate in Section B. Consequently, the entire optimization problem outlined in (2) and (3) transforms into a convex optimization problem.

Then, using the method of Lagrange multipliers, (2) is transformed into

$$\min_q \max_l L_\alpha(q, l), \quad L_\alpha(q, l) := \mathbb{E} [q^2(s, a)/2 + \{\gamma \Omega_{\alpha, \pi_b}(q)(s') + r - q(s, a)\}l(s, a)]. \quad (4)$$

Being motivated by the above formulation, our MSQP algorithm, specified in Algorithm 1, approximates this formulation by replacing expectations with sample averages and restricting optimization to function classes with bounded complexity.

Remark 1 (Computation). *Although minimax optimization is generally difficult to solve, it is computationally feasible when we choose RKHS or linear function classes for \mathcal{L} . In this case, we can solve the inner maximization problem analytically in closed form, as the objective function is linear in l . As a result, the minimax optimization problem reduces to empirical risk minimization.*

3.2 Minimax Q^* -learning with Penalization

Next, we examine a policy learning algorithm utilizing Q^* -functions. To learn Q^* , our objective function is derived from the constrained optimization problem:

$$\arg \min_{q \in \mathcal{Q}^{*'}} 0.5 \mathbb{E}_{(s, a) \sim P_b} [q^2(s, a)] \quad (6)$$

where $\mathcal{Q}^{*'}$ consists of all functions $q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ satisfying

$$\forall (s, a) \in (\mathcal{S} \times \mathcal{A})_b; \mathbb{E}_{s' \sim P(\cdot | s, a)} [\gamma \max_{a' \in \mathcal{A}} q(s', a') + r - q(s, a) | s, a] = 0.$$

Next, again using the method of Lagrange multipliers, (6) is transformed into

$$\min_q \max_l L_0(q, l), \quad L_0(q, l) := \mathbb{E} [q^2(s, a)/2 + \{\gamma \max_{a'} q(s', a') + r - q(s, a)\}l(s, a)]. \quad (7)$$

Note L_0 is the limit of L_α as $\alpha \rightarrow 0$.

Our MQP algorithm, specified in Algorithm 2, similarly approximates this formulation by replacing expectations with sample averages and restricting optimization to function classes with bounded complexity. Our final policy is greedy with respect to the learned Q -function but restricting to the support of the offline data in order to avoid exploiting regions not covered by the offline data.

Remark 2 (Prominent differences). *There exist several other minimax estimators for Q^* including BRM (Antos et al., 2008) and MABO (Xie and Jiang, 2020). Although these ensure convergence guarantees in terms of Bellman residual errors, they do not ensure the guarantee in terms of L^2 -errors, which is our focus. Our minimax objective function differs significantly from that of the aforementioned approaches, and its unique design plays a pivotal role in enabling L^2 -rates.*

4 L^2 -convergence Rates for Soft Q -functions and Q^* -functions

To analyze our Q -estimators we first establish conditions that ensure $q_\alpha^* = \arg \min_{q \in \mathcal{Q}} \max_{l \in \mathcal{L}} L_\alpha(q, l)$ on the support $(\mathcal{S} \times \mathcal{A})_b$. Building on this, we prove L^2 -convergence rates for \hat{q}_α and \hat{q}_0 . These L^2 -convergence guarantees are subsequently translated into performance guarantees of the policies we output in [Section 5](#).

4.1 Identification of Soft Q -functions

Consider an L^2 -space \mathcal{H} where the inner product is define as $\langle h_1, h_2 \rangle = \mathbb{E}_{(s,a) \sim P_b} [h_1(s, a)h_2(s, a)]$. Then we define two operators and a key function:³

$$\begin{aligned} P_\alpha^* : \mathcal{H} &\ni f \mapsto \mathbb{E}_{s' \sim P(s,a), a' \sim \pi_\alpha^*} [f(s', a') \mid (s, a) = \cdot] \in \mathcal{H}, \\ \{P_\alpha^*\}^\top : \mathcal{H} &\ni f \mapsto \int P(\cdot \mid s, a) \pi_\alpha^*(\cdot \mid \cdot) f(s, a) d(s, a) \in \mathcal{H}, \\ l_\alpha^*(s, a) &:= \begin{cases} \frac{(I - \gamma \{P_\alpha^*\}^\top)^{-1} (P_b(s, a) q_\alpha^*(s, a))}{P_b(s, a)} & (s, a) \in (\mathcal{S} \times \mathcal{A})_b, \\ 0 & (s, a) \notin (\mathcal{S} \times \mathcal{A})_b. \end{cases} \end{aligned}$$

These satisfy a key adjoint property, which we leverage to show (q_α^*, l_α^*) is a saddle point of $L_\alpha(q, l)$.

Lemma 1. $\forall q \in \mathcal{H}$, we have $\langle l_\alpha^*, (I - \gamma P_\alpha^*)q \rangle_{\mathcal{H}} = \langle q_\alpha^*, q \rangle_{\mathcal{H}}$.

Our first assumption ensures that l_α^* exists.

Assumption 1. Suppose $\|d_{\pi_\alpha^*, P_b}/P_b\|_\infty < \infty$. Note the infinity norm $\|\cdot\|_\infty$ is over $\mathcal{S} \times \mathcal{A}$.

Proposition 1. Under [Assumption 1](#), we have $\|l_\alpha^*\|_\infty < \infty$.

[Proposition 1](#) is immediate noting that $(I - \gamma \{P_\alpha^*\}^\top)^{-1} (P_b(\cdot) q_\alpha^*(\cdot)) = \sum_{t=0}^{\infty} \gamma^t (\{P_\alpha^*\}^\top)^t (P_b q_\alpha^*)$ and recalling the discounted occupancy measure under π_α^* with initial distribution μ_0 is written as $d_{\pi_\alpha^*, \mu_0} = (1 - \gamma)(I - \gamma \{P_\alpha^*\}^\top)^{-1}(\mu_0)$. Hence, $\|l_\alpha^*\|_\infty \leq (1 - \gamma)^{-1} R_{\max} \|d_{\pi_\alpha^*, \mu_0}/P_b\|_\infty$. Note that $\|d_{\pi_\alpha^*, P_b}/P_b\|_\infty$ crucially differs with the standard density-ratio-based concentrability coefficient $\|d_{\pi_\alpha^*, \mu_0}/P_b\|_\infty$ in offline RL. Unlike $\|d_{\pi_\alpha^*, P_b}/P_b\|_\infty$, the value of $\|d_{\pi_\alpha^*, \mu_0}/P_b\|_\infty$ can be infinite when the initial distribution μ_0 is not covered by offline data P_b as the practical motivating example is explained in the footnote in [Section 1](#) and [Example 2](#).

Our next assumption ensures $q_\alpha^* \geq 0$, which also guarantees that $l_\alpha^* \geq 0$.

Assumption 2. Suppose $\alpha \log \|\pi_\alpha^*/\pi_b\|_\infty \leq R_{\min}$.

[Assumption 2](#) can be satisfied by rescaling reward (*i.e.*, rescaling R_{\min}) as long as $\|\pi_\alpha^*/\pi_b\|_\infty$ is finite. Hence, it is very mild. Putting [Lemma 1](#) together with our assumptions we have the following.

Lemma 2. Suppose [Assumptions 1](#) and [2](#) hold. Then, (q_α^*, l_α^*) is a saddle point of $L_\alpha(q, l)$ over $q \in \mathcal{H}, l \in \mathcal{H}$, *i.e.*, $L_\alpha(q, l_\alpha^*) \geq L_\alpha(q_\alpha^*, l_\alpha^*) \geq L_\alpha(q_\alpha^*, l) \forall q \in \mathcal{H}, \forall l \in \mathcal{H}$.

Recall that a point (\tilde{q}, \tilde{l}) is a saddle point if and only if the strong duality holds, and $\tilde{q} \in \arg \min_{q \in \mathcal{H}} \sup_{l \in \mathcal{H}} L_\alpha(q, l)$, $\tilde{l} \in \arg \max_{l \in \mathcal{H}} \inf_{q \in \mathcal{H}} L_\alpha(q, l)$ using the general characterization ([Bertsekas, 2009](#)). Hence, [Lemma 2](#) ensures $q_\alpha^* \in \arg \min_{q \in \mathcal{H}} \max_{l \in \mathcal{H}} L_\alpha(q, l)$.

Next, we consider the constrained optimization problem when we use function classes $\mathcal{Q} \subset \mathcal{H}, \mathcal{L} \subset \mathcal{H}$. As long as the saddle point is included in $(\mathcal{Q}, \mathcal{L})$, we can prove that q_α^* is a unique minimaxer.

Lemma 3. Suppose [Assumptions 1](#) and [2](#) hold, $q_\alpha^* \in \mathcal{Q}$, and $l_\alpha^* \in \mathcal{L}$. Then, we have that $q_\alpha^* = \arg \min_{q \in \mathcal{Q}} \sup_{l \in \mathcal{L}} L_\alpha(q, l)$ on the support $(\mathcal{S} \times \mathcal{A})_b$.

This establishes that realizability ($q_\alpha^* \in \mathcal{Q}, l_\alpha^* \in \mathcal{L}$) is sufficient to identify q_α^* on the offline data distribution. At a high level, $q_\alpha^* \in \arg \min_{q \in \mathcal{Q}} \sup_{l \in \mathcal{L}} L_\alpha(q, l)$ is established through the invariance of saddle points, *i.e.*, saddle points over original sets remain saddle points over restricted sets. Its uniqueness is verified by the strong convexity in q of $L_\alpha(q, l)$ induced by $\mathbb{E}_{(s,a) \sim P_b} [q^2(s, a)]$.

³We use the notation \cdot^\top because $P_\alpha^{*\top}$ is interpreted as the adjoint operator in the non-weighted L^2 -space.

4.2 L^2 -convergence Rate for Soft Q-estimators

Based on the population-level results in [Section 4.1](#), we give a finite-sample error analysis of \hat{q}_α

Assumption 3 (Realizability of soft Q-function). *Suppose $q_\alpha^* \in \mathcal{Q}$ and $\|q\|_\infty \leq B_{\mathcal{Q}} \forall q \in \mathcal{Q}$.*

Assumption 4 (Realizability of Lagrange multiplier). *Suppose $l_\alpha^* \in \mathcal{L}$ and $\|l\|_\infty \leq B_{\mathcal{L}} \forall l \in \mathcal{L}$.*

It is natural to set $B_{\mathcal{Q}} = (1 - \gamma)^{-1} R_{\max}$ and $B_{\mathcal{L}} = (1 - \gamma)^{-1} R_{\max} \|d_{\pi_\alpha^*, P_b} / P_b\|_\infty$, but letting these be arbitrary offers further flexibility to our results.

Theorem 1 (L^2 -convergence of soft Q-estimators). *Suppose Assumptions 1, 2, 3, and 4 hold. Then, with probability $1 - \delta$, the L^2 -error $\|\hat{q}_\alpha - q_\alpha^*\|_2$ is upper-bounded by*

$$c (B_{\mathcal{Q}}^2 + B_{\mathcal{Q}} B_{\mathcal{L}} \{\alpha + \ln(|\mathcal{A}|)\})^{1/2} (\ln(|\mathcal{Q}||\mathcal{L}|/\delta)/n)^{1/4}.$$

Our result is significant as it relies on realizability-type conditions rather than Bellman closedness. Since the majority of existing works focus on non-regularized Q-functions, we postpone the comparison to these existing works to the next section. Note when \mathcal{Q} and \mathcal{L} are infinite, we can easily replace $|\mathcal{Q}|, |\mathcal{L}|$ with their L^∞ -covering numbers following [Uehara et al. \(2021\)](#). Details are given in the appendix.

4.3 L^2 -convergence Rate for Q^* -functions

Next, we give analogous finite-sample error analysis of \hat{q}_0 leveraging the same reasoning.

Assumption 5 (Realizability of Q^* -functions). *Suppose $q^* \in \mathcal{Q}$ and $\|q\|_\infty \leq B_{\mathcal{Q}} \forall q \in \mathcal{Q}$.*

Next, we define the Lagrange multiplier:

$$\begin{aligned} \{P^*\}^\top : \mathcal{H} \ni f &\mapsto \int P(\cdot | s, a) \pi^*(\cdot | \cdot) f(s, a) d\mu(s, a) \in \mathcal{H}, \\ l^* &:= \{(I - \gamma \{P^*\}^\top)^{-1} (q^* P_{\pi_b})\} / P_{\pi_b}. \end{aligned}$$

While l^* involves the density ratio, this is always well-defined as long as $\|d_{\pi^*, P_b} / P_b\|_\infty < \infty$. Then, it can be similarly established that (q^*, l^*) is a saddle point of $L_0(q, l)$ over $q \in \mathcal{H}, l \in \mathcal{H}$ as we show in [Lemma 2](#). We lastly require its realizability.

Assumption 6 (Realizability of Lagrange multiplier). *Suppose $\|d_{\pi^*, P_b} / P_b\|_\infty < \infty$ and $l^* \in \mathcal{L}$. Further suppose $\|l\|_\infty \leq B_{\mathcal{L}} \forall l \in \mathcal{L}$.*

Theorem 2 (L^2 -convergence of Q-estimators). *Suppose Assumptions 5 and 6 hold. Then, with probability $1 - \delta$, the L^2 -error $\|\hat{q}_0 - q^*\|_2$ is upper-bounded by*

$$c (B_{\mathcal{Q}}^2 + B_{\mathcal{L}} B_{\mathcal{Q}})^{1/2} (\ln(|\mathcal{Q}||\mathcal{L}|/\delta)/n)^{1/4}.$$

To the best of our knowledge, this is the first guarantee on L^2 errors for learning q^* using *general function approximation without relying on Bellman completeness*. This is highly nontrivial, and we have carefully crafted our algorithm to obtain this guarantee. Existing results are often specific to particular models, such as linear models ([Shi et al., 2022a](#)), or they require Bellman completeness ([Chen and Jiang, 2019](#); [Chen and Qi, 2022](#)), or they are limited to offline policy evaluation scenarios ([Huang and Jiang, 2022](#)) (*i.e.*, cases involving linear Bellman operators, but nonlinear Bellman operators). Actually, it seems that even under the assumption of Bellman completeness, obtaining an L^2 guarantee *without strong coverage assumptions* remains unclear. A detailed comparison among these different approaches is presented in [Section A](#).

5 Finite Sample Guarantee of MSQP

In this section, we present our primary sample complexity guarantee for our MSQP algorithm under the assumptions of realizability of q_α^* and l_α^* and partial coverage. We first show the learned policy $\hat{\pi}_\alpha$ can compete with π_α^* . Finally we show $\hat{\pi}_\alpha$ can compete with π^* by selecting α properly.

We first introduce the flattened behavior policy π_b^\diamond , which is uniform on the support of π_b . We use it as a technical device to define a model-free concentrability coefficient following [Xie et al. \(2021\)](#).

Definition 1 (Model-free concentrability coefficient). *Define*

$$C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} := \sup_{q \in \mathcal{Q}} \frac{\mathbb{E}_{s \sim d_{\pi_\alpha^*, \mu_0}, a \sim \pi_b^\diamond(a|s)} [\|q(s, a) - q_\alpha^*(s, a)\|_2^2]}{\mathbb{E}_{(s, a) \sim P_b} [\|q(s, a) - q_\alpha^*(s, a)\|_2^2]}$$

where $\pi_b^\diamond(\cdot | s) = \begin{cases} 0 & \pi_b(\cdot | s) = 0 \\ 1/|\{a \in \mathcal{A} | \pi_b(a | s) > 0\}| & \pi_b(\cdot | s) > 0 \end{cases}$ is the flattened behavior policy.

Clearly, $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}}$ is smaller than density-ratio-based concentrability coefficient, in other words,

$$C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} \leq \max_{(s, a)} \frac{d_{\pi_\alpha^*, \mu_0}(s) \pi_b^\diamond(a | s)}{P_b(s) \pi_b(a | s)}.$$

Here, we always have $\|\pi_b^\diamond/\pi_b\| < \infty$ even if $\pi_b(a | s)$ is 0 for some (s, a) . In the special case where $\pi_b(a | s) \geq 1/C'$ for any (s, a) , we have $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} \leq C' \|d_{\pi_\alpha^*, \mu_0}/P_b\|_\infty$. The coefficient $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}}$ is a refined concentrability coefficient, which adapts to a function class \mathcal{Q} . For example, in linear MDPs, it reduces to a relative condition number as follows. Similar properties are obtained in related works (Xie et al., 2021; Uehara and Sun, 2021).

Example 1 (Linear MDPs). *A linear MDP is one such that, for a known feature vector $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, the true density satisfies $P(s' | s, a) = \langle \mu^*(s'), \phi(s, a) \rangle$ for some $\mu^* : \mathcal{S} \rightarrow \mathbb{R}^d$ and the reward function satisfies $\mathbb{E}[r | s, a] = \langle \theta_r, \phi(s, a) \rangle$ for some $\theta_r \in \mathbb{R}^d$.*

In linear MDPs, q_α^ is clearly linear in $\phi(s, a)$. Hence, the natural function class is $\mathcal{Q} = \{(\theta, \phi(s, a)) | \|\theta\| \leq B\}$ for a certain $B \in \mathbb{R}^+$. Then, we have*

$$C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} = \sup_{x \neq 0} \frac{x^\top \mathbb{E}_{s \sim d_{\pi_\alpha^*, \mu_0}, a \sim \pi_b^\diamond(a|s)} [\phi(s, a) \phi(s, a)^\top] x}{x^\top \mathbb{E}_{(s, a) \sim P_b} [\phi(s, a) \phi(s, a)^\top] x}.$$

We are now prepared to present our main result, which states that given the realizability of the soft Q-function q_α^* and Lagrange multiplier l_α^* , it is possible to compete with π_α^* under the coverage condition $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} < \infty, \|d_{\pi_\alpha^*, P_b}/P_b\|_\infty < \infty$.

Theorem 3 ($\hat{\pi}_\alpha$ can compete with π_α^*). *Fix $\alpha > 0$. Suppose Assumptions 1, 2, 3, and 4 hold. With probability $1 - \delta$, the regret $J(\pi_\alpha^*) - J(\hat{\pi}_\alpha)$ is upper-bounded by*

$$n^{-1/4} \text{Poly} \left(|\mathcal{A}|, \mathcal{B}_{\mathcal{Q}}, \mathcal{B}_{\mathcal{L}}, C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}}, \ln \left(\frac{|\mathcal{Q}| |\mathcal{L}|}{\delta} \right), \frac{1}{1 - \gamma}, R_{\max} \right)$$

The proof mainly consists of two steps: (1) obtaining L^2 -errors of \hat{q}_α as previously demonstrated in Theorem 1, (2) translating this error into the error of $\hat{\pi}_\alpha$. In the second step, the Lipschitz continuity of the softmax function plays a crucial role. If there is no regularization ($\alpha = 0$) and the greedy policy of q_0^* is utilized, the second step does not proceed (without any further additional assumptions).

Our ultimate goal is to compete with π^* . Theorem 3 serves as the primary foundation for this goal. The remaining task is to analyze the approximation error $J(\pi^*) - J(\pi_\alpha^*)$. Fortunately, this term can be controlled through α and the density ratio between π^* and π_b . Then, by properly controlling α , we can obtain the following sample complexity result.

Theorem 4 (PAC guarantee of $\hat{\pi}_\alpha$). *Fix any $\epsilon > 0$. Suppose Assumptions 1, 2, 3, and 4 hold for $\alpha = c/n^{1/8}$ and $\|\pi_0^*/\pi_b\|_\infty \leq C_0, C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} < \infty$. Then, if n is at least*

$$\epsilon^{-8} \text{Poly}(|\mathcal{A}|, \mathcal{B}_{\mathcal{Q}}, \mathcal{B}_{\mathcal{L}}, C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}}, \ln(|\mathcal{Q}| |\mathcal{L}| / \delta), (1 - \gamma)^{-1}, \ln(C_0), R_{\max}),$$

with probability at least $1 - \delta$, we can ensure $J(\pi^) - J(\hat{\pi}_\alpha) \leq \epsilon$.*

In summary, the realizability of q_α^*, l_α^* , per-step coverage $\|\pi_0^*/\pi_b\|_\infty < \infty$ and partial coverage $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} < \infty, \|d_{\pi_\alpha^*, P_b}/P_b\|_\infty < \infty$ are sufficient to compete with π^* . This is a novel and attractive result. Firstly, if we solely use the naïve FQI or Bellman residual minimization, existing PAC results require the global coverage $\|d_{\pi, \mu_0}/P_b\|_\infty < \infty$ for any possible policy π (Munos and Szepesvári, 2008; Antos et al., 2008). Our result only requires coverage under a single policy π_α^* (near-optimal policy). Secondly, we only require the realizability of two functions, and we do not necessitate

realizability-type conditions for all policies in the policy class or Bellman completeness, unlike existing works with partial coverage (Xie et al., 2021; Jiang and Huang, 2020).

The most similar result is Zhan et al. (2022). However, our guarantee possesses a certain advantage over their guarantee as follows. They demonstrate the realizability of certain functions \tilde{w}_α^* , v_α^* and partial coverage $\|d_{\tilde{\pi}_\alpha, \mu_0}/P_b\| < \infty$ are sufficient conditions in offline RL, where $\tilde{w}_\alpha^* = d_{\tilde{\pi}_\alpha, \mu_0}/P_b$ ($\tilde{\pi}_\alpha$ is a certain regularized optimal policy, but fundamentally distinct from π_α^*) and v_α^* is a near-optimal regularized value function parameterized by α . Here, we have $\tilde{w}_0^* = w^*$, $v_0^* = v^*$. Our guarantee has a similar flavor in the sense that it roughly illustrates realizability and partial coverage are sufficient conditions. However, the meanings of realizability and partial coverage are significantly different. In particular, by employing our algorithm, we can ensure PAC guarantees under the boundedness of the refined concentrability coefficient $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} < \infty$ (and $\|d_{\pi_\alpha^*, P_b}/P_b\|_\infty$ through $\mathcal{B}_{\mathcal{L}}$). As a result, the L^∞ -norm of the density-ratio-based concentrability coefficient $\|d_{\pi_\alpha^*, \mu_0}/P_b\|_\infty$ can even be infinite. More specifically, we can permit situations where $\max_s \mu_0(s)/P_b(s) = \infty$ as we will see the practical example soon. Conversely, Zhan et al. (2022) excludes this possibility since the algorithm explicitly estimates the density ratio \tilde{w}_α^* .

Example 2 (Contextual bandit under external validity). *We consider the contextual bandit setting where we want to optimize $J(\pi) = \mathbb{E}_{s \sim \mu_0, a \sim \pi(s), r \sim P_r(s, a)}[r]$ using offline data $s \sim P_b, a \sim \pi_b(s), r \sim P(s, a)$. This is the simplest RL setting with $\gamma = 0$. Here, note μ_0 could be different from P_b . This case often happens in practice as discussed in the literature on causal inference related to external validity (Pearl and Bareinboim, 2014; Dahabreh et al., 2019; Uehara et al., 2020), which refers to the shift between the target population and the offline data. Here, our PAC guarantee does not require that $\mu_0(s)$ is covered by $P_b(s)$ in terms of the density ratio as long as the relative condition number is upper-bounded when we use linear models. On the other hand, Zhan et al. (2022) excludes this possibility.*

Despite the aforementioned advantage of our approach, unfortunately, our sample complexity of $O(1/\epsilon^8)$ is slower compared to that of $O(1/\epsilon^6)$ in Zhan et al. (2022). In the following, we demonstrate that MQP, which is a special version of MSQP when $\alpha \rightarrow 0$, can achieve a faster rate of $O(1/\epsilon^2)$.

6 Finite Sample Guarantee of MQP

In this section, building upon the convergence result of \hat{q}_0 , we demonstrate the finite sample guarantee of our MQP algorithm under partial coverage. We first introduce the soft margin.

Assumption 7 (Soft margin). *For any $a' \in \mathcal{A}$, there exists $t_0 \in \mathbb{R}^+$, $\beta \in (0, \infty]$ such that*

$$\mathbb{P}_{s \sim d_{\pi_\alpha^*, \mu_0}}(0 < |q^*(s, \pi^*(s)) - q^*(s, a')| < t) \leq (t/t_0)^\beta$$

for any $t > 0$. Here, we use the convention $x^\infty = 0$ if $0 < x < 1$ and $x^\infty = \infty$ if $x > 1$.

In the extreme case, if there exists a gap in q^* (also known as a hard margin) so that the best action is always better than the second-best by some lower bounded amount, then the soft margin is satisfied with $\beta = \infty$. Thus, the soft margin is more general than the gap condition used in Simchowitz and Jamieson (2019); Wu et al. (2022). Crucially, a gap generally does *not* exist in continuous state spaces unless Q-functions are discontinuous or one action is always option, or a gap involves a large t_0 constant in discrete state spaces with bad dependence on the number of states. In contrast, a soft margin with some $\beta > 0$ generally holds (see, e.g., lemma 4 in Hu et al., 2021). The soft margin is widely used in the literature on classification, decision making, and RL (Audibert and Tsybakov, 2007; Perchet and Rigollet, 2013; Luedtke and Chambaz, 2020; Hu et al., 2021, 2022).

Theorem 5 (PAC guarantee of $\hat{\pi}_0$). *Suppose Assumptions 5, 6, and 7 hold and $\|\pi^*/\pi_b\|_\infty \leq C_0$. Fix any $\epsilon > 0$. Then, if n is at least*

$$\left\{ \frac{|\mathcal{A}|}{\epsilon} \right\}^{\frac{4+2\beta}{\beta}} \text{Poly} \left(t_0^{-1}, |\mathcal{A}|, \mathcal{B}_{\mathcal{Q}}, \mathcal{B}_{\mathcal{L}}, C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}}, \ln \left(\frac{|\mathcal{Q}||\mathcal{L}|}{\delta} \right), (1-\gamma)^{-1}, \ln(C_0), R_{\max} \right)$$

with probability at least $1 - \delta$, we can ensure $J(\pi^*) - J(\hat{\pi}_0) \leq \epsilon$.

The proof mainly consists of two steps: (1) obtaining L^2 -errors of \hat{q}_0 as demonstrated in Theorem 2, (2) translating this error into the error of $\hat{\pi}_0$. In the second step, the soft margin plays a crucial role.

These theorems indicate that the realizability of the Q-function q^* and Lagrange multiplier l^* , and the soft margin are sufficient for the PAC guarantee under partial coverage $C_{\mathcal{Q}, d_{\pi_\alpha^*, \mu_0}} < \infty$,

$\|d_{\pi^*, P_b}/P_b\|_\infty < \infty$. Our algorithm is *agnostic* to β and operates under any value of β . In particular, when there is a gap ($\beta = \infty$), we can achieve sample complexity of $O(1/\epsilon^2)$ ⁴. In comparison to [Theorem 3](#), although we additionally use the soft margin, the realizability in [Theorem 5](#) is more appealing since it is imposed on the standard Q-function q^* . The closest guarantee to our work can be found in [Chen and Jiang \(2022\)](#), which demonstrates that the existence of the gap in q^* , the realizability of q^* , $w^*(:= d_{\pi^*, \mu_0}/P_{\pi_b})$, and partial coverage $\|w^*\|_\infty < \infty$ are sufficient conditions. A similar comparison is made in [Ozdaglar et al. \(2023\)](#). In comparison to their work, we use the soft margin, which is significantly less stringent.

7 Conclusions

We propose two value-based algorithms, MSQP and MQP, that operate under realizability of certain functions and partial coverage (*i.e.*, single-policy-coverage). Notably, our guarantee does not require Bellman completeness and uniform-type realizability over the policy class. While guarantees with similar flavors are obtained in [Zhan et al. \(2022\)](#); [Chen and Jiang \(2022\)](#), MSQP can potentially relax the density-ratio-based partial coverage regarding the initial distribution as opposed to [Zhan et al. \(2022\)](#), and MQP can operate under the soft-margin, which is less stringent than the hard margin imposed in [Chen and Jiang \(2022\)](#). Moreover, both algorithms work on Q-functions, which are more commonly used in practice.

Acknowledgements

NK acknowledges support by the National Science Foundation under Grant Nos. IIS 1846210, CCF 2002272, IIS 2107304, CIF 2212262, IIS 2144994.

JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0304, the Sloan Research Fellowship, NSF CCF 2002272, NSF IIS 2107304, NSF CIF 2212262, ONR Young Investigator Award, and NSF CAREER Award 2144994.

WS acknowledges support by the National Science Foundation under Grant Nos. IIS-2154711.

References

- Andras Antos, Csaba Szepesvari, and Remi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- Jean-Yves Audibert and Alexandre B Tsybakov. Fast learning rates for plug-in classifiers. *The Annals of statistics*, 35(2):608–633, 2007.
- Chenjia Bai, Lingxiao Wang, Zhuoran Yang, Zhihong Deng, Animesh Garg, Peng Liu, and Zhaoran Wang. Pessimistic bootstrapping for uncertainty-driven offline reinforcement learning. *arXiv preprint arXiv:2202.11566*, 2022.
- Dimitri Bertsekas. *Convex optimization theory*, volume 1. Athena Scientific, 2009.
- Jonathan Chang, Masatoshi Uehara, Dhruv Sreenivas, Rahul Kidambi, and Wen Sun. Mitigating covariate shift in imitation learning via offline data with partial coverage. *Advances in Neural Information Processing Systems*, 34:965–979, 2021.
- Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. PMLR, 2019.
- Jinglin Chen and Nan Jiang. Offline reinforcement learning under value and density-ratio realizability: the power of gaps. *arXiv preprint arXiv:2203.13935*, 2022.
- Xiaohong Chen and Zhengling Qi. On well-posedness and minimax optimal rates of nonparametric q-function estimation in off-policy evaluation. *arXiv preprint arXiv:2201.06169*, 2022.

⁴Similar to the findings in [Wang et al. \(2022\)](#), in general offline RL, we may potentially achieve a result of $O(1/\epsilon)$. We leave room for further enhancements in future research.

- Stephen R Cole and Elizabeth A Stuart. Generalizing evidence from randomized clinical trials to target populations: the actg 320 trial. *American journal of epidemiology*, 172(1):107–115, 2010.
- Issa J Dahabreh, Sarah E Robertson, Eric J Tchetgen, Elizabeth A Stuart, and Miguel A Hernán. Generalizing causal inferences from individuals in randomized trials to all trial-eligible individuals. *Biometrics*, 75(2):685–694, 2019.
- Rasool Fakoor, Jonas W Mueller, Kavosh Asadi, Pratik Chaudhari, and Alexander J Smola. Continuous doubly constrained batch reinforcement learning. *Advances in Neural Information Processing Systems*, 34:11260–11273, 2021.
- Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.
- Bolin Gao and Laca Pavel. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*, 2017.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- Yichun Hu, Nathan Kallus, and Masatoshi Uehara. Fast rates for the regret of offline reinforcement learning. *arXiv preprint arXiv:2102.00479*, 2021.
- Yichun Hu, Nathan Kallus, and Xiaojie Mao. Fast rates for contextual linear optimization. *Management Science*, 68(6):4236–4245, 2022.
- Audrey Huang and Nan Jiang. Beyond the return: Off-policy function estimation under user-specified error-measuring distributions. In *Neurips*, 2022.
- Nan Jiang and Jiawei Huang. Minimax value interval for off-policy evaluation and policy optimization. *Advances in Neural Information Processing Systems*, 33:2747–2758, 2020.
- Ying Jin, Zhuoran Yang, and Zhaoran Wang. Is pessimism provably efficient for offline rl? *arXiv preprint arXiv:2012.15085*, 2020.
- Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21810–21823. Curran Associates, Inc., 2020.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Gene Li, Cong Ma, and Nathan Srebro. Pessimism for offline linear contextual bandits using lp confidence sets. *arXiv preprint arXiv:2205.10671*, 2022.
- Alex Luedtke and Antoine Chambaz. Performance guarantees for policy learning. In *Annales de l’IHP Probabilités et statistiques*, volume 56, page 2162. NIH Public Access, 2020.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(5), 2008.
- Ofir Nachum and Bo Dai. Reinforcement learning via fenchel-rockafellar duality. *arXiv preprint arXiv:2001.01866*, 2020.
- Gergely Neu, Anders Jonsson, and Vicenç Gómez. A unified view of entropy-regularized markov decision processes. *arXiv preprint arXiv:1705.07798*, 2017.
- Thanh Nguyen-Tang, Ming Yin, Sunil Gupta, Svetha Venkatesh, and Raman Arora. On instance-dependent bounds for offline reinforcement learning with linear function approximation. *arXiv preprint arXiv:2211.13208*, 2022.

- Asuman E Ozdaglar, Sarath Pattathil, Jiawei Zhang, and Kaiqing Zhang. Revisiting the linear-programming framework for offline rl with general function approximation. In *International Conference on Machine Learning*, pages 26769–26791. PMLR, 2023.
- Judea Pearl and Elias Bareinboim. External validity: From do-calculus to transportability across populations. In *Statistical Science*, pages 451–482. 2014.
- Vianney Perchet and Philippe Rigollet. The multi-armed bandit problem with covariates. *The Annals of Statistics*, 41(2):693–721, 2013.
- Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *arXiv preprint arXiv:2103.12021*, 2021.
- Paria Rashidinejad, Hanlin Zhu, Kunhe Yang, Stuart Russell, and Jiantao Jiao. Optimal conservative offline rl with general function approximation via augmented lagrangian. *arXiv preprint arXiv:2211.00716*, 2022.
- John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.
- Chengchun Shi, Sheng Zhang, Wenbin Lu, and Rui Song. Statistical inference of the value function for reinforcement learning in infinite-horizon settings. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3):765–793, 2022a.
- Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Pessimistic q-learning for offline reinforcement learning: Towards optimal sample complexity. *arXiv preprint arXiv:2202.13890*, 2022b.
- Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems*, 32, 2019.
- Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. *ICLR 2022*, 2021.
- Masatoshi Uehara, Masahiro Kato, and Shota Yasui. Off-policy evaluation and learning for external validity under a covariate shift. *Advances in Neural Information Processing Systems*, 33:49–61, 2020.
- Masatoshi Uehara, Masaaki Imaizumi, Nan Jiang, Nathan Kallus, Wen Sun, and Tengyang Xie. Finite sample analysis of minimax offline reinforcement learning: Completeness, fast rates and first-order efficiency. *arXiv preprint arXiv:2102.02981*, 2021.
- Xinqi Wang, Qiwen Cui, and Simon S Du. On gap-dependent bounds for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35:14865–14877, 2022.
- Jingfeng Wu, Vladimir Braverman, and Lin Yang. Gap-dependent unsupervised exploration for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 4109–4131. PMLR, 2022.
- Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- Tengyang Xie and Nan Jiang. Q* approximation schemes for batch reinforcement learning: A theoretical comparison. In *Conference on Uncertainty in Artificial Intelligence*, pages 550–559. PMLR, 2020.
- Tengyang Xie, Ching-An Cheng, Nan Jiang, Paul Mineiro, and Alekh Agarwal. Bellman-consistent pessimism for offline reinforcement learning. *Advances in neural information processing systems*, 34:6683–6694, 2021.
- Mengjiao Yang, Ofir Nachum, Bo Dai, Lihong Li, and Dale Schuurmans. Off-policy evaluation via the regularized lagrangian. *Advances in Neural Information Processing Systems*, 33:6551–6561, 2020.

- Ming Yin and Yu-Xiang Wang. Towards instance-optimal offline reinforcement learning with pessimism. *Advances in neural information processing systems*, 34:4065–4078, 2021.
- Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 14129–14142, 2020.
- Andrea Zanette and Martin J Wainwright. Bellman residual orthogonalization for offline reinforcement learning. *arXiv preprint arXiv:2203.12786*, 2022.
- Wenhao Zhan, Baihe Huang, Audrey Huang, Nan Jiang, and Jason Lee. Offline reinforcement learning with realizability and single-policy concentrability. In *Conference on Learning Theory*, pages 2730–2775. PMLR, 2022.
- Xuezhou Zhang, Yiding Chen, Xiaojin Zhu, and Wen Sun. Corruption-robust offline reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 5757–5773. PMLR, 2022.
- Hanlin Zhu, Paria Rashidinejad, and Jiantao Jiao. Importance weighted actor-critic for optimal conservative offline reinforcement learning. *arXiv preprint arXiv:2301.12714*, 2023.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- Brian D Ziebart, J Andrew Bagnell, and Anind K Dey. Modeling interaction via the principle of maximum causal entropy. In *ICML*, 2010.