
Sparse Feature Selection Makes Batch Reinforcement Learning More Sample Efficient

Botao Hao¹ Yaqi Duan² Tor Lattimore¹ Csaba Szepesvári^{1,3} Mengdi Wang^{2,1}

Abstract

This paper provides a statistical analysis of high-dimensional batch reinforcement learning (RL) using sparse linear function approximation. When there is a large number of candidate features, our result sheds light on the fact that sparsity-aware methods can make batch RL more sample efficient. We first consider the off-policy evaluation problem. To evaluate a new target policy, we analyze a Lasso fitted Q-evaluation method and establish a finite-sample error bound that has no polynomial dependence on the ambient dimension. To reduce the Lasso bias, we further propose a post model-selection estimator that applies fitted Q-evaluation to the features selected via group Lasso. Under an additional signal strength assumption, we derive a sharper instance-dependent error bound that depends on a divergence function measuring the distribution mismatch between the data distribution and occupancy measure of the target policy. Further, we study the Lasso fitted Q-iteration for batch policy optimization and establish a finite-sample error bound depending on the ratio between the number of relevant features and restricted minimal eigenvalue of the data's covariance. In the end, we complement the results with minimax lower bounds for batch-data policy evaluation/optimization that nearly match our upper bounds. The results suggest that having well-conditioned data is crucial for sparse batch policy learning.

1. Introduction

We consider batch reinforcement learning (RL), where the problem is to evaluate a target policy or to learn a good policy based on a given dataset (Szepesvári, 2010; Lange et al., 2012; Levine et al., 2020). While in online RL the central question is how to sequentially interact with the environment to balance exploration and exploitation, in batch RL the dataset is given a priori and the focus is typically on learning a near-optimal policy or evaluating a given target policy.

To handle RL systems with large or even infinite state spaces, we focus on the use of linear function approximation (Bellman et al., 1963; Schweitzer and Seidmann, 1985; Bertsekas and Tsitsiklis, 1996), that is, using a weighted linear combination of available features (aka basis functions) to represent high-dimensional transition/value functions. Results from the supervised learning literature show that the sample size needed to get accurate policy evaluations or near-optimal policies must scale at least linearly with d , the number of features (e.g., Example 15.14 of Wainwright, 2019).

We leverage the idea of sparse approximation and focus on situations when a smaller number of *unknown* relevant features is sufficient for solving the RL problem. Sparse regression has proved to be a powerful method for high-dimensional statistical learning with limited data (Tibshirani, 1996; Chen et al., 2001; Bunea et al., 2007; Bickel et al., 2009; Rish and Grabarnik, 2014), and we will borrow techniques from the sparse learning literature to improve the sample efficiency of batch RL, an idea with a considerable history in RL as witnessed by our literature review that follows below.

Contribution. We make several contributions:

- First, we consider the problem of off-policy evaluation (OPE). To promote sparse solutions, we iteratively fit state-action value functions using linear regression with ℓ_1 -regularization which we call this procedure as *Lasso fitted Q-evaluation* (Algorithm 1). We estab-

¹Deepmind ²Princeton University ³University of Alberta. Correspondence to: Botao Hao <haobotao000@gmail.com>, Mengdi Wang <mengdiw@princeton.edu>.

lish the first finite-sample error bound for Lasso fitted Q-evaluation that depends linearly on the number of relevant features and the restricted minimal eigenvalue. Importantly, the bound has no polynomial dependence on d . This appears to be the first theoretical bound for sparse batch policy evaluation.

- Second, to reduce the Lasso bias, we propose an improved *post model-selection estimator* (Algorithm 2) that applies fitted Q-evaluation with a smaller feature set that is selected using group Lasso. Under an additional separability assumption, we derive a sharper and nearly minimax-optimal error bound that is instance-dependent. The error bound is determined by a divergence function measuring the distribution mismatch, *restricted over the reduced feature space*, between the data distribution and the occupancy distribution of the target policy. This divergence defined over the reduced feature space is significantly smaller than its counterpart over the full d -dimensional space. In other words, *sparse feature selection reduces the distribution mismatch*. We also provide a nearly-matching lower bound, and these two results sharply characterize the statistical limits of sparse off-policy evaluation.
- Third, we extend our analysis to the batch policy optimization problem. We analyze the *Lasso fitted Q-iteration* (Algorithm 3) and show that the ℓ_∞ -norm of policy error depends linearly on the ratio between the number of relevant features and the restricted minimal eigenvalue of the dataset’s covariance matrix. Finally, we establish a minimax lower bound for sparse batch policy learning and show that the lower bound also depends on the aforementioned ratio. This is the first lower bound result, to the authors’ best knowledge, demonstrating the critical role played by the minimal eigenvalue of the dataset’s covariance matrix and the construction is highly non-trivial. The upper and lower bounds validate the belief that well-conditioned data is crucial for sample-efficient policy learning.

2. Preliminaries

2.1. Problem definition

A finite, infinite-horizon discounted Markov decision process (MDP) can be described by the tuple $M = (\mathcal{X}, \mathcal{A}, P, r, \gamma)$. Here, $\mathcal{X}$ is a finite set of states, $\mathcal{A}$ is a finite set of actions, $P : \mathcal{X} \times \mathcal{A} \rightarrow \Delta_{\mathcal{X}}$ is the transition probability function, $r : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function and $\gamma \in (0, 1)$ is the so-called discount factor. In this

paper, for the sake of simplicity, we stick to finite DMDPs. However, our results can be extended to more general cases with routine work.

We define a (stationary) policy π as a $\mathcal{X} \rightarrow \Delta_{\mathcal{A}}$ map, mapping states to distributions over actions. A policy, a distribution ξ_0 over $\mathcal{X}$ and a DMDP M together give rise to a probability measure $\mathbb{P}^\pi$ over the set of infinitely long state-action histories: for a history of form $x_1, a_1, r_1, x_2, a_2, r_2, \dots$, $\mathbb{P}^\pi(x_1) = \xi_0(x_1)$ and for $t = 1, 2, \dots$, $\mathbb{P}^\pi(a_t | x_1, a_1, \dots, x_{t-1}, a_{t-1}, x_t) = \pi(a_t | x_t)$ and $\mathbb{P}^\pi(x_{t+1} | x_1, a_1, \dots, x_{t-1}, a_{t-1}, x_t, a_t) = P(x_{t+1} | x_t, a_t)$. We denote by $\mathbb{E}^\pi$ the corresponding expectation operator. The value of policy π given the initial state distribution ξ_0 is $v_{\xi_0}^\pi := \mathbb{E}^\pi[\sum_{t=0}^{\infty} \gamma^t r(x_t, a_t)]$, where $\mathbb{E}^\pi$ depends on ξ_0 but just this dependence is not shown. We define $v^\pi : \mathcal{X} \rightarrow \mathbb{R}$, the *value function* of policy π , by letting $v^\pi(x)$ denote the value of π when it is started from state x .

A nonstationary policy is a sequence of maps from histories to probability distributions over actions and such a policy similarly induces a probability measure $\mathbb{P}^\pi$ and an underlying expectation operator $\mathbb{E}^\pi$. An *optimal policy* π^* maximizes the value from every state: $v^{\pi^*}(x) = \max_{\pi} v^\pi(x)$ for any $x \in \mathcal{X}$ where the maximum is over all policies, including nonstationary ones. As is well known, for finite DMDPs, a stationary optimal policy always exist (e.g., Szepesvári, 2010). The value function shared by optimal policies is called the *optimal value function* and is denoted by v^* .

We consider the following learning and optimization problems. The learner knows the state space $\mathcal{X}$ and action space $\mathcal{A}$. The reward function r is given in the form of a black box, which the learner can use to evaluate $r(x, a)$ for any pair of $(x, a) \in \mathcal{X} \times \mathcal{A}$. The only unknown is the transition probability function P . The learner is given a random dataset $\mathcal{D} = \{(x_n, a_n, x'_n)\}_{n=1}^N$ generated by using a (possibly nonstationary and unknown) *behavior policy* $\bar{\pi}$ in the DMDP M starting from some initial distribution which may be different from ξ_0 . We study two fundamental batch RL tasks:

- *Off-policy policy evaluation*: given $\mathcal{D}$ and black box access to a *target policy* π , ξ_0 and r , estimate the value, $v_{\xi_0}^\pi$, of π ;
- *Batch policy optimization*: given $\mathcal{D}$ and black box access to r , find an optimal policy.

Bellman operators. By slightly abusing the notation, we will view the transition probability kernel P as a left linear

operator, mapping from $\mathbb{R}^{\mathcal{X}}$ to $\mathbb{R}^{\mathcal{X} \times \mathcal{A}}$:

$$(Pv)(x, a) := \sum_{x'} P(x'|x, a)v(x').$$

The *Bellman optimality operator* $\mathcal{T} : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ is defined as

$$[\mathcal{T}v](x) = \max_{a \in \mathcal{A}} (r(x, a) + \gamma Pv(x, a)), \forall x \in \mathcal{X}.$$

The function $v^* \in \mathbb{R}^{\mathcal{X}}$ is the unique solution to the Bellman optimality equation $v = \mathcal{T}v$. The state-action value function of a policy π is defined as

$$Q^\pi(x, a) = r(x, a) + \gamma Pv^\pi(x, a).$$

We also introduce the Bellman operator $\mathcal{T}_\pi : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ for policy π as:

$$[\mathcal{T}_\pi v](x) = \sum_{a \in \mathcal{A}} \pi(a|x) (r(x, a) + \gamma Pv(x, a)), \forall x \in \mathcal{X}.$$

The function v^π is the unique solution to the Bellman equation $v^\pi = \mathcal{T}_\pi v^\pi$.

Linear function approximation. Let $\phi : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$ be a feature map which assigns to each state-action pair a d -dimensional feature vector. A feature map combined with a parameter vector $w \in \mathbb{R}^d$ gives rise to the linear function $g_w : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ defined by $g_w(x, a) = \phi(x, a)^\top w$, $(x, a) \in \mathcal{X} \times \mathcal{A}$, and the linear function class $\mathcal{G}_\phi = \{g_w : w \in \mathbb{R}^d\}$. When little a priori information is available on how to choose the features, agnostic choices often lead to dimensions which can be as large as the number of samples n . However, effective learning with many more features than the sample-size is possible when only $s \ll d$ features are relevant. Thus, we further consider a sparse linear function class $\mathcal{G}_{\phi, \mathcal{K}, s} = \{g_w : w \in \mathbb{R}^d, w_{\mathcal{K}^c} = 0, |\mathcal{K}| \leq s\}$.

Fix a feature map ϕ . For $f : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$, let $f^\pi : \mathcal{X} \rightarrow \mathbb{R}$ be defined by $f^\pi(x) = \sum_a \pi(a|x) f(x, a)$. We present a representation condition for $\mathcal{G}_{\phi, \mathcal{K}, s}$ that ensures sample-efficient estimation.

Assumption 2.1 (Policy completeness). For some feature map ϕ , $\mathcal{G}_{\phi, \mathcal{K}, s}$ is closed under the operators $f \mapsto Pf^\pi$ for any policy π of the DMDP.

Policy completeness requires function class $\mathcal{G}_{\phi, \mathcal{K}, s}$ can well capture Bellman operator. It is crucial for estimation consistency of fitted-Q-iteration type algorithm (Le et al., 2019; Duan and Wang, 2020) and implies the realizability condition regarding Q^π . Recently, Wang et al. (2020) has derived a lower bound to show that Q^π realizability condition only

is not sufficient for a sample-efficient estimation of OPE with linear function approximation. More discussions regarding those conditions in batch policy learning could refer Chen and Jiang (2019).

Our second assumption concerns the dataset:

Assumption 2.2. The dataset $\mathcal{D}$ consists of $N = KL$ samples from K independent episodes $\tau_1, \dots, \tau_K$. Each episode τ_k has L consecutive transitions generated by some unknown behavior policy $\bar{\pi}_k$ giving rise to a sample path $\tau_k = (x_0^{(k)}, a_0^{(k)}, x_0^{(k)'}, \dots, x_{L-1}^{(k)}, a_{L-1}^{(k)}, x_{L-1}^{(k)'})$.

3. Sparsity-Aware Off-Policy Evaluation

In this section we consider OPE problem, i.e., to estimate the value of a target policy π from logged experiences $\mathcal{D}$ generated using unknown behavior policies. We propose two sparsity-aware algorithms to approximate state-action functions using sparse parameters.

3.1. Lasso fitted Q-evaluation

We propose a straightforward modification of the fitted Q-evaluation method (Le et al., 2019; Duan and Wang, 2020) to account for sparsity. For $t = 1, 2, \dots, T-1$, the algorithm produces $\hat{w}_{t+1}$ using Lasso-regularized regression. To make the errors of different steps independent, the dataset $\mathcal{D}$ is split into T nonoverlapping folds $\mathcal{D}_1, \dots, \mathcal{D}_T$, so that all folds have $R = K/T$ episodes in them and only data from fold t is used in step t . To define the algorithm, it is useful to introduce $Q_w = r + \gamma g_w$ where $w \in \mathbb{R}^d$. For $a < b$, we also define the operator $\Pi_{[a, b]} : \mathbb{R} \rightarrow [a, b]$ that projects its input to $[a, b]$, i.e., $\Pi_{[a, b]}(x) = \max(\min(x, b), a)$. The pseudocode is given as Algorithm 1. In the last step, $m = N$ samples are used to produce the final output to guarantee that the error introduced by the Monte-Carlo averaging is negligible compared to the rest.

Remark 3.1. The last step Monte Carlo averaging is only for numerical integration, where the samples are newly drawn inside the algorithm (independent from batch data), so there is no bias here. We set $m = N$ to simplify the theory but it could be much larger than N for a more accurate approximation.

3.2. Post model-selection fitted Q-evaluation

Sparse regularization is known to induce a small bias in regression. However, this bias could get compounded through the iterative procedure of Algorithm 1. To avoid such bias and improve the accuracy, we aim to identify the set of rel-

Algorithm 1 Lasso fitted Q-evaluation

1: **Input:** Initial distribution ξ_0 , target policy π , T folds of dataset $\{\mathcal{D}_t\}_{t=1}^T$ of overall size N , regularization parameter $\lambda_1 > 0$, $m := N$, $\hat{w}_0 = 0 \in \mathbb{R}^d$.

2: **for** $t = 1, 2, \dots, T$ **do**

3: Calculate regression targets for $(x_i, a_i, x'_i) \in \mathcal{D}_t$:

$$y_i = \sum_a \pi(a|x'_i) Q_{\hat{w}_{t-1}}(x'_i, a).$$

4: Fit $\hat{w}_t$ through sparse linear regression $\hat{w}_t =$

$$\underset{w}{\operatorname{argmin}} \left\{ \lambda_1 \|w\|_1 + \frac{1}{|\mathcal{D}_t|} \sum_{(x_i, a_i, x'_i) \in \mathcal{D}_t} (\Pi_{[0, 1/(1-\gamma)]} y_i - \phi(x_i, a_i)^\top w)^2 \right\}.$$

5: **end for**

6: **Output:** $\hat{v}_{\hat{w}_T}^\pi = \frac{1}{m} \sum_{u=1}^m \Pi_{[0, 1/(1-\gamma)]}(Q_{\hat{w}_T}(\tilde{x}_u, \tilde{a}_u))$ with $\tilde{x}_u \sim \xi_0$, $\tilde{a}_u \sim \pi(\cdot|\tilde{x}_u)$.

evant features $\mathcal{K}$ before evaluating the policy based on the following proposition.

Proposition 3.2. Under Assumption 2.1, there exists a $d \times d$ matrix K^π such that $\forall (x, a) \in \mathcal{X} \times \mathcal{A}$

$$\mathbb{E}_{x' \sim P(\cdot|x, a)}[\phi^\pi(x')^\top | x, a] = \phi(x, a)^\top K^\pi, \quad (3.1)$$

where $\phi^\pi(x) = \sum_a \pi(a|x) \phi(x, a)$ and all but s rows of K^π are identically zero.

Thus we propose to estimate the set of relevant features $\mathcal{K}$ using group lasso (Yuan and Lin, 2006). Once the relevant feature set is identified, any regular policy evaluation method can be used over the learned feature set $\hat{\mathcal{K}}$. In Algorithm 2, for the ease of comparability with the previous method, we consider vanilla fitted Q-evaluation.

Remark 3.3. One may wonder whether it is necessary to refit the iterative regression and why not simply use the estimated $\hat{K}^\pi$ to get a plug-in estimator. This is because refitting typically performs strictly better than direct regularized learning and has less bias, as long as the feature selection succeeds (Belloni et al., 2013).

4. Performance Bounds For Sparse Off-Policy Evaluation

We study the finite-sample estimation error of Algorithms 1, 2. All the technical proofs are deferred to the Appendix. Let Σ be the expected uncentered covariance matrix of the

Algorithm 2 Post model-selection fitted Q-evaluation

1: **Input:** initial distribution ξ_0 , target policy π , dataset $\mathcal{D}$ of size N , $m := N$, number of iterations T , $\lambda_2, \lambda_3 > 0$.

2: Estimate $\hat{K}^\pi$ through: $\hat{K}^\pi =$

$$\underset{K \in \mathbb{R}^{d \times d}}{\operatorname{argmin}} \left\{ \lambda_2 \sum_{j=1}^d \|K_{j\cdot}\|_2 + \frac{1}{Nd} \sum_{n=1}^N \|\phi^\pi(x'_n)^\top - \phi(x_n, a_n)^\top K\|_2^2 \right\}. \quad (3.2)$$

3: Find $\hat{\mathcal{K}} = \{j \in [d] : \|\hat{K}_{j\cdot}^\pi\|_2 \neq 0\}$, the estimated set of relevant features. Here, $\hat{K}_{j\cdot}^\pi$ refers to the j th row of $\hat{K}^\pi$.

4: **Initialize:** $\hat{w}_0 = 0 \in \mathbb{R}^{|\hat{\mathcal{K}}|}$.

5: **for** $t = 1, 2, \dots, T$ **do**

6: Calculate regression targets for $n \in [N]$:

$$y_n = r(x_n, a_n) + \gamma \sum_a \pi(a|x'_n) [\phi(x'_n, a)_{\hat{\mathcal{K}}}]^\top \hat{w}_{t-1}.$$

7: Update $\hat{w}_t$ through $\hat{w}_{t+1} =$

$$\underset{w \in \mathbb{R}^{|\hat{\mathcal{K}}|}}{\operatorname{argmin}} \left\{ \frac{1}{N} \sum_{n=1}^N (y_n - [\phi(x_n, a_n)_{\hat{\mathcal{K}}}]^\top w)^2 + \lambda_3 \|w\|_2^2 \right\}.$$

8: **end for**

9: **Output:** $\hat{v}_{\hat{w}_T}^\pi = \frac{1}{m} \sum_{u=1}^m \Pi_{[0, 1/(1-\gamma)]}(Q_{\hat{w}_T}(\tilde{x}_u, \tilde{a}_u))$ with $\tilde{x}_u \sim \xi_0$, $\tilde{a}_u \sim \pi(\cdot|\tilde{x}_u)$.

batch data, given by

$$\Sigma := \mathbb{E} \left[\frac{1}{L} \sum_{h=0}^{L-1} \phi(x_h^{(1)}, a_h^{(1)}) \phi(x_h^{(1)}, a_h^{(1)})^\top \right], \quad (4.1)$$

where L is the length of one episode.

Remark 4.1. Our analysis applies to multiple behavior policies and our algorithms do not require the knowledge of the behavior policy. For example, if we have M behavior policies $\{\pi_1, \dots, \pi_M\}$ and one policy is selected following distribution p_π at the beginning of each episode to collect data, we could define Σ similarly but the expectation will take also w.r.t the randomness with p_π .

We need a notion of restricted eigenvalue that is common in high-dimensional statistics (Bickel et al., 2009; Bühlmann and Van De Geer, 2011).

Definition 4.2 (Restricted eigenvalue). Given a positive semi-definite matrix $Z \in \mathbb{R}^{d \times d}$ and integer $s \geq 1$, define

the restricted minimum eigenvalue of Z as $C_{\min}(Z, s) :=$

$$\min_{S \subset [d], |S| \leq s} \min_{\beta \in \mathbb{R}^d} \left\{ \frac{\langle \beta, Z\beta \rangle}{\|\beta_S\|_2^2} : \|\beta_{S^c}\|_1 \leq 3\|\beta_S\|_1 \right\}.$$

4.1. Finite-sample error bounds of Algorithm 1

We first provide a statistical error bound for the Lasso fitted Q-evaluation (Algorithm 1).

Theorem 4.3 (OPE error bound for Algorithm 1). Suppose Assumptions 2.1, 2.2 hold and $C_{\min}(\Sigma, s) > 0$. Let Algorithm 1 take N samples satisfying $N \gtrsim s^2 \log(d/\delta) L(1 - \gamma)^{-1} / C_{\min}(\Sigma, s)$. Set the number of iterations $T = \Theta(\log(N/(1 - \gamma)) / (1 - \gamma))$ and $\lambda_1 = (1 - \gamma)^{-1} \sqrt{T \log(2d/\delta) / N}$. Then, with probability at least $1 - 2\delta$,

$$|\hat{v}_{\hat{w}_T}^\pi - v^\pi| \lesssim \frac{1}{C_{\min}(\Sigma, s)} \sqrt{\frac{s^2 \log(d/\delta)}{N(1 - \gamma)^5}}. \quad (4.2)$$

Theorem 4.3 shows that the OPE error depends linearly on $s/C_{\min}(\Sigma, s)$. For comparison, when the sparsity is not considered, Duan and Wang (2020) proved the error bound (using our notations) of the form

$$|\hat{v}^\pi - v^\pi| \lesssim \sqrt{d / (C_{\min}(\Sigma, d) N (1 - \gamma)^4)}.$$

From the Definition 4.2, $C_{\min}(\Sigma, d) < C_{\min}(\Sigma, s)$. Comparing the two results (setting $\delta = 1/N$), we expect the new error bound to be significantly tighter, i.e., $C_{\min}(\Sigma, d) \frac{s^2 \log(dN)}{1 - \gamma} \ll C_{\min}^2(\Sigma, s) d$, when there is a high level of sparsity ($s \ll d$).

Remark 4.4. The restricted eigenvalue $C_{\min}(\Sigma, s)$ quantifies how well the behavior policy covers the state-action feature space. We need $C_{\min}(\Sigma, s) > 0$ meaning that the data is well-conditioned or the behavior policy provides good coverage over relevant features. In the tabular case, where $\phi(x, a)$ is a basis vector in $\mathbb{R}^{|\mathcal{X}| \times |\mathcal{A}|}$, then $C_{\min} = \min_{x,a} \mu^\pi(x, a) > 0$ implies that the behavior policy visits all state-action pairs with positive probability, where $\mu^\pi(x, a)$ is the visitation distribution under policy π and initial distribution ξ_0 : $\mu^\pi(x, a) = \frac{1}{L} \sum_{h=0}^{L-1} \mathbb{E}^\pi[\mathbb{I}((x_h, a_h) = (x, a))]$.

To ensure the success of policy evaluation/optimization with linear function approximation, similar assumptions regarding Σ in RL literature also appear in Abbasi-Yadkori et al. (2019), Duan and Wang (2020), Lazic et al. (2020). In addition, Agarwal et al. (2020) showed how to design a policy that can collect data that satisfies this condition.

4.2. Finite-sample error bounds of Algorithm 2

Next, we give a result for the post-selection model estimator for OPE (Algorithm 2). We will show that this algorithm provides a more accurate estimate under the additional condition that every relevant feature plays a nontrivial role.

Assumption 4.5 (Minimal signal strength). For some given $\delta > 0$, the *minimum signal strength* satisfies

$$\min_{j \in \mathcal{K}} \|K_{j \cdot}^\pi\|_2 / \sqrt{d} \geq \frac{64\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{2 \log(2d^2/\delta)}{N}},$$

where $K_{j \cdot}^\pi$ is the j th row of K^π defined in Eq. (3.1).

Then we provide a critical lemma showing that the group lasso step in Algorithm 2 is guaranteed to identify a sufficiently sparse feature set including all the relevant features with high probability.

Lemma 4.6 (Feature screening). Suppose Assumptions 2.1, 4.5 hold and $C_{\min}(\Sigma, s) > 0$. Set the regularization parameter $\lambda_2 = 4\sqrt{2 \log(2d^2/\delta) / (Nd)}$ for some $\delta > 0$ and let the sample size satisfy $N \gtrsim Ls^2 \log(d/\delta) / C_{\min}^2(\Sigma, s)$. Then with probability at least $1 - \delta$, the size of learned relevant feature set $\hat{\mathcal{K}}$ satisfies $|\hat{\mathcal{K}}| \lesssim s$ and $\hat{\mathcal{K}} \supseteq \mathcal{K}$ where $\mathcal{K}$ is the true relevant feature set of M .

The detailed proof of Lemma 4.6 is deferred to Appendix B.2. This lemma implies the model group lasso selected is sufficiently sparse and can identify all the important features with high probability. Thus, we have $|\hat{\mathcal{K}}| \asymp s$.

Now we analyze the policy evaluation error of Algorithm 2. According to Cramer-Rao lower bound for tabular OPE (Jiang and Li, 2016) and the minimax lower bound for OPE with linear function approximation (Duan and Wang, 2020), we expect the optimal OPE error to depend on the *distribution mismatch* between the target policy and the behavior policy that generated the data. To define the notion of distribution mismatch, we first need the notion of occupancy measures:

Definition 4.7 (Occupancy measures). Let $\bar{\mu}$ be the expected occupancy measure of observations $\{(x_n, a_n)\}_{n=1}^N$: $\bar{\mu}(x, a) = \sum_{n=1}^N \mathbb{P}(x_n = x, a_n = a) / N$ and μ^π be the discounted occupancy distribution of (x_h, a_h) under policy π and initial distribution ξ_0 : $\mu^\pi(x, a) = (1 - \gamma) \mathbb{E}^\pi[\sum_{h=0}^{\infty} \gamma^h \mathbb{I}(x_h = x, a_h = a)]$, $\forall x, a$.

Inspired by Theorem 5 of Duan and Wang (2020), we will measure the distribution mismatch using *restricted chi-square divergences* between $\bar{\mu}$ and μ^π .

Definition 4.8 (Restricted chi-square divergence). Let $\mathcal{G}$ be a set of real-valued functions over $\mathcal{X}$ and let p_1 and p_2 be

probability distributions over $\mathcal{X}$. We define the $\mathcal{G}$ -restricted chi-square divergence (or $\chi_{\mathcal{G}}^2$ -divergence) between p_1 and p_2 as

$$\chi_{\mathcal{G}}^2(p_1, p_2) := \sup_{f \in \mathcal{G}} \frac{\mathbb{E}_{p_1}[f(x)]^2}{\mathbb{E}_{p_2}[f(x)^2]} - 1.$$

By using the feature screening Lemma 4.6, and a similar analysis as by Duan and Wang (2020), we obtain the following instance-dependent error bound for sparse off-policy evaluation.

Theorem 4.9 (Instance-dependent error bound for sparse OPE). Suppose Assumptions 2.1, 2.2 and 4.5 hold and $C_{\min}(\Sigma, s) > 0$. Let $\delta \in (0, 1)$ and assume that Algorithm 2 is fed with N samples satisfying $N \gtrsim L \log(d^2/\delta)s^2/C_{\min}^2(\Sigma, s) + \gamma^2 L \log(s/\delta)s/(1-\gamma)^2$. Set $\lambda_2 = 4\sqrt{2 \log(2d^2/\delta)/Nd}$, $\lambda_3 = \lambda_{\min}(\Sigma) \log(12|\hat{\mathcal{K}}|/\delta)L|\hat{\mathcal{K}}|$. Letting the number of iterations $T \rightarrow \infty$, the following holds with probability at least $1 - 3\delta$,

$$|\hat{v}_{\hat{\omega}}^{\pi} - v^{\pi}| \lesssim \sqrt{1 + \chi_{\mathcal{G}(\hat{\mathcal{K}})}^2(\mu^{\pi}, \bar{\mu})} \sqrt{\frac{\log(1/\delta)}{N(1-\gamma)^4}}, \quad (4.3)$$

where $\bar{\mu}$ is the data generating distribution, $\mathcal{G}(\hat{\mathcal{K}})$ is the reduced feature space.

Remark 4.10 (Reduced distribution mismatch via sparse feature selection.). The OPE error bound of Theorem 4.9 depends on the statistics $\chi_{\mathcal{G}(\hat{\mathcal{K}})}^2(\mu^{\pi}, \bar{\mu})$ that quantifies the distribution mismatch between data and the target policy. This result implies the uncertainty for evaluating a new policy from batch data crucially and jointly depends on the two distributions as well as the function class used for fitting. When $\hat{\mathcal{K}}$ is a small subset of $[d]$, we have $\chi_{\mathcal{G}(\hat{\mathcal{K}})}^2 \ll \chi_{\mathcal{G}([d])}^2$. Therefore our instance-dependent error bound is expected to be significantly smaller than its counterpart that does not exploit sparsity.

4.3. Minimax lower bound for sparse OPE

To complete the picture, we provide a minimax lower bound of off-policy evaluation using sparse linear function approximation. The proof is an adaptation of the respective lower bound proof for linear case (Theorem 3 in Duan and Wang (2020)). It implies the bound in Theorem 4.9 is nearly minimax-optimal.

Theorem 4.11. There exists a DMDP instance with feature map $\phi(\cdot, \cdot) \in \mathbb{R}^d$ satisfying $\|\phi(x, a)\|_{\infty} \leq 1$ for all $(x, a) \in \mathcal{X} \times \mathcal{A}$ and Assumption 2.1 such that if $N \gtrsim sL(1-\gamma)^{-1}$,

then

$$\inf_{\hat{v}^{\pi}} \sup_{\phi, M \in \mathcal{M}_{\phi, s}(\mathcal{X}, \mathcal{A}, \gamma)} \mathbb{P}_M \left(|\hat{v}^{\pi}(\mathcal{D}) - v^{\pi}| \gtrsim \frac{1}{(1-\gamma)^2} \sqrt{1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^{\pi}, \bar{\mu})} \sqrt{\frac{1}{N}} \right) \geq \frac{1}{6},$$

where $\phi \in (\mathbb{R}^d)^{\mathcal{X} \times \mathcal{A}}$, $\mathbb{P}_M$ is the probability measure under the DMDP instance M and $\hat{v}^{\pi}(\cdot)$ sweeps through all algorithms that estimate values based on data $\mathcal{D}$.

Remark 4.12. It is worth to mention that in Theorem 4.11, the distribution mismatch term $1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^{\pi}, \bar{\mu})$ may also contain a $1 - \gamma$ term in the worse case. Thus the lower bound of sparse OPE also has a $\sqrt{1/(1-\gamma)^3}$ dependency in the worse case that matches the result for the lower bound of sparse batch policy optimization in Theorem 5.2. This matches the result in Yin and Wang (2020) (Theorem 3.1&Remark 3.3) since they define n as the number of episodes while we define N as the number of total sample pairs.

5. Sparsity-Aware Batch Policy Optimization

We extend our analysis to batch policy learning problem with sparse linear function approximation. Consider the Lasso fitted Q-iteration (see Algorithm 3) that has been studied in Calandriello et al. (2014) as a special case of an algorithm for sparse multi-task RL. It resembles Algorithm 1 except for that it calculates the regression target with an additional “max” operation.

Algorithm 3 Lasso-regularized fitted Q-iteration

- 1: **Input:** T folds of dataset $\{\mathcal{D}_t\}_{t=1}^T$, regularization parameter λ_1 , $\hat{w}_0 = 0 \in \mathbb{R}^d$.
 - 2: **Repeat:**
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: Calculate regression targets: for $(x_i, a_i, x'_i) \in \mathcal{D}_t$,
 $y_i = \max_{a \in \mathcal{A}} Q_{\hat{w}_{t-1}}(x'_i, a)$.
 - 5: Based on $\{(\Pi_{[0,1/(1-\gamma)]} y_i, \phi(x_i, a_i))\}_{(x_i, a_i, x'_i) \in \mathcal{D}_t}$,
 fit $\hat{w}_t$ through Lasso as in Algorithm 1.
 - 6: **end for**
 - 7: **Output:** policy $\hat{\pi}_T(\cdot|x) = \max_{a \in \mathcal{A}} Q_{\hat{w}_T}(x, a)$, $\forall x \in \mathcal{X}$.
-

The next theorem proves the approximate optimality of the learned policy using Lasso fitted Q-iteration.

Theorem 5.1. Suppose Assumptions 2.1-2.2 holds and $C_{\min}(\Sigma, s) > 0$. Let $N \gtrsim s^2 L(1-\gamma)^{-1}/C_{\min}(\Sigma, s)$. Let Algorithm 3 take $T = \Theta(\log(N/(1-\gamma))/(1-\gamma))$ and

$\lambda_1 = (1 - \gamma)^{-1} \sqrt{T \log(2d/\delta)/N}$. Then, with probability at least $1 - \delta$,

$$\|v^{\hat{\pi}_T} - v^*\|_\infty \lesssim \frac{s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(d/\delta)}{N(1 - \gamma)^7}}. \quad (5.1)$$

Theorem 5.1 suggests that the sample size needed to get a good policy depends mainly on the number of relevant features s , instead of the large ambient dimension d , provided that the data is well-conditioned. This result is not surprising: Calandriello et al. (2014) gave a similar upper bound for sparse FQI for the setting of generative model. Le et al. (2019) provided a generalization theory for policy evaluation/learning with a general function class and their error bound depends on the VC-dimension of the class, but it requires a stronger coefficient concentration condition.

5.1. Minimax lower bound for policy optimization

In the end, we study the fundamental limits of sparse batch policy learning. We establish an information-theoretic minimax lower bound that nearly match the aforementioned upper bound.

Theorem 5.2. Let $\gamma \geq \frac{2}{3}$ and $\hat{\pi}$ denote an algorithm that maps dataset $\mathcal{D}$ to a policy $\hat{\pi}(\mathcal{D})$. If $N \gtrsim sL(1 - \gamma)^{-1}$, then for any $\hat{\pi}$, there always exists a DMDP instance M with feature map $\phi(\cdot, \cdot) \in \mathbb{R}^d$ satisfying $\|\phi(x, a)\|_\infty \leq 1$ for all $(x, a) \in \mathcal{X} \times \mathcal{A}$ and Assumption 2.1, such that

$$\mathbb{P}_M \left(v_{\xi_0}^* - v_{\xi_0}^{\hat{\pi}(\mathcal{D})} \gtrsim \sqrt{\frac{1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu})}{N(1 - \gamma)^3}} \right) \geq \frac{1}{6}, \quad (5.2)$$

where μ^* is the discounted state-action occupancy measure of π^* . In addition, we have

$$\mathbb{P}_M \left(v_{\xi_0}^* - v_{\xi_0}^{\hat{\pi}(\mathcal{D})} \gtrsim \sqrt{\frac{s/C_{\min}(\Sigma, s)}{N(1 - \gamma)^3}} \right) \geq \frac{1}{6}. \quad (5.3)$$

Theorems 5.1, 5.2 show that the statistical error of batch policy learning is fundamentally determined by the ratio $s/C_{\min}(\Sigma, s)$. Note that there remains a gap $\sqrt{s/C_{\min}(\Sigma, s)}$ between Theorems 5.1 and 5.2, due to the nature of Lasso regression.

Remark 5.3. Minimax sample complexity lower bound for solving MDP has been studied in the setting with a generative model that allows querying any (s, a) for independent samples. Azar et al. (2013) constructed a hard instance of tabular MDP and, by reducing policy optimization to a testing a Bernoulli distribution, proved a lower bound

$SA/(1 - \gamma)^3$ which is known to be sharp. Yang and Wang (2019) extended the construction to linear MDP and show that the sample complexity lower bound is $d/(1 - \gamma)^3$ under a generative model. There also exists matching upper bound in the same setting.

Our Theorem 5.2 applies to the setting of batch episodic data where are highly dependent. Due to this major difference, we have to use a more intricate proof based on likelihood test to establish a minimax lower bound. Further, Theorem 5.2 characterizes for the first time that the lower bound depends on the minimal eigenvalue of the data's population covariance.

Remark 5.4. Earlier results such as those of (Munos and Szepesvári, 2008; Antos et al., 2008; Le et al., 2019) require stronger forms of concentration condition that the state-action occupancy measure (or a ratio involving this measure) is entrywisely bounded across all policies. Such entrywise bound can be very large if the state-action space $\mathcal{X}$ is large. In contrast, our results only require that the data's covariance Σ is well-conditioned on restricted supports, which is a much weaker assumption.

Further, one can use the empirical minimal eigenvalue to get a rough error estimate. Theorem 5.2 further validates that the minimal eigenvalue indeed determines the statistical limit of batch policy optimization. The result is the first of its kind to our best knowledge.

6. Related Work

Off-policy evaluation (OPE). OPE often serves the starting point of batch RL. A direct approach was to fit value function from data using approximate dynamic programming, e.g., the policy evaluation analog of fitted Q-iteration (Ernst et al., 2005; Munos and Szepesvári, 2008; Le et al., 2019) or least square policy iteration (Lagoudakis and Parr, 2003). Another popular class of OPE methods used importance sampling to get unbiased value estimate of a new policy (Precup et al., 2000) and improved by doubly-robust technique to reduce the variance (Jiang and Li, 2016; Thomas and Brunskill, 2016). To alleviate the curse of horizon (Li et al., 2015; Jiang and Li, 2016; Yin and Wang, 2020), marginalized importance sampling was suggested by estimating state marginal importance ratio without reweighting the entire trajectory (Hallak and Mannor, 2017; Liu et al., 2018; Xie et al., 2019). In general, estimating marginalized importance ratio could be sample-expensive and even intractable. Recently, practical duality-inspired methods were developed for estimating this ratio using function ap-

proximation (Nachum et al., 2019; Uehara and Jiang, 2019; Zhang et al., 2020a;b; Yang et al., 2020).

On the theoretical side, Uehara and Jiang (2019); Yin and Wang (2020); Kallus and Uehara (2020) established asymptotic optimality and efficiency for OPE in the tabular setting. Duan and Wang (2020) showed that fitted Q-evaluation with linear function approximation is minimax optimal and provided matching upper and lower bounds that depend on a distribution mismatch term. Another closely related work was by Le et al. (2019) who studied batch policy evaluation and optimization with more general function approximation. They showed the complexity of batch RL depends on the complexity of the function class, assuming a ‘‘concentration coefficient’’ condition (Munos and Szepesvári, 2008) that the state-action visitation density is bounded entrywisely across policies. More recently, Uehara and Jiang (2019) provided theoretical investigations into OPE using general function approximators for marginalized importance weights and value functions but did not show the statistical optimality.

Sparse learning in RL. The use of feature selection by regularization in RL has been explored in a number of prior works. Kolter and Ng (2009); Geist and Scherrer (2011); Hoffman et al. (2011); Painter-Wakefield and Parr (2012) studied *on-policy* evaluation with ℓ_1 -regularization for temporal-difference (TD) learning but none of them come with a theoretical analysis. Liu et al. (2012) studied off-policy evaluation by regularized TD learning but only provided algorithmic convergence guarantee without statistical error analysis.

Ghavamzadeh et al. (2011); Geist et al. (2012) proposed Lasso-TD with finite-sample statistical analysis for estimating the value function in *Markov reward process*. In particular, they derived in-sample *prediction error* bound $\mathcal{O}((s \log(d)/\psi n)^{1/2})$ under ψ -minimum eigenvalue condition on the empirical feature gram matrix. Although this bound also has no polynomial dependency on d , in-sample prediction error generally can not be translated to the estimation error of target policy in the OPE problem and their bound can not characterize the distribution mismatch between behavior policy and target policy. On the other hand, no minimax lower bound has been investigated so far.

In addition, Farahmand et al. (2008; 2016) considered ℓ_2 -regularization in fitted Q-iteration/policy iteration for policy optimization in a reproducing kernel Hilbert space, and finite-sample performance bounds for these algorithms were proved built on a coefficient concentration condition. Calan-


driello et al. (2014) developed Lasso fitted Q-iteration for sparse multi-task reinforcement learning and assumed a generative model (Kakade et al., 2003) for sampling transitions. Ibrahimi et al. (2012) derived a $\mathcal{O}(p\sqrt{T})$ regret bound in high-dimensional sparse linear quadratic systems where p is the dimension of the state space. More recently, Hao et al. (2020) studied the hardness of online reinforcement learning in fixed horizon, sparse linear Markov decision process.

Sparse linear regression. Sparse regression receives considerable attention in high-dimensional statistics in the past decade. Lasso (Tibshirani, 1996), is arguably the most widely used method to conduct sparse regression. Theoretical analysis of Lasso is well-studied in Zhao and Yu (2006); Bickel et al. (2009); Wainwright (2009). For a thorough review of Lasso as well as high-dimensional statistics, we refer the readers to Hastie et al. (2015); Wainwright (2019); Bühlmann and Van De Geer (2011). However, extending existing analysis from regression to batch RL is much more involved due to the complex optimization structure, non-i.i.d data collection, and covariate shift.

7. Experiment

We conducted experiments with a mountain car example. We use 800 radial basis functions for linear value function approximation. The number of episodes collected by behavior policies ranges from 2 to 100. We compare our Lasso-FQE with the standard FQE with ℓ_2 ridge regularization. For each algorithm we report the performance for the best regularization parameter λ in the range $\{0.02, 0.05, 0.1, 0.2, 0.5\}$. We pick a near-optimal policy as the target, and we measure the estimation error by $|\hat{v} - v^*|/|v^*|$. All the results are averaged over 20 runs.

The left panel of Figure 7 shows that our Lasso-FQE clearly has smaller estimation error compared with FQE, proving the sparse feature selection is effective in a practical RL example. The right panel of Figure 7 demonstrates how the distribution mismatch (χ^2 -divergence term) affects OPE error (with sample size fixed). We constructed multiple behavior policies with varying levels of ϵ -greedy noise, and plot their OPE error against their (restricted) χ^2 -divergence from the target policy. The results confirm our theorems (Theorem 4.9&4.11) that the (restricted) chi-square divergence sharply determines the (sparse) OPE error.



8. Conclusion

In this work we focus on high-dimensional batch RL using sparse linear function approximation. While previous work in RL recognized the possibility of bringing tools from sparse learning to RL, they lacked a clean theoretical framework and formal results. By building on the strength of the linear DMDP framework, our result show that learning and planning in linear DMDPs can be done in the “feature space” even in the presence of sparsity and when only batch data is available.

Acknowledgements

Csaba Szepesvári gratefully acknowledges funding from the Canada CIFAR AI Chairs Program, Amii and NSERC. Mengdi Wang gratefully acknowledges funding from the U.S. National Science Foundation (NSF) grant CMMI1653435, Air Force Office of Scientific Research (AFOSR) grant FA9550-19-1-020, and C3.ai DTL.

References

- Yasin Abbasi-Yadkori, Peter Bartlett, Kush Bhatia, Nevena Lazic, Csaba Szepesvári, and Gellért Weisz. Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*, pages 3692–3702, 2019.
- Alekh Agarwal, Mikael Henaff, Sham Kakade, and Wen Sun. Pc-pg: Policy cover directed exploration for provable policy gradient learning. *arXiv preprint arXiv:2007.08459*, 2020.
- András Antos, Csaba Szepesvári, and Rémi Munos. Fitted Q-iteration in continuous action-space MDPs. In *Advances in neural information processing systems*, pages 9–16, 2008.
- Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J Kappen. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3):325–349, 2013.
- I. R. Bellman, R. Kalaba, and B. Kotkin. Polynomial approximation – a new computational technique in dynamic programming. *Math. Comp.*, 17(8):155–161, 1963.
- Alexandre Belloni, Victor Chernozhukov, et al. Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547, 2013.
- Dimitri P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific, 1995.
- Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- Peter J Bickel, Ya’acov Ritov, Alexandre B Tsybakov, et al. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- Peter Bühlmann and Sara Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- F. Bunea, A. Tsybakov, and M. Wegkamp. Sparsity oracle inequalities for the Lasso. *Electronic Journal of Statistics*, 1:169–194, 2007.
- Daniele Calandriello, Alessandro Lazaric, and Marcello Restelli. Sparse multi-task reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 819–827, 2014.
- Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051, 2019.
- Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- Yaqi Duan and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. *International Conference on Machine Learning*, 2020.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.
- A.m. Farahmand, M. Ghavamzadeh, Cs. Szepesvári, and S. Mannor. Regularized policy iteration with nonparametric function spaces. *JMLR*, 17:1–66, 2016.
- Amir massoud Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor. Regularized fitted Q-iteration: Application to planning. In *European Workshop on Reinforcement Learning*, pages 55–68. Springer, 2008.
- Matthieu Geist and Bruno Scherrer. ℓ^1 -penalized projected Bellman residual. In *European Workshop on Reinforcement Learning*, pages 89–101. Springer, 2011.
- Matthieu Geist, Bruno Scherrer, Alessandro Lazaric, and Mohammad Ghavamzadeh. A Dantzig selector approach to temporal difference learning. In *Proceedings of the 29th International Conference on Machine Learning*, pages 347–354, 2012.
- Mohammad Ghavamzadeh, Alessandro Lazaric, Rémi Munos, and Matthew Hoffman. Finite-sample analysis of Lasso-TD. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1177–1184, 2011.

- Assaf Hallak and Shie Mannor. Consistent on-line off-policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1372–1383. JMLR. org, 2017.
- Botao Hao, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Online sparse reinforcement learning. *arXiv preprint arXiv:2011.04018*, 2020.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC press, 2015.
- Matthew W Hoffman, Alessandro Lazaric, Mohammad Ghavamzadeh, and Rémi Munos. Regularized least squares temporal difference learning with nested ℓ^2 and ℓ^1 penalization. In *European Workshop on Reinforcement Learning*, pages 102–114. Springer, 2011.
- Morteza Ibrahimi, Adel Javanmard, and Benjamin V Roy. Efficient reinforcement learning for high dimensional linear quadratic systems. In *Advances in Neural Information Processing Systems*, pages 2636–2644, 2012.
- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661, 2016.
- Sham Machandranath Kakade et al. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- Nathan Kallus and Masatoshi Uehara. Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *Journal of Machine Learning Research*, 21(167):1–63, 2020.
- J Zico Kolter and Andrew Y Ng. Regularization and feature selection in least-squares temporal difference learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 521–528, 2009.
- Michail G Lagoudakis and Ronald Parr. Least-squares policy iteration. *Journal of machine learning research*, 4(Dec):1107–1149, 2003.
- Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.
- Nevena Lazic, Dong Yin, Mehrdad Farajtabar, Nir Levine, Dilan Gorur, Chris Harris, and Dale Schuurmans. A maximum-entropy approach to off-policy evaluation in average-reward mdps. *Conference on Neural Information Processing Systems*, 2020.
- Hoang Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *International Conference on Machine Learning*, pages 3703–3712, 2019.
- 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.
- Lihong Li, Remi Munos, and Csaba Szepesvari. Toward minimax off-policy value estimation. In *Artificial Intelligence and Statistics*, pages 608–616, 2015.
- Bo Liu, Sridhar Mahadevan, and Ji Liu. Regularized off-policy TD-learning. In *Advances in Neural Information Processing Systems*, pages 836–844, 2012.
- Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, pages 5356–5366, 2018.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(May): 815–857, 2008.
- Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. DualDICE: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, pages 2315–2325, 2019.
- Christopher Painter-Wakefield and Ronald Parr. Greedy algorithms for sparse reinforcement learning. In *Proceedings of the 29th International Conference on Machine Learning*, pages 867–874, 2012.
- Doina Precup, Richard S Sutton, and Satinder Singh. Eligibility traces for off-policy policy evaluation. In *ICML’00 Proceedings of the Seventeenth International Conference on Machine Learning*, 2000.
- Irina Rish and Genady Grabarnik. *Sparse Modeling: Theory, Algorithms, and Applications*. CRC Press, 2014.
- Paul J Schweitzer and Abraham Seidmann. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110(2):568–582, 1985.
- Csaba Szepesvári. *Algorithms for Reinforcement Learning*. Morgan and Claypool, 2010.
- Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.
- Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Masatoshi Uehara and Nan Jiang. Minimax weight and Q-function learning for off-policy evaluation. *arXiv preprint arXiv:1910.12809*, 2019.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Martin J Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE transactions on information theory*, 55(5): 2183–2202, 2009.

Martin J Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.

Ruosong Wang, Dean P Foster, and Sham M Kakade. What are the statistical limits of offline rl with linear function approximation? *arXiv preprint arXiv:2010.11895*, 2020.

Tengyang Xie, Yifei Ma, and Yu-Xiang Wang. Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems*, pages 9665–9675, 2019.

Lin Yang and Mengdi Wang. Sample-optimal parametric Q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004, 2019.

Mengjiao Yang, Ofir Nachum, Bo Dai, Lihong Li, and Dale Schuurmans. Off-policy evaluation via the regularized lagrangian. *arXiv preprint arXiv:2007.03438*, 2020.

Ming Yin and Yu-Xiang Wang. Asymptotically efficient off-policy evaluation for tabular reinforcement learning. *arXiv preprint arXiv:2001.10742*, 2020.

Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. GenDICE: Generalized offline estimation of stationary values. *arXiv preprint arXiv:2002.09072*, 2020a.

Shangdong Zhang, Bo Liu, and Shimon Whiteson. GradientDICE: Rethinking generalized offline estimation of stationary values. *arXiv preprint arXiv:2001.11113*, 2020b.

Peng Zhao and Bin Yu. On model selection consistency of lasso. *Journal of Machine learning research*, 7(Nov):2541–2563, 2006.

The appendix is organized as follows.

- Appendix A. Proof of Proposition 3.2.
- Appendix B. Proofs of results about off-policy evaluation in Section 4.
 - Appendix B.1. Proof of Theorem 4.3.
 - Appendix B.2. Proof of Lemma 4.6.
 - Appendix B.3. Proof of Theorem 4.9.
- Appendix C. Proof of Theorem 5.1: upper bound of batch policy optimization.
- Appendix D. Proof of Theorem 5.2: minimax lower bound of batch policy optimization.
- Appendix E. Proofs of auxiliary lemmas.
- Appendix F. Some supporting lemmas.

A. Proof of Proposition 3.2

Proof of Proposition 3.2. Fix a policy π . Since M is a linear MDP, for every $i \in [d]$ there exist $w_i \in \mathbb{R}^d$ such that $P\phi_i^\pi = g_{w_i}$. Thus, for any $(x, a) \in \mathcal{X} \times \mathcal{A}$,

$$\begin{aligned} \mathbb{E}_{x' \sim P(\cdot|x,a)}[\phi^\pi(x')^\top | x, a] &= (P\phi_1^\pi, \dots, P\phi_d^\pi) = (\phi(x, a)^\top w_1, \dots, \phi(x, a)^\top w_d) \\ &= \phi(x, a)^\top (w_1, \dots, w_d). \end{aligned}$$

Thus, Eq. (3.1) holds if we choose

$$K^\pi = (w_1, \dots, w_d).$$

Under the sparsity assumption, Assumption 2.1, there exists $\mathcal{K} \subset [d]$ such that $w_{ij} = 0$ when $j \notin \mathcal{K}$. This shows that all but $|\mathcal{K}|$ rows of K^π are identically zero, finishing the proof. $\square$

B. Proofs of Off-Policy Evaluation

B.1. Proof of Theorem 4.3: Lasso fitted Q-evaluation

Recall that we split the whole dataset $\mathcal{D}$ into T folds and each fold consists of R episodes or RL sample transitions. At t th phase, only the fresh fold of dataset $\mathcal{D}_t = \{(x_i^{(t)}, a_i^{(t)}, x_i^{(t)'})\}_{i=1}^{RL}$ is used.

Step 1: Approximate value iteration. We first show that the execution of Algorithm 1 is equivalent to approximate value iteration. Denote a Lasso estimator with respect to a function V at t th phase:

$$\hat{w}_t(V) = \operatorname{argmin}_{w \in \mathbb{R}^d} \left(\frac{1}{RL} \sum_{i=1}^{RL} \left(V(x_i^{(t)'}) - \phi(x_i^{(t)}, a_i^{(t)})^\top w \right)^2 + \lambda_1 \|w\|_1 \right). \quad (\text{B.1})$$

Note that $\hat{w}_t(\cdot)$ only depends data collected at the t th phase. Define the parameterized value function as

$$V_w^\pi(x) = \sum_a \pi(a|x) (r(x, a) + \gamma \phi(x, a)^\top w).$$

Define an approximate Bellman operator for target policy π , i.e. $\hat{\mathcal{T}}_\pi^{(t)} : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ as:

$$[\hat{\mathcal{T}}_\pi^{(t)} V](x) := V_{\hat{w}_t(V)}^\pi(x) = \sum_a \pi(a|x) (r(x, a) + \gamma \phi(x, a)^\top \hat{w}_t(V)). \quad (\text{B.2})$$

Note this $\widehat{\mathcal{T}}_\pi^{(t)}$ is a randomized operator that only depends data in the t th fold. It is easy to see that if $(\widehat{w}_t)_{t=1}^T$ is the sequence of weights computed in Algorithm 1 then $\widehat{w}_t = \widehat{w}_t(\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi)$ and also

$$V_{\widehat{w}_t}^\pi = \widehat{\mathcal{T}}_\pi^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi. \quad (\text{B.3})$$

Step 2: Linear representation of Bellman operator. Recall that the true Bellman operator for target policy π , i.e. $\mathcal{T}_\pi : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ is defined as

$$[\mathcal{T}_\pi V](x) := \sum_a \pi(a|x) \left(r(x, a) + \gamma \sum_{x'} P(x'|x, a) V(x') \right). \quad (\text{B.4})$$

We first verify for each phase $t \in [T]$, $\mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_t}^\pi$ has a linear representation. From Assumption 2.1, there exists a vector $\bar{w}_t \in \mathbb{R}^d$ with $\bar{w}_{t,k} = 0$ if $k \notin \mathcal{K}$ such that

$$[\mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_t}^\pi](x) = \sum_a \pi(a|x) \left(r(x, a) + \gamma \phi(x, a)^\top \bar{w}_t \right). \quad (\text{B.5})$$

It shows that $\mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_t}^\pi$ has a linear representation if the reward could also be linearly represented. For notation simplicity, we drop the superscript of $x_i^{(t)}$ and $a_i^{(t)}$ for the following derivations when there is no ambiguity.

Step 3: Sparse linear regression. We interpret $\bar{w}_t$ as the ground truth of the lasso estimator in Algorithm 1 at phase t , in terms of the following sparse linear regression:

$$\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi(x'_i) = \phi(x_i, a_i)^\top \bar{w}_t + \varepsilon_i, i = 1 \dots, RL, \quad (\text{B.6})$$

where $\varepsilon_i = \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi(x'_i) - \phi(x_i, a_i)^\top \bar{w}_t$. Define a filtration $\{\mathcal{F}_i\}_{i=1, \dots, RL}$ with $\mathcal{F}_i$ generated by $\{(x_1, a_1), \dots, (x_i, a_i)\}$. By the definition of $V_{\widehat{w}_{t-1}}^\pi$, we have

$$\begin{aligned} \mathbb{E}[\varepsilon_i | \mathcal{F}_i] &= \mathbb{E}[\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi(x'_i) | \mathcal{F}_i] - \phi(x_i, a_i)^\top \bar{w}_t \\ &= \sum_{x'} [\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi](x') P(x' | x_i, a_i) - \phi(x_i, a_i)^\top \bar{w}_t \\ &= \sum_{k \in \mathcal{K}} \phi_k(x_i, a_i) \sum_{x'} [\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}^\pi](x') \psi_k(x') - \phi(x_i, a_i)^\top \bar{w}_t = 0. \end{aligned}$$

Therefore, $\{\varepsilon_i\}_{i=1}^{RL}$ is a sequence of martingale difference noises and $|\varepsilon_i| \leq 1/(1-\gamma)$ due to the truncation operator $\Pi_{[0,1/(1-\gamma)]}$. The next lemma bounds the difference between $\widehat{w}_t$ and $\bar{w}_t$ within ℓ_1 -norm. The proof is deferred to Appendix E.2.

Lemma B.1. Consider the sparse linear regression described in Eq. (B.6). Suppose the restricted minimum eigenvalue of Σ satisfy $C_{\min}(\Sigma, s) > 0$ and the number of episodes used in phase t satisfies

$$R \geq \frac{C_1 \log(3d^2/\delta) s^2}{C_{\min}(\Sigma, s)},$$

for some absolute constant $C_1 > 0$. With the choice of $\lambda_1 = (1-\gamma)^{-1} \sqrt{\log(2d/\delta)/(RL)}$, the following holds with probability at least $1 - \delta$,

$$\|\widehat{w}_t - \bar{w}_t\|_1 \leq \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2d/\delta)}{RL}}. \quad (\text{B.7})$$

Note that the samples we use between phases are mutually independent. Thus, Eq. (B.7) uniformly holds for all $t \in [T]$ with probability at least $1 - T\delta$.

Step 4: Error decomposition. Recall that $\widehat{v}_{\widehat{w}_T}^\pi = \frac{1}{m} \sum_{u=1}^m \Pi_{[0,1/(1-\gamma)]} (Q_{\widehat{w}_T}(\tilde{x}_u, \tilde{a}_u))$ and we denote $\bar{v}_{\widehat{w}_T}^\pi = \sum_x V_{\widehat{w}_T}^\pi(x) \xi_0(x)$. According to Eq. (B.3), we decompose the policy evaluation error by Monte Carlo error, estimation error

and approximation error as follows:

$$\begin{aligned}
 |\hat{v}_{\hat{w}_T}^\pi - v^\pi| &= \left| \hat{v}_{\hat{w}_T}^\pi - \bar{v}_{\bar{w}_T}^\pi + \sum_x \left(V_{\hat{w}_T}^\pi(x) - v^\pi(x) \right) \xi_0(x) \right| \\
 &= \left| \hat{v}_{\hat{w}_T}^\pi - \bar{v}_{\bar{w}_T}^\pi + \sum_x \left([\hat{\mathcal{T}}_\pi^{(T)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi](x) - [\mathcal{T}_\pi v^\pi](x) \right) \xi_0(x) \right| \\
 &\leq \underbrace{\left| \sum_x \left([\hat{\mathcal{T}}_\pi^{(T)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi](x) - [\mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi](x) \right) \xi_0(x) \right|}_{\text{estimation error}} \\
 &\quad + \underbrace{\left| \sum_x [\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi - v^\pi)](x) \xi_0(x) \right|}_{\text{approximation error}} + \underbrace{\left| \hat{v}_{\hat{w}_T}^\pi - \bar{v}_{\bar{w}_T}^\pi \right|}_{\text{Monte Carlo error}}.
 \end{aligned} \tag{B.8}$$

Since $\tilde{x}_u, \tilde{a}_u$ is i.i.d sampled from ξ_0 and π , standard Hoeffding's inequality shows that

$$\left| \hat{v}_{\hat{w}_T}^\pi - \bar{v}_{\bar{w}_T}^\pi \right| \leq \sqrt{\frac{\log(1/\delta)}{m(1-\gamma)^2}} = \sqrt{\frac{\log(1/\delta)}{N(1-\gamma)^2}}, \tag{B.9}$$

with probability at least $1 - \delta$.

Recall that $\nu_t^\pi = \mathbb{E}^\pi[\phi(x_t, a_t) | x_0 \sim \xi_0]$. To bound the estimation error, combining Eqs. (B.2) and (B.5) together, we have

$$\begin{aligned}
 &\sum_x \left([\hat{\mathcal{T}}_\pi^{(T)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi](x) - [\mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi](x) \right) \xi_0(x) \\
 &= \gamma \sum_x \sum_a \pi(a|x) \phi(x, a)^\top (\hat{w}_{T-1} - \bar{w}_{T-1}) \xi_0(x) \\
 &= \gamma (\nu_0^\pi)^\top (\hat{w}_{T-1} - \bar{w}_{T-1}).
 \end{aligned} \tag{B.10}$$

To bound approximation error, we expand it by Eq. (B.4):

$$\begin{aligned}
 &\sum_x [\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi - v^\pi)](x) \xi_0(x) \\
 &= \gamma \sum_x \left(\sum_a \pi(a|x) \sum_{x'} P(x'|x, a) (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi - v^\pi)(x') \right) \xi_0(x).
 \end{aligned}$$

According to Eq. (B.10), we decompose

$$\begin{aligned}
 &(\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi - v^\pi)(x) \leq (V_{\hat{w}_{T-1}}^\pi - v^\pi)(x) \\
 &= [\hat{\mathcal{T}}_\pi^{(T-1)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - \mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi + \mathcal{T}_\pi \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - \mathcal{T}_\pi v^\pi](x) \\
 &= \gamma \sum_a \pi(a|x) \phi(x, a)^\top (\hat{w}_{T-2} - \bar{w}_{T-2}) + [\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - v^\pi)](x).
 \end{aligned}$$

This implies

$$\begin{aligned}
 &\sum_x [\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-1}}^\pi - v^\pi)](x) \xi_0(x) \\
 &\leq \gamma^2 \sum_x \left(\sum_a \pi(a|x) \gamma \sum_{x'} P(x'|x, a) \left(\sum_a \pi(a|x') \phi(x', a)^\top (\hat{w}_{T-2} - \bar{w}_{T-2}) \right) \right) \xi_0(x) \\
 &\quad + \sum_x \left(\sum_a \pi(a|x) \gamma \sum_{x'} [\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - v^\pi)](x') P(x'|x, a) \right) \xi_0(x) \\
 &= \gamma^2 \mathbb{E}^\pi[\phi(x_1, a_1) | x \sim \xi_0]^\top (\hat{w}_{T-2} - \bar{w}_{T-2}) \\
 &\quad + \gamma \mathbb{E}^\pi[\mathcal{T}_\pi (\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - v^\pi)(x_1) | x_0 \sim \xi_0].
 \end{aligned} \tag{B.11}$$

Combining Eqs. (B.8), (B.10) and (B.11) together, we have

$$\begin{aligned} |\bar{v}_{\hat{w}_T}^\pi - v^\pi| &\leq |\gamma(\nu_0^\pi)^\top (\hat{w}_{T-1} - \bar{w}_{T-1})| + |\gamma^2(\nu_1^\pi)^\top (\hat{w}_{T-2} - \bar{w}_{T-2})| \\ &\quad + \left| \gamma \mathbb{E}^\pi \left[[\mathcal{T}_\pi(\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{T-2}}^\pi - v^\pi)](x_1) | x_0 \sim \xi_0 \right] \right|. \end{aligned}$$

Iteratively implementing the above decomposition, we have

$$\begin{aligned} |\bar{v}_{\hat{w}_T}^\pi - v^\pi| &\leq \sum_{t=0}^{T-1} \gamma^{t+1} |(\nu_t^\pi)^\top (\hat{w}_{T-t-1} - \bar{w}_{T-t-1})| + \gamma^T \left| \mathbb{E}^\pi \left[(\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_0}^\pi - v^\pi)(x_T) | x_0 \sim \xi_0 \right] \right| \\ &\leq \sum_{t=0}^{T-1} \gamma^{t+1} |(\nu_t^\pi)^\top (\hat{w}_{T-t-1} - \bar{w}_{T-t-1})| + \frac{2\gamma^T}{1-\gamma} \\ &\leq \sum_{t=0}^{T-1} \gamma^{t+1} \|\nu_t^\pi\|_\infty \|\hat{w}_{T-t-1} - \bar{w}_{T-t-1}\|_1 + \frac{2\gamma^T}{1-\gamma}. \end{aligned}$$

Since we assume $\|\phi(x, a)\|_\infty \leq 1$, then $\|\nu_t^\pi\|_\infty \leq 1$ as well. Using the fact that $\sum_{t=0}^{T-1} \gamma^t \leq 1/(1-\gamma)$, we have

$$|\bar{v}_{\hat{w}_T}^\pi - v^\pi| \leq \frac{1}{1-\gamma} \max_{t=0, \dots, T-1} \|\hat{w}_t - \bar{w}_t\|_1 + \frac{2\gamma^T}{1-\gamma}.$$

Suppose the sample size satisfies

$$N \geq \frac{C_1 \log(3d^2/\delta) s^2}{C_{\min}(\Sigma, s)} \frac{L}{1-\gamma} \log(N/(1-\gamma)),$$

for a sufficient large constant $C_1 > 0$. Applying Lemma B.1 over $t = 0, \dots, T-1$, it implies

$$|\bar{v}_{\hat{w}_T}^\pi - v^\pi| \leq \frac{1}{(1-\gamma)^2} \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(2d/\delta)}{RL}} + \frac{2\gamma^T}{1-\gamma},$$

holds with probability at least $1 - T\delta$. By elementary change of base formula and Taylor expansion, we have

$$\log_{1/\gamma}(x) = \frac{\log(x)}{\log(1/\gamma)} \approx \frac{\log(x)\lambda}{1-\gamma}.$$

By properly choosing $T = \Theta(\log(N/(1-\gamma))/(1-\gamma))$, we have with probability at least $1 - \delta$,

$$|\bar{v}_{\hat{w}_T}^\pi - v^\pi| \leq \frac{1}{(1-\gamma)^{5/2}} \frac{32\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(N/(1-\gamma)) \log(2dT/\delta)}{N}},$$

where we use $N = TRL$. Combining with Monte Carlo approximation error Eq. (B.9) This ends the proof. $\blacksquare$

B.2. Proof of Lemma 4.6: Feature selection

We study the feature screening and sparsity properties of the model selected by the regularized estimator $\hat{K}^\pi$. Recall that from the identity Eq. (3.1), we solve the following multivariate regression problem:

$$\phi^\pi(x'_n)^\top = \phi(x_n, a_n)^\top K^\pi + \varepsilon_n, \quad n \in [N], \quad (\text{B.12})$$

where $x'_n \sim P(\cdot | x_n, a_n)$ and $\varepsilon_n = \phi^\pi(x'_n)^\top - \mathbb{E}[\phi^\pi(x'_n)^\top] \in \mathbb{R}^d$ is the noise vector. Define a filtration $\{\mathcal{F}_n\}_{n=1, \dots, N}$ with $\mathcal{F}_n$ generated by $\{(x_1, a_1), \dots, (x_n, a_n)\}$. It is easy to see $\mathbb{E}[\varepsilon_n | \mathcal{F}_n] = 0$ such that $\{\varepsilon_n\}_{n=1}^N$ are martingale difference vectors. We introduce some notations for simplicities:

- Denote the response $Y_j = (\phi_j^\pi(x'_1), \dots, \phi_j^\pi(x'_N))^\top$ for $j \in [d]$ where $\phi_j^\pi(\cdot)$ is j th coordinate of $\phi^\pi(\cdot)$. And $Y = (Y_1^\top, \dots, Y_d^\top)^\top \in \mathbb{R}^{Nd \times 1}$.
- Denote the noise $W_j = (\varepsilon_{1j}, \dots, \varepsilon_{Nj})^\top$ where ε_{nj} is the j th coordinate of ε_n , and $W = (W_1^\top, \dots, W_d^\top)^\top \in \mathbb{R}^{Nd \times 1}$.
- Denote the design matrix as

$$\Phi = \begin{pmatrix} \phi(x_1, a_1)^\top \\ \vdots \\ \phi(x_N, a_N)^\top \end{pmatrix} \in \mathbb{R}^{N \times d}, X = \begin{pmatrix} \Phi & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Phi \end{pmatrix} \in \mathbb{R}^{Nd \times d^2}.$$

Note that X is a block diagonal matrix.

- Let β_j^* as the j th column of K^π and $\beta^* = (\beta_1^{*\top}, \dots, \beta_d^{*\top})^\top \in \mathbb{R}^{d^2 \times 1}$ as the regression coefficient.
- For every $\beta \in \mathbb{R}^{d^2}$, we define $\beta^j = (\beta_{j+(i-1)d} : i \in [d])^\top$ as the vector formed by the coefficients corresponding to the j th variable. For instance, β^{*j} is the j th row of K^π . If $J \subseteq [d]$, denote $\beta_J \in \mathbb{R}^{d^2}$ by stacking the vectors $\beta^j \mathbb{1}\{j \in J\}$. Write $\mathcal{S}(\beta) = \{j : \beta^j \neq 0\}$ as the relevant feature set of β .
- For a vector $\beta \in \mathbb{R}^{d^2}$, define the $\ell_{2,p}$ -norm for $1 \leq p < \infty$ as:

$$\|\beta\|_{2,p} = \left(\sum_{j=1}^d \left(\sum_{i=1}^d \beta_{j+(i-1)d}^2 \right)^{p/2} \right)^{1/p},$$

and the $\ell_{2,0}$ -norm as:

$$\|\beta\|_{2,0} = \sum_{j=1}^d \mathbb{1}\{\|\beta^j\|_2 \neq 0\}.$$

Therefore, we can rewrite Eq. (B.12) into an ordinary linear regression form with group sparse structure on the regression coefficients β^* :

$$Y = X\beta^* + W.$$

Note that $\mathcal{S}(\beta^*) = \mathcal{K}$ where $\mathcal{K}$ is defined in Assumption 2.1 since K^π is row-sparse. The corresponding group lasso estimator defined in Eq. (3.2) can be rewritten into:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \left\{ \frac{1}{Nd} \|Y - X\beta\|_2^2 + \lambda_2 \sum_{j=1}^d \|\beta^j\|_2 \right\}, \quad (\text{B.13})$$

and $\mathcal{S}(\hat{\beta}) = \hat{\mathcal{K}}$. The regularization parameter is chosen as

$$\lambda_2 = 4\sqrt{\frac{2\log(2d^2/\delta)}{Nd}}, \quad (\text{B.14})$$

for some $\delta > 0$.

Now we study the feature screening property of $\hat{\beta}$ in four steps.

Step 1. By the optimality of $\hat{\beta}$, we have

$$\frac{1}{Nd} \|Y - X\hat{\beta}\|_2^2 + \lambda_2 \sum_{j=1}^d \|\hat{\beta}^j\|_2 \leq \frac{1}{Nd} \|Y - X\beta^*\|_2^2 + \lambda_2 \sum_{j=1}^d \|\beta^{*j}\|_2.$$

Plugging in $Y = X\beta^* + W$,

$$\begin{aligned} \frac{1}{Nd} \|X(\hat{\beta} - \beta^*)\|_2^2 &\leq \frac{2}{Nd} W^\top X(\hat{\beta} - \beta^*) + \lambda_2 \sum_{j=1}^d \left(\|\beta^{*j}\|_2 - \|\hat{\beta}^j\|_2 \right) \\ &\leq \frac{2}{Nd} \max_{j \in [d]} \|(X^\top W)^j\|_2 \sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 + \lambda_2 \sum_{j=1}^d \left(\|\beta^{*j}\|_2 - \|\hat{\beta}^j\|_2 \right) \end{aligned} \quad (\text{B.15})$$

where the last inequality is from Hölder's inequality.

Step 2. Next, we will bound the noise term: $\|(X^\top W)^j\|_2$. From the definitions of X and W , we write it explicitly as

$$\|(X^\top W)^j\|_2 = \sqrt{\sum_{i=1}^d \left(\sum_{n=1}^N \phi_j(x_n, a_n) \varepsilon_{ni} \right)^2}.$$

It is easy to verify that $\{\phi_j(x_n, a_n) \varepsilon_{ni}\}_{n=1}^N$ is also a martingale difference sequence for any $i, j \in [d]$ and $|\phi_j(x_n, a_n) \varepsilon_{ni}| \leq 1$ since we assume $\|\phi(x, a)\|_\infty \leq 1$ for any state-action pair. According to Azuma-Hoeffding inequality (Lemma F.2), for all $\tilde{\delta} > 0$,

$$\mathbb{P}\left(\left| \sum_{n=1}^N \phi_j(x_n, a_n) \varepsilon_{ni} \right| \geq \tilde{\delta}\right) \leq 2 \exp\left(-\frac{\tilde{\delta}^2}{2N}\right).$$

Using the union bound twice, the following holds,

$$\mathbb{P}\left(\max_{j \in [d]} \sqrt{\sum_{i=1}^d \left(\sum_{n=1}^N \phi_j(x_n, a_n) \varepsilon_{ni} \right)^2} \geq \sqrt{d\tilde{\delta}}\right) \leq 2d^2 \exp\left(-\frac{\tilde{\delta}^2}{2N}\right). \quad (\text{B.16})$$

Letting $\delta = 2d^2 \exp(-\tilde{\delta}^2/2N)$, we have with probability at least $1 - \delta$,

$$\frac{1}{Nd} \max_{j \in [d]} \|(X^\top W)^j\|_2 \leq \frac{\lambda_2}{4}.$$

Define an event $\mathcal{A}$ as

$$\mathcal{A} = \left\{ \frac{1}{Nd} \max_{j \in [d]} \|(X^\top W)^j\|_2 \leq \frac{\lambda_2}{4} \right\}.$$

Then we have $\mathbb{P}(\mathcal{A}) \geq 1 - \delta$.

Step 3. According to Karush–Kuhn–Tucker (KKT) condition, the solution $\hat{\beta}$ of the optimization problem Eq. (B.13) satisfies

$$\begin{cases} (X^\top (Y - X\hat{\beta}))^j / (Nd) = \lambda_2 \hat{\beta}^j / \|\hat{\beta}^j\|_2, & \text{if } \hat{\beta}^j \neq 0, \\ \|(X^\top (Y - X\hat{\beta}))^j\|_2 / (Nd) \leq \lambda_2, & \text{if } \hat{\beta}^j = 0. \end{cases}$$

Under event $\mathcal{A}$ and using KKT condition, we have if $\hat{\beta}^j \neq 0$, then

$$\begin{aligned} \lambda_2 &= \left\| \frac{1}{Nd} (X^\top (Y - X\hat{\beta}))^j \right\|_2 = \left\| \frac{1}{Nd} (X^\top X(\beta^* - \hat{\beta}))^j + \frac{1}{Nd} (X^\top W)^j \right\|_2 \\ &\leq \left\| \frac{1}{Nd} (X^\top X(\beta^* - \hat{\beta}))^j \right\|_2 + \left\| \frac{1}{Nd} (X^\top W)^j \right\|_2 \\ &\leq \left\| \frac{1}{Nd} (X^\top X(\beta^* - \hat{\beta}))^j \right\|_2 + \frac{1}{4} \lambda_2, \end{aligned}$$

which implies

$$\begin{aligned} \frac{1}{(Nd)^2} \|(X^\top X(\hat{\beta} - \beta^*))_{\mathcal{S}(\hat{\beta})}\|_2^2 &= \frac{1}{(Nd)^2} \sum_{j \in \mathcal{S}(\hat{\beta})} \|(X^\top X(\hat{\beta} - \beta^*))^j\|_2^2 \\ &\geq |\mathcal{S}(\hat{\beta})| \frac{9}{16} \lambda_2^2. \end{aligned} \quad (\text{B.17})$$

We define a notation of restricted maximum eigenvalue with respect to $\mathcal{S}(\beta^*)$ and X :

$$\tilde{C}_{\max}(m) = \max_{\beta \in \mathbb{R}^{d^2 \times d^2} : \|\beta_{\mathcal{S}(\beta^*)^c}\|_{2,0} \leq m} \frac{\beta^\top X^\top X \beta}{N \|\beta\|_2^2}. \quad (\text{B.18})$$

Denote $\hat{m} = |\mathcal{S}(\hat{\beta}) \setminus \mathcal{S}(\beta^*)|$. Then we have

$$\begin{aligned} \|(X^\top X(\hat{\beta} - \beta^*))_{\mathcal{S}(\hat{\beta})}\|_2 &\leq \sup_{\|\alpha_{\mathcal{S}(\beta^*)^c}\|_{2,0} \leq \hat{m}} \frac{|\alpha^\top X^\top X(\hat{\beta} - \beta^*)|}{\|\alpha\|_2} \\ &\leq \sup_{\|\alpha_{\mathcal{S}(\beta^*)^c}\|_{2,0} \leq \hat{m}} \frac{\|\alpha^\top X^\top\|_2 \|X(\hat{\beta} - \beta^*)\|_2}{\|\alpha\|_2} \\ &= \sup_{\|\alpha_{\mathcal{S}(\beta^*)^c}\|_{2,0} \leq \hat{m}} \frac{\sqrt{|\alpha^\top X^\top X \alpha|}}{\|\alpha\|_2} \|X(\hat{\beta} - \beta^*)\|_2 \\ &\leq \sqrt{N C_{\max}(\hat{m})} \|X(\hat{\beta} - \beta^*)\|_2. \end{aligned} \quad (\text{B.19})$$

Combining Eqs. (B.17) and (B.19) together, we have

$$|\mathcal{S}(\hat{\beta})| \leq \frac{16 C_{\max}(\hat{m})}{9 N d^2 \lambda_2^2} \|X(\hat{\beta} - \beta^*)\|_2^2, \quad (\text{B.20})$$

holds with probability at least $1 - \delta$.

Step 4. It remains to control the in-sample prediction error $\|X(\hat{\beta} - \beta^*)\|_2^2$. Under event $\mathcal{A}$, using Eq. (B.15) implies

$$\frac{1}{Nd} \|X(\hat{\beta} - \beta^*)\|_2^2 \leq \frac{\lambda_2}{2} \sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 + \lambda_2 \sum_{j=1}^d \left(\|\beta^{*j}\|_2 - \|\hat{\beta}^j\|_2 \right).$$

Adding $\sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 \lambda_2 / 2$ to both sides and using the fact that $\|\beta^{*j} - \beta^{*j}\|_2 + \|\hat{\beta}^j\|_2 - \|\beta^{*j}\|_2 = 0$ for $j \neq \mathcal{S}(\beta^*)$, we have

$$\begin{aligned} &\frac{1}{Nd} \|X(\hat{\beta} - \beta^*)\|_2^2 + \frac{\lambda_2}{2} \sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 \\ &\leq \lambda_2 \sum_{j \in \mathcal{S}(\beta^*)} \left(\|\hat{\beta}^j - \beta^{*j}\|_2 + \|\beta^{*j}\|_2 - \|\hat{\beta}^j\|_2 \right) \\ &\leq 2 \lambda_2 \sqrt{s} \|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_2, \end{aligned} \quad (\text{B.21})$$

where the last inequality is from Cauchy-Schwarz inequality. Recall that the expected uncentered covariance matrix is defined as

$$\Sigma = \mathbb{E} \left[\frac{1}{L} \sum_{h=0}^{L-1} \phi(x_h^{(1)}, a_h^{(1)}) \phi(x_h^{(1)}, a_h^{(1)})^\top \right],$$

and we define the empirical uncentered covariance matrix as

$$\hat{\Sigma} = \frac{1}{N} \sum_{n=1}^N \phi(x_n, a_n) \phi(x_n, a_n)^\top = \frac{1}{K} \sum_{k=1}^K \frac{1}{L} \sum_{h=0}^{L-1} \phi(x_h^{(k)}, a_h^{(k)}) \phi(x_h^{(k)}, a_h^{(k)})^\top,$$

with $N = KL$. Denote the expected and empirical uncentered covariance matrices for the multivariate linear regression as

$$\hat{\Psi} = \begin{pmatrix} \hat{\Sigma} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \hat{\Sigma} \end{pmatrix} \in \mathbb{R}^{d^2 \times d^2}; \Psi = \begin{pmatrix} \Sigma & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Sigma \end{pmatrix} \in \mathbb{R}^{d^2 \times d^2}.$$

We introduce a generalization of restricted eigenvalue condition (Definition 4.2) for multivariate linear regression.

Definition B.2 ($\ell_{2,1}$ -restricted eigenvalue). Given a symmetric matrix $H \in \mathbb{R}^{d^2 \times d^2}$ and integer $s \geq 1$, the restricted eigenvalue of H is defined as

$$\tilde{C}_{\min}(H, s) := \min_{S \subset [d], |S| \leq s} \min_{\beta \in \mathbb{R}^{d^2}} \left\{ \frac{\langle \beta, H\beta \rangle}{\|\beta_S\|_2^2} : \beta \in \mathbb{R}^{d^2}, \|\beta_{S^c}\|_{2,1} \leq 3\|\beta_S\|_{2,1} \right\}.$$

Next lemma provides a lower bound for $\tilde{C}_{\min}(\hat{\Phi}, s)$. The proof is deferred to Appendix E.1.

Lemma B.3. Assume the $\ell_{2,1}$ -restricted eigenvalue of Ψ satisfies $\tilde{C}_{\min}(\Psi, s) > 0$ for some $\delta > 0$. Suppose the sample size satisfies $N \geq 32^2 L \log(3d^2/\delta) s^2 / \tilde{C}_{\min}(\Psi, s)^2$. Then the $\ell_{2,1}$ -restricted eigenvalue of $\hat{\Psi}$ satisfies $\tilde{C}_{\min}(\hat{\Psi}, s) > \tilde{C}_{\min}(\Psi, s)/2$ with probability at least $1 - \delta$.

On the other hand, from Eq. (B.21), we know that

$$\frac{1}{2} \sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 \leq 2 \sum_{j \in \mathcal{S}(\beta^*)} \|\hat{\beta}^j - \beta^{*j}\|_2,$$

and thus,

$$\sum_{j \in \mathcal{S}(\beta^*)^c} \|\hat{\beta}^j - \beta^{*j}\|_2 \leq 3 \sum_{j \in \mathcal{S}(\beta^*)} \|\hat{\beta}^j - \beta^{*j}\|_2. \quad (\text{B.22})$$

This implies that $\|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)^c}\|_{2,1} \leq 3\|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_{2,1}$. Applying Lemma B.3, the following holds with probability at least $1 - \delta$,

$$\frac{\|X(\hat{\beta} - \beta^*)\|_2^2}{N\|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_2^2} = \frac{(\hat{\beta} - \beta^*)^\top \hat{\Psi}(\hat{\beta} - \beta^*)}{\|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_2^2} \geq \tilde{C}_{\min}(\hat{\Psi}, s) \geq \tilde{C}_{\min}(\Psi, s)/2. \quad (\text{B.23})$$

Plugging the above bound into Eq. (B.21),

$$\begin{aligned} \frac{1}{Nd} \|X(\hat{\beta} - \beta^*)\|_2^2 &\leq 2\lambda_2 \sqrt{s} \|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_2 \\ &\leq \frac{1}{\sqrt{N}} \|X(\hat{\beta} - \beta^*)\|_2 \frac{4\sqrt{s}\lambda_2}{\sqrt{\tilde{C}_{\min}(\Psi, s)}}. \end{aligned}$$

Combining with Eq. (B.20) and the choice of λ_2 in Eq. (B.14), we reach

$$|\mathcal{S}(\hat{\beta})| \leq \frac{16\tilde{C}_{\max}}{9Nd^2\lambda_2^2} \frac{16\lambda_2^2 Nd^2 s}{\tilde{C}_{\min}(\Psi, s)} = \frac{256\tilde{C}_{\max}(\hat{m})s}{9\tilde{C}_{\min}(\Psi, s)},$$

with probability at least $1 - \delta$, as long as $N \geq 32^2 L \log(3d^2/\delta) s^2 / \tilde{C}_{\min}(\Psi, s)^2$. If the vanilla restricted eigenvalue (Definition 4.2) of Σ satisfies $C_{\min}(\Sigma, s) > 0$, then we have for any $\mathcal{S} \subset [d], |\mathcal{S}| \leq s$, and any $\beta_j \in \mathbb{R}^d$ satisfying $\|(\beta_j)_{\mathcal{S}^c}\|_1 \leq 3\|(\beta_j)_{\mathcal{S}}\|_1$,

$$\frac{\beta_j^\top \Sigma \beta_j}{\|(\beta_j)_{\mathcal{S}}\|_2^2} \geq C_{\min}(\Sigma, s) > 0.$$

Consider a sequence of vectors $\beta_1, \dots, \beta_d$ satisfying $\|(\beta_j)_{\mathcal{S}^c}\|_1 \leq 3\|(\beta_j)_{\mathcal{S}}\|_1$. Then for $\beta = (\beta_1^\top, \dots, \beta_d^\top)^\top$, we have

$$\beta^\top \Psi \beta = \sum_{j=1}^d \beta_j^\top \Sigma \beta_j \geq C_{\min}(\Sigma, s) \sum_{j=1}^d \|(\beta_j)_{\mathcal{S}}\|_2^2 = C_{\min}(\Sigma, s) \|\beta_{\mathcal{S}}\|_2^2.$$

Therefore, we conclude $\tilde{C}_{\min}(\Psi, s) \geq C_{\min}(\Sigma, s) > 0$ such that

$$|\mathcal{S}(\hat{\beta})| \leq \frac{256\tilde{C}_{\max}(\hat{m})s}{9C_{\min}(\Sigma, s)}, \quad (\text{B.24})$$

with probability at least $1 - \delta$, as long as $N \geq 32^2 L \log(3d^2/\delta) s^2 / C_{\min}^2(\Sigma, s)$.

For any $\beta = (\beta_1^\top, \dots, \beta_d^\top)^\top$ satisfying $\|\beta_{\mathcal{S}(\beta^*)^c}\|_{2,0} \leq m_0$, we have $\|(\beta_j)_{\mathcal{S}(\beta^*)^c}\|_0 \leq m_0$ for any $j \in [d]$. Given a positive semi-definite matrix $Z \in \mathbb{R}^{d \times d}$ and integer $s \geq 1$, the restricted maximum eigenvalue of Z are defined as

$$C_{\max}(Z, s) := \max_{\mathcal{S} \subset [d], |\mathcal{S}| \leq s} \max_{\beta \in \mathbb{R}^d} \left\{ \frac{\langle \beta, Z\beta \rangle}{\|\beta\|_2^2} : \beta \in \mathbb{R}^d, \|\beta_{\mathcal{S}^c}\|_0 \leq s \right\}.$$

Using the definition of $C_{\max}(\hat{\Sigma}, s)$, it holds that

$$\beta_j^\top \hat{\Sigma} \beta_j \leq C_{\max}(\hat{\Sigma}, s) \|\beta_j\|_2^2, \text{ for any } j \in [d].$$

Summing the above inequality from 1 to d ,

$$\beta^\top \frac{X^\top X}{N} \beta = \sum_{j=1}^d \beta_j^\top \hat{\Sigma} \beta_j \leq C_{\max}(\hat{\Sigma}, s) \sum_{j=1}^d \|\beta_j\|_2^2 = C_{\max}(\hat{\Sigma}, s) \|\beta\|_2^2.$$

This implies $\tilde{C}_{\max}(\hat{m}) \leq C_{\max}(\hat{\Sigma}, \hat{m})$. As shown in the Lemma 1 in [Belloni et al. \(2013\)](#), we have $C_{\max}(\hat{\Sigma}, m) \leq 4C_{\max}(\Sigma, m)$ for any $m + s \leq \log(n)$ as long as $n \gtrsim s$.

Step 5. Recall that $\hat{m} = |\mathcal{S}(\hat{\beta}) \setminus \mathcal{S}(\beta^*)|$ and denote

$$\mathcal{M} = \left\{ m \in \mathbb{N}^+ : m > \frac{2048sC_{\max}(\Sigma, m)}{9C_{\min}(\Sigma, s)} \right\}.$$

Suppose there is a $m_0 \in \mathcal{M}$ such that $\hat{m} > m_0$. From Eq. (B.24), we know that

$$\hat{m} \leq |\mathcal{S}(\hat{\beta})| \leq C_{\max}(\Sigma, m) \frac{1024s}{9C_{\min}(\Sigma, s)}.$$

According to Lemma 3 in [Belloni et al. \(2013\)](#) for the sublinearity of sparse maximum eigenvalues, we have

$$C_{\max}(\Sigma, \hat{m}) \leq \lceil \hat{m}/m_0 \rceil C_{\max}(\Sigma, m_0) \leq 2C_{\max}(\Sigma, m_0) \hat{m}/m_0.$$

where the last inequality we use $\lceil \kappa \rceil \leq 2\kappa$. Putting the above two results together, we have

$$m_0 \leq \frac{2048sC_{\max}(\Sigma, m_0)}{9C_{\min}(\Sigma, s)}.$$

This leads a contradiction with the definition of $\mathcal{M}$. Therefore, $\hat{m} \leq m_0$ for all $m_0 \in \mathcal{M}$. This implies

$$\begin{aligned} |\mathcal{S}(\hat{\beta})| &\leq \min_{m_0 \in \mathcal{M}} C_{\max}(\Sigma, m_0) \frac{1024s}{9C_{\min}(\Sigma, s)} \\ &= \left\lceil \frac{1024 \min_{m_0 \in \mathcal{M}} C_{\max}(\Sigma, m_0)}{9C_{\min}(\Sigma, s)} \right\rceil s \lesssim s. \end{aligned}$$

The term $\min_{m_0 \in \mathcal{M}} C_{\max}(\Sigma, m_0) / C_{\min}(\Sigma, s)$ essentially characterizes the condition number of Σ on a restricted support and is upper bounded by the condition number defined in the full support. Now we finish the proof of the first part of Lemma 4.6 and start to prove the second part of Lemma 4.6 under separability condition.

According to Eq. (B.22), under event $\mathcal{A}$ we have

$$\begin{aligned} \|\hat{\beta} - \beta^*\|_{2,1} &= \sum_{j=1}^d \|\hat{\beta}^j - \beta^{*j}\|_2 = \sum_{j \in \mathcal{S}(\beta^*)} \|\hat{\beta}^j - \beta^{*j}\|_2 + \sum_{j \in \mathcal{S}(\beta^*)^c} \|\hat{\beta}^j - \beta^{*j}\|_2 \\ &\leq 4 \sum_{j \in \mathcal{S}(\beta^*)} \|\hat{\beta}^j - \beta^{*j}\|_2. \end{aligned}$$

From Eq. (B.23),

$$\begin{aligned} \sum_{j \in \mathcal{S}(\beta^*)} \|\hat{\beta}^j - \beta^{*j}\|_2 &\leq \sqrt{s} \|(\hat{\beta} - \beta^*)_{\mathcal{S}(\beta^*)}\|_2 \\ &\leq \sqrt{\frac{2s}{\tilde{C}_{\min}(\Sigma, s)}} \frac{1}{\sqrt{N}} \|X(\hat{\beta} - \beta^*)\|_2 \\ &\leq \frac{4\sqrt{2}sd\lambda_2}{\tilde{C}_{\min}(\Sigma, s)} \leq \frac{4\sqrt{2}sd\lambda_2}{C_{\min}(\Sigma, s)}. \end{aligned}$$

Combining the above two inequality together and plugging in the choice of λ_2 , we can bound

$$\|\hat{\beta} - \beta^*\|_{2,1} \leq \frac{64\sqrt{2}s\sqrt{d}}{C_{\min}(\Sigma, s)} \sqrt{\frac{2\log(2d^2/\delta)}{N}}.$$

with probability at least $1 - \delta$. Under Assumption 4.5, the following holds that with probability at least $1 - \delta$,

$$\min_{j \in \mathcal{S}(\beta^*)} \|\beta^{*j}\|_2 > \|\hat{\beta} - \beta^*\|_{2,1} \geq \|\hat{\beta} - \beta^*\|_{2,\infty}.$$

If there is a $j \in \mathcal{S}(\beta^*)$ but $j \notin \mathcal{S}(\hat{\beta})$, we have

$$\|\hat{\beta}^j - \beta^{*j}\|_2 = \|\beta^{*j}\|_2 > \|\hat{\beta} - \beta^*\|_{2,\infty}.$$

On the other hand,

$$\|\hat{\beta}^j - \beta^{*j}\|_2 \leq \|\hat{\beta} - \beta^*\|_{2,\infty},$$

which leads a contradiction. Now we conclude that $\mathcal{S}(\hat{\beta}) \supseteq \mathcal{K}$. This ends the proof. $\blacksquare$

B.3. Proof of Theorem 4.9: Instance-dependent upper bound

We restate the instance-dependent error bound of vanilla fitted Q-evaluation algorithm on the full support.

Theorem B.4 (Theorem 5 in Duan and Wang (2020)). Suppose Assumption 2.1 holds. Suppose $\phi(x, a)^\top \Sigma^{-1} \phi(x, a) \lesssim d$ for any pair of (x, a) . Let $\delta \in (0, 1)$ and Algorithm 2 without feature selection stage takes N samples satisfying

$$N \gtrsim \frac{\gamma^2 \log(d/\delta)d}{(1-\gamma)^3}.$$

Set regularization parameter $\lambda_3 = \lambda_{\min}(\Sigma) \log(12d/\delta)C_1 d/(1-\gamma)$. Letting the number of iteration $T \rightarrow \infty$, the following holds with probability at least $1 - \delta$,

$$\hat{v}_w^\pi - v^\pi \leq \frac{1}{1-\gamma} \sum_{t=0}^{\infty} \gamma^{t+1} \sqrt{(\nu_t^\pi)^\top \Sigma^{-1} (\nu_t^\pi)} \sqrt{\frac{\log(1/\delta)}{N}} + \frac{\gamma \ln(12d/\delta)d}{N(1-\gamma)^{3.5}},$$

where $\nu_t^\pi = \mathbb{E}^\pi[\phi(x_t, a_t)|x \sim \xi_0]$.

If the true relevant feature set $\mathcal{K}$ is known in an oracle case, we could directly run the algorithm on $\mathcal{K}$ such that all the dependency on d can be reduced to s and the instance-dependent term turns to be defined in the $\mathcal{K}$ that is much sharper than the original one. Fortunately, Lemma 4.6 implies $\hat{\mathcal{K}} \supseteq \mathcal{K}$ and $|\hat{\mathcal{K}}| \lesssim s$. Suppose

$$N \gtrsim \frac{\gamma^2 \log(s/\delta)s}{(1-\gamma)^3} \gtrsim \frac{\gamma^2 \log(|\hat{\mathcal{K}}|/\delta)|\hat{\mathcal{K}}|}{(1-\gamma)^3}.$$

Rewriting Theorem B.4 with respect to $\hat{\mathcal{K}}$, we have

$$|\hat{v}_w^\pi - v^\pi| \lesssim \frac{1}{1-\gamma} \sum_{t=0}^{\infty} \gamma^{t+1} \sqrt{(\tilde{\nu}_t^\pi)^\top \tilde{\Sigma}^{-1} (\tilde{\nu}_t^\pi)} \sqrt{\frac{\log(1/\delta)}{N}},$$

where $\tilde{\nu}_t^\pi = [\nu_t^\pi]_{\hat{\mathcal{K}}}$ and $\tilde{\Sigma} = \Sigma_{\hat{\mathcal{K}} \times \hat{\mathcal{K}}}$. The corresponding condition $\phi(x, a)^\top \Sigma^{-1} \phi(x, a) \lesssim |\hat{\mathcal{K}}|$ can be satisfied due to $C_{\min}(\Sigma, s) > 0$ and $\|\phi(x, a)\|_\infty \leq 1$. From Definitions 4.7, 4.8 and Lemma B.2 in (Duan and Wang, 2020), we have

$$\begin{aligned} \sqrt{1 + \chi_{\mathcal{G}(\hat{\mathcal{K}})}^2(\mu^\pi, \bar{\mu})} &= \sum_{t=0}^{\infty} \gamma^t \sup_{f \in \mathcal{G}(\hat{\mathcal{K}})} \frac{(1 - \gamma) \mathbb{E}^\pi[f(x_t, a_t) | x_0 \sim \xi_0]}{\sqrt{\mathbb{E}[\frac{1}{L} \sum_{h=0}^{L-1} f^2(x_{1h}, a_{1h})]}} \\ &= (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \sqrt{(\tilde{\nu}_t^\pi)^\top \tilde{\Sigma}^{-1} (\tilde{\nu}_t^\pi)}. \end{aligned}$$

Now we end the proof. ■

C. Proof of Theorem 5.1: Lasso Fitted Q-iteration

C.1. Proof of Theorem 5.1

The main structure of this proof is similar to the proof of Theorem 4.3 in Appendix B.1 but we need to utilize the contraction property of Bellman optimality operator. Recall that we split the whole dataset into T folds and each fold consists of R episodes or RL sample transitions. The overall sample size is $N = TRL$.

Step 1. We verify that the execution of Algorithm 3 is equivalent to the approximate value iteration. Recall that a generic Lasso estimator with respect to a function V at t th phase is defined in Eq. (B.1) as

$$\hat{w}_t(V) = \underset{w \in \mathbb{R}^d}{\operatorname{argmin}} \left(\frac{1}{RL} \sum_{i=1}^{RL} \left(\Pi_{[0,1/(1-\gamma)]} V(x_i^{(t)}) - \phi(x_i^{(t)}, a_i^{(t)})^\top w \right)^2 + \lambda_1 \|w\|_1 \right).$$

Define $V_w(x) = \max_{a \in \mathcal{A}} (r(x, a) + \gamma \phi(x, a)^\top w)$. For simplicity, we write $\hat{w}_t := \hat{w}_t(V_{\hat{w}_{t-1}})$ for short. Define an approximate Bellman optimality operator $\hat{\mathcal{T}}^{(t)} : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ as:

$$[\hat{\mathcal{T}}^{(t)} V](x) := \max_a \left[r(x, a) + \gamma \phi(x, a)^\top \hat{w}_t(V) \right]. \quad (\text{C.1})$$

Note this $\hat{\mathcal{T}}^{(t)}$ is a randomized operator that only depends data collected at t th phase. Algorithm 3 is equivalent to the following approximate value iteration:

$$[\hat{\mathcal{T}}^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}}](x) = \max_a \left[r(x, a) + \gamma \phi(x, a)^\top \hat{w}_t \right] = \max_a Q_{\hat{w}_t}(x, a) = V_{\hat{w}_t}(x). \quad (\text{C.2})$$

Recall that the true Bellman optimality operator $\mathcal{T} : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$ is defined as

$$[\mathcal{T} V](x) := \max_a \left[r(x, a) + \gamma \sum_{x'} P(x' | x, a) V(x') \right]. \quad (\text{C.3})$$

Step 2. We verify that the true Bellman operator on $\Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}}$ can also be written as a linear form. From Assumption 2.1, there exists a vector $\bar{w}_t \in \mathbb{R}^d$ with $\bar{w}_{t,k} = 0$ if $k \notin \mathcal{K}$ such that

$$[\mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}}](x) = \max_a \left[r(x, a) + \gamma \phi(x, a)^\top \bar{w}_t \right]. \quad (\text{C.4})$$

Step 3. We start to bound $\|V_{\hat{w}_t} - v^*\|_\infty$ for each phase t . By the approximate value iteration form Eq. (C.2) and the definition of optimal value function,

$$\begin{aligned} \|V_{\hat{w}_t} - v^*\|_\infty &= \|\hat{\mathcal{T}}^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}} - \mathcal{T} v^*\|_\infty \\ &= \|\hat{\mathcal{T}}^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}} - \mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}}\|_\infty \\ &\quad + \|\mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\hat{w}_{t-1}} - \mathcal{T} v^*\|_\infty. \end{aligned} \quad (\text{C.5})$$

The first term mainly captures the error between approximate Bellman optimality operator and true Bellman optimality operator while the second term can be bounded by the contraction of true Bellman operator. From linear forms Eqs. (C.2) and (C.4), it holds for any $x \in \mathcal{X}$,

$$\begin{aligned}
 & [\widehat{\mathcal{T}}^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}](x) - [\mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}](x) \\
 &= \max_a \left[r(x, a) + \gamma \phi(x, a)^\top \widehat{w}_t \right] - \max_a \left[r(x, a) + \gamma \phi(x, a)^\top \bar{w}_t \right] \\
 &\leq \gamma \max_a |\phi(x, a)^\top (\widehat{w}_t - \bar{w}_t)| \\
 &\leq \gamma \max_{a, x} \|\phi(x, a)\|_\infty \|\widehat{w}_t - \bar{w}_t\|_1.
 \end{aligned} \tag{C.6}$$

Applying Lemma B.1, with the choice of $\lambda_1 = (1-\gamma)^{-1} \sqrt{\log(2d/\delta)/RL}$, the following error bound holds with probability at least $1 - \delta$,

$$\|\widehat{w}_t - \bar{w}_t\|_1 \leq \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2d/\delta)}{RL}}, \tag{C.7}$$

where R satisfies $R \geq C_1 \log(3d^2/\delta) s^2 / C_{\min}(\Sigma, s)$.

Note that the samples we use between phases are mutually independent. Thus Eq. (C.7) uniformly holds for all $t \in [T]$ with probability at least $1 - T\delta$. Plugging it into Eq. (C.6), we have for any phase $t \in [T]$,

$$\|\widehat{\mathcal{T}}^{(t)} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}} - \mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}}\|_\infty \leq \gamma \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2d/\delta)}{RL}}, \tag{C.8}$$

holds with probability at least $1 - \delta$.

To bound the second term in Eq. (C.5), we use the contraction property of true Bellman operator such that

$$\|\mathcal{T} \Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}} - \mathcal{T} v^*\|_\infty \leq \gamma \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}} - v^*\|_\infty. \tag{C.9}$$

Plugging Eqs. (C.8) and (C.9) into Eq. (C.5), it holds that

$$\|V_{\widehat{w}_t} - v^*\|_\infty \leq \gamma \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2dT/\delta)}{RL}} + \gamma \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{t-1}} - v^*\|_\infty, \tag{C.10}$$

with probability at least $1 - \delta$. Recursively using Eq. (C.10), the following holds with probability $1 - \delta$,

$$\begin{aligned}
 & \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{T-1}} - v^*\|_\infty \leq \|V_{\widehat{w}_{T-1}} - v^*\|_\infty \\
 &= \gamma \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2dT/\delta)}{RL}} + \gamma \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{T-2}} - v^*\|_\infty \\
 &\leq \gamma^{T-1} \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_0} - v^*\|_\infty + \sum_{t=1}^{T-1} \gamma^t \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \frac{1}{1-\gamma} \sqrt{\frac{\log(2dT/\delta)}{RL}} \\
 &\leq \frac{2\gamma^{T-1}}{1-\gamma} + \frac{1}{(1-\gamma)^2} \frac{16\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(2dT/\delta)}{RL}},
 \end{aligned}$$

where the first inequality is due to that $\Pi_{[0,1/(1-\gamma)]}$ can only make error smaller and the last inequality is from $\sum_{t=1}^{T-1} \gamma^t \leq 1/(1-\gamma)$. By properly choosing $T = \Theta(\log(N/(1-\gamma))/(1-\gamma))$, it implies

$$\|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{T-1}} - v^*\|_\infty \leq \frac{1}{(1-\gamma)^{5/2}} \frac{32\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(2dT/\delta) \log(N/(1-\gamma))}{N}},$$

holds with probability at least $1 - \delta$. From Proposition 2.14 in Bertsekas (1995),

$$\|v^{\widehat{\pi}_T} - v^*\|_\infty \leq \frac{1}{1-\gamma} \|Q_{\widehat{w}_T} - Q^*\|_\infty \leq \frac{2}{1-\gamma} \|\Pi_{[0,1/(1-\gamma)]} V_{\widehat{w}_{T-1}} - v^*\|_\infty. \tag{C.11}$$

Putting the above together, we have with probability at least $1 - \delta$,

$$\|v^{\hat{\pi}_T} - v^*\|_\infty \leq \frac{64\sqrt{2}s}{C_{\min}(\Sigma, s)} \sqrt{\frac{\log(2dT/\delta) \log(N/(1-\gamma))}{N(1-\gamma)^7}},$$

when the sample size N satisfies

$$N \geq \frac{C_1 s^2 L \log(3d^2/\delta) T}{C_{\min}(\Sigma, s)},$$

for some sufficiently large constant C_1 . This ends the proof. $\blacksquare$

D. Proof of Theorem 5.2: Minimax Lower Bound of Policy Optimization

D.1. Reducing to likelihood test

We prove the minimax lower bound by conducting likelihood test. Similar to Lemma C.1 in [Duan and Wang \(2020\)](#), we have Lemma D.1 below.

Lemma D.1. Let M_α and M_β be two MDP instances with transition kernels $p_\alpha(x' | x, a)$ and $p_\beta(x' | x, a)$. Suppose Assumption 2.2 holds. Define likelihood functions

$$\mathcal{L}_i(\mathcal{D}) := \prod_{k=1}^K \bar{\xi}_0(x_0^{(k)}) \prod_{h=0}^{L-1} \bar{\pi}(a_h^{(k)} | x_h^{(k)}) p_i(x_{h+1}^{(k)} | x_h^{(k)}, a_h^{(k)}), \quad i = \alpha, \beta.$$

Denote $\mathbb{P}_\alpha$ the probability space generated by running M_α following the behavioral policy $\bar{\pi}$. If $\mathbb{P}_\alpha\left(\frac{\mathcal{L}_\beta(\mathcal{D})}{\mathcal{L}_\alpha(\mathcal{D})} \geq \frac{1}{2}\right) \geq \frac{1}{2}$ and there exist scalars $\rho_\alpha, \rho_\beta \geq 0$ such that

$$\{\text{policy } \pi \mid v_{M_\alpha, \xi_0}^* - v_{M_\alpha, \xi_0}^\pi \geq \rho_\alpha\} \cap \{\text{policy } \pi \mid v_{M_\beta, \xi_0}^* - v_{M_\beta, \xi_0}^\pi \geq \rho_\beta\} = \emptyset, \quad (\text{D.1})$$

then for any policy learning algorithm $\hat{\pi}$,

$$\mathbb{P}_\alpha\left(v_{M_\alpha, \xi_0}^* - v_{M_\alpha, \xi_0}^{\hat{\pi}(\mathcal{D})} \geq \rho_\alpha\right) \geq \frac{1}{6} \quad \text{or} \quad \mathbb{P}_\beta\left(v_{M_\beta, \xi_0}^* - v_{M_\beta, \xi_0}^{\hat{\pi}(\mathcal{D})} \geq \rho_\beta\right) \geq \frac{1}{6}. \quad (\text{D.2})$$

We learn from Lemma D.1 that as long as $\mathbb{P}_\alpha\left(\frac{\mathcal{L}_\beta(\mathcal{D})}{\mathcal{L}_\alpha(\mathcal{D})} \geq \frac{1}{2}\right) \geq \frac{1}{2}$ and (D.1) hold, the lower bound is achieved at model M_α or M_β . In the following, we construct MDP models and analyze these two conditions separately.

D.2. Constructing MDP instances

We assume without loss of generality that the number of active features s is even. We consider a simplest case where the MDP only consists of two states, *i.e.* $\mathcal{X} = \{\bar{x}, \underline{x}\}$. At each state, the agent chooses from $\frac{s}{2} + s(d-s)$ actions $\mathcal{A} = \{a_1, a_2, \dots, a_{\frac{s}{2}}\} \cup \{\bar{a}_{i,k} \mid i = 1, 2, \dots, \frac{s}{2}, k = \pm 1, \pm 2, \dots, \pm(d-s)\}$. Here, we only use $\bar{a}_{i,j}$ in collecting the dataset $\mathcal{D}$.

We first introduce Lemma D.2, which will be used in the construction of feature mapping $\phi : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$.

Lemma D.2. For any $s \in \mathbb{Z}_+$, there exists an s -by- s orthogonal matrix $\Theta \in \mathbb{R}^{s \times s}$ satisfying

$$\sqrt{s} \cdot |\Theta_{i,j}| \leq \sqrt{2}, \quad \text{for } i, j = 1, 2, \dots, s. \quad (\text{D.3})$$

Proof. Consider the discrete cosine transform (DCT) matrix $\Theta \in \mathbb{R}^{s \times s}$, given by

$$\Theta_{i,1} = \frac{1}{\sqrt{s}}, \quad i = 1, 2, \dots, s, \quad \Theta_{i,j} = \sqrt{\frac{2}{s}} \cos \frac{(2i-1)(j-1)\pi}{2s}, \quad i = 1, 2, \dots, s, \quad j = 2, 3, \dots, s.$$

Θ is orthogonal and satisfies (D.3). $\square$

Let $\Theta \in \mathbb{R}^{s \times s}$ be the orthogonal matrix given in Lemma D.2. We fix $\mathcal{K} \subseteq [d]$ to be the active feature set and denote by $\phi_{\mathcal{K}}$ the corresponding coordinates of ϕ . We now construct $\phi_{\mathcal{K}} : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^s$ as follows: For $i = 1, 2, \dots, \frac{s}{2}$, $k = \pm 1, \pm 2, \dots, \pm(d-s)$, let

$$\phi_{\mathcal{K}}(\bar{x}, a_i) := \sqrt{\frac{s}{2}} \cdot \Theta \begin{pmatrix} 0 & 0 & \cdots & 1 & 0 & \cdots & 0 & 0 \end{pmatrix}^{\top} \in \mathbb{R}^s,$$

$$\phi_{\mathcal{K}}(\bar{x}, \bar{a}_{i,k}) := \sqrt{\frac{s}{2}} \cdot \Theta \begin{pmatrix} 0 & 0 & \cdots & 1 - \varsigma_1 & \varsigma_1 & \cdots & 0 & 0 \end{pmatrix}^{\top} \in \mathbb{R}^s,$$

$$\phi_{\mathcal{K}}(\underline{x}, a_i) = \phi_{\mathcal{K}}(\underline{x}, \bar{a}_{i,k}) := \sqrt{\frac{s}{2}} \cdot \Theta \begin{pmatrix} 0 & 0 & \cdots & \varsigma_2 & 1 - \varsigma_2 & \cdots & 0 & 0 \end{pmatrix}^{\top} \in \mathbb{R}^s,$$

where $\varsigma_1, \varsigma_2 \in (0, 1)$ will be determined later. By construction, we have $\|\phi_{\mathcal{K}}(x, a)\|_{\infty} \leq 1$ for any $(x, a) \in \mathcal{X} \times \mathcal{A}$. Note that $\phi_{\mathcal{K}}$ abstracts all the dynamic information for state-action pairs, and $\phi_{\mathcal{K}^c}$ does not affect the transition model or reward function. Therefore, it is sufficient for us to use $\phi_{\mathcal{K}}$ when identifying the optimal policy or calculate value functions.

We propose $\frac{s}{2}$ MDP models $M_1, M_2, \dots, M_{\frac{s}{2}}$, where M_i has transition kernel $p_i(x' | x, a) = \phi_{\mathcal{K}}(x, a)^{\top} \psi_i(x')$ given by

$$\psi_i(\bar{x}) = \sqrt{\frac{2}{s}} \cdot \Theta \begin{pmatrix} 1 - \delta_1 & \delta_2 & \cdots & 1 & 0 & \cdots & 1 - \delta_1 & \delta_2 \end{pmatrix}^{\top} \in \mathbb{R}^s,$$

$$\psi_i(\underline{x}) = \sqrt{\frac{2}{s}} \cdot \Theta \begin{pmatrix} \delta_1 & 1 - \delta_2 & \cdots & 0 & 1 & \cdots & \delta_1 & 1 - \delta_2 \end{pmatrix}^{\top} \in \mathbb{R}^s.$$

Here, $\delta_1, \delta_2 \in [0, 2(1 - \gamma))$ are parameters reflecting the small differences among actions.

The reward functions are the same for all models and are chosen as

$$r(\bar{x}, a_i) = r(\bar{x}, \bar{a}_{i,k}) = 1, \quad r(\underline{x}, a_i) = r(\underline{x}, \bar{a}_{i,k}) = 0,$$

for $i = 1, 2, \dots, \frac{s}{2}$, $j = \pm 1, \pm 2, \dots, \pm(d-s)$.

D.3. Analyzing the concentration of the likelihood ratio

We devise a behavior policy $\bar{\pi}$ and verify the likelihood ratio condition under the data collecting scheme in Assumption 2.2. We start from an initial distribution $\bar{\xi}_0$ and take a behavior policy $\bar{\pi}(\bar{a}_{i,k} | \bar{x}) = \bar{\pi}(\bar{a}_{i,k} | \underline{x}) = \frac{1}{s(d-s)}$ for any i, k . Under this specific $\bar{\pi}$, due to symmetry, all MDP models $M_1, M_2, \dots, M_{\frac{s}{2}}$ have the same marginal distribution at each time step

$l = 0, 1, \dots, L-1$, which we denote by $\bar{\xi}_l = \begin{pmatrix} \bar{\xi}_l(\bar{x}) \\ \bar{\xi}_l(\underline{x}) \end{pmatrix} \in \mathbb{R}^2$. Define the average distribution as $\bar{\xi} := \frac{1}{L} \sum_{l=0}^{L-1} \bar{\xi}_l \in \mathbb{R}^2$.

Take

$$p_{\min} := \min \left\{ p_i(\bar{x} | \bar{x}, \bar{a}_{i,k}), p_i(\underline{x} | \bar{x}, \bar{a}_{i,k}), p_i(\bar{x} | \underline{x}, \bar{a}_{i,k}), p_i(\underline{x} | \underline{x}, \bar{a}_{i,k}) \right\}$$

$$i = 1, 2, \dots, \frac{s}{2}, k = \pm 1, \pm 2, \dots, \pm(d-s)$$

and

$$\Sigma^{\circ} := \bar{\xi}(\bar{x}) \begin{pmatrix} 1 - \varsigma_1 \\ \varsigma_1 \end{pmatrix} \begin{pmatrix} 1 - \varsigma_1 & \varsigma_1 \end{pmatrix} + \bar{\xi}(\underline{x}) \begin{pmatrix} \varsigma_2 \\ 1 - \varsigma_2 \end{pmatrix} \begin{pmatrix} \varsigma_2 & 1 - \varsigma_2 \end{pmatrix}.$$

Parallel to Lemma C.3 in (Duan and Wang, 2020), we provide concentration results of the likelihood ratio in Lemma D.3. The proof can be found in Appendix E.2.1.

Lemma D.3. If we take $\delta_1, \delta_2 \geq 0$ such that

$$(\delta_1 \quad -\delta_2) \Sigma^{\circ} \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix} \leq \frac{sp_{\min}}{100N}, \quad \delta_1 \vee \delta_2 \leq \frac{p_{\min}}{100\sqrt{L}}, \quad (\text{D.4})$$

then for any $i, j = 1, 2, \dots, s, i \neq j$, it holds that

$$\mathbb{P}_i \left(\frac{\mathcal{L}_j(\mathcal{D})}{\mathcal{L}_i(\mathcal{D})} \geq \frac{1}{2} \right) \geq \frac{1}{2}. \quad (\text{D.5})$$

Lemma D.3 suggests that as long as (D.4) is satisfied, the likelihood test in Lemma D.1 works for any pair of indices $(\alpha, \beta) = (i, j), i \neq j$.

D.4. Calculating the gap in values

For model M_i , the optimal policy is given by

$$\pi_i^*(\bar{x}) = a_i \quad \text{and} \quad \pi_i^*(\underline{x}) = \begin{cases} a_j \text{ for any } j \neq i, & \text{if } (1 - \varsigma_2)\delta_2 > \varsigma_2\delta_1, \\ a_i, & \text{otherwise.} \end{cases}$$

For computational simplicity, we take initial distribution $\xi_0 := \begin{pmatrix} \xi_0(\bar{x}) \\ \xi_0(\underline{x}) \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \in \mathbb{R}^2$. In the following Lemma D.4, we provide an estimation for the difference between values of optimal and sub-optimal policies. See Appendix E.2.2 for the proof.

Lemma D.4. If $\delta_1 \leq \frac{1-\gamma}{\gamma}$, $\delta_2 \leq \varsigma_2$, then for any policy π such that $\pi(\bar{x}) \neq a_i$, it holds that

$$v_{M_i, \xi_0}^* - v_{M_i, \xi_0}^\pi \geq \frac{\gamma\delta_1}{2(1-\gamma)} \cdot \frac{1}{1-\gamma+2\gamma\varsigma_2}.$$

According to Lemma D.4, if we take

$$\rho_i = \rho' := \frac{\gamma\delta_1}{2(1-\gamma)} \cdot \frac{1}{1-\gamma+2\gamma\varsigma_2}, \quad i = 1, 2, \dots, \frac{s}{2}, \quad (\text{D.6})$$

then condition (D.1) in Lemma D.1 holds for any $(\alpha, \beta) = (i, j), i \neq j$.

D.5. Choosing parameters

We now integrate Lemmas D.1, D.3 and D.4. Specifically, we choose parameters $\varsigma_1, \varsigma_2, \delta_1, \delta_2$ and $\bar{\xi}$ that maximize ρ' in (D.6) under the constraint (D.4).

We first consider the optimization problem

$$\text{maximize } \delta_1, \quad \text{subject to } (\delta_1 \quad -\delta_2) \Sigma^\circ \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix} \leq \frac{sp_{\min}}{100N}.$$

It has solution

$$\delta_1 = \sqrt{\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)}} \sqrt{\frac{sp_{\min}}{100N}}, \quad \delta_2 = \sqrt{\frac{(\Sigma_{12}^\circ)^2}{\Sigma_{22}^\circ \det(\Sigma^\circ)}} \sqrt{\frac{sp_{\min}}{100N}}. \quad (\text{D.7})$$

Plugging (D.7) into (D.6) and assuming that $p_{\min} \geq \frac{\varsigma_2}{2}$, we have

$$\rho' \geq \frac{\gamma}{2(1-\gamma)} \cdot \frac{\sqrt{\varsigma_2}}{1-\gamma+2\gamma\varsigma_2} \cdot \sqrt{\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)}} \sqrt{\frac{s}{200N}}. \quad (\text{D.8})$$

We maximize the right hand side of (D.8) over ς_2 , and obtain

$$\varsigma_2 = \frac{1-\gamma}{2\gamma}, \quad \rho' \geq \frac{\sqrt{\gamma}}{80} \sqrt{\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)}} \sqrt{\frac{s}{N(1-\gamma)^3}}.$$

We further let $\varsigma_1 \in [\frac{1-\gamma}{2\gamma}, 1 - \frac{1-\gamma}{2\gamma})$ and suppose the sample size

$$N \geq \frac{(\Sigma_{22}^\circ \vee \Sigma_{12}^\circ)^2}{\Sigma_{22}^\circ \det(\Sigma^\circ)} \frac{400sL}{1-\gamma}. \quad (\text{D.9})$$

In this case, $p_{\min} \geq \frac{s}{2}$ and $\delta_1 \vee \delta_2 \leq \frac{p_{\min}}{100\sqrt{L}} \leq \varsigma_2 \leq \frac{1-\gamma}{\gamma}$.

In summary, if the sample size N satisfies (D.9) and we take

$$\gamma \geq \frac{1}{2}, \quad \varsigma_1 \in \left[\frac{1-\gamma}{2\gamma}, 1 - \frac{1-\gamma}{2\gamma}\right), \quad \varsigma_2 = \frac{1-\gamma}{2\gamma} \quad \text{and} \quad \delta_1, \delta_2 \text{ in (D.7),}$$

then the conditions in Lemmas D.3 and D.4 are satisfied and (D.2) holds for

$$\rho := \frac{1}{80\sqrt{2}} \sqrt{\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)}} \sqrt{\frac{s}{N(1-\gamma)^3}}. \quad (\text{D.10})$$

Remark that under this construction, we still have the flexibility to take $\varsigma_1 \nearrow 1 - \frac{1-\gamma}{2\gamma}$ so that Σ° is very ill-conditioned. For instance, if we take $\varsigma_1 = 1 - \frac{1-\gamma}{\gamma}$, then $\det(\Sigma^\circ)$ or $\lambda_{\min}(\Sigma^\circ)$ at least has the order of $(1-\gamma)^3$.

In order that condition (D.9) is as weak as possible, we take $\gamma \geq \frac{2}{3}$, $\varsigma_1 = \frac{1-\gamma}{2\gamma}$ and $\bar{\xi}(\bar{x}) = \bar{\xi}(\underline{x}) = \frac{1}{2}$. In this setting, if $N \geq 2000sL(1-\gamma)^{-1}$ then (D.9) holds.

D.6. Relating to mismatch terms

In this part, we relate $\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)}$ in (D.10) to mismatch terms $\chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu})$ and $C_{\min}(\Sigma, s)$.

D.6.1. RESTRICTED χ^2 -DIVERGENCE

According to Lemma B.2 in Duan and Wang (2020), we have

$$1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu}) = \sup_{f \in \mathcal{G}(\mathcal{K})} \frac{\mathbb{E}_i[f(x, a) \mid (x, a) \sim \mu^*]^2}{\mathbb{E}_i[f^2(x, a) \mid (x, a) \sim \bar{\mu}]} = (\nu_{\mathcal{K}}^*)^\top \Sigma_{\mathcal{K}}^{-1} \nu_{\mathcal{K}}^*,$$

where

$$\nu_{\mathcal{K}}^* := \mathbb{E}_i[\phi_{\mathcal{K}}(x, a) \mid (x, a) \sim \mu^*] = \frac{1}{1-\gamma} \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_i[\phi_{\mathcal{K}}(x_t, a_t) \mid x_0 \sim \xi_0, \pi_i^*] \in \mathbb{R}^s$$

and

$$\begin{aligned} \Sigma_{\mathcal{K}} &:= \mathbb{E}_i[\phi_{\mathcal{K}}(x, a) \phi_{\mathcal{K}}(x, a)^\top \mid (x, a) \sim \bar{\mu}] \\ &= \mathbb{E}_i\left[\frac{1}{L} \sum_{h=0}^{L-1} \phi_{\mathcal{K}}(x_h^{(k)}, a_h^{(k)}) \phi_{\mathcal{K}}(x_h^{(k)}, a_h^{(k)})^\top \mid x_0^{(k)} \sim \bar{\xi}_0, \bar{\pi}\right] \in \mathbb{R}^{s \times s}. \end{aligned}$$

For model M_i , $\bar{x}$ is an absorbing state under the optimal policy π_i^* . Therefore, $\mu^*(\bar{x}) = 1$ and $\nu_{\mathcal{K}}^* = \phi_{\mathcal{K}}(\bar{x}, a_i)$. Under our proposed behavior policy $\bar{\pi}$, we have

$$\Sigma_{\mathcal{K}} = \Theta \begin{pmatrix} \Sigma^\circ & 0 & \cdots & 0 \\ 0 & \Sigma^\circ & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \Sigma^\circ \end{pmatrix} \Theta^\top. \quad (\text{D.11})$$

It follows that

$$1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu}) = (\nu_{\mathcal{K}}^*)^\top \Sigma_{\mathcal{K}}^{-1} \nu_{\mathcal{K}}^* = \frac{s}{2} ((\Sigma^\circ)^{-1})_{1,1} = \frac{s \Sigma_{22}^\circ}{2 \det(\Sigma^\circ)}.$$

To this end, we have

$$\rho = \frac{1}{80} \sqrt{1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu})} \frac{1}{\sqrt{N(1-\gamma)^3}}.$$

This implies there always exists an DMDP instance $M \in \mathcal{M}_{\phi,s}(\mathcal{X}, \mathcal{A}, \gamma)$ with feature $\phi \in (\mathbb{R}^d)^{\mathcal{X} \times \mathcal{A}}$ satisfying $\|\phi(x, a)\|_\infty \leq 1$ for all $(x, a) \in \mathcal{X} \times \mathcal{A}$, then

$$\mathbb{P}_M \left(v_{\xi_0}^* - v_{\xi_0}^{\hat{\pi}(\mathcal{D})} \gtrsim \frac{1}{(1-\gamma)^{\frac{3}{2}}} \sqrt{1 + \chi_{\mathcal{G}(\mathcal{K})}^2(\mu^*, \bar{\mu})} \sqrt{\frac{1}{N}} \right) \geq \frac{1}{6},$$

where μ^* is the discounted state-action occupancy measure of π^* .

D.6.2. RESTRICTED MINIMUM EIGENVALUE

The uncentered covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ is given by

$$\Sigma := \mathbb{E}_i \left[\frac{1}{L} \sum_{h=0}^{L-1} \phi(x_h^{(k)}, a_h^{(k)}) \phi(x_h^{(k)}, a_h^{(k)})^\top \middle| x_0^{(k)} \sim \bar{\xi}_0, \bar{\pi} \right] \in \mathbb{R}^{s \times s}.$$

In the following, we specify the choice of $\phi_{\mathcal{K}^c}(\bar{x}, \bar{a}_{i,k})$ and $\phi_{\mathcal{K}^c}(\underline{x}, \bar{a}_{i,k})$ and show that if

$$\bar{\xi}(\bar{x})\varsigma_1^2 + \bar{\xi}(\underline{x})(1 - \varsigma_2)^2 \geq \bar{\xi}(\bar{x})(1 - \varsigma_1)^2 + \bar{\xi}(\underline{x})\varsigma_2^2, \quad (\text{D.12})$$

then

$$\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)} \geq \frac{1}{2C_{\min}(\Sigma, s)}. \quad (\text{D.13})$$

Under condition (D.12), it holds that $\Sigma_{22}^\circ \geq \Sigma_{11}^\circ$, therefore, $\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)} \geq \frac{\text{Tr}(\Sigma^\circ)}{2\det(\Sigma^\circ)}$. In addition, for the 2-by-2 matrix Σ° , we have $\lambda_{\min}(\Sigma^\circ) + \lambda_{\max}(\Sigma^\circ) = \text{Tr}(\Sigma^\circ)$ and $\lambda_{\min}(\Sigma^\circ)\lambda_{\max}(\Sigma^\circ) = \det(\Sigma^\circ)$. It follows that

$$\frac{\Sigma_{22}^\circ}{\det(\Sigma^\circ)} \geq \frac{\text{Tr}(\Sigma^\circ)}{2\det(\Sigma^\circ)} \geq \frac{\lambda_{\max}(\Sigma^\circ)}{2\lambda_{\max}(\Sigma^\circ)\lambda_{\min}(\Sigma^\circ)} = \frac{1}{2\lambda_{\min}(\Sigma^\circ)}.$$

We next relate $\lambda_{\min}(\Sigma^\circ)$ to $C_{\min}(\Sigma, s)$.

Let $\bar{\Theta} \in \mathbb{R}^{(d-s) \times (d-s)}$ be an orthogonal matrix given by Lemma D.2. We take

$$\phi_{\mathcal{K}^c}(\bar{x}, \bar{a}_{i,k}) = \phi_{\mathcal{K}^c}(\underline{x}, \bar{a}_{i,k}) := \text{sign}(k) \sqrt{\frac{d-s}{2}} \cdot \text{col}_k(\bar{\Theta}),$$

for $i = 1, 2, \dots, \frac{s}{2}$, $k = \pm 1, \pm 2, \dots, \pm(d-s)$. It holds that $\|\phi_{\mathcal{K}^c}(\bar{x}, \bar{a}_{i,k})\|_\infty \leq 1$ and $\|\phi_{\mathcal{K}^c}(\underline{x}, \bar{a}_{i,k})\|_\infty \leq 1$. For notational simplicity, let $\mathcal{K} = [s]$. Under our proposed behavior policy $\bar{\pi}(\bar{a}_{i,k} | \bar{x}) = \bar{\pi}(\bar{a}_{i,k} | \underline{x}) = \frac{1}{s(d-s)}$, we have

$$\Sigma = \begin{pmatrix} \Sigma_{\mathcal{K}} & 0 \\ 0 & \frac{1}{2} I_{d-s} \end{pmatrix}. \quad (\text{D.14})$$

By (D.11), $\lambda_{\min}(\Sigma_{\mathcal{K}}) = \lambda_{\min}(\Sigma^\circ)$. We also note that $\text{Tr}(\Sigma^\circ) = \bar{\xi}(\bar{x})\|(1 - \varsigma_1, \varsigma_1)\|_2^2 + \bar{\xi}(\underline{x})\|(\varsigma_2, 1 - \varsigma_2)\|_2^2 \leq 1$, and therefore

$$\lambda_{\min}(\Sigma^\circ) \leq \frac{\text{Tr}(\Sigma^\circ)}{2} \leq \frac{1}{2}.$$

It follows that $\lambda_{\min}(\Sigma) = \lambda_{\min}(\Sigma^\circ)$, which further implies $C_{\min}(\Sigma, s) \geq \lambda_{\min}(\Sigma) = \lambda_{\min}(\Sigma^\circ)$. On the other hand, the eigenvector of Σ corresponding to $\lambda_{\min}(\Sigma^\circ)$ has support set $\mathcal{K}$ and is s -sparse. Hence, $\lambda_{\min}(\Sigma^\circ) \geq C_{\min}(\Sigma, s)$. In this way, we have proved $C_{\min}(\Sigma, s) = \lambda_{\min}(\Sigma^\circ)$ for Σ defined in (D.14).

In the special case where $\varsigma_1 = \varsigma_2 = \frac{1-\gamma}{2\gamma}$ and $\bar{\xi}(\bar{x}) = \bar{\xi}(\underline{x}) = \frac{1}{2}$, condition (D.12) holds. Plugging (D.13) into (D.10), we finish our proof of Theorem 5.2.

E. Proofs of Auxiliary Results

E.1. Proof of Lemma B.3

We prove if the population covariance matrix satisfies the restricted eigenvalue condition, the empirical covariance matrix satisfies it as well with high probability. Recall that

$$\widehat{\Sigma} = \frac{1}{K} \sum_{k=1}^K \frac{1}{L} \sum_{h=0}^{L-1} \phi(x_h^{(k)}, a_h^{(k)}) \phi(x_h^{(k)}, a_h^{(k)})^\top,$$

and

$$\widehat{\Psi} = \begin{pmatrix} \widehat{\Sigma} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \widehat{\Sigma} \end{pmatrix}.$$

For any $i, j \in [d]$, define

$$v_{ij}^{(k)} = \frac{1}{L} \sum_{h=0}^{L-1} \phi_i(x_h^{(k)}, a_h^{(k)}) \phi_j(x_h^{(k)}, a_h^{(k)}) - \Sigma_{ij}.$$

It is easy to verify $\mathbb{E}[v_{ij}^{(k)}] = 0$ and $|v_{ij}^{(k)}| \leq 1$ since we assume $\|\phi(x, a)\|_\infty \leq 1$. Note that from the data collection process, samples between different episodes are independent. This implies $v_{ij}^{(1)}, \dots, v_{ij}^{(K)}$ are independent. By standard Hoeffding's inequality (Proposition 5.10 in [Vershynin \(2010\)](#)), we have

$$\mathbb{P}\left(\left|\sum_{k=1}^K v_{ij}^{(k)}\right| \geq \delta\right) \leq 3 \exp\left(-\frac{C_0 \delta^2}{K}\right),$$

for some absolute constant $C_0 > 0$. Applying the union bound over $i, j \in [d]$, we have

$$\begin{aligned} \mathbb{P}\left(\max_{i,j} \left|\sum_{k=1}^K v_{ij}^{(k)}\right| \geq \delta\right) &\leq 3d^2 \exp\left(-\frac{C_0 \delta^2}{K}\right) \\ \Rightarrow \mathbb{P}\left(\|\widehat{\Sigma} - \Sigma\|_\infty \geq \delta\right) &\leq 3d^2 \exp\left(-\frac{C_0 \delta^2}{K}\right). \end{aligned}$$

Since the blocks of Ψ are the same, the following holds with probability $1 - \delta$.

$$\|\widehat{\Psi} - \Psi\|_\infty \leq \sqrt{\frac{\log(3d^2/\delta)}{K}}.$$

Therefore, when the number of episodes $K \geq 32^2 \log(3d^2/\delta) s^2 / \widetilde{C}_{\min}(\Psi, s)^2$, the following holds with probability at least $1 - \delta$,

$$\|\widehat{\Psi} - \Psi\|_\infty \leq \frac{\widetilde{C}_{\min}(\Psi, s)}{32s}.$$

Next lemma shows that if the restricted eigenvalue condition holds for one positive semi-definite block diagonal matrix Σ_0 , then it holds with high probability for another positive semi-definite block diagonal matrix Σ_1 as long as Σ_0 and Σ_1 are close enough in terms of entry-wise max norm.

Lemma E.1 (Corollary 6.8 in [Bühlmann and Van De Geer \(2011\)](#)). Let Σ_0 and Σ_1 be two positive semi-definite block diagonal matrices. Suppose that the restricted eigenvalue of Σ_0 satisfies $\widetilde{C}_{\min}(\Sigma_0, s) > 0$ and $\|\Sigma_1 - \Sigma_0\|_\infty \leq \widetilde{C}_{\min}(\Sigma_0, s)/(32s)$. Then the restricted eigenvalue of Σ_1 satisfies $\widetilde{C}_{\min}(\Sigma_1, s) > \widetilde{C}_{\min}(\Sigma_0, s)/2$.

Applying Lemma E.1 with $\widehat{\Psi}$ and Ψ , we have the restricted eigenvalue of $\widehat{\Psi}$ satisfies $\widetilde{C}_{\min}(\widehat{\Psi}, s) > \widetilde{C}_{\min}(\Psi, s)/2$ with probability at least $1 - \delta$, as long as the sample size $N \geq 32^2 L \log(3d^2/\delta) s^2 / \widetilde{C}_{\min}(\Psi, s)^2$. This ends the proof. $\blacksquare$

E.2. Proof of Lemma B.1

We prove the ℓ_1 -norm bound of estimating $\bar{w}_t$ using a fresh fold of batch data. We overload the notation $\hat{\Sigma}$ to denote

$$\hat{\Sigma} = \frac{1}{RL} \sum_{i=1}^{RL} \phi(x_i, a_i) \phi(x_i, a_i)^\top = \frac{1}{R} \sum_{r=1}^R \frac{1}{L} \sum_{h=1}^L \phi(x_h^{(r)}, a_h^{(r)}) \phi(x_h^{(r)}, a_h^{(r)})^\top.$$

Similar to the proof of Lemma B.3 in Appendix E.1, we can have with probability at least $1 - \delta$,

$$\|\hat{\Sigma} - \Sigma\|_\infty \leq \sqrt{\frac{C_1}{R} \log\left(\frac{3d^2}{\delta}\right)},$$

where C_1 is an absolute constant. When $R \geq C_1 32^2 \log(3d^2/\delta) s^2 / C_{\min}(\Sigma, s)$, we have

$$\|\hat{\Sigma} - \Sigma\|_\infty \leq \frac{C_{\min}(\Sigma, s)}{32s}.$$

Applying Lemma E.1, we have $C_{\min}(\hat{\Sigma}, s) > C_{\min}(\Sigma, s)/2$ with probability at least $1 - \delta$. Note that $\{\varepsilon_i \phi_j(x_i, a_i)\}_{i=1}^{RL}$ is a martingale difference sequence and $|\varepsilon_i \phi_j(x_i, a_i)| \leq 1/(1 - \gamma)$. Similar to the proof of Eq. (B.16) by Azuma-Hoeffding inequality,

$$\mathbb{P}\left(\max_{j \in [d]} \left| \frac{1}{RL} \sum_{i=1}^{RL} \varepsilon_i \phi_j(x_i, a_i) \right| \leq \frac{1}{1 - \gamma} \sqrt{\frac{\log(2d/\delta)}{RL}}\right) \geq 1 - \delta.$$

Denote event $\mathcal{E}$ as

$$\mathcal{E} = \left\{ \max_{j \in [d]} \left| \frac{1}{RL} \sum_{i=1}^{RL} \varepsilon_i \phi_j(x_i, a_i) \right| \leq \lambda_1 \right\}.$$

Then $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$. Under event $\mathcal{E}$, applying (B.31) in Bickel et al. (2009), we have

$$\|\hat{w}_t - \bar{w}_t\|_1 \leq \frac{16\sqrt{2}s\lambda_1}{C_{\min}(\Sigma, s)},$$

holds with probability at least $1 - 2\delta$. This ends the proof. ■

E.2.1. PROOF OF LEMMA D.3

Proof of Lemma D.3. It is easy to see that

$$\ln \frac{\mathcal{L}_j(\mathcal{D})}{\mathcal{L}_i(\mathcal{D})} = \sum_{k=1}^K \sum_{l=0}^{L-1} \ln \frac{p_j(x_{l+1}^{(k)} | x_l^{(k)}, a_l^{(k)})}{p_i(x_{l+1}^{(k)} | x_l^{(k)}, a_l^{(k)})} = \sum_{k=1}^K \sum_{l=0}^{L-1} \ln(1 - \Lambda_l^{(k)}),$$

where

$$\Lambda_l^{(k)} := \frac{p_i(x_{l+1}^{(k)} | x_l^{(k)}, a_l^{(k)}) - p_j(s_{l+1}^{(k)} | s_l^{(k)}, a_l^{(k)})}{p_i(s_{l+1}^{(k)} | s_l^{(k)}, a_l^{(k)})} = \frac{\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(s_{l+1}^{(k)}) - \psi_j(s_{l+1}^{(k)}))}{p_i(s_{l+1}^{(k)} | s_l^{(k)}, a_l^{(k)})}.$$

If we take $\delta_1 \vee \delta_2 \leq \frac{p_{\min}}{2}$, then $|\Lambda_l^{(k)}| \leq \frac{1}{2}$ and

$$\ln \frac{\mathcal{L}_j(\mathcal{D})}{\mathcal{L}_i(\mathcal{D})} \geq - \underbrace{\sum_{k=1}^K \sum_{l=0}^{L-1} \Lambda_l^{(k)}}_{E_1} - \underbrace{\sum_{k=1}^K \sum_{l=0}^{L-1} (\Lambda_l^{(k)})^2}_{E_2}. \quad (\text{E.1})$$

Since $\mathbb{E}_i[\Lambda_l^{(k)} | s_l^{(k)}, a_l^{(k)}] = 0$, we apply Freedman's inequality to analyze E_1 . The conditional variances satisfy

$$\begin{aligned} & \mathbb{E}_i[(\Lambda_l^{(k)})^2 | s_l^{(k)}, a_l^{(k)}] \\ &= p_i(\bar{x} | s_l^{(k)}, a_l^{(k)}) \left(\frac{\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(\bar{x}) - \psi_j(\bar{x}))}{p_i(\bar{x} | s_l^{(k)}, a_l^{(k)})} \right)^2 \\ & \quad + p_i(\underline{x} | s_l^{(k)}, a_l^{(k)}) \left(\frac{\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(\underline{x}) - \psi_j(\underline{x}))}{p_i(\underline{x} | s_l^{(k)}, a_l^{(k)})} \right)^2 \\ &= \frac{(\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(\bar{x}) - \psi_j(\bar{x})))^2}{p_i(\bar{x} | s_l^{(k)}, a_l^{(k)}) p_i(\underline{x} | s_l^{(k)}, a_l^{(k)})} \leq \frac{1}{p_{\min}(1 - p_{\min})} (\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(\bar{x}) - \psi_j(\bar{x})))^2. \end{aligned}$$

Denote $\Xi_k := \frac{1}{L} \sum_{l=0}^{L-1} (\phi(s_l^{(k)}, a_l^{(k)})^\top (\psi_i(\bar{x}) - \psi_j(\bar{x})))^2$. Note that

$$\psi_i(\bar{x}) - \psi_j(\bar{x}) = \sqrt{\frac{2}{s}} \cdot \Theta \begin{pmatrix} 2i-1 & 2i & & 2j-1 & 2j \\ 0 & 0 & \cdots & \delta_1 & -\delta_2 & \cdots & -\delta_1 & \delta_2 & \cdots & 0 & 0 \end{pmatrix}^\top \in \mathbb{R}^s,$$

therefore,

$$\mathbb{E}[\Xi_k] = \frac{4}{s} (\delta_1 \quad -\delta_2) \Sigma^\circ \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix} \quad \text{and} \quad |\Xi_k| \leq (\delta_1 \vee \delta_2)^2.$$

By Bernstein inequality and the independence of trajectories $\tau_1, \tau_2, \dots, \tau_K$, we have with $\mathbb{P}_i$ -probability at least $\frac{5}{6}$,

$$\frac{1}{K} \sum_{k=1}^K \Xi_k \leq \left(\sqrt{\mathbb{E}[\Xi_k]} + (\delta_1 \vee \delta_2) \sqrt{\frac{2 \ln 6}{3K}} \right)^2 =: \sigma^2. \quad (\text{E.2})$$

Since $|\Lambda_l^{(k)}| \leq p_{\min}^{-1} (\delta_1 \vee \delta_2)$, by Freedman's inequality, with $\mathbb{P}_i$ -probability at least $\frac{5}{6}$,

$$\frac{1}{N} \sum_{k=1}^K \sum_{l=0}^{L-1} \Lambda_l^{(k)} \leq \frac{\sigma}{\sqrt{p_{\min}(1 - p_{\min})}} \sqrt{\frac{2 \ln 6}{N}} + p_{\min}^{-1} (\delta_1 \vee \delta_2) \frac{2 \ln 6}{3N} \quad \text{and} \quad (\text{E.2}) \text{ holds.} \quad (\text{E.3})$$

Combining (E.2) and (E.3), we use union bound and derive that with $\mathbb{P}_i$ -probability at least $\frac{2}{3}$,

$$\sum_{k=1}^K \sum_{l=0}^{L-1} \Lambda_l^{(k)} \leq 4\sqrt{\ln 6} \sqrt{(\delta_1 \quad -\delta_2) \Sigma^\circ \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix}} \sqrt{\frac{N}{p_{\min}s}} + \frac{2 \ln 6}{3} (\delta_1 \vee \delta_2) \left(\sqrt{\frac{6L}{p_{\min}}} + \frac{1}{p_{\min}} \right). \quad (\text{E.4})$$

As for E_2 , since

$$\mathbb{E}_i \left[\frac{1}{L} \sum_{l=0}^{L-1} (\Lambda_l^{(k)})^2 \right] \leq \frac{4}{s p_{\min}(1 - p_{\min})} (\delta_1 \quad -\delta_2) \Sigma^\circ \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix}, \quad \left| \frac{1}{L} \sum_{l=0}^{L-1} (\Lambda_l^{(k)})^2 \right| \leq \frac{(\delta_1 \vee \delta_2)^2}{p_{\min}^2},$$

by Bernstein's inequality, with $\mathbb{P}_i$ -probability at least $\frac{5}{6}$,

$$\sum_{k=1}^K \sum_{l=0}^{L-1} (\Lambda_l^{(k)})^2 \leq \left(2\sqrt{2} \sqrt{(\delta_1 \quad -\delta_2) \Sigma^\circ \begin{pmatrix} \delta_1 \\ -\delta_2 \end{pmatrix}} \sqrt{\frac{N}{s p_{\min}}} + \frac{\delta_1 \vee \delta_2}{p_{\min}} \sqrt{\frac{2L \ln 6}{3}} \right)^2. \quad (\text{E.5})$$

Plugging (E.4) and (E.5) into (E.1) and applying condition (D.4), we obtain (D.5).

□

E.2.2. PROOF OF LEMMA D.4

Proof of Lemma D.4. We consider another policy π'_i such that $\pi'_i(\bar{x}) = a_j$ for some $a_j \neq a_i$ and $\pi'_i(\underline{x}) = \pi_i^*(\underline{x})$. It holds that

$$v_{M_i, \xi_0}^{\pi'_i} = \sup \{v_{M_i, \xi_0}^\pi \mid \pi(\bar{x}) \neq a_i\}.$$

For any π , denote its corresponding transition matrix by $P_i^\pi = \begin{pmatrix} p_i^\pi(\bar{x}|\bar{x}) & p_i^\pi(\underline{x}|\bar{x}) \\ p_i^\pi(\bar{x}|\underline{x}) & p_i^\pi(\underline{x}|\underline{x}) \end{pmatrix} \in \mathbb{R}^{2 \times 2}$. Define $v_{M_i}^* := \begin{pmatrix} v_{M_i}^*(\bar{x}) \\ v_{M_i}^*(\underline{x}) \end{pmatrix} \in \mathbb{R}^2$. We have the following decomposition,

$$v_{M_i, \xi_0}^* - v_{M_i, \xi_0}^{\pi'_i} = \sum_{t=0}^{\infty} \gamma^{t+1} \xi_0^\top (P_i^{\pi'_i})^t (P_i^{\pi_i^*} - P_i^{\pi'_i}) v_{M_i}^* = \gamma \xi_0^\top (I - \gamma P_i^{\pi'_i})^{-1} (P_i^{\pi_i^*} - P_i^{\pi'_i}) v_{M_i}^*. \quad (\text{E.6})$$

Under model M_i , when $\delta_2 \leq \varsigma_2$, π_i^* and π'_i satisfy

$$\begin{aligned} p_i^{\pi_i^*}(\bar{x}|\bar{x}) &= 1, & p_i^{\pi'_i}(\bar{x}|\bar{x}) &= 1 - \delta_1, & p_i^{\pi_i^*}(\underline{x}|\bar{x}) &= 0, & p_i^{\pi'_i}(\underline{x}|\bar{x}) &= \delta_1, \\ p_i^{\pi_i^*}(\bar{x}|\underline{x}) &= p_i^{\pi'_i}(\bar{x}|\underline{x}) \leq 2\varsigma_2, & p_i^{\pi_i^*}(\underline{x}|\underline{x}) &= p_i^{\pi'_i}(\underline{x}|\underline{x}) \geq 1 - 2\varsigma_2, \end{aligned}$$

therefore,

$$(P_i^{\pi_i^*} - P_i^{\pi'_i}) v_{M_i}^* = \begin{pmatrix} \delta_1 & -\delta_1 \\ 0 & 0 \end{pmatrix} v_{M_i}^* = \begin{pmatrix} \delta_1 \\ 0 \end{pmatrix} \cdot (v_{M_i}^*(\bar{x}) - v_{M_i}^*(\underline{x})).$$

To this end, we reduce (E.6) into

$$v_{M_i, \xi_0}^* - v_{M_i, \xi_0}^{\pi'_i} = \gamma \xi_0^\top (I - \gamma P_i^{\pi'_i})^{-1} \begin{pmatrix} \delta_1 \\ 0 \end{pmatrix} \cdot (v_{M_i}^*(\bar{x}) - v_{M_i}^*(\underline{x})). \quad (\text{E.7})$$

Note that

$$v_{M_i}^*(\bar{x}) = (1 - \gamma)^{-1} \quad \text{and} \quad v_{M_i}^*(\underline{x}) = \frac{\gamma}{1 - \gamma} \cdot \frac{p_i^{\pi_i^*}(\underline{x}|\underline{x})}{1 - \gamma(1 - p_i^{\pi_i^*}(\underline{x}|\underline{x}))}.$$

Therefore,

$$v_{M_i}^*(\bar{x}) - v_{M_i}^*(\underline{x}) = \frac{1}{1 - \gamma} - \frac{\gamma}{1 - \gamma} \cdot \frac{p_i^{\pi_i^*}(\underline{x}|\underline{x})}{1 - \gamma(1 - p_i^{\pi_i^*}(\underline{x}|\underline{x}))} = \frac{1}{1 - \gamma p_i^{\pi_i^*}(\underline{x}|\underline{x})} \geq \frac{1}{1 - \gamma(1 - 2\varsigma_2)}. \quad (\text{E.8})$$

In addition, we have

$$\xi_0^\top (I - \gamma P_i^{\pi'_i})^{-1} \begin{pmatrix} \delta_1 \\ 0 \end{pmatrix} = \frac{\delta_1}{1 - \gamma} \cdot \frac{1 - \gamma p_i^{\pi_i^*}(\underline{x}|\underline{x})}{1 + \gamma \delta_1 - \gamma p_i^{\pi_i^*}(\underline{x}|\underline{x})}$$

Under the condition $\delta_1 \leq \frac{1-\gamma}{\gamma}$, we have $\gamma \delta_1 \leq 1 - \gamma \leq 1 - \gamma p_i^{\pi_i^*}(\underline{x}|\underline{x})$, therefore,

$$\xi_0^\top (I - \gamma P_i^{\pi'_i})^{-1} \begin{pmatrix} \delta_1 \\ 0 \end{pmatrix} \geq \frac{\delta_1}{2(1 - \gamma)}. \quad (\text{E.9})$$

Plugging (E.8) and (E.9) into (E.7), we finish our proof. $\square$

F. Supporting Lemmas

Lemma F.1. Let $Z_1, \dots, Z_n$ be random, positive-semidefinite adaptively chosen matrices with dimension d . Suppose $\lambda_{\max}(Z_i) \leq R^2$ almost surely for all i . Let $Z^+ = \sum_{i=1}^n Z_i$ and $W = \sum_{i=1}^n \mathbb{E}[Z_i | Z_1, \dots, Z_{i-1}]$. Then for any μ and any $\alpha \in (0, 1)$,

$$\mathbb{P}\left(\lambda_{\min}(Z^+) \leq (1 - \alpha)\mu \text{ and } \lambda_{\min}(W) \geq \mu\right) d \left(\frac{1}{e^\alpha(1 - \alpha)^{1-\alpha}}\right)^{\mu/R^2}$$

Lemma F.2 (Azuma-Hoeffding's inequality). Let $\mathcal{F}_n = \sigma(x_1, \dots, x_n)$ be a sequence of σ -fields known as a filtration. Let $\{(x_n, \mathcal{F}_n)\}_{n=1}^\infty$ be a martingale difference sequence for which there are constants $\{(a_k, b_k)_{k=1}^n\}$ such that $x_k \in [a_k, b_k]$ almost surely for $k = 1, \dots, n$. Then for all $t \geq 0$,

$$\mathbb{P}\left(\left|\sum_{k=1}^n x_k\right| \geq t\right) \leq 2 \exp\left(-\frac{2t^2}{\sum_{k=1}^n (b_k - a_k)^2}\right).$$