

THEORETICAL STUDY OF PROVABLY EFFICIENT OFF-LINE REINFORCEMENT LEARNING WITH TRAJECTORY-WISE REWARD

Anonymous authors

Paper under double-blind review

ABSTRACT

The remarkable success of reinforcement learning (RL) heavily relies on observing the reward of every visited state-action pair. In many real world applications, however, an agent can observe only a score that represents the quality of the whole trajectory, which is referred to as the *trajectory-wise reward*. In such a situation, it is difficult for standard RL methods to well utilize trajectory-wise reward, and large bias and variance errors can be incurred in policy evaluation. In this work, we propose a novel offline RL algorithm, called Pessimistic Value Iteration with Reward Decomposition (PARTED), which decomposes the trajectory return into per-step proxy rewards via least-squares-based reward redistribution, and then performs pessimistic value iteration based on the learned proxy reward. To ensure the value functions constructed by PARTED are always pessimistic with respect to the optimal ones, we design a new penalty term to offset the uncertainty of the proxy reward. For general episodic MDPs with large state space, we show that PARTED with overparameterized neural network function approximation achieves an $\tilde{O}(D_{\text{eff}}H^2/\sqrt{N})$ suboptimality, where H is the length of episode, N is the total number of samples, and D_{eff} is the effective dimension of the neural tangent kernel matrix. To further illustrate the result, we show that PARTED achieves an $\tilde{O}(dH^3/\sqrt{N})$ suboptimality with linear MDPs, where d is the feature dimension, which matches with that with neural network function approximation, when $D_{\text{eff}} = dH$. To the best of our knowledge, PARTED is the first offline RL algorithm that is provably efficient in general MDP with trajectory-wise reward.

1 INTRODUCTION

Reinforcement learning (RL) aims at searching for an optimal policy in an unknown environment [Sutton & Barto \(2018\)](#). To achieve this goal, an instantaneous reward is typically required at every step so that RL algorithms can maximize the cumulative reward of a Markov Decision Process (MDP). In recent years, RL has achieved remarkable empirical success with a high quality reward function [Mnih et al. \(2015\)](#); [Levine et al. \(2016\)](#); [Silver et al. \(2017\)](#); [Senior et al. \(2020\)](#); [Maei \(2011\)](#). However, in many real-world scenarios, instantaneous rewards are hard or impossible to be obtained. For example, in the autonomous driving task [Shalev-Shwartz et al. \(2016\)](#), it is very costly and time consuming to score every state-action pair that the agent (car) visits. In contrast, it is fairly easy to score the entire trajectory after the agent completing the task [Chatterji et al. \(2021\)](#). Therefore, in practice, it becomes more reasonable to adopt trajectory-wise reward schemes, in which only a return signal that represents the quality of the entire trajectory is revealed to the agent in the end. In recent years, trajectory-wise rewards have become prevalent in many real-world applications [Gong et al. \(2019\)](#); [Olivecrona et al. \(2017\)](#); [Lin et al. \(2018\)](#); [Hein et al. \(2017\)](#); [Rahmandad et al. \(2009\)](#).

Although trajectory-wise rewards are convenient to be obtained, it is often challenging for standard RL algorithms to utilize such a type of rewards well due to the high bias and variance it can introduce in the policy evaluation process [Arjona-Medina et al. \(2019\)](#), which leads to unsatisfactory policy optimization results. To address such an issue, [Chatterji et al. \(2021\)](#); [Pacchiano et al. \(2021\)](#) proposed to encode the whole trajectory and search for a non-Markovian trajectory-dependent optimal policy using the contextual bandit method. Although this type of approaches have promising theoretical

guarantees, they are difficult to be implemented in practice due to the difficulty of searching the large trajectory-dependent policy space whose dimension increases exponentially with the horizon length. Another type of approaches widely adopted in practice is called *reward redistribution*, which learns a reward function by allocating the trajectory-wise reward to every visited state-action pairs based on their contributions [Arjona-Medina et al. \(2019\)](#); [Liu et al. \(2019\)](#); [Gangwani et al. \(2020\)](#); [Ren et al. \(2021b\)](#); [Efroni et al. \(2021\)](#). Since the reward function in reward redistribution is typically learned via solving a supervised learning problem, such an approach is sample-efficient and can be integrated into the existing RL frameworks easily. However, most of existing reward redistribution approaches do not have theoretical performance guarantee. So far, only [Efroni et al. \(2021\)](#) proposes a provably efficient reward redistribution algorithm, but is only applicable to tabular episodic MDP and requires both reward and transition kernel to be horizon-independent.

Despite the superior performance of the reward redistribution method, all previous algorithms considered only the *online* setting, which are not applicable to many critical domains where offline sampling is preferred (or can be required), as interactive data collection could be very costly and risky [Shalev-Shwartz et al. \(2016\)](#); [Gottesman et al. \(2019\)](#). *How to design reward redistribution in offline RL for trajectory-wise rewards is an important but fully unexplored problem.* For such a problem, designing reward redistribution algorithms can be hard due to the insufficient sample coverage issue [Wang et al. \(2020a\)](#) in offline RL. Further challenges can be encountered when we try to design *provably* efficient reward distribution algorithms for general MDPs with large state space and horizon-dependent rewards and transition kernels, which has not been studied in online setting.

Thus, the goal of this work is to design an offline RL algorithm with reward redistribution for trajectory-wise rewards, which has provable efficiency guarantee for general episodic MDPs.

1.1 MAIN CONTRIBUTIONS

In this paper, we consider episodic MDP with possibly infinity state space and horizon-dependent reward function and transition kernel. The trajectory-wise reward adopts a standard sum-form as considered previously in [Han et al. \(2021\)](#); [Zheng et al. \(2018\)](#); [Klissarov & Precup \(2020\)](#); [Oh et al. \(2018\)](#); [Ren et al. \(2021b\)](#); [Efroni et al. \(2021\)](#), in which only the summation of rewards over the visited state-action pairs is revealed at the end of each episode.

We propose a novel Pessimistic vAlue iteRaTion with rEward Decomposition (PARTED) algorithm for offline RL with trajectory-wise rewards, which incorporates a least-square-based reward redistribution into the pessimistic value iteration (PEVI) algorithm [Jin et al. \(2021\)](#); [Yin et al. \(2021; 2020; 2022\)](#); [Yin & Wang \(2021b\)](#). Differently from the standard PEVI with instantaneous reward, in PARTED, reward need to be learned separately from the value function by training a regression model to decompose the trajectory return into per-step proxy rewards. In order to capture the reward and value function for a large state space, we adopt overparameterized neural networks for function approximation. Moreover, to offset the estimation error of proxy rewards, we design a penalty function by transferring the uncertainty from the covariance matrix of trajectory features to step-wise proxy rewards via an "one-block-hot" vector, which is new in the literature.

We show that our proposed new penalty term ensures that the value functions constructed by PARTED are always pessimistic with respect to the optimal ones. Furthermore, with overparameterized neural network function approximation, we show that PARTED achieves an $\tilde{O}(D_{\text{eff}}H^2/\sqrt{N})$ suboptimality, where H is the length of episode, N is the total number of samples, and D_{eff} is the effective dimension of neural tangent kernel matrix. To further illustrate our result, we show that PARTED achieves an $\tilde{O}(dH^3/\sqrt{N})$ suboptimality in the linear MDP setting, where d is the feature dimension, which matches that in the neural network function approximation setting when $D_{\text{eff}} = dH$. To the best of our knowledge, PARTED is the first-known offline RL algorithm that is provably efficient in general episodic MDPs with trajectory-wise rewards.

1.2 RELATED WORKS

Trajectory-Wise Reward RL. Policy optimization with trajectory-rewards is extremely difficult. A variety of practical strategies have been proposed to resolve this technical challenge by redistributing trajectory rewards to step-wise rewards. RUDDER [Arjona-Medina et al. \(2019\)](#) trains a return

predictor of state-action sequence with LSTM Hochreiter & Schmidhuber (1997), and the reward at each horizon is then assigned by the difference between the predications of two adjacent sub-trajectories. Later, Liu et al. (2019) improves RUDDER and utilizes a Transformer Vaswani et al. (2017) for better reward learning. IRCR Gangwani et al. (2020) assigns the proxy reward of a state-action pair as the normalized value of trajectory returns that contain the correspondingly state-action pair. RRD Ren et al. (2021b) learns a proxy reward function by solving a supervised learning problem together with a Monte-Carlo sampling strategy. Although those methods have achieved great empirical success, they all lack overall theoretical performance guarantee.

Differently from empirical studies, existing theoretical works of trajectory-wise reward RL are rare and focus only on the online setting. One line of research assumes trajectory reward being non-Markovian, and thus focuses on searching for a non-Markovian, trajectory-dependent optimal policy. Chatterji et al. (2021) assumes that trajectory-wise reward is a binary signal generated by a logistic classifier with trajectory embedding as the input. In this setting, the policy optimization problem is reduced to a linear contextual bandit problem in which the trajectory embedding is the contextual vector. Pacchiano et al. (2021) considers a similar setting as Chatterji et al. (2021) but assumes only having access to a binary preference score between two trajectories instead of an absolute reward of a trajectory. Another line of research assumes that the trajectory-wise reward is the summation of underlying step-wise Markovian rewards. The goal of this line of work is to search for an optimal Markovian policy. Cohen et al. (2021) adopted a mirror descent approach so that the summation of rewards alone is sufficient to perform the policy optimization. This approach relies on the on-policy unbiased sampling of trajectory rewards, and can hardly be extended to the offline setting. Efroni et al. (2021) proposed to recover the reward by solving a least-squared regression problem that fits the summation of reward estimation toward the trajectory reward.

To our best knowledge, offline RL with trajectory-wise rewards (where no interaction with the environment is allowed) has not been studied before, and our work develops the first-known algorithm for such a setting with provable sample efficiency guarantee. Further, although our reward redistribution approach applies the least-square based method, which has also been adopted in Efroni et al. (2021), our algorithm is designed for general MDPs with possibly infinite state and horizon-dependent rewards and transition kernels, which is very different from that in Efroni et al. (2021) designed for tabular MDPs with time-independent rewards and transition kernels.

Offline RL. The major challenge in offline RL is the insufficient sample coverage in the pre-collected dataset, which arises from the lack of exploration Wang et al. (2020a); Liu et al. (2020). To address such a challenge, two types of algorithms have been studied: (1) regularized approaches, which prevent the policy from visiting states and actions that are less covered by the dataset Dadashi et al. (2021); Fujimoto et al. (2019b;a); Wang et al. (2020b); Fujimoto & Gu (2021); (2) pessimistic approach, which penalize the estimated values of the less-covered state-action pairs Buckman et al. (2020); Kumar et al. (2020). So far, a number of provably efficient pessimistic offline RL algorithm have been proposed in both tabular MDP setting Yin et al. (2020); Shi et al. (2022); Yan et al. (2022); Li et al. (2022); Yin & Wang (2021a); Ren et al. (2021a); Xie et al. (2021b); Yin & Wang (2021b); Rashidinejad et al. (2021) and linear MDP setting Jin et al. (2021); Xie et al. (2021a); Zanette et al. (2021); Wang et al. (2020a); Zanette (2021); Foster et al. (2021); Yin et al. (2022). However, the efficiency of all those works relies on both the availability of instantaneous reward and special structures of MDP, which can hardly be satisfied in practical settings. In this work, we take a first step towards relaxing those two assumptions by proposing PARTED, which is provably efficient in general episodic MDPs with trajectory-wise rewards.

2 PRELIMINARY AND PROBLEM FORMULATION

2.1 EPISODIC MARKOV DECISION PROCESS

An episodic Markov decision process (MDP) is defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$, where \mathcal{S} and \mathcal{A} are the state and action spaces, $H > 0$ is the length of each episode, and $\mathbb{P} = \{\mathbb{P}_h\}_{h \in [H]}$ and $r = \{r_h\}_{h \in [H]}$ are the transition kernel and reward, respectively, where $[n] = \{1, 2, \dots, n\}$ for integer $n \geq 1$. We assume \mathcal{S} is a measurable space of possibly infinite cardinality and \mathcal{A} is a finite set. For each $h \in [H]$, $\mathbb{P}_h(\cdot | s, a)$ denotes the transition probability when action a is taken at state s at timestep h , and $r_h(s, a) \in [0, 1]$ is a **random** reward that is observed with state-action pair (s, a) at

timestep h . We denote the mean of the reward as $R_h(s, a) = \mathbb{E}[r_h(s, a)|s, a]$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. For any policy $\pi = \{\pi_h\}_{h \in [H]}$, we define the state value function $V_h^\pi(\cdot) : \mathcal{S} \rightarrow \mathbb{R}$ and state-action value function $Q_h^\pi(\cdot) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ at each timestep h as

$$V_h^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=h}^H r_t(s_t, a_t) \middle| s_h = s \right], \quad Q_h^\pi(s, a) = \mathbb{E}_\pi \left[\sum_{t=h}^H r_t(s_t, a_t) \middle| (s_h, a_h) = (s, a) \right],$$

where the expectation \mathbb{E}_π is taken with respect to the randomness of the trajectory induced by policy π , which is obtained by taking action $a_t \sim \pi_t(\cdot|s_t)$ and transiting to the next state $s_{t+1} \sim \mathbb{P}_t(\cdot|s_t, a_t)$ at timestep $t \in [H]$. At each timestep $h \in [H]$, for any $f : \mathcal{S} \rightarrow \mathbb{R}$, we define the transition operator as $(\mathbb{P}_h f)(s, a) = \mathbb{E}[f(s_{h+1})|(s_h, a_h) = (s, a)]$ and the Bellman operator as $(\mathbb{B}_h f)(s, a) = R_h(s, a) + (\mathbb{P}_h f)(s, a)$. For episodic MDP $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$, we have

$$Q_h^\pi(s, a) = (\mathbb{B}_h V_{h+1}^\pi)(s, a), \quad V_h^\pi(s) = \langle Q_h^\pi(s, \cdot), \pi_h(\cdot|s) \rangle_{\mathcal{A}}, \quad V_{H+1}^\pi(s) = 0,$$

where $\langle \cdot, \cdot \rangle_{\mathcal{A}}$ denotes the inner product over \mathcal{A} . We define the optimal policy π^* as the policy that yields the optimal value function, i.e., $V_h^{\pi^*}(s) = \sup_{\pi} V_h^\pi(s)$ for all $s \in \mathcal{S}$ and $h \in [H]$. For simplicity, we denote $V_h^{\pi^*}$ and $Q_h^{\pi^*}$ as V_h^* and Q_h^* , respectively. The Bellman optimality equation is given as follows

$$Q_h^*(s, a) = (\mathbb{B}_h V_h^*)(s, a), \quad V_h^*(s) = \operatorname{argmax}_{a \in \mathcal{A}} Q_h^*(s, a), \quad V_{H+1}^*(s) = 0, \quad (1)$$

The goal of reinforcement learning is to learn the optimal policy π^* . For any fixed π , we define the performance metric as

$$\text{SubOpt}(\pi, s) = V_1^*(s) - V_1^\pi(s),$$

which is the suboptimality of the policy π given the initial state $s_1 = s$.

2.2 TRAJECTORY-WISE REWARD AND OFFLINE RL

In the trajectory-wise reward setting, the transition of the environment is still Markovian and the agent can still observe and interact with the environment instantly as in standard MDPs. However, unlike standard MDPs in which the agent can receive an instantaneous reward $r_h(s, a)$ for every visited state-action pair x at each timestep h , in the trajectory-wise reward setting, only a reward that is associated with the whole trajectory can be observed at the end of the episode, i.e., $r(\tau)$ where $\tau = \{(s_1^T, a_1^T), \dots, (s_H^T, a_H^T)\}$ denotes a trajectory and (s_h^T, a_h^T) is the h -th state-action pair in trajectory τ , which is called "trajectory reward" in the sequel. In this work, we consider the setting in which the trajectory reward is the summation of the underlying instantaneous reward in the trajectory of MDP $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$, i.e., $r(\tau) = \sum_{h=1}^H r_h(s_h^T, a_h^T)$. We denote the mean of the trajectory reward as $R(\tau) = \mathbb{E}[r(\tau)|\tau] = \sum_{h=1}^H R_h(s_h^T, a_h^T)$. Such a sum-form reward has been commonly considered in both theoretical [Efroni et al. \(2021\)](#) and empirical studies [Han et al. \(2021\)](#); [Zheng et al. \(2018\)](#); [Klissarov & Precup \(2020\)](#); [Oh et al. \(2018\)](#); [Ren et al. \(2021b\)](#). It models the situations where the agent's goal is captured by a certain metric with additive properties, e.g., the energy cost of a car during driving, the click rate of advertisements during a time interval, or the distance of a robot's running. Such a form of reward can be more general than the standard RL feedback and is expected to be more common in practical scenarios. Note that RL problems under trajectory-wise rewards is very challenging, as traditional policy optimization approach typically fails due the obscured feedback received from the environment, which causes large value function evaluation error [Han et al. \(2021\)](#).

We consider the offline RL setting, in which a learner has access only to a pre-collected dataset \mathcal{D} consisting of N trajectories $\{\tau_i, r(\tau_i)\}_{i,h=1}^{N,H}$ rolled out from some possibly unknown behavior policy μ , where τ_i and $r(\tau_i)$ are the i -th trajectory and the observed trajectory reward of τ_i , respectively. Given this batch data \mathcal{D} with only trajectory-wise rewards and a target accuracy ϵ , our goal is to find a policy π such that $\text{SubOpt}(\pi, s) \leq \epsilon$ for all $s \in \mathcal{S}$.

2.3 OVERPARAMETERIZED NEURAL NETWORK

In this paper, we consider the function approximation setting, in which the state-action value function is approximated by a two-layer neural network. To simplify the notation, we denote $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ and

view it as a subset of \mathbb{R}^d . We further assign a feature vector $x \in \mathcal{X}$ to represent a state-action pair (s, a) . Without loss of generality, we assume that $\|x\|_2 = 1$ for all $x \in \mathcal{X}$. We also allow $x = 0$ to represent a null state-action pair. We now define a two-layer neural network $f(\cdot, b, w) : \mathcal{X} \rightarrow \mathbb{R}$ with $2m$ neurons and weights (b, w) as

$$f(x; b, w) = \frac{1}{\sqrt{2m}} \sum_{r=1}^{2m} b_j \cdot \sigma(w_r^\top x), \quad \forall x \in \mathcal{X}, \quad (2)$$

where $\sigma(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function, $b_r \in \mathbb{R}$ and $w_r \in \mathbb{R}^d$ for all $r \in [2m]$, and $b = (b_1, \dots, b_{2m})^\top \in \mathbb{R}^{2m}$ and $w = (w_1^\top, \dots, w_{2m}^\top)^\top \in \mathbb{R}^{2md}$. We make the following assumption for $\sigma(\cdot)$, which can be satisfied by a number of activation functions such as ReLU and tanh(\cdot).

Assumption 1. For all $x \in \mathcal{X}$, we have $|\sigma'(x)| \leq C_\sigma < +\infty$ and $\sigma'(0) = 0$.

We initialize b and w via a symmetric initialization scheme Gao et al. (2019); Bai & Lee (2019): for any $1 \leq r \leq m$ we set $b_{0,r} \sim \text{Unif}(\{-1, 1\})$ and $w_{0,r} \sim N(0, I_d/d)$, where I_d is the identity matrix in \mathbb{R}^d , and for any $m+1 \leq r \leq 2m$, we set $b_{0,r} = -b_{0,r-m}$ and $w_{0,r} = w_{0,r-m}$. Under such an initialization, the initial neural network is a zero function, i.e. $f(x; b_0, w_0) = 0$ for all $x \in \mathcal{X}$, where $b_0 = [b_{0,1}, \dots, b_{0,2m}]^\top$ and $w_0 = [w_{0,1}^\top, \dots, w_{0,2m}^\top]^\top$ are initialization parameters. During training, we fix the value of b at its initial value and only optimize w . To simplify the notation, we denote $f(x; b, w)$ as $f(x; w)$ and $\nabla_w f(x, w)$ as $\phi(x, w)$.

Notations. We use $\tilde{\mathcal{O}}(X)$ to refer to a quantity that is upper bounded by X , up to poly-log factors of d, H, N, m and $(1/\delta)$. Furthermore, we use $\mathcal{O}(X)$ to refer to a quantity that is upper bounded by X up to constant multiplicative factors. We use I_d as the identity matrix in dimension d . Similarly, we denote by $\mathbf{0}_d \in \mathbb{R}^d$ as the vector whose components are zeros. For any square matrix M , we let $\|M\|_2$ denote the operator norm of M . Finally, for any positive definite matrix $M \in \mathbb{R}^{d \times d}$ and any vector $x \in \mathbb{R}^d$, we define $\|x\|_M = \sqrt{x^\top M x}$.

3 ALGORITHM

In this section, we propose a Pessimistic vAlue iteRaTion with rEward Decomposition (PARTED) algorithm based on the neural network function approximation. PARTED shares a similar structure as that of pessimistic value iteration (PEVI) Jin et al. (2021); Xie et al. (2021b); Yin et al. (2020), but has a very different design due to trajectory-wise rewards. In PEVI, a pessimistic estimator of the value function is constructed from the dataset \mathcal{D} and the Bellman optimality equation is then iterated based such an estimator. Since instantaneous rewards are available in PEVI, given a function class \mathcal{G} , PEVI constructs an estimated Bellman backup of value function $(\hat{\mathbb{B}}_h \hat{V}_{h+1})$ by solving the following regression problem for all $h \in [H]$ in the backward direction:

$$(\hat{\mathbb{B}}_h \hat{V}_{h+1}) = \underset{g_h \in \mathcal{G}}{\text{argmin}} L_{\text{PEVI}}^h(g_h) = \sum_{\tau \in \mathcal{D}} \left(r_h(x_h^\tau) + \hat{V}_{h+1}(s_{h+1}^\tau) - g_h(x_h^\tau) \right)^2 + \lambda \cdot \text{Reg}(g_h). \quad (3)$$

In eq. (3), $\hat{V}_{h+1}(\cdot)$ is the pessimistic estimator of optimal value function constructed for horizon $h+1$, $\lambda > 0$ is a regularization parameter and $\text{Reg}(\cdot)$ is the regularization function. The optimal state-action value function can then be estimated as $\hat{Q}_h(\cdot) = \min\{(\hat{\mathbb{B}}_h \hat{V}_{h+1})(\cdot) - \Gamma_h(\cdot), H\}^+$, where $-\Gamma_h$ is a negative penalty used to offset the uncertainty in $(\hat{\mathbb{B}}_h \hat{V}_{h+1})(\cdot)$ and guarantee the pessimism of \hat{Q}_h .

However, in PARTED (see Algorithm 1) designed for trajectory-wise rewards, since instantaneous reward $r_h(\cdot)$ is not available, we can no longer obtain $(\hat{\mathbb{B}}_h \hat{V}_{h+1})$ in the same way as PEVI by solving the regression problem in eq. (3). To overcome such an issue, in PARTED, we construct two estimators \hat{r}_h and $(\hat{\mathbb{P}}_h \hat{V}_{h+1})$ for instantaneous reward r_h and transition value function $(\mathbb{P}_h \hat{V}_{h+1})$, respectively. The estimated Bellman backup can then be formulated as $(\hat{\mathbb{B}}_h \hat{V}_{h+1})(\cdot) = \hat{r}_h(\cdot) + (\hat{\mathbb{P}}_h \hat{V}_{h+1})(\cdot)$.

Reward Redistribution. In order to estimate the instantaneous rewards from the trajectory-wise reward, we use a neural network $f(\cdot, \theta_h)$ given in eq. (2) to represent per-step mean reward $R_h(\cdot)$ for all $h \in [H]$, where $\theta_h \in \mathbb{R}^{2md}$ is the parameter. We further assume, for simplicity, that all the neural networks share the same initial weights denoted by $\theta_0 \in \mathbb{R}^{2md}$. We define the following loss function $L_r(\cdot) : \mathbb{R}^{2mdH} \rightarrow \mathbb{R}$ for reward redistribution as

$$L_r(\Theta) = \sum_{\tau \in \mathcal{D}} \left[\sum_{h=1}^H f(x_h^\tau, \theta_h) - r(\tau) \right]^2 + \lambda_1 \cdot \sum_{h=1}^H \|\theta_h - \theta_0\|_2^2, \quad (4)$$

Algorithm 1 Neural Pessimistic Value Iteration with Reward Decomposition (PARTED)

Input: Dataset $\mathcal{D} = \{\tau_i, r(\tau_i)\}_{i,h=1}^{N,H}$
Initialization: Set \widehat{V}_{H+1} as zero function
Obtain \widehat{R}_h and $\widehat{\Theta}$ according to eq. (5)
for $h = H, H-1, \dots, 1$ **do**
 Obtain $\widehat{\mathbb{P}}_h \widehat{V}_{h+1}$ and \widehat{w}_h according to eq. (7)
 Obtain $\Gamma_h(\cdot, \widehat{\Theta}, \widehat{w}_h)$ according to eq. (10)
 Obtain $\widehat{Q}_h(\cdot)$ and $\widehat{V}_h(\cdot)$ according to eq. (11) and let $\widehat{\pi}_h(\cdot|s) = \operatorname{argmax}_{\pi_h} \langle \widehat{Q}_h(s, \cdot), \pi_h(\cdot|s) \rangle$
end for

where $\Theta = [\theta_1^\top, \dots, \theta_H^\top]^\top \in \mathbb{R}^{2mdH}$ and $\lambda_1 > 0$ is a regularization parameter. Then, the per-step proxy reward $\widehat{R}_h(\cdot)$ is obtained by solving the following optimization problem

$$\widehat{R}_h(\cdot) = f(\cdot, \widehat{\theta}_h), \quad \text{where } \widehat{\Theta} = \operatorname{argmin}_{\Theta \in \mathbb{R}^{2mdH}} L_r(\Theta) \text{ and } \widehat{\Theta} = [\widehat{\theta}_1^\top, \dots, \widehat{\theta}_H^\top]^\top. \quad (5)$$

Transition Value Function Estimation. Similarly, we use H neural networks given in eq. (2) with parameter $\{w_h\}_{h \in [H]}$ to estimate $\{(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)\}_{h \in [H]}$, where $w_h \in \mathbb{R}^{2md}$ is the parameter of the h -th network. Specifically, for each $h \in [H]$, we define the loss function $L_v^h(w_h): \mathbb{R}^{2md} \rightarrow \mathbb{R}$ as

$$L_v^h(w_h) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - f(x_h^\tau, w_h) \right)^2 + \lambda_2 \cdot \|w_h - w_0\|_2^2, \quad (6)$$

where $\lambda_2 > 0$ is a regularization parameter and w_0 is the initialization shared by all neural networks. The estimated transition value function $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ can be obtained by solving the following optimization problem

$$(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = f(\cdot, \widehat{w}_h), \quad \text{where } \widehat{w}_h = \operatorname{argmin}_{w_h \in \mathbb{R}^{2md}} L_v^h(w_h). \quad (7)$$

Penalty Term Construction. It remains to construct the penalty term Γ_h to offset the uncertainties in \widehat{R}_h and $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})$. First consider the penalty of $\widehat{R}_h(\cdot)$ for each $h \in [H]$. For any $\tau \in \mathcal{D}$ and $\Theta \in \mathbb{R}^{2mdH}$, we define a trajectory feature $\Phi(\tau, \Theta) = [\phi(x_1^\tau, \theta_1)^\top, \dots, \phi(x_H^\tau, \theta_H)^\top]^\top$. Based on $\Phi(\tau, \Theta)$, the trajectory feature covariance matrix $\Sigma(\Theta) \in \mathbb{R}^{2mdH \times 2mdH}$ is then defined as

$$\Sigma(\Theta) = \lambda_1 \cdot I_{2mdH} + \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta) \Phi(\tau, \Theta)^\top.$$

We also define an "one-block-hot" vector $\Phi_h(x, \Theta) = [\mathbf{0}_{2md}^\top, \dots, \phi(x, \theta_h)^\top, \dots, \mathbf{0}_{2md}^\top]^\top$ for all $x \in \mathcal{X}$, where $\Phi_h(x, \Theta) \in \mathbb{R}^{2mdH}$ is a vector in which $[\Phi_h(x, \Theta)]_{2md(h-1)+1:2mdh} = \phi(x, \theta_h)$ and the rest entries are zero. The penalty term of reward for a given $\Theta \in \mathbb{R}^{2mdH}$ is defined as:

$$b_{r,h}(x, \Theta) = [\Phi_h(x, \Theta)^\top \Sigma^{-1}(\Theta) \Phi_h(x, \Theta)]^{1/2}, \quad \forall x \in \mathcal{X}. \quad (8)$$

Note that the reward penalty term $b_{r,h}(x, \Theta)$ is new and first proposed in this work. By constructing $b_{r,h}(x, \Theta)$ in this way, we can capture the effect of uncertainty caused by solving the trajectory-wise regression problem in eq. (4), which is contained in the covariance matrix $\Sigma(\Theta)$, on the proxy reward $f(\cdot, \widehat{\theta}_h)$ at each step $h \in [H]$, via the "one-block-hot" vector $\Phi_h(\cdot, \Theta)$.

Next, we consider the penalty of $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ for each $h \in [H]$. We define the per-step feature covariance matrix $\Lambda_h(w_h) \in \mathbb{R}^{2md \times 2md}$ as

$$\Lambda_h(w) = \lambda_2 \cdot I_{2md} + \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w) \phi(x_h^\tau, w)^\top.$$

Then, the penalty term of $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ for a given $w \in \mathbb{R}^{2md}$ is defined as

$$b_{v,h}(x, w) = [\phi(x, w)^\top \Lambda_h(w)^{-1} \phi(x, w)^\top]^{1/2}, \quad \forall x \in \mathcal{X}. \quad (9)$$

Finally, combining eqs. (8) and (9), the penalty term for $\widehat{\mathbb{B}}_h \widehat{V}_{h+1}(\cdot)$ is constructed as

$$\Gamma_h(x, \Theta, w) = \beta_1 b_{r,h}(x, \Theta) + \beta_2 b_{v,h}(x, w), \quad (10)$$

where $\beta_1, \beta_2 > 0$ are parameters. The estimator of $Q_h(\cdot)$ and $V_h(\cdot)$ can then be obtained as

$$\widehat{Q}_h(\cdot) = \min\{\widehat{R}_h(\cdot) + (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) - \Gamma_h(\cdot, \widehat{\Theta}, \widehat{w}_h), H\}^+, \quad \widehat{V}_h(\cdot) = \operatorname{argmax}_{a \in \mathcal{A}} \widehat{Q}_h(\cdot, a). \quad (11)$$

Furthermore, for any $h \in [H]$, we denote $\mathcal{V}_h(x, R_{\beta_1}, R_{\beta_2}, \lambda_1, \lambda_2)$ as the class of functions that takes the form $\bar{V}_h(\cdot) = \max_{a \in \mathcal{A}} \bar{Q}_h(\cdot, a)$, where

$$\begin{aligned} \bar{Q}_h(x) = & \min\{\langle \phi(x, \theta_0), \theta - \theta_0 \rangle + \langle \phi(x, w_0), w - w_0 \rangle \\ & - \beta_1 \cdot \sqrt{\Phi_h(x, \theta_0)^\top \Sigma^{-1} \Phi_h(x, \theta_0)} - \beta_2 \cdot \sqrt{\phi(x, w_0)^\top \Lambda^{-1} \phi(x, w_0)}, H\}^+, \end{aligned}$$

in which $\|\theta - \theta_0\|_2 \leq H \sqrt{N/\lambda_1}$, $\|w - w_0\|_2 \leq H \sqrt{N/\lambda_2}$, $\beta_1 \in [0, R_{\beta_1}]$, $\beta_2 \in [0, R_{\beta_2}]$, $\|\Sigma\|_2 \geq \lambda_1$ and $\|\Lambda\|_2 \geq \lambda_2$. To this end, for any $\epsilon > 0$, we define $\mathcal{N}_{\epsilon, h}^v$ as the ϵ -covering number of $\mathcal{V}_h(x, R_{\beta_1}, R_{\beta_2}, \lambda_1, \lambda_2)$ with respect to the ℓ_∞ -norm on \mathcal{X} , and we let $\mathcal{N}_\epsilon^v = \max_{h \in [H]} \{\mathcal{N}_{\epsilon, h}^v\}$.

4 MAIN RESULTS

4.1 SUBOPTIMALITY OF PARTED FOR GENERAL MDPs

In the overparameterized scheme, the neural network width $2m$ is considered to be much larger than the number of trajectories N and horizon length H . Under such a scheme, the training process of neural networks can be captured by the framework of neural tangent kernel (NTK) [Jacot et al. \(2018\)](#). Specifically, conditioning on the realization of w_0 , we define a kernel $K(x, x') : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ as

$$K(x, x') = \langle \phi(x, w_0), \phi(x', w_0) \rangle = \frac{1}{2m} \sum_{r=1}^{2m} \sigma'(w_{0,r}^\top x) \sigma'(w_{0,r}^\top x') x^\top x', \quad \forall (x, x') \in \mathcal{X} \times \mathcal{X},$$

where $\sigma'(\cdot)$ is the derivative of the action function $\sigma(\cdot)$. It can be shown that $f(\cdot, w)$ is close to its linearization at w_0 when m is sufficiently large and w is not too far away from w_0 , i.e.,

$$f(x, w) \approx f_0(x, w) = f(x, w_0) + \langle \phi(x, w_0), w - w_0 \rangle = \langle \phi(x, w_0), w - w_0 \rangle, \quad \forall x \in \mathcal{X}.$$

Note that $f_0(x, w)$ belongs to a reproducing kernel Hilbert space (RKHS) with kernel $K(\cdot, \cdot)$. Similarly, consider the sum of H neural networks $f(\tau, \Theta) = \sum_{h=1}^H f(x_h^\tau, \theta_h)$ with the same initialization θ_0 for each neural network, where $\tau = [x_1^\top, \dots, x_H^\top]^\top$ and $\Theta = [\theta_1^\top, \dots, \theta_H^\top]^\top$. If θ_h is not too far away from θ_0 for all $h \in [H]$ and m is sufficiently large, it can be shown that the dynamics of $f(\tau, \Theta)$ belong to a RKHS with kernel K_H defined as $K_H(\tau, \tau') = \sum_{h=1}^H K(x_h^\tau, x_h^{\tau'})$. We further define \mathcal{H}_K and \mathcal{H}_{K_H} as the RKHS induced by $K(\cdot, \cdot)$ and $K_H(\cdot, \cdot)$, respectively.

Based on the kernel $K(\cdot, \cdot)$ and $K_H(\cdot, \cdot)$, we define the Gram matrix $K_r, K_{v, h} \in \mathbb{R}^{N \times N}$ as

$$K_r = [K_H(\tau_i, \tau_j)]_{i, j \in [N]}, \quad \text{and} \quad K_{v, h} = [K(x_h^{\tau_i}, x_h^{\tau_j})]_{i, j \in [N]}.$$

We further define a function class as follows

$$\mathcal{F}_{B_1, B_2} = \left\{ f_\ell(x) = \int_{\mathbb{R}^d} \sigma'(w^\top x) \cdot x^\top \ell(w) dp(w) : \sup_w \|\ell(w)\|_2 \leq B_1, \sup_w \frac{\|\ell(w)\|_2}{p(w)} \leq B_2 \right\},$$

where $\ell : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a mapping, B_1, B_2 are positive constants, and p is the density of $N(0, I_d/d)$. We then make the following assumption regarding the expressive power of the above function class.

Assumption 2. *We assume that there exist $a_1, a_2, A_1, A_2 > 0$ such that $R_h(\cdot) \in \mathcal{F}_{a_1, a_2}$ and $(\widehat{\mathbb{P}}_h f)(\cdot) \in \mathcal{F}_{A_1, A_2}$ for any $f(\cdot) : \mathcal{X} \rightarrow [0, H]$.*

Assumption 2 ensures that both $R_h(\cdot)$ and $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ can be captured by an infinite width neural network. Note that Assumption 2 is mild since \mathcal{F}_{B_1, B_2} is an expressive function class as shown in Lemma C.1 of [Gao et al. \(2019\)](#). Similar assumptions have also been adopted in many previous works that consider neural network function approximation [Yang et al. \(2020\)](#); [Cai et al. \(2019\)](#); [Wang et al. \(2019\)](#); [Xu et al. \(2021\)](#); [Qiu et al. \(2021\)](#). Additionally, we assume that the data collection process explores the state-action space and trajectory space well. Note that similar assumptions have also been adopted in [Duan et al. \(2020\)](#); [Yin et al. \(2022\)](#); [Jin et al. \(2021\)](#).

Assumption 3 (Well-Explored Dataset). *Suppose the N trajectories in dataset \mathcal{D} are independently and identically induced by a fixed behaviour policy μ . There exist absolute constants $C_\sigma > 0$ and $C_\zeta > 0$ such that*

$$\lambda_{\min}(\overline{M}(\Theta_0)) \geq C_\sigma \quad \text{and} \quad \lambda_{\min}(\overline{m}_h(w_0)) \geq C_\zeta \quad \forall h \in [H],$$

where

$$\overline{M}(\Theta_0) = \mathbb{E}_\mu [\Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top] \quad \text{and} \quad \overline{m}_h(w_0) = \mathbb{E}_\mu [\phi(x_h^\tau, w_0) \phi(x_h^\tau, w_0)^\top].$$

We can now present the suboptimality of the policy $\hat{\pi}$ obtained via Algorithm 1.

Theorem 1. *Consider Algorithm 1. Suppose Assumption 1-3 hold. Let $\lambda_1 = \lambda_2 = 1 + 1/N$, $\beta_1 = R_{\beta_1}$ and $\beta_2 = R_{\beta_2}$, in which R_{β_1} and R_{β_2} satisfy*

$$R_{\beta_1} \geq H \left(4a_2^2 \lambda_1 / d + 2 \log \det(I + K_r / \lambda_1) + 10 \log(NH^2) \right)^{1/2},$$

$$R_{\beta_2} \geq H \left(8A_2^2 \lambda_2 / d + 4 \max_{h \in [H]} \{ \log \det(I + K_{v,h} / \lambda_2) \} + 6C_\epsilon + 16 \log(NH^2 N_\epsilon^v) \right)^{1/2},$$

where $\epsilon = \sqrt{\lambda_2 C_\epsilon H} / (2NC_\phi)$, $C_\epsilon \geq 1$ is an adjustable parameter, and $C_\phi > 0$ is an absolute constant. In addition, let m be sufficiently large. Then, with probability at least $1 - (N^2 H^4)^{-1}$, we have

$$\text{SubOpt}(\hat{\pi}, s) \leq \tilde{\mathcal{O}} \left(\frac{H \max\{\beta_1, \beta_2\}}{\sqrt{N}} \right) + \varepsilon_1,$$

where

$$\varepsilon_1 = \max\{\beta_1 H^{5/3}, \beta_2 H^{7/6}\} \tilde{\mathcal{O}} \left(\frac{N^{1/12}}{m^{1/12}} \right) + \tilde{\mathcal{O}} \left(\frac{H^{17/6} N^{5/3}}{m^{1/6}} \right).$$

Theorem 1 shows that Algorithm 1 can find an ϵ -optimal policy with $\tilde{\mathcal{O}}(H^2 \max\{\beta_1, \beta_2\}^2 / \epsilon^2)$ episodes of offline data in the trajectory-wise reward setting up to a function approximation error ε_1 , which vanishes as the neural network width $2m$ increases. Note that the dependence of ε_1 on the network width, which is $\mathcal{O}(m^{-1/12})$, matches that of the approximation error in the previous work of value iteration algorithm with neural network function approximation Yang et al. (2020).

Discussion of Proof of Theorem 1. Comparing to the analysis of PEVI for linear MDP with instantaneous reward, which has been extensively studied in offline RL Jin et al. (2021); Yin et al. (2022; 2021), our analysis needs to address the following two new challenges: (1) In instantaneous reward setting, both $R_h(\cdot)$ and $(\mathbb{P}_h \hat{V}_{h+1})(\cdot)$ can be learned together by solving a single regression problem in per-step scale. However, in our trajectory-wise reward setting, $R_h(\cdot)$ and $(\mathbb{P}_h \hat{V}_{h+1})(\cdot)$ need to be learned separately by solving two regression problems (eqs. (5) and (7)) in different scales, i.e., eq. (5) is in trajectory scale and eq. (7) is in per-step scale. In order to apply union concentrations to bound the Bellman estimation error $|(\mathbb{B}_h \hat{V}_h)(\cdot) - (\mathbb{B}_h \hat{V}_h)(\cdot)|$, we need to develop new techniques to handle the mismatch between eqs. (5) and (7) in terms of scale. (2) In linear MDP, both $R_h(\cdot)$ and $(\mathbb{P}_h \hat{V}_{h+1})(\cdot)$ can be captured exactly by linear functions. However, in the more general MDP that we consider, we need to develop new analysis to bound the estimation error that caused by the insufficient expressive power of neural networks in order to characterize the optimality of $\hat{\theta}_h$ and \hat{w}_h in eqs. (5) and (7), respectively.

To obtain a more concrete suboptimality bound for Algorithm 1, we impose an assumption on the spectral structure of kernels K_H and K .

Assumption 4 (Finite Spectrum NTK Yang et al. (2020)). *Conditioned on the randomness of (b_0, w_0) , let T_{K_H} and T_K be the integral operator induced by K_H and K (see Appendix H for definition of T_{K_H} and T_K), respectively, and let $\{\omega_j\}_{j \geq 1}$ and $\{v_j\}_{j \geq 1}$ be eigenvalues of T_{K_H} and T_K , respectively. We have $\omega_j = 0$ for all $j \geq D_1 + 1$ and $v_j = 0$ for all $v_j \geq D_2 + 1$, where D_1, D_2 are positive integers.*

Assumption 4 implies that \mathcal{H}_{K_H} and \mathcal{H}_K are D_1 -dimensional and D_2 -dimensional, respectively. For concrete examples of neural networks that satisfy Assumption 4, please refer to Section B.3 in Yang et al. (2020). Note that such an assumption is in parallel to the "effective dimension" assumption in Zhou et al. (2020); Valko et al. (2013).

Corollary 1. *Consider Algorithm 1. Suppose Assumption 1-4 hold. Let $\lambda_1 = \lambda_2 = 1 + 1/N$, $\beta_1 = \tilde{\mathcal{O}}(HD_1)$ and $\beta_2 = \tilde{\mathcal{O}}(H \max\{D_1, D_2\})$. Then, with probability at least $1 - (N^2 H^4)^{-1}$, we have*

$$\text{SubOpt}(\hat{\pi}, s) = \tilde{\mathcal{O}} \left(D_{\text{eff}} H^2 / \sqrt{N} \right) + \varepsilon_2,$$

where $D_{\text{eff}} = \max\{D_1, D_2\}$ denotes the effective dimension and

$$\varepsilon_2 = \max \left\{ \sqrt{H}, \max\{D_1, D_2\}, \frac{H^{5/3} N^{19/12}}{m^{1/12}} \right\} \tilde{\mathcal{O}} \left(\frac{H^{13/6} N^{1/12}}{m^{1/12}} \right).$$

Corollary 1 states that when β_1 and β_2 are chosen properly according to the dimension of \mathcal{H}_{K_H} and \mathcal{H}_K , the suboptimality of the policy $\hat{\pi}$ incurred by Algorithm 1 converges to an ϵ -optimal policy with $\tilde{\mathcal{O}}(D_{\text{eff}}^2 H^4 / \epsilon^2)$ episodes of offline data up to a function approximation error ε_2 .

4.2 SUBOPTIMALITY OF PARTED UNDER LINEAR MDPs

In this section, we briefly illustrate our result by instantiating PARTED to simpler linear MDPs with trajectory-wise rewards. We further provide a detailed treatment of the linear MDP setting with trajectory-wise rewards in Appendix C.

With an abuse of notation, we define the linear MDP as follows.

Definition 1 (Linear MDP Jin et al. (2020); Yang & Wang (2019)). *We say an episodic MDP $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$ is a linear MDP with a known feature map $\phi(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^d$ if there exist an unknown vector $w_h^*(s) \in \mathbb{R}^d$ over \mathcal{S} and an unknown vector $\theta_h^* \in \mathbb{R}^d$ such that*

$$\mathbb{P}_h(s'|s, a) = \langle \phi(s, a), w_h^*(s') \rangle, \quad R_h(s, a) = \langle \phi(s, a), \theta_h^* \rangle, \quad (12)$$

for all $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ at each step $h \in [H]$. Here we assume $\|\phi(x)\|_2 \leq 1$ for all $x \in \mathcal{X}$ and $\max\{\|w_h^*(\mathcal{S})\|_2, \|\theta_h^*\|_2\} \leq \sqrt{d}$ at each step $h \in [H]$, where $\|w_h^*(\mathcal{S})\|_2 = \int_{\mathcal{S}} \|w_h^*(s)\|_2 ds$.

In linear MDPs, it has been shown that both reward $R_h(\cdot)$ and transition value function $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ are linear functions with respect to $\phi(\cdot)$ Agarwal et al. (2019); Jin et al. (2020). Thus, with only trajectory-wise rewards, we can construct the proxy reward $\widehat{R}_h(\cdot)$ and estimated transition value function $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ by solving two linear regression problems that take similar forms as eq. (4) and eq. (6) with feature $\Phi(\tau) = [\phi(x_1^\top), \dots, \phi(x_h^\top)]^\top$ and $\phi(x)$, respectively. For a detailed description of the algorithm, please see Algorithm 2 in Appendix C.1. The following theorem characterizes the suboptimality of Algorithm 2.

Theorem 2 (Informal). *Consider PARTED for linear MDP in Algorithm 2. Under appropriate hyperparameter setting and dataset coverage assumption, we have $\text{SubOpt}(\widehat{\pi}, s) \leq \widetilde{O}(dH^3/\sqrt{N})$ holds with high probability.*

Note that linear function with feature $\Phi(\tau)$ and $\phi(x)$ belongs to RKHS with kernel $K'_H(\tau, \tau') = \langle \Phi(\tau), \Phi(\tau') \rangle$ and $K'(x, x') = \langle \phi(x), \phi(x') \rangle$, respectively. Thus, $\mathcal{H}_{K'_H}$ is dH -dimensional and $\mathcal{H}_{K'}$ is d -dimensional. The suboptimality of PARTED for linear MDP in Theorem 2 will match that in Corollary 1 if we let $D_1 = dH$ and $D_2 = d$ (which implies $D_{\text{eff}} = dH$), where the dynamic of neural networks can be approximately captured by RKHSs \mathcal{H}_{K_H} and \mathcal{H}_K defined in Section 4.1.

To highlight why trajectory-wise reward RL is more challenging than instantaneous reward RL, we observe that Theorem 2 with *trajectory-wise rewards* has an additional dependence on the horizon H , compared to the suboptimality $\widetilde{O}(dH^2/\sqrt{N})$ (Jin et al., 2021, Corollary 4.5) of PEVI for linear MDP with *instantaneous rewards*. This additional dependence on H is caused by the reward redistribution process, in which PARTED needs to solve a trajectory-level regression problem with feature $\Phi(\tau) \in \mathbb{R}^{dH}$, which inevitably introduces large uncertainty in the regression solution used to construct the per-step proxy reward.

5 CONCLUSION

In this paper, we propose a novel offline RL algorithm, called PARTED, to handle the episodic RL problem with trajectory-wise rewards. PARTED uses a least-square-based reward redistribution method for reward estimation and incorporates a new penalty term to offset the uncertainty of proxy reward. Under the neural network function approximation, we prove that PARTED achieves an $\widetilde{O}(D_{\text{eff}}H^2/\sqrt{N})$ suboptimality, which matches the order $\widetilde{O}(dH^3/\sqrt{N})$ of linear MDP (that we further establish) when the effective dimension satisfies $D_{\text{eff}} = dH$. To the best of our knowledge, this is the first offline RL algorithm that is provably efficient in general episodic MDP setting with trajectory-wise rewards. As a future direction, it is interesting to incorporate the randomized return decomposition in Ren et al. (2021b) to improve the scalability of PARTED in the long horizon scenario.

REFERENCES

Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep.*, 2019.

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *Proc. International Conference on Machine Learning (ICML)*, pp. 242–252. PMLR, 2019.
- Jose A Arjona-Medina, Michael Gillhofer, Michael Widrich, Thomas Unterthiner, Johannes Brandstetter, and Sepp Hochreiter. Rudder: return decomposition for delayed rewards. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- Yu Bai and Jason D Lee. Beyond linearization: on quadratic and higher-order approximation of wide neural networks. *arXiv preprint arXiv:1910.01619*, 2019.
- Jacob Buckman, Carles Gelada, and Marc G Bellemare. The importance of pessimism in fixed-dataset policy optimization. *arXiv preprint arXiv:2009.06799*, 2020.
- Qi Cai, Zhuoran Yang, Jason D Lee, and Zhaoran Wang. Neural temporal-difference and q-learning provably converge to global optima. *arXiv preprint arXiv:1905.10027*, 2019.
- Niladri Chatterji, Aldo Pacchiano, Peter Bartlett, and Michael Jordan. On the theory of reinforcement learning with once-per-episode feedback. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proc. International Conference on Machine Learning (ICML)*, pp. 844–853. PMLR, 2017.
- Alon Cohen, Haim Kaplan, Tomer Koren, and Yishay Mansour. Online markov decision processes with aggregate bandit feedback. In *Proc. Conference on Learning Theory (COLT)*, pp. 1301–1329. PMLR, 2021.
- Robert Dadashi, Shideh Rezaeifar, Nino Vieillard, Léonard Hussenot, Olivier Pietquin, and Matthieu Geist. Offline reinforcement learning with pseudometric learning. In *Proc. International Conference on Machine Learning (ICML)*, pp. 2307–2318. PMLR, 2021.
- Yaqi Duan, Zeyu Jia, and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *Proc. International Conference on Machine Learning (ICML)*, pp. 2701–2709. PMLR, 2020.
- Yonathan Efroni, Nadav Merlis, and Shie Mannor. Reinforcement learning with trajectory feedback. In *Proc. AAAI Conference on Artificial Intelligence*, volume 35, pp. 7288–7295, 2021.
- Dylan J Foster, Akshay Krishnamurthy, David Simchi-Levi, and Yunzong Xu. Offline reinforcement learning: fundamental barriers for value function approximation. *arXiv preprint arXiv:2111.10919*, 2021.
- Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Scott Fujimoto, Edoardo Conti, Mohammad Ghavamzadeh, and Joelle Pineau. Benchmarking batch deep reinforcement learning algorithms. *arXiv preprint arXiv:1910.01708*, 2019a.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *Proc. International Conference on Machine Learning (ICML)*, pp. 2052–2062. PMLR, 2019b.
- Tanmay Gangwani, Yuan Zhou, and Jian Peng. Learning guidance rewards with trajectory-space smoothing. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 33:822–832, 2020.
- Ruiqi Gao, Tianle Cai, Haochuan Li, Cho-Jui Hsieh, Liwei Wang, and Jason D Lee. Convergence of adversarial training in overparametrized neural networks. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- Yaobang Gong, Mohamed Abdel-Aty, Qing Cai, and Md Sharikur Rahman. Decentralized network level adaptive signal control by multi-agent deep reinforcement learning. *Transportation Research Interdisciplinary Perspectives*, 1:100020, 2019.

- Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nature medicine*, 25(1):16–18, 2019.
- Beining Han, Zhizhou Ren, Zuofan Wu, Yuan Zhou, and Jian Peng. Off-policy reinforcement learning with delayed rewards. *arXiv preprint arXiv:2106.11854*, 2021.
- Daniel Hein, Stefan Depeweg, Michel Tokic, Steffen Udluft, Alexander Hentschel, Thomas A Runkler, and Volkmar Sterzing. A benchmark environment motivated by industrial control problems. In *Proc. Symposium Series on Computational Intelligence (SSCI)*, pp. 1–8. IEEE, 2017.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Proc. Conference on Learning Theory (COLT)*, pp. 2137–2143. PMLR, 2020.
- Ying Jin, Zhuoran Yang, and Zhaoran Wang. Is pessimism provably efficient for offline rl? In *Proc. International Conference on Machine Learning (ICML)*, pp. 5084–5096. PMLR, 2021.
- Martin Klissarov and Doina Precup. Reward propagation using graph convolutional networks. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 33:12895–12908, 2020.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 33:1179–1191, 2020.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Gen Li, Laixi Shi, Yuxin Chen, Yuejie Chi, and Yuting Wei. Settling the sample complexity of model-based offline reinforcement learning. *arXiv preprint arXiv:2204.05275*, 2022.
- Kaixiang Lin, Renyu Zhao, Zhe Xu, and Jiayu Zhou. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proc. International Conference on Knowledge Discovery & Data Mining*, pp. 1774–1783, 2018.
- Yang Liu, Yunan Luo, Yuanyi Zhong, Xi Chen, Qiang Liu, and Jian Peng. Sequence modeling of temporal credit assignment for episodic reinforcement learning. *arXiv preprint arXiv:1905.13420*, 2019.
- Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Provably good batch reinforcement learning without great exploration. *arXiv preprint arXiv:2007.08202*, 2020.
- Hamid Rez Maei. *Gradient temporal-difference learning algorithms*. PhD thesis, University of Alberta, 2011.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- Junhyuk Oh, Yijie Guo, Satinder Singh, and Honglak Lee. Self-imitation learning. In *Proc. International Conference on Machine Learning (ICML)*, pp. 3878–3887. PMLR, 2018.
- Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1):1–14, 2017.

- Aldo Pacchiano, Aadirupa Saha, and Jonathan Lee. Dueling rl: reinforcement learning with trajectory preferences. *arXiv preprint arXiv:2111.04850*, 2021.
- Shuang Qiu, Jieping Ye, Zhaoran Wang, and Zhuoran Yang. On reward-free rl with kernel and neural function approximations: single-agent mdp and markov game. In *Proc. International Conference on Machine Learning (ICML)*, pp. 8737–8747. PMLR, 2021.
- Hazhir Rahmandad, Nelson Repenning, and John Sterman. Effects of feedback delay on learning. *System Dynamics Review*, 25(4):309–338, 2009.
- Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Tongzheng Ren, Jialian Li, Bo Dai, Simon S Du, and Sujay Sanghavi. Nearly horizon-free offline reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021a.
- Zhizhou Ren, Ruihan Guo, Yuan Zhou, and Jian Peng. Learning long-term reward redistribution via randomized return decomposition. *arXiv preprint arXiv:2111.13485*, 2021b.
- Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander WR Nelson, Alex Bridgland, et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.
- Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.
- Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Pessimistic q-learning for offline reinforcement learning: Towards optimal sample complexity. *arXiv preprint arXiv:2202.13890*, 2022.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Joel A Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.
- Michal Valko, Nathaniel Korda, Rémi Munos, Ilias Flaounas, and Nelo Cristianini. Finite-time analysis of kernelised contextual bandits. *arXiv preprint arXiv:1309.6869*, 2013.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Lingxiao Wang, Qi Cai, Zhuoran Yang, and Zhaoran Wang. Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*, 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*, 2020a.
- Ziyu Wang, Alexander Novikov, Konrad Zolna, Josh S Merel, Jost Tobias Springenberg, Scott E Reed, Bobak Shahriari, Noah Siegel, Caglar Gulcehre, Nicolas Heess, et al. Critic regularized regression. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 33:7768–7778, 2020b.
- Tengyang Xie, Ching-An Cheng, Nan Jiang, Paul Mineiro, and Alekh Agarwal. Bellman-consistent pessimism for offline reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021a.

- Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, and Yu Bai. Policy finetuning: Bridging sample-efficient offline and online reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021b.
- Tengyu Xu, Yingbin Liang, and Guanghui Lan. Crpo: A new approach for safe reinforcement learning with convergence guarantee. In *Proc. International Conference on Machine Learning (ICML)*, pp. 11480–11491. PMLR, 2021.
- Yuling Yan, Gen Li, Yuxin Chen, and Jianqing Fan. The efficacy of pessimism in asynchronous q-learning. *arXiv preprint arXiv:2203.07368*, 2022.
- Lin Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *Proc. International Conference on Machine Learning (ICML)*, pp. 6995–7004. PMLR, 2019.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. On function approximation in reinforcement learning: Optimism in the face of large state spaces. *arXiv preprint arXiv:2011.04622*, 2020.
- Ming Yin and Yu-Xiang Wang. Optimal uniform ope and model-based offline reinforcement learning in time-homogeneous, reward-free and task-agnostic settings. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021a.
- Ming Yin and Yu-Xiang Wang. Towards instance-optimal offline reinforcement learning with pessimism. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021b.
- Ming Yin, Yu Bai, and Yu-Xiang Wang. Near-optimal provable uniform convergence in offline policy evaluation for reinforcement learning. *arXiv preprint arXiv:2007.03760*, 2020.
- Ming Yin, Yu Bai, and Yu-Xiang Wang. Near-optimal offline reinforcement learning via double variance reduction. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Ming Yin, Yaqi Duan, Mengdi Wang, and Yu-Xiang Wang. Near-optimal offline reinforcement learning with linear representation: Leveraging variance information with pessimism. *arXiv preprint arXiv:2203.05804*, 2022.
- Andrea Zanette. Exponential lower bounds for batch reinforcement learning: batch rl can be exponentially harder than online rl. In *Proc. International Conference on Machine Learning (ICML)*, pp. 12287–12297. PMLR, 2021.
- Andrea Zanette, Martin J Wainwright, and Emma Brunskill. Provable benefits of actor-critic methods for offline reinforcement learning. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.
- Zeyu Zheng, Junhyuk Oh, and Satinder Singh. On learning intrinsic rewards for policy gradient methods. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.
- Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *Proc. International Conference on Machine Learning (ICML)*, pp. 11492–11502. PMLR, 2020.

Supplementary Materials

A PROOF FLOW OF THEOREM 1

In this section, we present the main proof flow of Theorem 1. We first decompose the suboptimality $\text{SubOpt}(\pi, s)$, and then present the two main results of Lemma 1 and Lemma 2 to bound the evaluation error and summation of penalty terms, respectively. The detailed proof of Lemma 1 and Lemma 2 can be found at Appendix D and Appendix E.

We define the evaluation error at each step $h \in [H]$ as

$$\delta_h(s, a) = (\mathbb{B}_h \widehat{V}_{h+1})(s, a) - \widehat{Q}_h(s, a), \quad (13)$$

where \mathbb{B}_h is the Bellman operator defined in Section 2.1 and \widehat{V}_h and \widehat{Q}_h are estimation of state- and state-action value functions, respectively. To proceed the proof, we first decompose the suboptimality into three parts as follows via the standard technique (see Section A in Jin et al. (2021)).

$$\begin{aligned} \text{SubOpt}(\pi, s) &= - \sum_{h=1}^H \mathbb{E}_\pi [\delta_h(s_h, a_h) | s_1 = s] + \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) | s_1 = s] \\ &\quad + \sum_{h=1}^H \mathbb{E}_{\pi^*} [\langle \widehat{Q}_h(s_h, \cdot), \pi_h^*(\cdot | s_h) - \widehat{\pi}_h(\cdot | s_h) \rangle | s_1 = s]. \end{aligned} \quad (14)$$

In Algorithm 1, the output policy at each horizon $\widehat{\pi}_h$ is greedy with respect to the estimated Q-value \widehat{Q}_h . Thus, we have

$$\langle \widehat{Q}_h(s_h, \cdot), \pi_h^*(\cdot | s_h) - \widehat{\pi}_h(\cdot | s_h) \rangle \leq 0, \quad \forall h \in [H], \quad \forall s_h \in \mathcal{S}.$$

According to eq. (14), we have the following holds for the suboptimality of $\widehat{\pi} = \{\widehat{\pi}_h\}_{h=1}^H$

$$\text{SubOpt}(\widehat{\pi}, s) = - \sum_{h=1}^H \mathbb{E}_{\widehat{\pi}} [\delta_h(s_h, a_h) | s_1 = s] + \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) | s_1 = s]. \quad (15)$$

In the following lemma, we provide the first main technical result for the proof, which bounds the evaluation error $\delta_h(s, a)$. Recall that we use \mathcal{X} to represent the joint state-action space $\mathcal{S} \times \mathcal{A}$ and use x to represent a state action pair (s, a) .

Lemma 1. *Let $\lambda_1, \lambda_2 = 1 + 1/N$. Suppose Assumption 2 holds. With probability at least $1 - \mathcal{O}(N^{-2}H^{-4})$, it holds for all $h \in [H]$ and $x \in \mathcal{X}$ that*

$$-\varepsilon_b \leq \delta_h(x) \leq 2 \left[\beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b \right], \quad \forall x \in \mathcal{X}, \quad \forall h \in [H],$$

where

$$\begin{aligned} \varepsilon_b &= \max\{\beta_1 H^{2/3}, \beta_2 H^{1/6}\} \mathcal{O} \left(\frac{N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right) + \mathcal{O} \left(\frac{H^{17/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right), \\ \beta_1 &= H \left(\frac{4a_2^2 \lambda_1}{d} + 2 \log \det \left(I + \frac{K_N^r}{\lambda_1} \right) + 10 \log(NH^2) \right)^{1/2}, \\ \beta_2 &= H \left(\frac{8A_2^2 \lambda_2}{d} + 4 \max \left\{ \log \det \left(I + \frac{K_{N,h}^v}{\lambda_2} \right) \right\} + 6C_\epsilon + 16 \log(NH^2 \mathcal{N}_\epsilon^v) \right)^{1/2}, \\ \epsilon &= \sqrt{\lambda_2 C_\epsilon H / (2NC_\phi)}, \text{ where } C_\epsilon \geq 1. \end{aligned}$$

Proof. The main technical development of the proof lies in handling the uncertainty caused by redistributing the trajectory-wise reward via solving a trajectory-level regression problem and analyzing the dynamics of neural network optimization. The detailed proof is provided in Appendix D. \square

Applying Lemma 1 to eq. (15) yields

$$\begin{aligned} \text{SubOpt}(\hat{\pi}, s) &= -\sum_{h=1}^H \mathbb{E}_{\hat{\pi}} [\delta_h(s_h, a_h) | s_1 = s] + \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) | s_1 = s] \\ &\leq 3H\varepsilon_b + 2\beta_1 \cdot \sum_{h=1}^H b_{r,h}(x, \hat{\Theta}) + 2\beta_2 \cdot \sum_{h=1}^H b_{v,h}(x, \hat{w}_h). \end{aligned} \quad (16)$$

The following lemma captures the second main technical result for the proof, which bounds the summation of the penalty terms $\beta_1 \cdot \sum_{h=1}^H b_{r,h}(x, \hat{\Theta}) + \beta_2 \cdot \sum_{h=1}^H b_{v,h}(x, \hat{w}_h)$.

Lemma 2. *Suppose Assumption 2&3 hold. We have the following holds with probability $1 - \mathcal{O}(N^{-2}H^{-4})$*

$$\begin{aligned} &\beta_1 \cdot \sum_{h=1}^H b_{r,h}(x, \hat{\Theta}) + \beta_2 \cdot \sum_{h=1}^H b_{v,h}(x, \hat{w}_h) \\ &\leq \left(\frac{\beta_1}{\sqrt{C_\sigma}} + \frac{\beta_2}{\sqrt{C_\zeta}} \right) \frac{\sqrt{2}HC_\phi}{\sqrt{N}} + \max\{\beta_1 H^{5/3}, \beta_2 H^{7/6}\} \cdot \mathcal{O} \left(\frac{N^{1/12}(\log m)^{1/4}}{m^{1/12}} \right). \end{aligned}$$

Proof. The proof develops new analysis to characterize the summation of the penalty term $b_{r,h}$ constructed by trajectory features, which is unique in the trajectory-wise reward setting. The detailed proof is provided in Appendix E. \square

Applying Lemma 2 to eq. (16), we have

$$\begin{aligned} &\text{SubOpt}(\hat{\pi}, s) \\ &\leq 3H\varepsilon_b + \left(\frac{\beta_1}{\sqrt{C_\sigma}} + \frac{\beta_2}{\sqrt{C_\zeta}} \right) \frac{2\sqrt{2}HC_\phi}{\sqrt{N}} + \max\{\beta_1 H^{5/3}, \beta_2 H^{7/6}\} \cdot \mathcal{O} \left(\frac{N^{1/12}(\log m)^{1/4}}{m^{1/12}} \right) \\ &\leq 4H\varepsilon_b + \left(\frac{\beta_1}{\sqrt{C_\sigma}} + \frac{\beta_2}{\sqrt{C_\zeta}} \right) \frac{2\sqrt{2}HC_\phi}{\sqrt{N}}, \end{aligned} \quad (17)$$

which completes the proof.

B PROOF OF COROLLARY 1

To provide a concrete bound for $\text{SubOpt}(\hat{\pi}, s)$ defined in eq. (17), we first need to bound the penalty coefficients β_1, β_2 under Assumption 4. Recalling the properties of β_1, β_2 in Theorem 1, we have

$$\begin{aligned} &H \left(\frac{4a_2^2\lambda_1}{d} + 2 \log \det \left(I + \frac{K_N^r}{\lambda_1} \right) + 10 \log(NH^2) \right)^{1/2} \leq R_{\beta_1} = \beta_1, \quad (18) \\ &H \left(\frac{8A_2^2\lambda_2}{d} + 4 \max_{h \in [H]} \left\{ \log \det \left(I + \frac{K_{N,h}^v}{\lambda_2} \right) \right\} + 6C_\epsilon + 16 \log(NH^2\mathcal{N}_\epsilon^v) \right)^{1/2} \leq R_{\beta_2} = \beta_2. \end{aligned} \quad (19)$$

Recall that we use \mathcal{X} to represent the joint state-action space $\mathcal{S} \times \mathcal{A}$ and use x to represent a state action pair (s, a) . We define the maximal information gain associated with RHKS with kernels K_N^r and $K_{N,h}^v$ as follows

$$\Gamma_{K_N^r}(N, \lambda_1) = \sup_{\mathcal{D} \subset \mathcal{D}_\tau} \{1/2 \cdot \log \det(I_{2dmH} + \lambda_1^{-1} \cdot K_N^r)\}, \quad (20)$$

$$\Gamma_{K_{N,h}^v}(N, \lambda_2) = \sup_{\mathcal{D} \subset \mathcal{D}_x} \{1/2 \cdot \log \det(I_{2dm} + \lambda_2^{-1} \cdot K_{N,h}^v)\}, \quad (21)$$

where \mathcal{D}_x and \mathcal{D}_τ are discrete subsets of state-action pair $x \in \mathcal{X}$ and trajectory $\tau \in \mathcal{X} \times \dots \times \mathcal{X}$ with cardinality no more than N , respectively. Applying Lemma 9 in Appendix H and Assumption 4, we have

$$\Gamma_{K_N^r}(N, \lambda_1) \leq C_{K_1} \cdot D_1 \cdot \log N \quad \text{and} \quad \Gamma_{K_{N,h}^v}(N, \lambda_2) \leq C_{K_2} \cdot D_2 \cdot \log N, \quad (22)$$

where C_{K_1}, C_{K_2} are absolute constants. Recall that $\mathcal{N}_{\epsilon, h}^v$ is the cardinality of the function class. Next, we proceed to bound the term $\mathcal{N}_\epsilon^v = \max_{h \in [H]} \{\mathcal{N}_{\epsilon, h}^v\}$.

$$\mathcal{V}_h(x, R_\theta, R_w, R_{\beta_1}, R_{\beta_2}, \lambda_1, \lambda_2) = \left\{ \max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a, \theta, w, \beta_1, \beta_2, \Sigma, \Lambda)\} : \mathcal{S} \rightarrow [0, H] \right. \\ \left. \text{with } \|\theta - \theta_0\|_2 \leq R_\theta, \|w - w_0\|_2 \leq R_w, \beta_1 \in [0, R_{\beta_1}], \beta_2 \in [0, R_{\beta_2}], \|\Sigma\|_2 \geq \lambda_1, \|\Lambda\|_2 \geq \lambda_2 \right\},$$

where $R_\theta = H\sqrt{N/\lambda_1}$, $R_w = H\sqrt{N/\lambda_2}$, and

$$\bar{Q}_h(x, \theta, w, \beta_1, \beta_2, \Sigma, \Lambda) = \min\{\langle \phi(x, \theta_0), \theta - \theta_0 \rangle + \langle \phi(x, w_0), w - w_0 \rangle \\ - \beta_1 \cdot \sqrt{\Phi_h(x, \theta_0)^\top \Sigma^{-1} \Phi_h(x, \theta_0)} - \beta_2 \cdot \sqrt{\phi(x, w_0)^\top \Lambda^{-1} \phi(x, w_0)}, H\}^+.$$

Note that

$$\begin{aligned} & \left| \max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a, \theta, w, \beta_1, \beta_2, \Sigma, \Lambda)\} - \max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a, \theta', w', \beta'_1, \beta'_2, \Sigma', \Lambda')\} \right| \\ & \leq \max_{a \in \mathcal{A}} \left| \bar{Q}_h(s, a, \theta, w, \beta_1, \beta_2, \Sigma, \Lambda) - \bar{Q}_h(s, a, \theta', w', \beta'_1, \beta'_2, \Sigma', \Lambda') \right| \\ & \stackrel{(i)}{\leq} \max_{a \in \mathcal{A}} |\langle \phi(x, \theta_0), \theta - \theta' \rangle| + \max_{a \in \mathcal{A}} |\langle \phi(x, w_0), w - w' \rangle| \\ & \quad + \max_{a \in \mathcal{A}} \left| (\beta_1 - \beta'_1) \cdot \sqrt{\Phi_h(x, \theta_0)^\top \Sigma^{-1} \Phi_h(x, \theta_0)} \right| \\ & \quad + \max_{a \in \mathcal{A}} \left| \beta'_1 \cdot \left[\sqrt{\Phi_h(x, \theta_0)^\top \Sigma^{-1} \Phi_h(x, \theta_0)} - \sqrt{\Phi_h(x, \theta_0)^\top \Sigma'^{-1} \Phi_h(x, \theta_0)} \right] \right| \\ & \quad + \max_{a \in \mathcal{A}} \left| (\beta_2 - \beta'_2) \cdot \sqrt{\phi(x, w_0)^\top \Lambda^{-1} \phi(x, w_0)} \right| \\ & \quad + \max_{a \in \mathcal{A}} \left| \beta'_2 \cdot \left[\sqrt{\phi(x, w_0)^\top \Lambda^{-1} \phi(x, w_0)} - \sqrt{\phi(x, w_0)^\top \Lambda'^{-1} \phi(x, w_0)} \right] \right| \\ & \stackrel{(ii)}{\leq} \max_{a \in \mathcal{A}} |\langle \Phi_h(x, \Theta_0), \Theta - \Theta' \rangle| + \max_{a \in \mathcal{A}} |\langle \phi(x, w_0), w - w' \rangle| + \frac{C_\phi}{\sqrt{\lambda_1}} |\beta_1 - \beta'_1| + \frac{C_\phi}{\sqrt{\lambda_2}} |\beta_2 - \beta'_2| \\ & \quad + R_{\beta_1} \max_{a \in \mathcal{A}} \left| \|\Phi_h(x, \theta_0)\|_{\Sigma^{-1}} - \|\Phi_h(x, \theta_0)\|_{\Sigma'^{-1}} \right| \\ & \quad + R_{\beta_2} \max_{a \in \mathcal{A}} \left| \|\phi(x, w_0)\|_{\Lambda^{-1}} - \|\phi(x, w_0)\|_{\Lambda'^{-1}} \right|, \end{aligned} \tag{23}$$

where (i) follows from contractive properties of operators $\min\{\cdot, H\}^+$ and $\max_{a \in \mathcal{A}}\{\cdot\}$ and the triangle inequality, and (ii) follows from the fact that $\|\phi(x, w_0)\|_2, \|\Phi_h(x, \theta_0)\|_2 \leq C_\phi$.

Following arguments similar to those in the proof of Corollaries 4.8, Corollaries 4.4 and Section D.1 in Yang et al. (2020), we have the followings hold for terms in the right hand side of eq. (23)

$$|\langle \Phi_h(x, \Theta_0), \Theta - \Theta' \rangle| = |g_1(x) - g_2(x)| \quad \text{where } \|g_i\|_{\mathcal{H}_{K_H}} \leq R_g = 2H\sqrt{\Gamma_{K_N^r}(N, \lambda_1)} \quad \forall i \in \{1, 2\}, \tag{24}$$

$$|\langle \phi(x, w_0), w - w' \rangle| = |h_1(x) - h_2(x)| \quad \text{where } \|h_i\|_{\mathcal{H}_K} \leq R_h = 2H\sqrt{\Gamma_{K_{N,h}^v}(N, \lambda_2)} \quad \forall i \in \{1, 2\}, \tag{25}$$

$$\begin{aligned} \|\Phi_h(x, \theta_0)\|_{\Sigma^{-1}} - \|\Phi_h(x, \theta_0)\|_{\Sigma'^{-1}} &= \|\Psi(x)\|_\Omega - \|\Psi(x)\|_{\Omega'}, \\ \|\phi(x, w_0)\|_{\Lambda^{-1}} - \|\phi(x, w_0)\|_{\Lambda'^{-1}} &= \|\psi(x)\|_\Upsilon - \|\psi(x)\|_{\Upsilon'}, \end{aligned}$$

where $g_1(\cdot), g_2(\cdot)$ are two functions in RKHS \mathcal{H}_{K_H} , $h_1(\cdot), h_2(\cdot)$ are two functions in RKHS \mathcal{H}_K , $\Psi(\cdot)$ and $\psi(\cdot)$ are feature mappings of RKHSs \mathcal{H}_{K_H} and \mathcal{H}_K , respectively, $\Omega, \Omega' : \mathcal{H}_{K_H} \rightarrow \mathcal{H}_{K_H}$ are self-adjoint operators with eigenvalues bounded in $[0, 1/\lambda_1]$, and $\Upsilon, \Upsilon' : \mathcal{H}_K \rightarrow \mathcal{H}_K$ are self-adjoint operators with eigenvalues bounded in $[0, 1/\lambda_2]$. We define the following two function classes

$$\mathcal{F}_1 = \{\|\Psi(\cdot)\|_\Omega : \|\Omega\|_2 \leq 1/\lambda_1\}, \tag{26}$$

and

$$\mathcal{F}_2 = \{\|\psi(\cdot)\|_{\Upsilon} : \|\Upsilon\|_2 \leq 1/\lambda_2\}. \quad (27)$$

For any $\epsilon > 0$, we denote $\mathcal{N}(\epsilon, \mathcal{H}, R)$ as the ϵ -covering of $\{f \in \mathcal{H} : \|f\|_{\mathcal{H}} \leq R\}$, denote $\mathcal{N}(\epsilon, \mathcal{F}_1, \lambda_1)$ as the ϵ -covering number of \mathcal{F}_1 in eq. (26), denote $\mathcal{N}(\epsilon, \mathcal{F}_2, \lambda_2)$ as the ϵ -covering number of \mathcal{F}_2 in eq. (27), and denote $\mathcal{N}(\epsilon, R)$ as the ϵ -covering number of the interval $[0, R]$ with respect to the Euclidean distance. Note that eq. (23) implies

$$\begin{aligned} \mathcal{N}_{\epsilon, h}^v &\leq \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m^H}, R_g) \cdot \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m}, R_h) \cdot \mathcal{N}(\epsilon/(6C_\phi), R_{\beta_1}) \cdot \mathcal{N}(\epsilon/(6C_\phi), R_{\beta_2}) \\ &\quad \cdot \mathcal{N}(\epsilon/(6R_{\beta_1}), \mathcal{F}_1, \lambda_1) \cdot \mathcal{N}(\epsilon/(6R_{\beta_2}), \mathcal{F}_2, \lambda_2). \end{aligned} \quad (28)$$

Based on Corollary 4.1.13 in Vershynin (2018), we have the followings hold for $\mathcal{N}(\epsilon/(6C_\phi), R_{\beta_1})$ and $\mathcal{N}(\epsilon/(6C_\phi), R_{\beta_2})$, respectively

$$\mathcal{N}(\epsilon/(6C_\phi), R_{\beta_1}) \leq 1 + 12C_\phi R_{\beta_1}/\epsilon \text{ and } \mathcal{N}(\epsilon/(6C_\phi), R_{\beta_2}) \leq 1 + 12C_\phi R_{\beta_2}/\epsilon. \quad (29)$$

Moreover, as shown in Lemma D.2 and Lemma D.3 in Yang et al. (2020), under the finite spectrum NTK assumption in Assumption 4, we have the followings hold

$$\log \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m^H}, R_g) \leq C_1 \cdot D_1 \cdot [\log(6R_g/\epsilon) + C_2], \quad (30)$$

$$\log \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m}, R_h) \leq C_3 \cdot D_2 \cdot [\log(6R_h/\epsilon) + C_4], \quad (31)$$

$$\log \mathcal{N}(\epsilon/(6R_{\beta_1}), \mathcal{F}_1, \lambda_1) \leq C_5 \cdot D_1^2 \cdot [\log(6R_{\beta_1}/\epsilon) + C_6], \quad (32)$$

$$\log \mathcal{N}(\epsilon/(6R_{\beta_2}), \mathcal{F}_2, \lambda_2) \leq C_7 \cdot D_2^2 \cdot [\log(6R_{\beta_2}/\epsilon) + C_8]. \quad (33)$$

where C_i ($i \in \{1, \dots, 8\}$) are absolute constants that do not rely on N , H or ϵ . Then, substituting eq. (29)-(33) into eq. (28), we have

$$\begin{aligned} \log \mathcal{N}_{\epsilon, h}^v &\leq \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m^H}, R_g) + \mathcal{N}(\epsilon/6, \mathcal{H}_{K_m}, R_h) + \mathcal{N}(\epsilon/(6C_\phi), R_{\beta_1}) + \mathcal{N}(\epsilon/(6C_\phi), R_{\beta_2}) \\ &\quad + \mathcal{N}(\epsilon/(6R_{\beta_1}), \mathcal{F}_1, \lambda_1) + \mathcal{N}(\epsilon/(6R_{\beta_2}), \mathcal{F}_2, \lambda_2) \\ &\leq \log(1 + 12C_\phi R_{\beta_1}/\epsilon) + \log(1 + 12C_\phi R_{\beta_2}/\epsilon) + C_1 D_1 [\log(6R_g/\epsilon) + C_2] \\ &\quad + C_3 D_2 [\log(6R_h/\epsilon) + C_4] + C_5 D_1^2 [\log(6R_{\beta_1}/\epsilon) + C_6] + C_7 D_2^2 [\log(6R_{\beta_2}/\epsilon) + C_8], \end{aligned} \quad (34)$$

We next proceed to show that there exists an absolute constant $R_{\beta_1} > 0$ such that eq. (18) holds. Substituting eq. (22) to eq. (18), we can obtain

$$\text{L.H.S of eq. (18)} \leq H \left(\frac{4a_2^2 \lambda_1}{d} + 4C_{K_1} D_1 \log N + 10 \log(NH^2) \right)^{1/2}.$$

If we let

$$R_{\beta_1} = C_{\beta_1} H \sqrt{D_1 \log(NH^2)}, \quad (35)$$

in which C_{β_1} is a sufficiently large constant, then we have the following holds

$$\text{L.H.S of eq. (18)} \leq R_{\beta_1}.$$

Note that eq. (34) directly implies that

$$\begin{aligned} \log \mathcal{N}_\epsilon^v &= \max_{h \in [H]} \{\log \mathcal{N}_{\epsilon, h}^v\} \\ &\leq \log(1 + 12C_\phi R_{\beta_1}/\epsilon) + \log(1 + 12C_\phi R_{\beta_2}/\epsilon) + C_1 D_1 [\log(6R_g/\epsilon) + C_2] \\ &\quad + C_3 D_2 [\log(6R_h/\epsilon) + C_4] + C_5 D_1^2 [\log(6R_{\beta_1}/\epsilon) + C_6] + C_7 D_2^2 [\log(6R_{\beta_2}/\epsilon) + C_8] \\ &\stackrel{(i)}{\leq} C'_1 D_1^2 \log(R_{\beta_1}/\epsilon) + C'_2 D_2^2 \log(R_{\beta_2}/\epsilon) + C'_3 D_1 \log(H \sqrt{D_1}/\epsilon) + C'_4 D_2 \log(H \sqrt{D_2}/\epsilon) \\ &\stackrel{(ii)}{\leq} C''_1 D_1^2 \log(NH^2 \sqrt{D_1}/\epsilon) + C''_2 D_2^2 \log(R_{\beta_2}/\epsilon) + C'_3 D_1 \log(H \sqrt{D_1}/\epsilon) \\ &\quad + C'_4 D_2 \log(H \sqrt{D_2}/\epsilon), \end{aligned} \quad (36)$$

where in (i) we let C'_1, C'_2, C'_3 and C'_4 be sufficiently large absolute constants, in (ii) we use eq. (35) and let C''_1 be sufficiently large. Then, we proceed to show that there exists an absolute constant

$R_{\beta_2} > 0$ such that eq. (19) holds. Using eq. (22) and eq. (36), the left hand side of eq. (19) can be bounded as follows

$$\begin{aligned}
& \text{L.H.S of eq. (19)} \\
& \leq H \left(\frac{8A_2^2 \lambda_2}{d} + 8C_{K_2} \cdot D_2 \cdot \log N + 20 \log(NH^2) + 6C_\epsilon + 16 \log \mathcal{N}_\epsilon^v \right)^{1/2} \\
& \stackrel{(i)}{\leq} HC_{\beta_2,1} \sqrt{D_2 \log(NH^2)} + HC_{\beta_2,2} \sqrt{\log \mathcal{N}_\epsilon^v} + HC_{\beta_2,3} \sqrt{C_\epsilon} \\
& \leq HC_{\beta_2,1} \sqrt{D_2 \log(NH^2)} + HC_{\beta_2,2} \left[D_1 \sqrt{C_1'' \log(NH^2 \sqrt{D_1}/\epsilon)} + D_2 \sqrt{C_2' \log(R_{\beta_2}/\epsilon)} \right. \\
& \quad \left. + \sqrt{C_3' D_1 \log(H \sqrt{D_1}/\epsilon)} + \sqrt{C_4' D_2 \log(H \sqrt{D_2}/\epsilon)} \right] + HC_{\beta_2,3} \sqrt{C_\epsilon}.
\end{aligned}$$

where (i) follows from the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and $C_{\beta_2,1}$, $C_{\beta_2,2}$ and $C_{\beta_2,3}$ are sufficiently large constants. Clearly, if we let

$$R_{\beta_2} = C_{\beta_2} H \max\{D_1, D_2\} \log(NH^2 \max\{D_1, D_2\}/\epsilon), \quad (37)$$

where C_{β_2} is a sufficiently large absolute constant, then we have

$$\text{L.H.S of eq. (19)} \leq R_{\beta_2}. \quad (38)$$

Finally, substituting the value of R_{β_1} in eq. (35) and value of R_{β_2} in eq. (37) into eq. (17) and letting $C_\epsilon = \max\{D_1, D_2\}^2$ (which implies $\epsilon = \sqrt{\lambda_2} \max\{D_1, D_2\} H / (2NC_\phi)$), we have

SubOpt($\hat{\pi}, s$)

$$\begin{aligned}
& \leq \left(\frac{\beta_1}{\sqrt{C_\sigma}} + \frac{\beta_2}{\sqrt{C_\zeta}} \right) \frac{2\sqrt{2}HC_\phi}{\sqrt{N}} + 4H\epsilon_b \\
& \leq \max\{R_{\beta_1}, R_{\beta_2}\} \mathcal{O} \left(\frac{H}{\sqrt{N}} \right) + \max\{R_{\beta_1} H^{5/3}, R_{\beta_2} H^{7/6}\} \mathcal{O} \left(\frac{N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right) \\
& \quad + \mathcal{O} \left(\frac{H^{23/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right) \\
& \leq \mathcal{O} \left(\frac{H^2 \max\{D_1, D_2\}}{\sqrt{N}} \log \left(\frac{NH^2 \max\{D_1, D_2\}}{\epsilon} \right) \right) \\
& \quad + \max \left\{ \sqrt{H}, \max\{D_1, D_2\} \right\} \mathcal{O} \left(\frac{H^{13/6} N^{1/12} (\log m)^{1/4}}{m^{1/12}} \log(N^2 H^2) \right) \\
& \quad + \mathcal{O} \left(\frac{H^{23/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right) \\
& \stackrel{(i)}{\leq} \mathcal{O} \left(\frac{H^2 \max\{D_1, D_2\}}{\sqrt{N}} \log(2C_\phi N^2 H) \right) \\
& \quad + \max \left\{ \sqrt{H}, \max\{D_1, D_2\}, \frac{H^{5/3} N^{19/12}}{m^{1/12}} \right\} \mathcal{O} \left(\frac{H^{13/6} N^{1/12} (\log m)^{1/4}}{m^{1/12}} \log(N^2 H^5 m) \right),
\end{aligned}$$

where (i) follows from the definition of ϵ and the fact that $\lambda_2 \geq 1$.

C LINEAR MDP WITH TRAJECTORY-WISE REWARD

In this section, we present the full details of our study on the offline RL in the linear MDP setting with trajectory-wise rewards.

C.1 LINEAR MDP AND ALGORITHM

We define the linear MDP Jin et al. (2020) as follows, where the transition kernel and expected reward function are linear in a feature map. We use \mathcal{X} to represent the joint state-action space $\mathcal{S} \times \mathcal{A}$ and use x to represent a state action pair.

Algorithm 2 Linear Pessimistic Value Iteration with Reward Decomposition (PARTED)

Input: Dataset $\mathcal{D} = \{\tau_i, r(\tau_i)\}_{i,h=1}^{N,H}$
Initialization: Set \widehat{V}_{H+1} as zero function
Obtain \widehat{R}_h and $\widehat{\Theta}$ according to eq. (41)
for $h = H, H-1, \dots, 1$ **do**
 Obtain $\widehat{\mathbb{P}}_h \widehat{V}_{h+1}$ and \widehat{w}_h according to eq. (43)
 Obtain $\Gamma_h(\cdot)$ according to eq. (46)
 $\widehat{Q}_h(\cdot) = \min\{\widehat{R}_h(\cdot) + \widehat{\mathbb{P}}_h \widehat{V}_{h+1}(\cdot) - \Gamma_h(\cdot), H-h+1\}^+$
 $\widehat{\pi}_h(\cdot|s) = \operatorname{argmax}_{\pi_h} \langle \widehat{Q}_h(s, \cdot), \pi_h(\cdot|s) \rangle$
 $\widehat{V}_h(\cdot) = \langle \widehat{Q}_h(\cdot, \cdot), \widehat{\pi}_h(\cdot|\cdot) \rangle_{\mathcal{A}}$
end for

Definition 2 (Linear MDP). We say an episodic MDP $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, H)$ is a linear MDP with a known feature map $\phi(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^d$ if there exist an unknown vector $w_h^*(s) \in \mathbb{R}^d$ over \mathcal{S} and an unknown vector $\theta_h^* \in \mathbb{R}^d$ such that

$$\mathbb{P}_h(s'|s, a) = \langle \phi(s, a), w_h^*(s') \rangle, \quad R_h(s, a) = \langle \phi(s, a), \theta_h^* \rangle, \quad (39)$$

for all $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ at each step $h \in [H]$. Here we assume $\|\phi(x)\|_2 \leq 1$ for all $x \in \mathcal{X}$ and $\max\{\|w_h^*(\mathcal{S})\|_2, \|\theta_h^*\|_2\} \leq \sqrt{d}$ at each step $h \in [H]$, where with an abuse of notation, we define $\|w_h^*(\mathcal{S})\|_2 = \int_{\mathcal{S}} \|w_h^*(s)\|_2 ds$.

We present our PARTED algorithm for linear MDPs with trajectory-wise rewards in Algorithm 2. Note that Algorithm 2 shares a structure similar to that of Algorithm 1. Specifically, we estimate each $R_h(\cdot)$ for all $h \in [H]$ using a linear function $\langle \phi(s, a), \theta_h \rangle$, where $\theta_h \in \mathbb{R}^d$ is a learnable parameter. We define the vector $\Theta = [\theta_1^\top, \dots, \theta_H^\top] \in \mathbb{R}^{dH}$ and the loss function $L_r : \mathbb{R}^{dH} \rightarrow \mathbb{R}$ for reward learning as

$$L_r(\Theta) = \sum_{\tau \in \mathcal{D}} \left[\sum_{h=1}^H \langle \phi(x_h^\tau), \theta_h \rangle - r(\tau) \right]^2 + \lambda_1 \cdot \sum_{h=1}^H \|\theta_h - \theta_0\|_2^2, \quad (40)$$

where $\lambda_1 > 0$ is a regularization parameter. We then define $\widehat{R}_h(\cdot)$ as

$$\widehat{R}_h(\cdot) = \langle \phi(\cdot), \widehat{\theta}_h \rangle, \quad \text{where } \widehat{\Theta} = \operatorname{argmin}_{\Theta \in \mathbb{R}^{2dmH}} L_r(\Theta) \text{ and } \widehat{\Theta} = [\widehat{\theta}_1^\top, \dots, \widehat{\theta}_H^\top]^\top. \quad (41)$$

Similarly, we also use linear function $\langle \phi(s, a), w_h \rangle$ to estimate transition value functions $\{(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot, \cdot)\}_{h \in [H]}$ for all $h \in [H]$, where $w_h \in \mathbb{R}^d$ is a learnable parameter. For each $h \in [H]$, we define the loss function $L_v^h(w_h) : \mathbb{R}^d \rightarrow \mathbb{R}$ as

$$L_v^h(w_h) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - \langle \phi(x_h^\tau), w_h \rangle \right)^2 + \lambda_2 \cdot \|w_h - w_0\|_2^2, \quad (42)$$

where $\lambda_2 > 0$ is a regularization parameter. We then define $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ as

$$(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = \langle \phi(\cdot), \widehat{w}_h \rangle, \quad \text{where } \widehat{w}_h = \operatorname{argmin}_{w_h \in \mathbb{R}^d} L_v^h(w_h). \quad (43)$$

It remains to construct the penalty term Γ_h . We first consider the penalty term that is used to offset the uncertainty raised from estimating the reward $R_h(\cdot)$ for each $h \in [H]$. We define the vectors $\Phi_h(x) = [\mathbf{0}_d^\top, \dots, \phi(x)^\top, \dots, \mathbf{0}_d^\top]^\top \in \mathbb{R}^{dH}$ and $\Phi(\tau) = [\phi(x_1^\tau), \dots, \phi(x_H^\tau)] \in \mathbb{R}^{dH}$, where $\Phi_h(x) \in \mathbb{R}^{dH}$ is a vector in which $[\Phi_h(x)]_{d(h-1)+1:dh} = \phi(x)$ and the rest entries are all zero. We define a matrix $\Sigma(\Theta) \in \mathbb{R}^{dH \times dH}$ as

$$\Sigma = \lambda_1 \cdot I_{dH} + \sum_{\tau \in \mathcal{D}} \Phi(\tau) \Phi(\tau)^\top.$$

The penalty term $b_{r,h}$ of the estimated reward is then defined as

$$b_{r,h}(x) = [\Phi_h(x)^\top \Sigma^{-1} \Phi_h(x)]^{1/2}, \quad \forall x \in \mathcal{X}. \quad (44)$$

Next, we consider the penalty term that is used to offset the uncertainty raised from estimating the transition value function $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ for each $h \in [H]$. We define a matrix $\Lambda_h \in \mathbb{R}^{d \times d}$ as

$$\Lambda_h = \lambda_2 \cdot I_d + \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau) \phi(x_h^\tau)^\top.$$

The penalty term $b_{v,h}$ of the estimated transition value function is then defined as

$$b_{v,h}(x) = [\phi(x)^\top \Lambda_h^{-1} \phi(x)]^{1/2}, \quad \forall x \in \mathcal{X}. \quad (45)$$

Finally, the penalty term for the estimated Bellman operation $\widehat{\mathbb{B}}_h \widehat{V}_{h+1}(\cdot)$ is obtained as

$$\Gamma_h(x) = \beta_1 b_{r,h}(x) + \beta_2 b_{v,h}(x), \quad (46)$$

where $\beta_1, \beta_2 > 0$ are constant factors.

C.2 MAIN RESULT

We consider the following dataset coverage assumption so that we can explicitly bound the suboptimality of Algorithm 2. Note that the following assumption has also been considered in Jin et al. (2021).

Assumption 5 (Well-Explored Dataset). *Suppose the N trajectories in dataset \mathcal{D} are independent and identically induced by a fixed behaviour policy μ . There exist absolute constants $C_\sigma > 0$ and $C_\varsigma > 0$ such that*

$$\lambda_{\min}(\overline{M}(\Theta_0)) \geq C_\sigma \quad \text{and} \quad \lambda_{\min}(\overline{m}_h(w_0)) \geq C_\varsigma \quad \forall h \in [H],$$

where

$$\overline{M} = \mathbb{E}_\mu [\Phi(\tau) \Phi(\tau)^\top] \quad \text{and} \quad \overline{m}_h(w_0) = \mathbb{E}_\mu [\phi(x_h^\tau) \phi(x_h^\tau)^\top].$$

We provide a formal statement of Theorem 2 as follows, which characterizes the suboptimality of Algorithm 2.

Theorem 3 (Formal Statement of Theorem 2). *Consider Algorithm 2. Let $\lambda_1 = \lambda_2 = 1$ and $\beta_1 = \mathcal{O}(H \sqrt{dH \log(N/\delta)})$ and $\beta_2 = \mathcal{O}(dH^2 \sqrt{\log(dH^3 N^{5/2}/\delta)})$. Then, with probability at least $1 - \delta$, we have*

$$\text{SubOpt}(\widehat{\pi}, s) \leq \mathcal{O} \left(\frac{dH^3}{\sqrt{N}} \sqrt{\log \left(\frac{dH^3 N^{5/2}}{\delta} \right)} \right).$$

C.3 PROOF FLOW OF THEOREM 2

In this section, we present the main proof flow of Theorem 2. Our main development is Lemma 3, the proof of which is presented in Appendix F.

Recalling the suboptimality of $\widehat{\pi} = \{\widehat{\pi}_h\}_{h=1}^H$ in eq. (15), we have

$$\text{SubOpt}(\widehat{\pi}, s) = - \sum_{h=1}^H \mathbb{E}_{\widehat{\pi}} [\delta_h(s_h, a_h) | s_1 = s] + \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) | s_1 = s],$$

where $\delta_h(\cdot)$ is the evaluation error defined as

$$\delta_h(s, a) = (\mathbb{B}_h \widehat{V}_{h+1})(s, a) - \widehat{Q}_h(s, a).$$

To characterize the suboptimality $\text{SubOpt}(\widehat{\pi}, s)$, we provide the following lemma to bound $\delta_h(\cdot)$ in the linear MDP setting.

Lemma 3. Let $\lambda_1, \lambda_2 = 1$, and let $\beta_1 = C_{\beta_1} H \sqrt{dH \log(N/\delta)}$ and $\beta_2 = C_{\beta_2} dH^2 \sqrt{\log(dH^3 N^{5/2}/\delta)}$, where C_{β_1}, C_{β_2} are two absolute constants. Suppose Assumption 2 holds. With probability at least $1 - \delta/2$, it holds for all $h \in [H]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$0 \leq \delta_h(x) \leq 2[\beta_1 \cdot b_{r,h}(x) + \beta_2 \cdot b_{v,h}(x)], \quad \forall x \in \mathcal{X}, \quad \forall h \in [H].$$

Proof. The main technical development here lies in handling additional challenges caused by the reward redistribution of trajectory-wise rewards, which are not present in linear MDPs with instantaneous rewards Jin et al. (2021). The detailed proof is provided in Appendix F. \square

Applying Lemma 3 to eq. (15), we can obtain

$$\text{SubOpt}(\hat{\pi}, s) \leq 2\beta_1 \cdot \sum_{h=1}^H b_{r,h}(x) + 2\beta_2 \cdot \sum_{h=1}^H b_{v,h}(x). \quad (47)$$

Then, following steps similar to those in Appendix E, we have the followings hold with probability at least $1 - \delta/2$

$$b_{r,h}(x) \leq \frac{C'}{\sqrt{N}} \quad \text{and} \quad b_{v,h}(x) \leq \frac{C''}{\sqrt{N}}, \quad (48)$$

where C' and C'' are absolute constants dependent only on C_σ, C_ς and $\log(1/\delta)$. Then, substituting eq. (48) into eq. (47), we have the following holds with probability $1 - \delta$

$$\begin{aligned} \text{SubOpt}(\hat{\pi}, s) &\leq 2C_{\beta_1} H^2 \sqrt{dH \log(N/\delta)} \cdot \frac{C'}{\sqrt{N}} + 2C_{\beta_2} dH^3 \sqrt{\log(dH^3 N^{5/2}/\delta)} \cdot \frac{C''}{\sqrt{N}} \\ &\leq \mathcal{O} \left(\frac{dH^3}{\sqrt{N}} \sqrt{\log \left(\frac{dH^3 N^{5/2}}{\delta} \right)} \right), \end{aligned}$$

which completes the proof.

D PROOF OF LEMMA 1

Recall that in Section 2.3 we let (b_0, w_0) be the initial value of network parameters obtained via the symmetric initialization scheme, which makes $f(\cdot; w_0)$ a zero function. We denote $(\mathbb{B}_h \hat{V}_{h+1})(\cdot) = \hat{R}_h(\cdot) + (\hat{\mathbb{P}}_h \hat{V}_{h+1})(\cdot)$ as the estimator of Bellman operator $(\mathbb{B}_h \hat{V}_{h+1})(\cdot) = R_h(\cdot) + (\mathbb{P}_h \hat{V}_{h+1})(\cdot)$. To prove Lemma 1, we show that $(\mathbb{B}_h \hat{V}_{h+1})(\cdot) - \beta_1 b_{r,h}(\cdot, \hat{\Theta}) - \beta_2 b_{v,h}(\cdot, \hat{w})$ is approximately a pessimistic estimator of $(\mathbb{B}_h \hat{V}_{h+1})(\cdot)$ up to a function approximation error. We consider m to be sufficiently large such that $m \geq NH^2$.

D.1 UNCERTAINTY OF ESTIMATED REWARD $\hat{R}_h(\cdot)$

In this step, we aim to bound the estimation error $|\hat{R}_h(\cdot) - R_h(\cdot)|$. Since $\hat{\Theta}$ is the global minimizer of the loss function L_r defined in eq. (4), we have

$$\begin{aligned} L_r(\hat{\Theta}) &= \sum_{\tau \in \mathcal{D}} \left[\sum_{h=1}^H f(x_h^\tau, \hat{\theta}_h) - r(\tau) \right]^2 + \lambda_1 \cdot \sum_{h=1}^H \|\hat{\theta}_h - \theta_0\|_2^2 \\ &\leq L_r(\Theta_0) = \sum_{\tau \in \mathcal{D}} \left[\sum_{h=1}^H f(x_h^\tau, \theta_0) - r(\tau) \right]^2 \stackrel{(i)}{=} \sum_{\tau \in \mathcal{D}} [r(\tau)]^2 \stackrel{(ii)}{\leq} NH^2, \end{aligned} \quad (49)$$

where (i) follows from the fact that $f(x, \theta_0) = 0$ for all $x \in \mathcal{X}$ and (ii) follows from the fact that $r(\tau) \leq H$ for any trajectory τ and we have total N trajectories in the offline sample set \mathcal{D} . We define the vector $\Theta_0 = [\theta_0^\top, \dots, \theta_0^\top]^\top \in \mathbb{R}^{2mdH}$. Note that eq. (49) implies

$$\|\hat{\theta}_h - \theta_0\|_2^2 \leq \|\hat{\Theta} - \Theta_0\|_2^2 = \sum_{h=1}^H \|\hat{\theta}_h - \theta_0\|_2^2 \leq NH^2/\lambda_1, \quad \forall h \in [H]. \quad (50)$$

Hence, each $\hat{\theta}_h$ belongs to the Euclidean ball $\mathcal{B}_\theta = \{\theta \in \mathbb{R}^{2md} : \|\theta - \theta_0\|_2 \leq H\sqrt{N/\lambda_1}\}$.

Since the radius of \mathcal{B}_θ does not depend on m , when m is sufficient large it can be shown that $f(\cdot, \theta)$ is close to its linearization at θ_0 , i.e.,

$$f(\cdot, \theta) \approx \langle \phi(\cdot, \theta_0), \theta - \theta_0 \rangle, \quad \forall \theta \in \mathcal{B}_\theta,$$

where $\phi(\cdot, \theta) = \nabla_\theta f(\cdot, \theta)$. Furthermore, according to Assumption 2, there exists a function $\ell_{a_1, a_2} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that the mean of the true reward function $R_h(\cdot) = \mathbb{E}[r_h(\cdot)]$ satisfies

$$R_h(x) = \int_{\mathbb{R}^d} \sigma'(\theta^\top x) \cdot x^\top \ell_r(\theta) dp(\theta), \quad (51)$$

where $\sup_\theta \|\ell_r(\theta)\|_2 \leq a_1$, $\sup_\theta (\|\ell_r(\theta)\|_2 / p(\theta)) \leq a_2$ and p is the density of the distribution $N(0, I_d/d)$. We then proceed to bound the difference between $\hat{R}_h(\cdot)$ and $R_h(\cdot)$.

Step I. In the first step, we show that with high probability the mean of the true reward $R_h(\cdot)$ can be well-approximated by a linear function with the feature vector $\phi(\cdot, \theta_0)$. Lemma 4 in Appendix G implies that that $R_h(\cdot)$ in eq. (51) can be well-approximated by a finite-width neural network, i.e., with probability at least $1 - N^{-2}H^{-4}$ over the randomness of initialization θ_0 , for all $h \in [H]$, there exists a function $\tilde{R}_h(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ satisfying

$$\sup_{x \in \mathcal{X}} \left| \tilde{R}_h(x) - R_h(x) \right| \leq \frac{2(L_\sigma a_2 + C_\sigma^2 a_2^2) \sqrt{\log(N^2 H^5)}}{\sqrt{m}}, \quad (52)$$

where $\tilde{R}_h(\cdot)$ can be written as

$$\tilde{R}_h(x) = \frac{1}{\sqrt{m}} \sum_{r=1}^m \sigma'(\theta_{0,r}^\top x) \cdot x^\top \ell_r,$$

where $\|\ell_r\|_2 \leq a_2/\sqrt{dm}$ for all $r \in [m]$ and $\theta_0 = [\theta_{0,1}, \dots, \theta_{0,m}]$ is generated via the symmetric initialization scheme. We next proceed to show that there exists a vector $\tilde{\theta}_h \in \mathbb{R}^{2md}$ such that $\tilde{R}_h(\cdot) = \langle \phi(\cdot, \theta_0), \tilde{\theta}_h - \theta_0 \rangle$. Let $\tilde{\theta}_h = [\tilde{\theta}_{h,1}^\top, \dots, \tilde{\theta}_{h,2m}^\top]^\top$, in which $\tilde{\theta}_{h,r}^\top = \theta_{0,r} + b_{0,r} \cdot \ell_r / \sqrt{2}$ for all $r \in \{1, \dots, m\}$ and $\tilde{\theta}_{h,r}^\top = \theta_{0,r} + b_{0,r} \cdot \ell_{r-m} / \sqrt{2}$ for all $r \in \{m+1, \dots, 2m\}$. Then, we have

$$\begin{aligned} \tilde{R}_h(x) &= \frac{1}{\sqrt{2m}} \sum_{r=1}^m \sqrt{2}(b_{0,r})^2 \cdot \sigma'(\theta_{0,r}^\top x) \cdot x^\top \ell_r \\ &= \frac{1}{\sqrt{2m}} \sum_{r=1}^m \frac{1}{\sqrt{2}} (b_{0,r})^2 \cdot \sigma'(\theta_{0,r}^\top x) \cdot x^\top \ell_r + \frac{1}{\sqrt{2m}} \sum_{r=1}^m \frac{1}{\sqrt{2}} (b_{0,r})^2 \cdot \sigma'(\theta_{0,r}^\top x) \cdot x^\top \ell_{r-m} \\ &= \frac{1}{\sqrt{2m}} \sum_{r=1}^{2m} b_{0,r} \cdot \sigma'(\theta_{0,r}^\top x) \cdot x^\top (\tilde{\theta}_{h,r} - \theta_{0,r}) \\ &= \phi(x, \theta_0)^\top (\tilde{\theta}_h - \theta_0). \end{aligned} \quad (53)$$

Thus, the true mean reward $R_h(\cdot)$ is approximately linear with the feature $\phi(\cdot, \theta_0)$. Since $\tilde{\theta}_{h,r} - \theta_{0,r} = b_{0,r} \cdot \ell_r / \sqrt{2}$ or $b_{0,r} \cdot \ell_{r-m} / \sqrt{2}$, we have

$$\left\| \tilde{\theta}_h - \theta_0 \right\|_2 \leq a_2 \sqrt{2dm}.$$

Step II. In this step, we show that $\hat{R}_h(\cdot)$ learned by neural network in Algorithm 1 can be well-approximated by its counterpart learned by a linear function with feature $\phi(\cdot, \theta_0)$.

Consider the following least-square loss function

$$\begin{aligned} \bar{L}_r(\Theta) &= \sum_{\tau \in \mathcal{D}} \left[\sum_{h=1}^H \langle \phi(x_h^\tau, \theta_0), \theta_h - \theta_0 \rangle - r(\tau) \right]^2 + \lambda_1 \cdot \sum_{h=1}^H \|\theta_h - \theta_0\|_2^2 \\ &= \sum_{\tau \in \mathcal{D}} [\langle \Phi(\tau, \Theta), \Theta - \Theta_0 \rangle - r(\tau)]^2 + \lambda_1 \cdot \|\Theta - \Theta_0\|_2^2. \end{aligned}$$

The global minimizer of $\bar{L}_r(\Theta)$ is defined as

$$\bar{\Theta} = \underset{\Theta \in \mathbb{R}^{2dmH}}{\operatorname{argmin}} \bar{L}_r(\Theta) \text{ and } \bar{\Theta} = [\bar{\theta}_1^\top, \dots, \bar{\theta}_H^\top]^\top. \quad (54)$$

We define $\bar{R}_h(\cdot) = \langle \phi(\cdot, \theta_0), \bar{\theta}_h - \theta_0 \rangle$ for all $h \in [H]$. We then proceed to bound the term $|\hat{R}_h(x) - \bar{R}_h(x)|$ as follows

$$\begin{aligned} |\hat{R}_h(x) - \bar{R}_h(x)| &= |f(x, \hat{\theta}_h) - \langle \phi(x, \theta_0), \bar{\theta}_h - \theta_0 \rangle| \\ &= |f(x, \hat{\theta}_h) - \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \Theta_0 \rangle| \\ &= |f(x, \hat{\theta}_h) - \langle \Phi_h(x, \Theta_0), \hat{\Theta} - \Theta_0 \rangle + \langle \Phi_h(x, \Theta_0), \hat{\Theta} - \bar{\Theta} \rangle| \\ &\leq |f(x, \hat{\theta}_h) - \langle \Phi_h(x, \Theta_0), \hat{\Theta} - \Theta_0 \rangle| + |\langle \Phi_h(x, \Theta_0), \hat{\Theta} - \bar{\Theta} \rangle| \\ &= |f(x, \hat{\theta}_h) - \langle \phi_h(x, \theta_0), \hat{\theta}_h - \theta_0 \rangle| + |\langle \Phi_h(x, \Theta_0), \hat{\Theta} - \bar{\Theta} \rangle| \\ &\leq |f(x, \hat{\theta}_h) - \langle \phi_h(x, \theta_0), \hat{\theta}_h - \theta_0 \rangle| + \|\Phi_h(x, \Theta_0)\|_2 \|\hat{\Theta} - \bar{\Theta}\|_2 \\ &= \underbrace{|f(x, \hat{\theta}_h) - \langle \phi_h(x, \theta_0), \hat{\theta}_h - \theta_0 \rangle|}_{(i)} + \underbrace{\|\phi(x, \theta_0)\|_2 \|\hat{\Theta} - \bar{\Theta}\|_2}_{(ii)}. \end{aligned}$$

According to Lemma 5 and the fact that $\|\hat{\theta}_h - \theta_0\|_2 \leq H\sqrt{N/\lambda_1}$, we have the followings hold with probability at least $1 - N^{-2}H^{-4}$

$$(i) \leq \mathcal{O}\left(C_\phi \left(\frac{N^2 H^4}{\lambda_1^2 \sqrt{m}}\right)^{1/3} \sqrt{\log m}\right), \quad (55)$$

$$(ii) \leq C_\phi \|\hat{\Theta} - \bar{\Theta}\|_2. \quad (56)$$

We then proceed to bound the term $\|\hat{\Theta} - \bar{\Theta}\|_2$. Consider the minimization problem defined in eq. (5) and eq. (54). By the first order optimality condition, we have

$$\lambda_1 (\hat{\Theta} - \Theta_0) = \sum_{\tau \in \mathcal{D}} \left(r(\tau) - \sum_{h=1}^H f(x_h^\tau, \hat{\theta}_h) \right) \Phi(\tau, \hat{\Theta}) \quad (57)$$

$$\lambda_1 (\bar{\Theta} - \Theta_0) = \sum_{\tau \in \mathcal{D}} (r(\tau) - \langle \Phi(\tau, \Theta_0), \bar{\Theta} - \Theta_0 \rangle) \Phi(\tau, \Theta_0). \quad (58)$$

Note that eq. (58) implies

$$\Sigma(\Theta_0) (\bar{\Theta} - \Theta_0) = \sum_{\tau \in \mathcal{D}} r(\tau) \Phi(\tau, \Theta_0). \quad (59)$$

Adding the term $\sum_{\tau \in \mathcal{D}} \langle \Phi(\tau, \Theta_0), \hat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0)$ on both sides of eq. (57) yields

$$\begin{aligned} \Sigma(\Theta_0) (\hat{\Theta} - \Theta_0) &= \sum_{\tau \in \mathcal{D}} r(\tau) \Phi(\tau, \hat{\Theta}) \\ &\quad + \sum_{\tau \in \mathcal{D}} \left[\langle \Phi(\tau, \Theta_0), \hat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \hat{\theta}_h) \right) \Phi(\tau, \hat{\Theta}) \right]. \quad (60) \end{aligned}$$

Then, subtracting eq. (59) from eq. (60), we can obtain

$$\begin{aligned} \Sigma(\Theta_0) (\hat{\Theta} - \bar{\Theta}) &= \sum_{\tau \in \mathcal{D}} r(\tau) \left(\Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0) \right) \\ &\quad + \sum_{\tau \in \mathcal{D}} \left[\langle \Phi(\tau, \Theta_0), \hat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \hat{\theta}_h) \right) \Phi(\tau, \hat{\Theta}) \right], \quad (61) \end{aligned}$$

which implies

$$\begin{aligned} \left\| \Sigma(\Theta_0)(\widehat{\Theta} - \bar{\Theta}) \right\|_2 &\leq \sum_{\tau \in \mathcal{D}} r(\tau) \sqrt{\sum_{h=1}^H \left\| \phi(x_h^\tau, \theta_0) - \phi(x_h^\tau, \widehat{\theta}_h) \right\|_2^2} \\ &\quad + \sum_{\tau \in \mathcal{D}} \left\| \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \right\|_2. \end{aligned} \quad (62)$$

To bound the term $\left\| \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \right\|_2$, we proceed as follows

$$\begin{aligned} &\langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \\ &= \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle (\Phi(\tau, \Theta_0) - \Phi(\tau, \widehat{\Theta})) - \left(\langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle - \sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \\ &= \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle (\Phi(\tau, \Theta_0) - \Phi(\tau, \widehat{\Theta})) - \left[\sum_{h=1}^H \left(\langle \phi(x_h^\tau, \theta_0), \widehat{\theta}_h - \theta_h^0 \rangle - f(x_h^\tau, \widehat{\theta}_h) \right) \right] \Phi(\tau, \widehat{\Theta}), \end{aligned}$$

which implies

$$\begin{aligned} &\left\| \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \right\|_2 \\ &\leq \|\Phi(\tau, \Theta_0)\|_2 \|\widehat{\Theta} - \Theta_0\|_2 \|\Phi(\tau, \Theta_0) - \Phi(\tau, \widehat{\Theta})\|_2 \\ &\quad + \left[\sum_{h=1}^H \left| \langle \phi(x_h^\tau, \theta_0), \widehat{\theta}_h - \theta_h^0 \rangle - f(x_h^\tau, \widehat{\theta}_h) \right| \right] \|\Phi(\tau, \widehat{\Theta})\|_2, \\ &= \sqrt{\sum_{h=1}^H \|\phi(x_h^\tau, \theta_h^0)\|_2^2} \|\widehat{\Theta} - \Theta_0\|_2 \sqrt{\sum_{h=1}^H \|\phi(x_h^\tau, \theta_0) - \phi(x_h^\tau, \widehat{\theta}_h)\|_2^2} \\ &\quad + \left[\sum_{h=1}^H \left| \langle \phi(x_h^\tau, \theta_0), \widehat{\theta}_h - \theta_h^0 \rangle - f(x_h^\tau, \widehat{\theta}_h) \right| \right] \sqrt{\sum_{h=1}^H \|\phi(x_h^\tau, \widehat{\theta}_h)\|_2^2}, \end{aligned} \quad (63)$$

where the last equality follows from the fact that $\|\Phi(\tau, \Theta)\|_2^2 = \sum_{h=1}^H \|\phi(x_h^\tau, \theta_h)\|_2^2$ for any $\Theta \in \mathbb{R}^{2mdH}$. According to Lemma 5 and the fact that $\|\widehat{\theta}_h - \theta_h^0\|_2 \leq H\sqrt{N/\lambda_1}$, we have the followings hold with probability at least $1 - N^{-2}H^{-4}$ for all $h \in [H]$ and $\tau \in \mathcal{D}$

$$\|\phi(x_h^\tau, \theta_0)\|_2 \leq C_\phi \quad \text{and} \quad \|\phi(x_h^\tau, \widehat{\theta}_h)\|_2 \leq C_\phi, \quad (64)$$

$$\left\| \phi(x_h^\tau, \theta_0) - \phi(x_h^\tau, \widehat{\theta}_h) \right\|_2 \leq \mathcal{O} \left(C_\phi \left(\frac{H\sqrt{N/\lambda_1}}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right), \quad (65)$$

$$\left| \langle \phi(x_h^\tau, \theta_0), \widehat{\theta}_h - \theta_h^0 \rangle - f(x_h^\tau, \widehat{\theta}_h) \right| \leq \mathcal{O} \left(C_\phi \left(\frac{H^4 N^2 / \lambda_1^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right). \quad (66)$$

Substituting eq. (64), eq. (65) and eq. (66) into eq. (63), we have

$$\left\| \langle \Phi(\tau, \Theta_0), \widehat{\Theta} - \Theta_0 \rangle \Phi(\tau, \Theta_0) - \left(\sum_{h=1}^H f(x_h^\tau, \widehat{\theta}_h) \right) \Phi(\tau, \widehat{\Theta}) \right\|_2$$

$$\begin{aligned}
&\leq (H^2 \sqrt{N/\lambda_1}) \mathcal{O} \left(C_\phi^2 \left(\frac{H \sqrt{N/\lambda_1}}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \mathcal{O} \left(C_\phi^2 H^{3/2} \left(\frac{H^4 N^2 / \lambda_1^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) \\
&\leq \mathcal{O} \left(\frac{C_\phi^2 H^{17/6} N^{2/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{2/3}} \right). \tag{67}
\end{aligned}$$

Then, substituting eq. (67) into eq. (62), we have the following holds with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned}
&\left\| \Sigma(\Theta_0)(\hat{\Theta} - \bar{\Theta}) \right\|_2 \\
&\leq NH \cdot \sqrt{H} \cdot \mathcal{O} \left(C_\phi \left(\frac{H \sqrt{N/\lambda_1}}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + N \cdot \mathcal{O} \left(\frac{C_\phi^2 H^{17/6} N^{2/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{2/3}} \right) \\
&\leq \mathcal{O} \left(\frac{C_\phi^2 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{2/3}} \right),
\end{aligned}$$

where we use the fact that $r(\tau) \leq H$. We then proceed to bound the term $\left\| \hat{\Theta} - \Theta_0 \right\|_2$ as follows

$$\begin{aligned}
\left\| \hat{\Theta} - \bar{\Theta} \right\|_2 &= \left\| \Sigma^{-1}(\Theta_0) \Sigma(\Theta_0)(\hat{\Theta} - \bar{\Theta}) \right\|_2 \\
&\leq \left\| \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \Sigma(\Theta_0)(\hat{\Theta} - \bar{\Theta}) \right\|_2 \\
&\leq \lambda_1^{-1} \left\| \Sigma(\Theta_0)(\hat{\Theta} - \bar{\Theta}) \right\|_2 \\
&\leq \mathcal{O} \left(\frac{C_\phi^2 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{5/3}} \right). \tag{68}
\end{aligned}$$

Substituting eq. (68) into eq. (56), we can bound (ii) as follows

$$(ii) \leq \mathcal{O} \left(\frac{C_\phi^3 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{5/3}} \right). \tag{69}$$

Taking summation of the upper bounds of (i) in eq. (55) and (ii) in eq. (69), we have

$$\begin{aligned}
\left| \hat{R}_h(x) - \bar{R}_h(x) \right| &\leq (i) + (ii) \\
&\leq \mathcal{O} \left(C_\phi \left(\frac{N^2 H^4}{\lambda_1^2 \sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \mathcal{O} \left(\frac{C_\phi^3 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{5/3}} \right) \\
&\leq \mathcal{O} \left(\frac{C_\phi^3 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{5/3}} \right). \tag{70}
\end{aligned}$$

Step III. In this step, we show that the bonus term $b_{r,h}(\cdot, \hat{\Theta})$ in Algorithm 1 can be well approximated by $b_{r,h}(\cdot, \Theta_0)$. According to the definition of $b_{r,h}(\cdot, \Theta)$, we have

$$\begin{aligned}
&\left| b_{r,h}(x, \hat{\Theta}) - b_{r,h}(x, \Theta_0) \right| \\
&= \left| \left[\Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) \right]^{1/2} - \left[\Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0) \right]^{1/2} \right| \\
&\leq \left| \Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0) \right|^{1/2}, \tag{71}
\end{aligned}$$

where the last inequality follows from the fact that $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$. We then proceed to bound the term $\left| \Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0) \right|$ as follows

$$\left| \Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0) \right|$$

$$\begin{aligned}
&= \left| [\Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)]^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) \right| + \left| \Phi_h(x, \Theta_0)^\top (\Sigma^{-1}(\hat{\Theta}) - \Sigma^{-1}(\Theta_0)) \Phi_h(x, \hat{\Theta}) \right| \\
&\quad + \left| \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) (\Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)) \right| \\
&\leq \left| \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0) \right|^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) + \left| \Phi_h(x, \Theta_0)^\top (\Sigma^{-1}(\hat{\Theta}) - \Sigma^{-1}(\Theta_0)) \Phi_h(x, \hat{\Theta}) \right| \\
&\quad + \left| \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) (\Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)) \right| \\
&\leq \left\| \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) \right\|_2 \left\| \Phi_h(x, \hat{\Theta}) \right\|_2 \\
&\quad + \left\| \Phi_h(x, \Theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) - \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \Phi_h(x, \hat{\Theta}) \right\|_2 \\
&\quad + \left\| \Phi_h(x, \Theta_0) \right\|_2 \left\| \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0) \right\|_2 \\
&= \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\
&\quad + \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) (\Sigma(\hat{\Theta}) - \Sigma(\Theta_0)) \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\
&\quad + \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \\
&\leq \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\
&\quad + \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\hat{\Theta}) \right\|_2 \left\| \Sigma(\hat{\Theta}) - \Sigma(\Theta_0) \right\|_2 \left\| \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\
&\quad + \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma^{-1}(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \\
&\leq \frac{1}{\lambda_1} \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 + \frac{1}{\lambda_1^2} \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma(\hat{\Theta}) - \Sigma(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\
&\quad + \frac{1}{\lambda_1} \left\| \phi(x, \theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2, \tag{72}
\end{aligned}$$

where the last inequality follows from the fact that $\|\Sigma(\Theta)\|_2 \geq \lambda_1$ for any $\Theta \in \mathbb{R}^{2mdH}$. For $\Sigma(\hat{\Theta}) - \Sigma(\Theta_0)$, we have

$$\begin{aligned}
\Sigma(\hat{\Theta}) - \Sigma(\Theta_0) &= \sum_{\tau \in \mathcal{D}} \left[\Phi(\tau, \hat{\Theta}) \Phi(\tau, \hat{\Theta})^\top - \Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top \right] \\
&= \sum_{\tau \in \mathcal{D}} \left[\Phi(\tau, \hat{\Theta}) (\Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0))^\top + (\Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0)) \Phi(\tau, \Theta_0)^\top \right],
\end{aligned}$$

which implies

$$\begin{aligned}
&\left\| \Sigma(\hat{\Theta}) - \Sigma(\Theta_0) \right\|_2 \\
&\leq \sum_{\tau \in \mathcal{D}} \left[\left\| \Phi(\tau, \hat{\Theta}) \right\|_2 \left\| \Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0) \right\|_2 + \left\| \Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0) \right\|_2 \left\| \Phi(\tau, \Theta_0) \right\|_2 \right]. \tag{73}
\end{aligned}$$

By definition of $\Phi(\tau, \Theta)$, we have the followings hold for any $\Theta, \tilde{\Theta} \in \mathbb{R}^{2mdH}$

$$\left\| \Phi(\tau, \tilde{\Theta}) \right\|_2 = \sqrt{\sum_{h \in [H]} \left\| \phi(x_h^\tau, \tilde{\theta}_h) \right\|_2^2}, \tag{74}$$

$$\left\| \Phi(\tau, \tilde{\Theta}) - \Phi(\tau, \Theta_0) \right\|_2 = \sqrt{\sum_{h \in [H]} \left\| \phi(x_h^\tau, \theta_h) - \phi(x_h^\tau, \tilde{\theta}_h) \right\|_2^2}. \tag{75}$$

Applying Lemma 5 to eq. (74) and eq. (75), we have the followings hold with probability at least $1 - N^{-2}H^{-4}$

$$\left\| \Phi(\tau, \tilde{\Theta}) \right\|_2 \leq C_\phi \sqrt{H},$$

$$\left\| \Phi(\tau, \hat{\Theta}) - \Phi(\tau, \Theta_0) \right\|_2 \leq \mathcal{O} \left(\frac{C_\phi H^{5/6} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_1^{1/6}} \right).$$

Substituting the above two inequalities into eq. (73) yields

$$\left\| \Sigma(\hat{\Theta}) - \Sigma(\Theta_0) \right\|_2 \leq \mathcal{O} \left(\frac{C_\phi^2 H^{4/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_1^{1/6}} \right). \quad (76)$$

Finally, combining eq. (76) and eq. (181) and eq. (182) in Lemma 5, we can bound the right hand side of eq. (72) as

$$\begin{aligned} & \left| \Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \bar{\Theta})^\top \Sigma^{-1}(\bar{\Theta}) \Phi_h(x, \bar{\Theta}) \right| \\ & \leq \frac{1}{\lambda_1} \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 + \frac{1}{\lambda_1^2} \left\| \phi(x, \theta_0) \right\|_2 \left\| \Sigma(\hat{\Theta}) - \Sigma(\Theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) \right\|_2 \\ & \quad + \frac{1}{\lambda_1} \left\| \phi(x, \theta_0) \right\|_2 \left\| \phi(x, \hat{\theta}_h) - \phi(x, \theta_0) \right\|_2 \\ & \leq \mathcal{O} \left(\frac{C_\phi^2 H^{1/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_1^{7/6}} \right) + \mathcal{O} \left(\frac{C_\phi^4 H^{4/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_1^{13/6}} \right) \\ & = \mathcal{O} \left(\frac{C_\phi^4 H^{4/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_1^{13/6}} \right). \end{aligned} \quad (77)$$

Substituting eq. (77) into eq. (71), we have the following holds with probability at least $1 - N^{-2} H^{-4}$

$$\begin{aligned} & \left| b_{r,h}(x, \hat{\Theta}) - b_{r,h}(x, \Theta_0) \right| \\ & \leq \left| \Phi_h(x, \hat{\Theta})^\top \Sigma^{-1}(\hat{\Theta}) \Phi_h(x, \hat{\Theta}) - \Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0) \right|^{1/2} \\ & \leq \mathcal{O} \left(\frac{C_\phi^2 H^{2/3} N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_1^{13/12}} \right). \end{aligned} \quad (78)$$

Step IV. In Steps I and II, we show that the mean of the real reward $R_h(\cdot)$ can be well approximated by a linear function $\tilde{R}_h(\cdot)$ with feature $\phi(\cdot, \theta_0)$ and our learned reward $\hat{R}_h(\cdot)$ can be well approximated by a linear function $\bar{R}_h(\cdot)$ with feature $\phi(\cdot, \theta_0)$. In this step, we want to show that the reward estimation error $\left| \hat{R}_h(\cdot) - R_h(\cdot) \right|$ is approximately $\beta_1 \cdot b_{r,h}(x, \Theta_0)$ with an approximately chosen β_1 .

Recall that $\tilde{R}_h(\cdot) = \langle \phi(\cdot, \theta_0), \tilde{\theta}_h - \theta_0 \rangle$ and $\bar{R}_h(\cdot) = \langle \phi(\cdot, \theta_0), \bar{\theta}_h - \theta_0 \rangle$. Considering the difference between $\bar{R}_h(\cdot)$ and $\tilde{R}_h(\cdot)$, we have

$$\bar{R}_h(x) - \tilde{R}_h(x) = \langle \phi(x, \theta_0), \bar{\theta}_h - \tilde{\theta}_h \rangle = \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \tilde{\Theta} \rangle, \quad (79)$$

where the last equality follows from the definition of $\Phi_h(\cdot, \Theta)$. By eq. (59), we have

$$\bar{\Theta} - \Theta_0 = \Sigma(\Theta_0)^{-1} \sum_{\tau \in \mathcal{D}} r(\tau) \Phi(\tau, \Theta_0). \quad (80)$$

By the definition of $\Sigma(\Theta)$, we have

$$\tilde{\Theta} - \Theta_0 = \Sigma(\Theta_0)^{-1} \left[\lambda_1 (\tilde{\Theta} - \Theta_0) + \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top \right) (\tilde{\Theta} - \Theta_0) \right]. \quad (81)$$

Subtracting eq. (81) from eq. (80), we have

$$\bar{\Theta} - \tilde{\Theta} = -\lambda_1 \Sigma(\Theta_0)^{-1} (\tilde{\Theta} - \Theta_0) + \Sigma(\Theta_0)^{-1} \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \left[r(\tau) - \langle \Phi(\tau, \Theta_0), \tilde{\Theta} - \Theta_0 \rangle \right]. \quad (82)$$

Taking inter product of both sides of eq. (82) with vector $\Phi_h(x, \Theta_0)$ and using the fact that $R(\tau) = \sum_{h \in [H]} R_h(x_h^\tau)$ and $\langle \Phi(\tau, \Theta_0), \tilde{\Theta} - \Theta_0 \rangle = \sum_{h \in [H]} \langle \phi(x_h^\tau, \theta_0), \tilde{\theta}_h - \theta_0 \rangle$, we have

$$\begin{aligned} & \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \tilde{\Theta} \rangle \\ &= -\lambda_1 \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \Sigma(\Theta_0)^{-1/2} (\tilde{\Theta} - \Theta_0) \\ & \quad + \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \Sigma(\Theta_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) (r(\tau) - R(\tau)) \right) \\ & \quad + \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \Sigma(\Theta_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \left[\sum_{h \in [H]} (R_h(x_h^\tau) - \langle \phi(x_h^\tau, \theta_0), \tilde{\theta}_h - \theta_0 \rangle) \right] \right). \end{aligned} \quad (83)$$

Recall that $\tilde{R}_h(x_h^\tau) = \langle \phi(x_h^\tau, \theta_0), \tilde{\theta}_h - \theta_0 \rangle$, and eq. (83) implies that

$$\begin{aligned} & \left| \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \tilde{\Theta} \rangle \right| \\ & \leq \sqrt{\lambda_1} \left\| \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \right\|_2 \left\| \tilde{\Theta} - \Theta_0 \right\|_2 \\ & \quad + \left\| \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \right\|_2 \left\| \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \varepsilon(\tau) \right\|_{\Sigma(\Theta_0)^{-1}} \\ & \quad + \frac{1}{\sqrt{\lambda_1}} \left\| \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \right\|_2 \left(\sum_{\tau \in \mathcal{D}} \|\Phi(\tau, \Theta_0)\|_2 \sum_{h \in [H]} |R_h(x_h^\tau) - \tilde{R}_h(x_h^\tau)| \right), \end{aligned} \quad (84)$$

where we denote $\varepsilon(\tau) = r(\tau) - R(\tau)$ and use the fact that $\|\Sigma(\Theta)^{-1/2}\|_2 \leq 1/\sqrt{\lambda_1}$ for any $\Theta \in \mathbb{R}^{2mdH}$. By the definition of $\tilde{\Theta}$ in Step I, we have

$$\begin{aligned} \left\| \tilde{\Theta} - \Theta_0 \right\|_2 &= \sqrt{\sum_{h \in [H], r \in [m]} \left\| \tilde{\theta}_{h,r} - \theta_{h,r}^0 \right\|_2^2} \\ &= \sqrt{\sum_{h \in [H], r \in [m]} \|\ell_r\|_2^2} \\ &\leq r_2 \sqrt{H/d}. \end{aligned} \quad (85)$$

By Lemma 5 and eq. (52), we have the followings hold with probability at least $1 - N^{-2}H^{-4}$

$$\|\Phi(\tau, \Theta_0)\|_2 \leq C_\phi \sqrt{H}, \quad (86)$$

$$\left| R_h(x_h^\tau) - \tilde{R}_h(x_h^\tau) \right| \leq \frac{2(L_\sigma a_2 + C_\sigma^2 a_2^2) \sqrt{\log N^2 H^5}}{\sqrt{m}}. \quad (87)$$

Substituting eq. (85), eq. (86) and eq. (87) into eq. (84) and using the fact that $b_{r,h}(x, \Theta_0) = \left\| \Phi_h(x, \Theta_0)^\top \Sigma(\Theta_0)^{-1/2} \right\|_2$, we have

$$\begin{aligned} & \left| \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \tilde{\Theta} \rangle \right| \\ & \leq \left(a_2 \sqrt{\frac{\lambda_1 H}{d}} + \frac{2(L_\sigma a_2 + C_\sigma^2 a_2^2) C_\phi N H^{3/2} \sqrt{\log H N}}{\sqrt{\lambda_1 m}} + \left\| \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \varepsilon(\tau) \right\|_{\Sigma(\Theta_0)^{-1}} \right) b_{r,h}(x, \Theta_0). \end{aligned} \quad (88)$$

Given that the events in eq. (86) and eq. (87) occur, applying eq. (199) in Lemma 8, we have the following holds with probability at least $1 - N^{-2}H^{-4}$

$$\left\| \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \varepsilon(\tau) \right\|_{\Sigma(\Theta_0)^{-1}}^2$$

$$\leq H^2 \log \det(I + K_N^r / \lambda_1) + H^2 N (\lambda_1 - 1) + 4H^2 \log(NH^2), \quad (89)$$

where $K_N^r \in \mathbb{R}^{N \times N}$ is the Gram matrix defined as

$$K_N^r = [K_H(\tau_i, \tau_j)]_{i,j \in [N]} \in \mathbb{R}^{N \times N}.$$

Combining eq. (88) and eq. (89) and letting $\lambda_1 = 1 + N^{-1}$ and m be sufficiently large such that

$$\frac{2(L_\sigma a_2 + C_\sigma^2 a_2^2) C_\phi N H^{3/2} \sqrt{\log HN}}{\sqrt{\lambda_1 m}} \leq a_2 \sqrt{\frac{\lambda_1 H}{d}},$$

we have the following holds with probability at least $1 - N^{-2} H^{-2}$

$$\begin{aligned} & \left| \langle \Phi_h(x, \Theta_0), \bar{\Theta} - \tilde{\Theta} \rangle \right| \\ & \leq \left(2a_2 \sqrt{\frac{\lambda_1 H}{d}} + \sqrt{H^2 \log \det \left(I + \frac{K_N^r}{\lambda_1} \right) + H^2 + 4H^2 \log(NH^2)} \right) b_{r,h}(x, \Theta_0) \\ & \leq H \underbrace{\left(\frac{4a_2^2 \lambda_1}{d} + 2 \log \det \left(I + \frac{K_N^r}{\lambda_1} \right) + 10 \log(NH^2) \right)^{1/2}}_{\beta_1} b_{r,h}(x, \Theta_0), \end{aligned} \quad (90)$$

where in the last inequality we use the fact that $a + b \leq \sqrt{2(a^2 + b^2)}$. Substituting eq. (90) into eq. (79), we have the following holds with probability at least $1 - N^{-2} H^{-4}$

$$\left| \bar{R}_h(x) - \tilde{R}_h(x) \right| \leq \beta_1 \cdot b_{r,h}(x, \Theta_0), \quad (91)$$

where

$$\beta_1 = H \left(\frac{4a_2^2 \lambda_1}{d} + 2 \log \det \left(I + \frac{K_N^r}{\lambda_1} \right) + 10 \log(NH^2) \right)^{1/2}.$$

Next, we proceed to bound the reward estimation error $|R_h(x) - \hat{R}_h(x)|$. By the triangle inequality, we have

$$\begin{aligned} \left| R_h(x) - \hat{R}_h(x) \right| &= \left| R_h(x) - \tilde{R}_h(x) + \tilde{R}_h(x) - \bar{R}_h(x) + \bar{R}_h(x) - \hat{R}_h(x) \right| \\ &\leq \left| R_h(x) - \tilde{R}_h(x) \right| + \left| \tilde{R}_h(x) - \bar{R}_h(x) \right| + \left| \bar{R}_h(x) - \hat{R}_h(x) \right| \\ &\stackrel{(i)}{\leq} \frac{2(L_\sigma a_2 + C_\sigma^2 a_2^2) \sqrt{\log(HN)}}{\sqrt{m}} + \mathcal{O} \left(\frac{C_\phi^3 H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_1^{5/3}} \right) \\ &\quad + \beta_1 \cdot b_{r,h}(x, \Theta_0) \\ &\stackrel{(ii)}{\leq} \mathcal{O} \left(\frac{H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6}} \right) + \beta_1 \cdot b_{r,h}(x, \Theta_0). \end{aligned} \quad (92)$$

where (i) follows from eq. (52) and eq. (70) and (ii) follows from the fact that $\lambda_1 = 1 + 1/N$ and $L_\sigma, C_\sigma, a_2, C_\phi = \mathcal{O}(1)$.

D.2 UNCERTAINTY OF ESTIMATED TRANSITION VALUE FUNCTION $(\hat{\mathbb{P}}_h \hat{V}_{h+1})(\cdot)$

In this subsection, we aim to bound the estimation error of the transition value function $\left| (\hat{\mathbb{P}}_h \hat{V})(\cdot) - (\mathbb{P}_h \hat{V})(\cdot) \right|$. For each $h \in [H]$, since \hat{w}_h is the global minimizer of the loss function $L_v^h(w_h)$ defined in eq. (6), we have

$$L_v^h(\hat{w}_h) = \sum_{\tau \in \mathcal{D}} \left(\hat{V}_{h+1}(s_{h+1}^\tau) - f(x_h^\tau, \hat{w}_h) \right)^2 + \lambda_2 \cdot \|\hat{w}_h - w_0\|_2^2$$

$$\leq L_v^h(\widehat{w}_0) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - f(x_h^\tau, w_0) \right)^2 \stackrel{(i)}{=} \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) \right)^2 \stackrel{(ii)}{\leq} NH^2, \quad (93)$$

where (i) follows from the fact that $f(x, w_0) = 0$ for all $x \in \mathcal{X}$ and (ii) follows from the fact that $\widehat{V}_h(s) \leq H$ for any $h \in [H]$, $s \in \mathcal{S}$, and $|\mathcal{D}| = N$. Note that eq. (93) implies

$$\|\widehat{w}_h - w_0\|_2^2 \leq NH^2/\lambda_2, \quad \forall h \in [H]. \quad (94)$$

Hence, each \widehat{w}_h belongs to the Euclidean ball $\mathcal{B}_w = \{w \in \mathbb{R}^{2md} : \|w - w_0\|_2 \leq H\sqrt{N/\lambda_2}\}$, where λ_2 does not depend on the network width m . Since the radius of \mathcal{B}_w does not depend on m , when m is sufficient large, it can be shown that $f(\cdot, w)$ is close to its linearization at w_0 , i.e.,

$$f(\cdot, w) \approx \langle \phi(\cdot, w_0), w - w_0 \rangle, \quad \forall w \in \mathcal{B}_w,$$

where $\phi(\cdot, w) = \nabla_w f(\cdot, w)$. Furthermore, according to Assumption 2, there exists a function $\ell_{A_1, A_2} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ satisfies

$$(\mathbb{P}_h \widehat{V}_{h+1})(x) = \int_{\mathbb{R}^d} \sigma'(\theta^\top x) \cdot x^\top \ell_v(w) dp(w), \quad (95)$$

where $\sup_w \|\ell_v(w)\|_2 \leq A_1$, $\sup_w (\|\ell_v(w)\|_2 / p(w)) \leq A_2$ and p is the density of $N(0, I_d/d)$. We then proceed to bound the difference between $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ and $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$.

Step I. In the first step, we show that the transition value function $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ can be well-approximated by a linear function with the feature vector $\phi(\cdot, \theta_0)$. Lemma 4 in Appendix G implies that with probability at least $1 - N^{-2}H^{-4}$ over the randomness of initialization w_0 , for all $h \in [H]$, there exists a function $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ satisfying

$$\sup_{x \in \mathcal{X}} \left| (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\mathbb{P}_h \widehat{V}_{h+1})(x) \right| \leq \frac{2(L_\sigma A_2 + C_\sigma^2 A_2^2) \sqrt{\log(N^2 H^5)}}{\sqrt{m}}, \quad (96)$$

where $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ is a finite-width neural network which can be written as

$$(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) = \frac{1}{\sqrt{m}} \sum_{r=1}^m \sigma'(w_{0,r}^\top x) \cdot x^\top \ell_r^v,$$

where $\|\ell_r^v\|_2 \leq A_2/\sqrt{dm}$ for all $r \in [m]$ and $w_0 = [w_{0,1}, \dots, w_{0,m}]$ is generated via the symmetric initialization scheme. Following steps similar to those in eq. (53), we can show that there exists a vector $\tilde{w}_h \in \mathbb{R}^{2md}$ such that

$$(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = \langle \phi(\cdot, w_0), \tilde{w}_h - w_0 \rangle,$$

where $\tilde{w}_h = [\tilde{w}_{h,1}^\top, \dots, \tilde{w}_{h,2m}^\top]^\top$, in which $\tilde{w}_{h,r}^\top = w_{0,r} + b_{0,r} \cdot \ell_r^v / \sqrt{2}$ for all $r \in \{1, \dots, m\}$ and $\tilde{w}_{h,r}^\top = w_{0,r} + b_{0,r} \cdot \ell_{r-m}^v / \sqrt{2}$ for all $r \in \{m+1, \dots, 2m\}$. Moreover, since $\tilde{w}_{h,r} - w_{0,r} = b_{0,r} \cdot \ell_r^v / \sqrt{2}$ or $b_{0,r} \cdot \ell_{r-m}^v / \sqrt{2}$, we have

$$\|\tilde{w}_h - w_0\|_2 \leq A_2 \sqrt{2dm}.$$

Step II. In the second step, we show that with high probability, the estimation of the transition value function $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ in Algorithm 1 can be well-approximated by its counterpart learned with a linear function with the feature $\phi(\cdot, \theta_0)$.

Consider the following least-square loss function

$$\bar{L}_v^h(w_h) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - \langle \phi(x_h^\tau, w_0), w_h - w_0 \rangle \right)^2 + \lambda_2 \cdot \|w_h - w_0\|_2^2. \quad (97)$$

The global minimizer of $\bar{L}_v^h(w_h)$ is defined as

$$\bar{w}_h = \operatorname{argmin}_{w \in \mathbb{R}^{2dm}} \bar{L}_v^h(w). \quad (98)$$

We define $(\bar{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = \langle \phi(\cdot, w_0), \bar{w}_h - w_0 \rangle$ for all $h \in [H]$. Then, in a manner similar to the construction of $\widehat{Q}_h(\cdot)$ in Algorithm 1, we combine $\bar{R}_h(\cdot)$ in eq. (54), $b_{r,h}(\cdot, \Theta_0)$, $\bar{\mathbb{P}}_h \widehat{V}_{h+1}(\cdot)$ and $b_{v,h}(\cdot, w_0)$ to construct $\bar{Q}_h(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ as

$$\bar{Q}_h(\cdot) = \min\{\bar{R}_h(\cdot) + (\bar{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) - \beta_1 \cdot b_{r,h}(\cdot, \Theta_0) - \beta_2 \cdot b_{v,h}(\cdot, w_0), H\}^+. \quad (99)$$

Moreover, we define the estimated optimal state value function as

$$\bar{V}_h(\cdot) = \max_{a \in \mathcal{A}} \bar{Q}_h(\cdot, a). \quad (100)$$

We then proceed to bound the estimation error $\left| (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) - \bar{\mathbb{P}}_h \widehat{V}_{h+1}(x) \right|$ as follows

$$\begin{aligned} \left| (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) - \bar{\mathbb{P}}_h \widehat{V}_{h+1}(x) \right| &= |f(x, \widehat{w}_h) - \langle \phi(x, w_0), \bar{w}_h - w_0 \rangle| \\ &= |f(x, \widehat{w}_h) - \langle \phi(x, w_0), \widehat{w}_h - w_0 \rangle - \langle \phi(\cdot, w_0), \widehat{w}_h - \bar{w}_h \rangle| \\ &\leq |f(x, \widehat{w}_h) - \langle \phi(x, w_0), \widehat{w}_h - w_0 \rangle| + |\langle \phi(\cdot, w_0), \widehat{w}_h - \bar{w}_h \rangle| \\ &\leq \underbrace{|f(x, \widehat{w}_h) - \langle \phi(\cdot, w_0), \widehat{w}_h - w_0 \rangle|}_{(i)} + \underbrace{\|\phi(x, w_0)\|_2 \|\widehat{w}_h - \bar{w}_h\|_2}_{(ii)}. \end{aligned}$$

We then bound the term (i) and term (ii) in the above inequality. According to Lemma 5 and the fact that $\|\widehat{w}_h - w_0\|_2 \leq H\sqrt{N/\lambda_2}$, we have the followings hold with probability at least $1 - N^{-2}H^{-4}$

$$(i) \leq \mathcal{O}\left(C_\phi \left(\frac{N^2 H^4}{\lambda_2^2 \sqrt{m}}\right)^{1/3} \sqrt{\log m}\right), \quad (101)$$

$$(ii) \leq C_\phi \|\widehat{w}_h - \bar{w}_h\|_2. \quad (102)$$

We then proceed to bound $\|\widehat{w}_h - \bar{w}_h\|_2$. Consider the minimization problem defined in eq. (5) and eq. (54). By the first order optimality condition, we have

$$\lambda_2 (\widehat{w}_h - w_0) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - f(x_h^\tau, \widehat{w}_h) \right) \phi(x_h^\tau, \widehat{w}_h) \quad (103)$$

$$\lambda_2 (\bar{w}_h - w_0) = \sum_{\tau \in \mathcal{D}} \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - \langle \phi(x_h^\tau, w_0), \bar{w}_h - w_0 \rangle \right) \phi(x_h^\tau, w_0). \quad (104)$$

Note that eq. (104) implies

$$\Lambda_h(w_0) (\bar{w}_h - w_0) = \sum_{\tau \in \mathcal{D}} \widehat{V}_{h+1}(s_{h+1}^\tau) \phi(x_h^\tau, w_0). \quad (105)$$

Adding the term $\sum_{\tau \in \mathcal{D}} \langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0)$ on both sides of eq. (103) yields

$$\begin{aligned} \Lambda_h(w_0) (\widehat{w}_h - w_0) &= \sum_{\tau \in \mathcal{D}} \widehat{V}_{h+1}(s_{h+1}^\tau) \phi(x_h^\tau, \widehat{w}_h) \\ &\quad + \sum_{\tau \in \mathcal{D}} [\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)]. \end{aligned} \quad (106)$$

Then, by subtracting eq. (105) from eq. (106), we have

$$\begin{aligned} \Lambda_h(w_0) (\widehat{w}_h - \bar{w}_h) &= \sum_{\tau \in \mathcal{D}} \widehat{V}_{h+1}(s_{h+1}^\tau) (\phi(x_h^\tau, \widehat{w}_h) - \phi(x_h^\tau, w_0)) \\ &\quad + \sum_{\tau \in \mathcal{D}} [\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)], \end{aligned} \quad (107)$$

which implies

$$\begin{aligned} \|\Lambda_h(w_0) (\widehat{w}_h - \bar{w}_h)\|_2 &= \sum_{\tau \in \mathcal{D}} \widehat{V}_{h+1}(s_{h+1}^\tau) \|\phi(x_h^\tau, \widehat{w}_h) - \phi(x_h^\tau, w_0)\|_2 \\ &\quad + \sum_{\tau \in \mathcal{D}} \|\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)\|_2. \end{aligned} \quad (108)$$

To bound the term $\|\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)\|_2$, we proceed as follows

$$\begin{aligned} & \langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h) \\ &= \langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle (\phi(x_h^\tau, w_0) - \phi(x_h^\tau, \widehat{w}_h)) - (\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle - f(x_h^\tau, \widehat{w}_h)) \phi(x_h^\tau, \widehat{w}_h) \end{aligned}$$

which implies

$$\begin{aligned} & \|\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)\|_2 \\ & \leq \|\phi(x_h^\tau, w_0)\|_2 \|\widehat{w}_h - w_0\|_2 \|\phi(x_h^\tau, w_0) - \phi(x_h^\tau, \widehat{w}_h)\|_2 \\ & \quad + |\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle - f(x_h^\tau, \widehat{w}_h)| \|\phi(x_h^\tau, \widehat{w}_h)\|_2. \end{aligned} \quad (109)$$

According to Lemma 5 and the fact that $\|\widehat{w}_h - w_0\|_2 \leq H\sqrt{N/\lambda_2}$, we have the followings hold for all $h \in [H]$ and $\tau \in \mathcal{D}$ with probability at least $1 - N^{-2}H^{-4}$

$$\|\phi(x_h^\tau, w_0)\|_2 \leq C_\phi \quad \text{and} \quad \|\phi(x_h^\tau, \widehat{w}_h)\|_2 \leq C_\phi, \quad (110)$$

$$\|\phi(x_h^\tau, w_0) - \phi(x_h^\tau, \widehat{w}_h)\|_2 \leq \mathcal{O}\left(C_\phi \left(\frac{H\sqrt{N/\lambda_2}}{\sqrt{m}}\right)^{1/3} \sqrt{\log m}\right), \quad (111)$$

$$|\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle - f(x_h^\tau, \widehat{w}_h)| \leq \mathcal{O}\left(C_\phi \left(\frac{H^4 N^2 / \lambda_2^2}{\sqrt{m}}\right)^{1/3} \sqrt{\log m}\right). \quad (112)$$

Substituting eq. (110), eq. (111) and eq. (112) into eq. (109), we can obtain

$$\begin{aligned} & \|\langle \phi(x_h^\tau, w_0), \widehat{w}_h - w_0 \rangle \phi(x_h^\tau, w_0) - f(x_h^\tau, \widehat{w}_h) \phi(x_h^\tau, \widehat{w}_h)\|_2 \\ & \leq (H\sqrt{N/\lambda_2}) \mathcal{O}\left(C_\phi^2 \left(\frac{H\sqrt{N/\lambda_2}}{\sqrt{m}}\right)^{1/3} \sqrt{\log m}\right) + \mathcal{O}\left(C_\phi^2 \left(\frac{H^4 N^2 / \lambda_2^2}{\sqrt{m}}\right)^{1/3} \sqrt{\log m}\right) \\ & \leq \mathcal{O}\left(\frac{C_\phi^2 H^{4/3} N^{2/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{2/3}}\right). \end{aligned} \quad (113)$$

Substituting eq. (113) into eq. (108), we have the following holds with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned} & \|\Lambda_h(w_0)(\widehat{w}_h - \bar{w}_h)\|_2 \\ & \leq NH \cdot \mathcal{O}\left(C_\phi \left(\frac{H\sqrt{N/\lambda_2}}{\sqrt{m}}\right)^{1/3} \sqrt{\log m}\right) + N \cdot \mathcal{O}\left(\frac{C_\phi^2 H^{4/3} N^{2/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{2/3}}\right) \\ & \leq \mathcal{O}\left(\frac{C_\phi^2 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{2/3}}\right) \end{aligned} \quad (114)$$

where we use the fact that $\widehat{V}_{h+1}(s) \leq H$ for any $s \in \mathcal{S}$. We then proceed to bound $\|\widehat{w}_h - \bar{w}_h\|_2$ as follows

$$\begin{aligned} \|\widehat{w}_h - \bar{w}_h\|_2 &= \|\Lambda^{-1}(w_0) \Lambda(w_0)(\widehat{w}_h - \bar{w}_h)\|_2 \\ &\leq \|\Lambda^{-1}(w_0)\|_2 \|\Lambda(w_0)(\widehat{w}_h - \bar{w}_h)\|_2 \\ &\leq \frac{1}{\lambda_2} \cdot \mathcal{O}\left(\frac{C_\phi^2 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{2/3}}\right) \\ &\leq \mathcal{O}\left(\frac{C_\phi^2 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{5/3}}\right). \end{aligned} \quad (115)$$

Substituting eq. (115) into eq. (102) yields

$$(ii) \leq \mathcal{O}\left(\frac{C_\phi^3 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{5/3}}\right). \quad (116)$$

Taking summation of the upper bounds of (i) in eq. (101) and (ii) in eq. (116), respectively, we have the following holds for all $x \in \mathcal{X}$ with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned} \left| (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) - \overline{\mathbb{P}}_h \widehat{V}_{h+1}(x) \right| &\leq (i) + (ii) \\ &\leq \mathcal{O} \left(C_\phi \left(\frac{N^2 H^4}{\lambda_2^2 \sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \mathcal{O} \left(\frac{C_\phi^3 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{5/3}} \right) \\ &\leq \mathcal{O} \left(\frac{C_\phi^3 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{5/3}} \right). \end{aligned} \quad (117)$$

Step III. In this step, we show that the bonus term $b_{v,h}(\cdot, \widehat{w}_h)$ in Algorithm 1 can be well approximated by $b_{v,h}(\cdot, w_0)$. By the definition of $b_{v,h}(\cdot, w)$, we have

$$\begin{aligned} &|b_{v,h}(x, \widehat{w}_h) - b_{v,h}(x, w_0)| \\ &= \left| [\phi_h(x, \widehat{w}_h)^\top \Lambda^{-1}(\widehat{w}_h) \phi_h(x, \widehat{w}_h)]^{1/2} - [\phi_h(x, w_0)^\top \Lambda^{-1}(w_0) \phi_h(x, w_0)]^{1/2} \right| \\ &\leq \left| \phi_h(x, \widehat{w}_h)^\top \Lambda^{-1}(\widehat{w}_h) \phi_h(x, \widehat{w}_h) - \phi_h(x, w_0)^\top \Lambda^{-1}(w_0) \phi_h(x, w_0) \right|^{1/2}, \end{aligned} \quad (118)$$

where the last inequality follows from the fact that $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$. Following steps similar to those in eq. (72), we can obtain

$$\begin{aligned} &|\phi_h(x, \widehat{w}_h)^\top \Lambda^{-1}(\widehat{w}_h) \phi_h(x, \widehat{w}_h) - \phi_h(x, w_0)^\top \Lambda^{-1}(w_0) \phi_h(x, w_0)| \\ &\leq \|\phi(x, \widehat{w}_h) - \phi(x, w_0)\|_2 \|\Lambda^{-1}(\widehat{w}_h)\|_2 \|\phi(x, \widehat{w}_h)\|_2 \\ &\quad + \|\phi(x, w_0)\|_2 \|\Lambda^{-1}(\widehat{w}_h)\|_2 \|\Lambda(\widehat{w}_h) - \Lambda(w_0)\|_2 \|\Lambda^{-1}(w_0)\|_2 \|\phi(x, \widehat{w}_h)\|_2 \\ &\quad + \|\phi(x, w_0)\|_2 \|\Lambda^{-1}(w_0)\|_2 \|\phi(x, \widehat{w}_h) - \phi(x, w_0)\|_2 \\ &\leq \frac{1}{\lambda_2} \|\phi(x, \widehat{w}_h) - \phi(x, w_0)\|_2 \|\phi(x, \widehat{w}_h)\|_2 + \frac{1}{\lambda_2} \|\phi(x, w_0)\|_2 \|\Lambda(\widehat{w}_h) - \Lambda(w_0)\|_2 \|\phi(x, \widehat{w}_h)\|_2 \\ &\quad + \frac{1}{\lambda_2} \|\phi(x, w_0)\|_2 \|\phi(x, \widehat{w}_h) - \phi(x, w_0)\|_2, \end{aligned} \quad (119)$$

where the last inequality follows from the fact that $\|\Lambda(w)\|_2 \geq \lambda_2$ for any $w \in \mathbb{R}^{2md}$. For $\Lambda(\widehat{w}) - \Lambda(w_0)$, by following steps similar to those in eq. (73), we can obtain

$$\begin{aligned} &\|\Lambda(\widehat{w}_h) - \Lambda(w_0)\|_2 \\ &\leq \sum_{\tau \in \mathcal{D}} [\|\phi(x_h^\tau, \widehat{w}_h)\|_2 \|\phi(x_h^\tau, \widehat{w}_h) - \phi(x_h^\tau, w_0)\|_2 + \|\phi(x_h^\tau, \widehat{w}_h) - \phi(x_h^\tau, w_0)\|_2 \|\phi(x_h^\tau, w_0)\|_2]. \end{aligned} \quad (120)$$

Applying Lemma 5 to eq. (120), we have the followings hold with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned} \|\phi(x_h^\tau, \widehat{w}_h)\|_2 &\leq C_\phi, \\ \|\phi(x_h^\tau, \widehat{w}_h) - \phi(x_h^\tau, w_0)\|_2 &\leq \mathcal{O} \left(\frac{C_\phi H^{1/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_2^{1/6}} \right). \end{aligned}$$

Substituting the above two inequalities into eq. (73) yields

$$\|\Lambda(\widehat{w}_h) - \Lambda(w_0)\|_2 \leq \mathcal{O} \left(\frac{C_\phi^2 H^{1/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_2^{1/6}} \right). \quad (121)$$

Finally, combining eq. (121) and eq. (181) and eq. (182) in Lemma 5, the right hand side of eq. (119) can be bounded by

$$\begin{aligned} &|\phi_h(x, \widehat{w}_h)^\top \Lambda^{-1}(\widehat{w}_h) \phi_h(x, \widehat{w}_h) - \phi_h(x, w_0)^\top \Lambda^{-1}(w_0) \phi_h(x, w_0)| \\ &\leq \frac{1}{\lambda_2} \|\phi(x, \widehat{w}_h) - \phi(x, w_0)\|_2 \|\phi(x, \widehat{w}_h)\|_2 + \frac{1}{\lambda_2} \|\phi(x, w_0)\|_2 \|\Lambda(\widehat{w}_h) - \Lambda(w_0)\|_2 \|\phi(x, \widehat{w}_h)\|_2 \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{\lambda_2} \|\phi(x, w_0)\|_2 \|\phi(x, \hat{w}_h) - \phi(x, w_0)\|_2 \\
& \leq \mathcal{O}\left(\frac{C_\phi^2 H^{1/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_2^{7/6}}\right) + \mathcal{O}\left(\frac{C_\phi^4 H^{1/3} N^{1/6} \sqrt{\log m}}{m^{1/6} \lambda_2^{13/6}}\right).
\end{aligned}$$

By eq. (118), we have the following holds with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned}
& |b_{v,h}(x, \hat{w}_h) - b_{v,h}(x, w_0)| \\
& \leq |\phi_h(x, \hat{w}_h)^\top \Lambda^{-1}(\hat{w}_h) \phi_h(x, \hat{w}_h) - \phi_h(x, w_0)^\top \Lambda^{-1}(w_0) \phi_h(x, w_0)|^{1/2} \\
& \leq \mathcal{O}\left(\frac{C_\phi^2 H^{1/6} N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_2^{13/12}}\right). \tag{122}
\end{aligned}$$

Step IV. In Steps I and II, we show that $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ can be well approximated by a linear function $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ with the feature $\phi(\cdot, \theta_0)$, and $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ can be well approximated by a linear function $(\overline{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ with the feature $\phi(\cdot, \theta_0)$. In this step, we want to show that the difference between $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot)$ and $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ is approximately $\beta_2 \cdot b_{v,h}(x, \Theta_0)$ with an approximately chosen β_2 .

Recall that $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = \langle \phi(\cdot, w_0), \tilde{w}_h - w_0 \rangle$ and $(\overline{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot) = \langle \phi(\cdot, w_0), \bar{w}_h - w_0 \rangle$. Consider the difference between $(\overline{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ and $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$. We have

$$(\overline{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) = \langle \phi(x, w_0), \bar{w}_h - \tilde{w}_h \rangle, \tag{123}$$

By eq. (104), we have

$$\bar{w} - w_0 = \Lambda(w_0)^{-1} \sum_{\tau \in \mathcal{D}} \widehat{V}_{h+1}(s_{h+1}^\tau) \phi(x_h^\tau, w_0). \tag{124}$$

By the definition of $\Lambda(w)$, we have

$$\tilde{w} - w_0 = \Lambda(w_0)^{-1} \left[\lambda_2 (\tilde{w} - w_0) + \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \phi(x_h^\tau, w_0)^\top \right) (\tilde{w} - w_0) \right]. \tag{125}$$

Subtracting eq. (125) from eq. (124), we have

$$\bar{w} - \tilde{w} = -\lambda_2 \Lambda(w_0)^{-1} (\tilde{w} - w_0) + \Lambda(w_0)^{-1} \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left[\widehat{V}_{h+1}(s_{h+1}^\tau) - \langle \phi(x_h^\tau, w_0), \tilde{w} - w_0 \rangle \right]. \tag{126}$$

Taking inter product of both sides of eq. (126) with vector $\phi(x_h^\tau, w_0)$ and using the fact that $(\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(s_{h+1}^\tau) = \langle \phi(x_h^\tau, w_0), \tilde{w}_h - w_0 \rangle$, we have

$$\begin{aligned}
& \langle \phi_h(x_h^\tau, w_0), \bar{w} - \tilde{w} \rangle \\
& = -\lambda_2 \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} (\tilde{w} - w_0) \\
& \quad + \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left(\widehat{V}_{h+1}(s_{h+1}^\tau) - (\mathbb{P}_h \widehat{V}_{h+1})(x_h^\tau) \right) \right) \\
& \quad + \phi_h(x, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left((\overline{\mathbb{P}}_h \widehat{V}_{h+1})(s_{h+1}^\tau) - \langle \phi(x_h^\tau, w_0), \tilde{w}_h - w_0 \rangle \right) \right) \\
& = -\lambda_2 \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} (\tilde{w} - w_0) \\
& \quad + \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left(\overline{V}_{h+1}(s_{h+1}^\tau) - (\mathbb{P}_h \overline{V}_{h+1})(x_h^\tau) \right) \right) \\
& \quad + \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left(\Delta V_{h+1}(s_{h+1}^\tau) - (\mathbb{P}_h \Delta V_{h+1})(x_h^\tau) \right) \right)
\end{aligned}$$

$$+ \phi_h(x, w_0)^\top \Lambda(w_0)^{-1/2} \Lambda(w_0)^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \left((\mathbb{P}_h \widehat{V}_{h+1})(x_h^\tau) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x_h^\tau) \right) \right), \quad (127)$$

where in the last equality we denote $\Delta V_h(s) := \widehat{V}_h(s) - \bar{V}_h(s)$. By the definition of $\widehat{V}_h(\cdot)$ in Algorithm 1 and $\bar{V}_h(\cdot)$ in eq. (100), we have

$$\begin{aligned} & \left| \widehat{V}_h(x) - \bar{V}_h(x) \right| \\ & \leq \sup_{x \in \mathcal{X}} \left| \widehat{Q}_h(x) - \bar{Q}_h(x) \right| \\ & \leq \left| f(x, \widehat{\theta}_h) - \langle \phi(x, \theta_0), \widehat{\theta}_h - \theta_0 \rangle \right| + \left| f(x, \widehat{w}_h) - \langle \phi(x, w_0), \widehat{w}_h - w_0 \rangle \right| \\ & \quad + \beta_1 \left| b_{r,h}(x, \widehat{\Theta}) - b_{r,h}(x, \Theta_0) \right| + \beta_2 \left| b_{v,h}(x, \widehat{w}) - b_{v,h}(x, w_0) \right| \\ & \stackrel{(i)}{\leq} \mathcal{O} \left(C_\phi \left(\frac{H^4 N^2 / \lambda_1^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \mathcal{O} \left(C_\phi \left(\frac{H^4 N^2 / \lambda_2^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) \\ & \quad + \beta_1 \cdot \mathcal{O} \left(\frac{C_\phi^2 N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_1^{13/12}} \right) + \beta_2 \cdot \mathcal{O} \left(\frac{C_\phi^2 H^{1/6} N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_2^{13/12}} \right) \\ & \stackrel{(ii)}{\leq} \mathcal{O} \left(C_\phi \left(\frac{H^4 N^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \max\{H^{2/3} \beta_1, H^{1/6} \beta_2\} \cdot \mathcal{O} \left(\frac{C_\phi^2 N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right), \end{aligned} \quad (128)$$

where (i) follows from eq. (183) in Lemma 5, eq. (78) and eq. (122), and (ii) follows from the fact that $\lambda_1, \lambda_2 > 1$. Denoting

$$\varepsilon_v = \mathcal{O} \left(C_\phi \left(\frac{H^4 N^2}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right) + \max\{H^{2/3} \beta_1, H^{1/6} \beta_2\} \cdot \mathcal{O} \left(\frac{C_\phi^2 N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right),$$

we then have the following holds for all $h \in [H]$ and $s \in \mathcal{S}$

$$|\Delta V_h(s)| \leq \varepsilon_v.$$

Equation (127) together with eq. (128) imply

$$\begin{aligned} & \left| \langle \phi_h(x_h^\tau, w_0), \bar{w} - \widetilde{w} \rangle \right| \\ & \leq \sqrt{\lambda_2} \left\| \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \right\|_2 \|\widetilde{w} - w_0\|_2 \\ & \quad + \left\| \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \right\|_2 \left\| \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \varepsilon_v(x_h^\tau) \right\|_{\Lambda(w_0)^{-1}} \\ & \quad + \frac{2\varepsilon_v}{\sqrt{\lambda_2}} \left\| \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \right\|_2 \sum_{\tau \in \mathcal{D}} \|\phi(x_h^\tau, w_0)\|_2 \\ & \quad + \frac{1}{\sqrt{\lambda_2}} \left\| \phi_h(x_h^\tau, w_0)^\top \Lambda(w_0)^{-1/2} \right\|_2 \left(\sum_{\tau \in \mathcal{D}} \|\phi(x_h^\tau, w_0)\|_2 \left| (\mathbb{P}_h \widehat{V}_{h+1})(x_h^\tau) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x_h^\tau) \right| \right), \end{aligned} \quad (129)$$

where we denote $\varepsilon_v(x_h^\tau) := \bar{V}_{h+1}(s_{h+1}^\tau) - (\mathbb{P}_h \bar{V}_{h+1})(x_h^\tau)$ and use the fact that $\|\Lambda(w)^{-1/2}\|_2 \leq 1/\sqrt{\lambda_2}$ for any $w \in \mathbb{R}^{2md}$. By the definition of \widetilde{w} in Step I, we have

$$\|\widetilde{w} - w_0\|_2 = \|\ell_v\|_2 \leq A_2 \sqrt{H/d}. \quad (130)$$

By Lemma 5 and eq. (52), we have the followings hold with probability at least $1 - N^{-2}H^{-4}$ over the randomness of initialization w_0

$$\|\phi(x_h^\tau, w_0)\|_2 \leq C_\phi, \quad (131)$$

$$\left| (\mathbb{P}_h \widehat{V}_{h+1})(x_h^\tau) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x_h^\tau) \right| \leq \frac{2(L_\sigma A_2 + C_\sigma^2 A_2^2) \sqrt{\log N^2 H^5}}{\sqrt{m}}. \quad (132)$$

Substituting eq. (130), eq. (131) and eq. (132) into eq. (129) and using the fact that $b_{v,h}(x, w_0) = \|\phi_h(x, w_0)^\top \Lambda(w_0)^{-1/2}\|_2$, we have

$$\begin{aligned} |\langle \phi_h(x_h^\tau, w_0), \bar{w} - \tilde{w} \rangle| &\leq \left(A_2 \sqrt{\frac{\lambda_2 H}{d}} + \frac{2(L_\sigma A_2 + C_\sigma^2 A_2^2) C_\phi N H^{3/2} \sqrt{\log H N}}{\sqrt{\lambda_2 m}} \right. \\ &\quad \left. + \left\| \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \varepsilon_v(x_h^\tau) \right\|_{\Lambda(w_0)^{-1}} \right) b_{v,h}(x_h^\tau, w_0). \end{aligned} \quad (133)$$

Given that the events in eq. (131) and eq. (132) occur, applying eq. (198) in Lemma 8, we have the following holds with probability at least $1 - N^{-2} H^{-4}$

$$\begin{aligned} &\left\| \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \varepsilon_v(x_h^\tau) \right\|_{\Lambda(w_0)^{-1}}^2 \\ &\leq 2H^2 \log \det(I + K_{N,h}^v / \lambda_2) + 2H^2 N (\lambda_2 - 1) + 4H^2 \log(\mathcal{N}_{\epsilon,h}^v / \delta) + 8N^2 C_\phi^2 \epsilon^2 / \lambda_2, \end{aligned} \quad (134)$$

where $K_{N,h}^v \in \mathbb{R}^{N \times N}$ is the Gram matrix defined as

$$K_{N,h}^v = [K(x_h^{\tau_i}, x_h^{\tau_j})]_{i,j \in [N]} \in \mathbb{R}^{N \times N},$$

and $\mathcal{N}_{\epsilon,h}^v$ is the cardinality of the following function class

$$\begin{aligned} \mathcal{V}_h(x, R_\theta, R_w, R_{\beta_1}, R_{\beta_2}, \lambda_1, \lambda_2) &= \{\max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a)\} : \mathcal{S} \rightarrow [0, H]\} \\ &\text{with } \|\theta\|_2 \leq R_\theta, \|w\|_2 \leq R_w, \beta_1 \in [0, R_{\beta_1}], \beta_2 \in [0, R_{\beta_2}], \|\Sigma\|_2 \geq \lambda_1, \|\Lambda\|_2 \geq \lambda_2, \end{aligned}$$

where $R_\theta = H \sqrt{N/\lambda_1}$, $R_w = H \sqrt{N/\lambda_2}$ and

$$\begin{aligned} \bar{Q}_h(x) &= \min\{\langle \phi(x, \theta_0), \theta - \theta_0 \rangle + \langle \phi(x, w_0), w - w_0 \rangle \\ &\quad - \beta_1 \cdot \sqrt{\Phi_h(x, \theta_0)^\top \Sigma^{-1} \Phi_h(x, \theta_0)} - \beta_2 \cdot \sqrt{\phi(x, w_0)^\top \Lambda^{-1} \phi(x, w_0)}, H\}^+. \end{aligned}$$

Combining eq. (133) and eq. (134), defining $\mathcal{N}_\epsilon^v = \max_{h \in [H]} \{\mathcal{N}_{\epsilon,h}^v\}$ and letting

$$\epsilon = \sqrt{\lambda_2 C_\epsilon H / (2N C_\phi)} \text{ where } C_\epsilon \geq 1 \text{ and } \lambda_2 = 1 + N^{-1}, \quad (135)$$

and m be sufficiently large such that

$$\frac{2(L_\sigma A_2 + C_\sigma^2 A_2^2) C_\phi N H^{3/2} \sqrt{\log H N}}{\sqrt{\lambda_2 m}} \leq A_2 \sqrt{\frac{\lambda_2 H}{d}},$$

we have the following holds with probability at least $1 - N^{-2} H^{-4}$

$$\begin{aligned} &|\langle \phi_h(x_h^\tau, w_0), \bar{w} - \tilde{w} \rangle| \\ &\leq \left(2A_2 \sqrt{\frac{\lambda_2 H}{d}} + \sqrt{2H^2 \log \det \left(I + \frac{K_{N,h}^v}{\lambda_2} \right) + 3C_\epsilon H^2 + 8H^2 \log(NH^2 \mathcal{N}_\epsilon^v)} \right) b_{v,h}(x, w_0) \\ &\leq H \underbrace{\left(\frac{8A_2^2 \lambda_2}{d} + 4 \max_{h \in [H]} \left\{ \log \det \left(I + \frac{K_{N,h}^v}{\lambda_2} \right) \right\} + 6C_\epsilon + 16 \log(NH^2 \mathcal{N}_\epsilon^v) \right)^{1/2}}_{\beta_2} b_{v,h}(x, w_0), \end{aligned} \quad (136)$$

where in the last inequality we use the fact that $a + b \leq \sqrt{2(a^2 + b^2)}$. Substituting eq. (136) into eq. (123), we conclude that the following holds with probability at least $1 - N^{-2} H^{-4}$

$$\left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| \leq \beta_2 \cdot b_{v,h}(x, w_0), \quad (137)$$

where

$$\beta_2 = H \left(\frac{8A_2^2\lambda_2}{d} + 4 \max_{h \in [H]} \left\{ \log \det \left(I + \frac{K_{N,h}^v}{\lambda_2} \right) \right\} + 22 \log(NH^2\mathcal{N}_\epsilon^v) \right)^{1/2}.$$

Next, we proceed to bound the term $\left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right|$. By the triangle inequality, we have

$$\begin{aligned} & \left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| \\ &= \left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) + (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\overline{\mathbb{P}}_h \widehat{V}_{h+1})(x) + (\overline{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| \\ &\leq \left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| + \left| (\widetilde{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\overline{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| + \left| (\overline{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| \\ &\stackrel{(i)}{\leq} \frac{2(L_\sigma A_2 + C_\sigma^2 A_2^2) \sqrt{\log(N^2 H^5)}}{\sqrt{m}} + \beta_2 \cdot b_{v,h}(x, w_0) + \mathcal{O} \left(\frac{C_\phi^3 H^{4/3} N^{5/3} \sqrt{\log(m)}}{m^{1/6} \lambda_2^{5/3}} \right) \\ &\stackrel{(ii)}{\leq} \mathcal{O} \left(\frac{H^{4/3} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right) + \beta_2 \cdot b_{v,h}(x, w_0), \end{aligned} \quad (138)$$

where (i) follows from eq. (96), eq. (117) and eq. (137) and (ii) follows from the fact that $\lambda_2 = 1 + 1/N$ and $L_\sigma, C_\sigma, A_2, C_\phi = \mathcal{O}(1)$.

D.3 UPPER AND LOWER BOUNDS ON EVALUATE ERROR $\delta_h(\cdot)$

By definition, we have the following holds with probability $1 - 2N^{-2}H^{-4}$

$$\begin{aligned} & \left| (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - (\mathbb{B}_h \widehat{V}_{h+1})(x) \right| \\ &= \left| \widehat{R}_h(x) + (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) - R_h(x) - (\mathbb{P}_h \widehat{V}_{h+1})(x) \right| \\ &\leq \left| \widehat{R}_h(x) - R_h(x) \right| + \left| (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) - (\mathbb{P}_h \widehat{V}_{h+1})(x) \right| \\ &\stackrel{(i)}{\leq} \beta_1 \cdot b_{r,h}(x, \Theta_0) + \beta_2 \cdot b_{v,h}(x, w_0) + \mathcal{O} \left(\frac{H^{4/3} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right) \\ &\quad + \mathcal{O} \left(\frac{H^{17/6} N^{5/3} \sqrt{\log(m)}}{m^{1/6}} \right) \\ &\leq \beta_1 \cdot b_{r,h}(x, \Theta_0) + \beta_2 \cdot b_{v,h}(x, w_0) + \mathcal{O} \left(\frac{H^{17/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}} \right), \end{aligned} \quad (139)$$

where (i) follows from eq. (92) and eq. (139). Moreover, by the triangle inequality, eq. (78) and eq. (122), we have the following holds with probability $1 - 2N^{-2}H^{-4}$

$$\begin{aligned} & \beta_1 \cdot b_{r,h}(x, \Theta_0) + \beta_2 \cdot b_{v,h}(x, w_0) \\ &\leq \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}) + \beta_1 \cdot \left| b_{r,h}(x, \widehat{\Theta}) - b_{r,h}(x, \Theta_0) \right| \\ &\quad + \beta_2 \cdot \left| b_{v,h}(x, \widehat{w}) - b_{v,h}(x, w_0) \right| \\ &\leq \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) \\ &\quad + \beta_1 \cdot \mathcal{O} \left(\frac{C_\phi^2 H^{2/3} N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_1^{13/12}} \right) + \beta_2 \cdot \mathcal{O} \left(\frac{C_\phi^2 H^{1/6} N^{1/12} (\log m)^{1/4}}{m^{1/12} \lambda_2^{13/12}} \right) \\ &\stackrel{(i)}{\leq} \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \max\{\beta_1 H^{2/3}, \beta_2 H^{1/6}\} \mathcal{O} \left(\frac{N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right), \end{aligned} \quad (140)$$

where (i) follows from the fact that $\lambda_1 = \lambda_2 = 1 + 1/N$ and $C_\phi = \mathcal{O}(1)$. Substituting eq. (140) into eq. (139), we can obtain

$$\left| (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - (\mathbb{B}_h \widehat{V}_{h+1})(x) \right|$$

$$\begin{aligned} &\leq \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) \\ &\quad + \max\{\beta_1 H^{2/3}, \beta_2 H^{1/6}\} \mathcal{O}\left(\frac{N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right) + \mathcal{O}\left(\frac{H^{17/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}}\right). \end{aligned}$$

Denoting

$$\varepsilon_b = \max\{\beta_1 H^{2/3}, \beta_2 H^{1/6}\} \mathcal{O}\left(\frac{N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right) + \mathcal{O}\left(\frac{H^{17/6} N^{5/3} \sqrt{\log(N^2 H^5 m)}}{m^{1/6}}\right),$$

we have

$$\left| (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - (\mathbb{B}_h \widehat{V}_{h+1})(x) \right| \leq \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b. \quad (141)$$

Up to this point, we characterize the uncertainty of $(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(\cdot)$. Next, we proceed to bound the suboptimality of Algorithm 1. Recalling the construction of $\widehat{Q}_h(x)$ in Algorithm 1, we have

$$\widehat{Q}_h(\cdot) = \min\{(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(\cdot) - \beta_1 \cdot b_{r,h}(\cdot, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(\cdot, \widehat{w}_h), H\}^+.$$

If $(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) < \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h)$, we have

$$\widehat{Q}_h(\cdot) = 0.$$

Note that $\widehat{V}_{h+1}(\cdot)$ is nonnegative. Recalling the definition of $\delta_h(x)$ in eq. (13), we have

$$\delta_h(x) = (\mathbb{B}_h \widehat{V}_{h+1})(x) - \widehat{Q}_h(x) = (\mathbb{B}_h \widehat{V}_{h+1})(x) > 0.$$

Otherwise, if $(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) > \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h)$, we have

$$\begin{aligned} \widehat{Q}_h(x) &= \min\{(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h), H\}^+ \\ &\leq (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h), \end{aligned}$$

which implies that

$$\begin{aligned} \delta_h(x) &\geq (\mathbb{B}_h \widehat{V}_{h+1})(x) - \left[(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) \right] \\ &= \left[(\mathbb{B}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) \right] + \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h). \end{aligned}$$

Note that eq. (141) implies the followings hold with probability $1 - 2N^{-2}H^{-4}$

$$(\mathbb{B}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) \geq -\beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) - \varepsilon_b. \quad (142)$$

$$(\mathbb{B}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) \leq \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b. \quad (143)$$

As a result, we have the following holds with probability $1 - 2N^{-2}H^{-4}$

$$\delta_h(x) \geq -\varepsilon_b. \quad (144)$$

It remains to establish the upper bound of $\delta_h(x)$. Considering the event in eq. (143) occurs, we have

$$\begin{aligned} &(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(\cdot) - \beta_1 \cdot b_{r,h}(\cdot, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(\cdot, \widehat{w}_h) \\ &\leq \left[(\mathbb{B}_h \widehat{V}_{h+1})(x) + \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b \right] - \beta_1 \cdot b_{r,h}(\cdot, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(\cdot, \widehat{w}_h) \\ &= (\mathbb{B}_h \widehat{V}_{h+1})(x) + \varepsilon_b \leq H + \varepsilon_b, \end{aligned}$$

where the last inequality follows from the fact that $R_h(x) \leq 1$ and $\widehat{V}_{h+1}(s) \leq H$ for all $x \in \mathcal{X}$ and $s \in \mathcal{S}$. Hence, we have

$$\begin{aligned} \widehat{Q}_h(x) &= \min\{(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h), H\}^+ \\ &\geq \min\{(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) - \varepsilon_b, H\}^+ \\ &= \max\{(\widehat{\mathbb{B}}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) - \varepsilon_b, 0\} \end{aligned}$$

$$\geq (\mathbb{B}_h \widehat{V}_{h+1})(x) - \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) - \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) - \varepsilon_b, \quad (145)$$

which by definition of $\delta_h(x)$ implies

$$\begin{aligned} \delta_h(x) &= (\mathbb{B}_h \widehat{V}_{h+1})(x) - \widehat{Q}_h(x) \\ &\leq (\mathbb{B}_h \widehat{V}_{h+1})(x) - (\mathbb{B}_h \widehat{V}_{h+1})(x) + \beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b \\ &\leq 2 \left[\beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b \right], \end{aligned} \quad (146)$$

where the last inequality follows from eq. (143). Combining eq. (144) and eq. (146), with probability $1 - 2N^{-2}H^{-4}$, we have

$$-\varepsilon_b \leq \delta_h(x) \leq 2 \left[\beta_1 \cdot b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot b_{v,h}(x, \widehat{w}_h) + \varepsilon_b \right], \quad \forall x \in \mathcal{X}, \quad \forall h \in [H],$$

which completes the proof.

E PROOF OF LEMMA 2

For $\sum_{h=1}^H b_{r,h}(x, \widehat{\Theta})$, we have the following holds with probability $1 - N^{-2}H^{-4}$

$$\begin{aligned} \sum_{h=1}^H b_{r,h}(x, \widehat{\Theta}) &\leq \sum_{h=1}^H b_{r,h}(x, \Theta_0) + \sum_{h=1}^H \left| b_{r,h}(x, \widehat{\Theta}) - b_{r,h}(x, \Theta_0) \right| \\ &\stackrel{(i)}{\leq} \sum_{h=1}^H b_{r,h}(x, \Theta_0) + \mathcal{O} \left(\frac{H^{5/3} N^{1/12} (\log m)^{1/4}}{m^{1/12}} \right), \end{aligned} \quad (147)$$

where (i) follows from eq. (78). We next proceed to bound the term $\sum_{h=1}^H b_{r,h}(x, \Theta_0)$. Recall that in Assumption 3 we define $\overline{M}(\Theta_0) = \mathbb{E}_\mu [\Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top]$. For all $\tau \in \mathcal{D}$, we define the following random matrix $\widehat{M}(\Theta_0)$

$$\widehat{M}(\Theta_0) = \sum_{\tau \in \mathcal{D}} A_\tau(\Theta_0), \quad \text{where } A_\tau(\Theta_0) = \Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top - \overline{M}(\Theta_0). \quad (148)$$

Note that eq. (64) implies $\|\Phi(\tau, \Theta_0)\|_2 \leq C_\phi \sqrt{H}$. By Jensen's inequality, we have

$$\|\overline{M}(\Theta_0)\|_2 \leq \mathbb{E}_\mu [\|\Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top\|_2] \leq C_\phi^2 H. \quad (149)$$

For any vector $v \in \mathbb{R}^{2mdH}$ with $\|v\|_2 = 1$, we have

$$\begin{aligned} \|A_\tau(\Theta_0)v\|_2 &\leq \|\Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top v\|_2 + \|\overline{M}(\Theta_0)v\|_2 \\ &\leq \|\Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top\|_2 \|v\|_2 + \|\overline{M}(\Theta_0)\|_2 \|v\|_2 \\ &\leq 2C_\phi^2 H \|v\|_2 = 2C_\phi^2 H, \end{aligned}$$

which implies

$$\|A_\tau(\Theta_0)\|_2 \leq 2C_\phi^2 H \text{ and } \|A_\tau(\Theta_0) A_\tau(\Theta_0)^\top\|_2 \leq \|A_\tau(\Theta_0)\|_2 \|A_\tau(\Theta_0)^\top\|_2 \leq 4C_\phi^4 H^2. \quad (150)$$

Since $\{A_\tau(\Theta_0)\}_{\tau \in \mathcal{D}}$ are i.i.d. and $\mathbb{E}[A_\tau(\Theta_0)] = 0$ for all τ , we have

$$\begin{aligned} \left\| \mathbb{E}_\mu [\widehat{M}(\Theta_0) \widehat{M}(\Theta_0)^\top] \right\|_2 &= \left\| \sum_{\tau \in \mathcal{D}} \mathbb{E}_\mu [A_\tau(\Theta_0) A_\tau(\Theta_0)^\top] \right\|_2 \\ &= N \cdot \left\| \mathbb{E}_\mu [A_{\tau_1}(\Theta_0) A_{\tau_1}(\Theta_0)^\top] \right\|_2 \\ &\stackrel{(i)}{\leq} N \cdot \mathbb{E}_\mu [\|A_{\tau_1}(\Theta_0) A_{\tau_1}(\Theta_0)^\top\|_2] \\ &\leq 4C_\phi^4 H^2 N, \end{aligned}$$

where (i) follows from Jensen's inequality. Similarly, we can also obtain

$$\left\| \mathbb{E}_\mu [\widehat{M}(\Theta_0)^\top \widehat{M}(\Theta_0)] \right\|_2 \leq 4C_\phi^4 H^2 N.$$

Applying Lemma 10 to $\widehat{M}(\Theta_0)$, for any fixed $h \in [H]$ and any $\xi_1 > 0$, we have

$$\mathbb{P} \left(\left\| \widehat{M}(\Theta_0) \right\|_2 \geq \xi_1 \right) \leq 4mdH \cdot \exp \left(-\frac{\xi_1^2/2}{4C_\phi^4 H^2 N + 2C_\phi^2 H/3 \cdot \xi_1} \right).$$

For any $\delta_1 \in (0, 1)$, let

$$\xi_1 = C_\phi^2 H \sqrt{10N \log \left(\frac{4mdH}{\delta_1} \right)} \quad \text{and} \quad N \geq \frac{40}{9} \log \left(\frac{4mdH}{\delta_1} \right).$$

Then, we have

$$\begin{aligned} \mathbb{P} \left(\left\| \widehat{M}(\Theta_0) \right\|_2 \geq \xi_1 \right) &\leq 4mdH \cdot \exp \left(-\frac{\xi_1^2/2}{4C_\phi^4 H^2 N + 2C_\phi^2 H/3 \cdot \xi_1} \right) \\ &\leq 4mdH \cdot \exp \left(-\frac{\xi_1^2}{10C_\phi^4 H^2 N} \right) = \delta_1, \end{aligned}$$

which implies that the following holds with probability at least $1 - \delta_1$ taken with respect to the randomness of \mathcal{D}

$$\begin{aligned} \left\| \widehat{M}(\Theta_0)/N \right\|_2 &= \left\| \frac{1}{N} \sum_{\tau \in \mathcal{D}} \Phi(\tau, \Theta_0) \Phi(\tau, \Theta_0)^\top - \overline{M}(\Theta_0) \right\|_2 \\ &\leq C_\phi^2 H \sqrt{\frac{10}{N} \log \left(\frac{4mdH}{\delta_1} \right)}. \end{aligned} \quad (151)$$

By the definition of $\Sigma(\Theta_0)$, we have

$$\widehat{M}(\Theta_0) = (\Sigma(\Theta_0) - \lambda_1 \cdot I_{2mdH}) - N \cdot \overline{M}(\Theta_0). \quad (152)$$

By Assumption 3, there exists an absolute constant $C_\sigma > 0$ such that $\lambda_{\min}(\overline{M}(\Theta_0)) \geq C_\sigma$, which implies that $\|\overline{M}(\Theta_0)^{-1}\|_2 \leq 1/C_\sigma$. Letting N be sufficiently large such that

$$N \geq \max \left\{ \frac{40C_\phi^4 H^2}{C_\sigma^2}, \frac{40}{9} \right\} \log \left(\frac{4mdH}{\delta_1} \right)$$

and combining eq. (151) and eq. (152), we have

$$\begin{aligned} \lambda_{\min}(\Sigma(\Theta_0)/N) &= \lambda_{\min}(\overline{M}(\Theta_0) + \widehat{M}(\Theta_0)/N + \lambda_1/N \cdot I_{2mdH}) \\ &\geq \lambda_{\min}(\overline{M}(\Theta_0)) - \left\| \widehat{M}(\Theta_0)/N \right\|_2 \\ &\geq C_\sigma - C_\phi^2 H \sqrt{\frac{10}{N} \log \left(\frac{4mdH}{\delta} \right)} \\ &\geq C_\sigma/2. \end{aligned}$$

Hence, the following holds with probability $1 - \delta_1$ with respect to randomness of \mathcal{D}

$$\left\| \Sigma(\Theta_0)^{-1} \right\|_2 \leq (N \cdot \lambda_{\min}(\Sigma(\Theta_0)/N))^{-1} \leq \frac{2}{NC_\sigma},$$

which implies the following holds for all $x \in \mathcal{X}$ and $h \in [H]$

$$\begin{aligned} b_{r,h}(x, \Theta_0) &= \sqrt{\Phi_h(x, \Theta_0)^\top \Sigma^{-1}(\Theta_0) \Phi_h(x, \Theta_0)} \\ &\leq \|\Phi_h(x, \Theta_0)\|_2 \cdot \left\| \Sigma^{-1}(\Theta_0) \right\|_2^{1/2} \leq \frac{\sqrt{2}C_\phi}{\sqrt{C_\sigma} \sqrt{N}}, \end{aligned} \quad (153)$$

where we use the fact that $\|\Phi_h(x, \Theta_0)\|_2 = \|\phi(x_h^\tau, \theta_0)\|_2 \leq C_\phi$. Substituting eq. (153) into eq. (147), we have

$$\sum_{h=1}^H b_{r,h}(x, \hat{\Theta}) \leq \frac{\sqrt{2}HC_\phi}{\sqrt{C_\sigma}\sqrt{N}} + \mathcal{O}\left(\frac{H^{5/3}N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right). \quad (154)$$

Next, we proceed to bound the term $\sum_{h=1}^H b_{v,h}(x, \hat{w}_h)$. According to eq. (122), we have the following holds with probability at least $1 - N^{-2}H^{-4}$

$$\begin{aligned} \sum_{h=1}^H b_{v,h}(x, \hat{w}_h) &\leq \sum_{h=1}^H b_{v,h}(x, w_0) + \sum_{h=1}^H |b_{v,h}(x, \hat{w}_h) - b_{v,h}(x, w_0)| \\ &\stackrel{(i)}{\leq} \sum_{h=1}^H b_{v,h}(x, w_0) + \mathcal{O}\left(\frac{H^{7/6}N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right). \end{aligned} \quad (155)$$

We then proceed to bound the summation of the penalty terms $\sum_{h=1}^H b_{v,h}(x, w_0)$. Recall that in Assumption 3 we define $\bar{m}_h(w_0) = \mathbb{E}_\mu [\phi(x_h^\tau, w_0)\phi(x_h^\tau, w_0)^\top]$. For all $h \in [H]$ and $\tau \in \mathcal{D}$, we define the following random matrix $\hat{m}_h(w_0)$

$$\hat{m}_h(w_0) = \sum_{\tau \in \mathcal{D}} B_h^\tau(w_0), \quad \text{where } B_h^\tau(w_0) = \phi(x_h^\tau, w_0)\phi(x_h^\tau, w_0)^\top - \bar{m}_h(w_0). \quad (156)$$

Note that eq. (64) implies $\|\phi(x_h^\tau, w_0)\|_2 \leq C_\phi$. By Jensen's inequality, we have

$$\|\bar{m}_h(w_0)\|_2 \leq \mathbb{E}_\mu [\|\phi(x_h^\tau, w_0)\phi(x_h^\tau, w_0)^\top\|_2] \leq C_\phi^2. \quad (157)$$

For any vector $v \in \mathbb{R}^{2md}$ with $\|v\|_2 = 1$, we have

$$\begin{aligned} \|B_h^\tau(w_0)v\|_2 &\leq \|\phi(x_h^\tau, w_0)\phi(x_h^\tau, w_0)^\top v\|_2 + \|\bar{m}_h(w_0)v\|_2 \\ &\leq \|\phi(x_h^\tau, w_0)\phi(x_h^\tau, w_0)^\top\|_2 \|v\|_2 + \|\bar{m}_h(w_0)\|_2 \|v\|_2 \\ &\leq 2C_\phi^2 \|v\|_2 = 2C_\phi^2, \end{aligned}$$

which implies

$$\|B_h^\tau(w_0)\|_2 \leq 2C_\phi^2 \text{ and } \|B_h^\tau(w_0)B_h^\tau(w_0)^\top\|_2 \leq \|B_h^\tau(w_0)\|_2 \|B_h^\tau(w_0)^\top\|_2 \leq 4C_\phi^4. \quad (158)$$

Since $\{B_h^\tau(w_0)\}_{\tau \in \mathcal{D}}$ are i.i.d. and $\mathbb{E}[B_h^\tau(w_0)] = 0$ for all τ , we have

$$\begin{aligned} \|E_\mu[\bar{m}_h(w_0)\bar{m}_h(w_0)^\top]\|_2 &= \left\| \sum_{\tau \in \mathcal{D}} E_\mu [B_h^\tau(w_0)B_h^\tau(w_0)^\top] \right\|_2 \\ &= N \cdot \|E_\mu [B_h^{\tau_1}(w_0)B_h^{\tau_1}(w_0)^\top]\|_2 \\ &\stackrel{(i)}{\leq} N \cdot E_\mu [\|B_h^{\tau_1}(w_0)B_h^{\tau_1}(w_0)^\top\|_2] \\ &\leq 4C_\phi^4 N, \end{aligned}$$

where (i) follows from Jensen's inequality. Similarly, we can also obtain

$$\|E_\mu[\bar{m}_h(w_0)^\top \bar{m}_h(w_0)]\|_2 \leq 4C_\phi^4 N.$$

Applying Lemma 10 to $\hat{m}_h(w_0)$, for any fixed $h \in [H]$ and any $\xi_2 > 0$, we have

$$\mathbb{P}(\|\hat{m}_h(w_0)\|_2 \geq \xi_2) \leq 4md \cdot \exp\left(-\frac{\xi_2^2/2}{4C_\phi^4 N + 2C_\phi^2/3 \cdot \xi_2}\right).$$

For any $\delta_2 \in (0, 1)$, let

$$\xi_2 = C_\phi^2 \sqrt{10N \log\left(\frac{4mdH}{\delta_2}\right)} \quad \text{and} \quad N \geq \frac{40}{9} \log\left(\frac{4mdH}{\delta_2}\right).$$

Then, we have

$$\begin{aligned} \mathbb{P}(\|\widehat{m}_h(w_0)\|_2 \geq \xi_2) &\leq 4md \cdot \exp\left(-\frac{\xi_2^2/2}{4C_\phi^4 N + 2C_\phi^2/3 \cdot \xi_2}\right) \\ &\leq 4md \cdot \exp\left(-\frac{\xi_2^2}{10C_\phi^4 N}\right) = \frac{\delta_2}{H}, \end{aligned}$$

which implies that we have the following holds with probability at least $1 - \delta_2/H$ taken with respect to the randomness of \mathcal{D}

$$\begin{aligned} \|\widehat{m}_h(w_0)/N\|_2 &= \left\| \frac{1}{N} \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau, w_0) \phi(x_h^\tau, w_0)^\top - \overline{m}_h(w_0) \right\|_2 \\ &\leq C_\phi^2 \sqrt{\frac{10}{N} \log\left(\frac{4mdH}{\delta_2}\right)}. \end{aligned} \quad (159)$$

By the definition of $\Lambda_h(w_0)$, we have

$$\widehat{m}_h(w_0) = (\Lambda_h(w_0) - \lambda_2 \cdot I_{2md}) - N \cdot \overline{m}_h(w_0). \quad (160)$$

By Assumption 3, there exists an absolute constant $C_\zeta > 0$ such that $\lambda_{\min}(\overline{m}_h(\Theta_0)) \geq C_\zeta$, which implies that $\|\overline{m}(w_0)^{-1}\|_2 \leq 1/C_\zeta$. Letting N be sufficiently large such that

$$N \geq \max\left\{\frac{40C_\phi^4}{C_\zeta^2}, \frac{40}{9}\right\} \log\left(\frac{4mdH}{\delta_2}\right)$$

and combining eq. (159) and eq. (160), we have

$$\begin{aligned} \lambda_{\min}(\Lambda_h(w_0)/N) &= \lambda_{\min}(\overline{m}(w_0) + \widehat{m}(w_0)/N + \lambda_1/N \cdot I_{2md}) \\ &\geq \lambda_{\min}(\overline{m}(w_0)) - \|\widehat{m}(w_0)/N\|_2 \\ &\geq C_\zeta - C_\phi^2 H \sqrt{\frac{10}{N} \log\left(\frac{4mdH}{\delta_2}\right)} \\ &\geq C_\zeta/2. \end{aligned}$$

Hence, the following holds with probability $1 - \delta_2/H$ with respect to randomness of \mathcal{D}

$$\|\Lambda_h(w_0)^{-1}\|_2 \leq (N \cdot \lambda_{\min}(\Lambda_h(w_0)/N))^{-1} \leq \frac{2}{NC_\zeta}. \quad (161)$$

Taking union bound of eq. (161) over $[H]$, we have the following holds for all $x \in \mathcal{X}$ and $h \in [H]$ with probability $1 - \delta_2$

$$\begin{aligned} b_{v,h}(x, w_0) &= \sqrt{\phi_h(x, w_0)^\top \Lambda_h^{-1}(w_0) \phi_h(x, w_0)} \\ &\leq \|\phi_h(x, w_0)\|_2 \cdot \|\Lambda_h(w_0)^{-1}\|_2^{1/2} \leq \frac{\sqrt{2}C_\phi}{\sqrt{C_\zeta}\sqrt{N}}, \end{aligned} \quad (162)$$

where we use the fact that $\|\phi(x_h^\tau, \theta_0)\|_2 \leq C_\phi$. Substituting eq. (162) into eq. (155), we have

$$\sum_{h=1}^H b_{v,h}(x, \widehat{w}) \leq \frac{\sqrt{2}HC_\phi}{\sqrt{C_\zeta}\sqrt{N}} + \mathcal{O}\left(\frac{H^{7/6}N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right). \quad (163)$$

Finally, letting $\delta_1 = N^{-2}H^{-4}/2$ and $\delta_2 = N^{-2}H^{-4}/2$ and combining eq. (154) and eq. (163), we have the following holds with probability $1 - N^{-2}H^{-4}$

$$\begin{aligned} &\beta_1 \cdot \sum_{h=1}^H b_{r,h}(x, \widehat{\Theta}) + \beta_2 \cdot \sum_{h=1}^H b_{v,h}(x, \widehat{w}) \\ &\leq \left(\frac{\beta_1}{\sqrt{C_\sigma}} + \frac{\beta_2}{\sqrt{C_\zeta}}\right) \frac{\sqrt{2}HC_\phi}{\sqrt{N}} + \max\{\beta_1 H^{5/3}, \beta_2 H^{7/6}\} \cdot \mathcal{O}\left(\frac{N^{1/12}(\log m)^{1/4}}{m^{1/12}}\right), \end{aligned}$$

which completes the proof.

F PROOF OF LEMMA 3

Similarly to the proof of Lemma 1, we first bound the uncertainty of the estimated reward $\widehat{R}_h(\cdot)$ in eq. (41) and then bound the uncertainty of the estimated transition value function $(\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$ in eq. (43).

F.1 UNCERTAINTY OF ESTIMATED REWARD $\widehat{R}_h(\cdot)$

Following steps similar to those in the proof of Lemma B.1 in Jin et al. (2021), we can obtain

$$\|\Theta^*\|_2 \leq H\sqrt{dH} \quad \text{and} \quad \|\widehat{\Theta}\|_2 \leq H\sqrt{dHN/\lambda_1}. \quad (164)$$

For simplicity, we denote $r(\tau) = \sum_{h \in [H]} r(x_h^\tau)$, $R(\tau) = \sum_{h \in [H]} R(x_h^\tau)$ and $\varepsilon(\tau) = R(\tau) - r(\tau)$. Consider the estimation error $R_h(\cdot) - \widehat{R}_h(\cdot)$. We have

$$\begin{aligned} & R_h(x) - \widehat{R}_h(x) \\ &= \langle \phi(x), \theta_h^* - \widehat{\theta}_h \rangle \\ &= \langle \Phi_h(x), \Theta^* - \widehat{\Theta} \rangle \\ &= \langle \Phi_h(x), \Theta^* \rangle - \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) r(\tau) \right) \\ &= \langle \Phi_h(x), \Theta^* \rangle - \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \Phi(\tau)^\top \Theta^* \right) + \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right) \\ &= \langle \Phi_h(x), \Theta^* \rangle - \Phi_h(x)^\top \Sigma^{-1} (\Sigma - \lambda_1 \cdot I_{dH}) \Theta^* + \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right) \\ &= -\lambda_1 \cdot \Phi_h(x)^\top \Sigma^{-1} \Theta^* + \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right). \end{aligned} \quad (165)$$

Applying the triangle inequality to eq. (165), we have

$$\left| R_h(x) - \widehat{R}_h(x) \right| \leq \underbrace{\lambda_1 \cdot \left| \Phi_h(x)^\top \Sigma^{-1} \Theta^* \right|}_{(i)} + \underbrace{\left| \Phi_h(x)^\top \Sigma^{-1} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right) \right|}_{(ii)}. \quad (166)$$

We then proceed to bound (i) and (ii) separately. For (i), we have

$$(i) = \lambda_1 \cdot \left| \Phi_h(x)^\top \Sigma^{-1/2} \Sigma^{-1/2} \Theta^* \right| \leq \lambda_1 \|\Phi_h(x)\|_{\Sigma^{-1}} \|\Theta^*\|_{\Sigma^{-1}} \stackrel{(i.1)}{\leq} H\sqrt{dH\lambda_1} \|\Phi_h(x)\|_{\Sigma^{-1}}, \quad (167)$$

where (i.1) follows from eq. (164) and the following inequality

$$\|\Theta^*\|_{\Sigma^{-1}} = \sqrt{\Theta^{*\top} \Sigma^{-1} \Theta^*} \leq \|\Sigma^{-1}\|_2^{1/2} \|\Theta^*\|_2 \leq H\sqrt{dH/\lambda_1}.$$

For (ii), we have

$$(ii) = \left| \Phi_h(x)^\top \Sigma^{-1/2} \Sigma^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right) \right| \leq \underbrace{\left\| \sum_{\tau \in \mathcal{D}} \Phi(\tau) \varepsilon(\tau) \right\|_{\Sigma^{-1}}}_{(iii)} \cdot \|\Phi_h(x)\|_{\Sigma^{-1}}. \quad (168)$$

Following steps similar to those in eq. (90) and Lemma B.2 in Jin et al. (2021), we have the following holds with probability at least $1 - \delta$

$$(iii) \leq H \cdot \sqrt{2 \log(1/\delta) + dH \cdot \log(1 + N/\lambda_1)},$$

which implies

$$(ii) \leq H \sqrt{2 \log(1/\delta) + dH \cdot \log(1 + N/\lambda_1)} \cdot \|\Phi_h(x)\|_{\Sigma^{-1}}. \quad (169)$$

Recalling that $b_{r,h}(x) = \|\Phi_h(x)\|_{\Sigma^{-1}}$ and substituting eq. (169) and eq. (167) into eq. (166), we can obtain

$$\left| R_h(x) - \widehat{R}_h(x) \right| \leq R_{\beta_1} \cdot b_{r,h}(x), \quad (170)$$

where R_{β_1} is an absolute constant satisfying

$$R_{\beta_1} \geq H \left(\sqrt{dH\lambda_1} + \sqrt{2 \log(1/\delta) + dH \cdot \log(1 + N/\lambda_1)} \right).$$

Letting $\lambda_1 = 1$ and $C_{\beta_1} > 0$ be a sufficiently large constant, we can verify that $R_{\beta_1} = C_{\beta_1} H \sqrt{dH \log(N/\delta)}$ satisfies the above inequality.

F.2 UNCERTAINTY OF ESTIMATED TRANSITION VALUE FUNCTION ($\widehat{\mathbb{P}}_h \widehat{V}_{h+1}(\cdot)$)

Following steps similar to those in the proof of Lemma B.1 in Jin et al. (2021), we can obtain

$$\|w^*\|_2 \leq H\sqrt{d} \quad \text{and} \quad \|\widehat{w}\|_2 \leq H\sqrt{dN/\lambda_2}. \quad (171)$$

Consider the estimation error $(\mathbb{P}_h \widehat{V}_{h+1})(\cdot) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(\cdot)$. For simplicity, we define $\varepsilon_v(x) = (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x)$ for all $x \in \mathcal{X}$. Following steps similar to those in eq. (165), we can obtain

$$\begin{aligned} & (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \\ & \leq -\lambda_2 \cdot \phi(x)^\top \Lambda_h^{-1} w_h^* + \phi(x)^\top \Lambda_h^{-1} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau) \varepsilon_v(x_h^\tau) \right). \end{aligned} \quad (172)$$

Applying the triangle inequality to eq. (172), we have

$$\begin{aligned} & \left| (\mathbb{P}_h \widehat{V}_{h+1})(x) - (\widehat{\mathbb{P}}_h \widehat{V}_{h+1})(x) \right| \\ & \leq \underbrace{\lambda_2 \cdot \left| \phi(x)^\top \Lambda_h^{-1} w_h^* \right|}_{(i)} + \underbrace{\left| \phi(x)^\top \Lambda_h^{-1} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau) \varepsilon_v(x_h^\tau) \right) \right|}_{(ii)}. \end{aligned} \quad (173)$$

Following steps similar to those in eq. (166), we can obtain

$$(i) \leq H\sqrt{d\lambda_2} \|\phi(x)\|_{\Lambda_h^{-1}}. \quad (174)$$

For (ii), we have

$$(ii) = \left| \phi(x)^\top \Lambda_h^{-1/2} \Lambda_h^{-1/2} \left(\sum_{\tau \in \mathcal{D}} \phi(x_h^\tau) \varepsilon_v(x_h^\tau) \right) \right| \leq \underbrace{\left\| \sum_{\tau \in \mathcal{D}} \phi(x_h^\tau) \varepsilon_v(x_h^\tau) \right\|_{\Lambda_h^{-1}}}_{(iii)} \|\phi(x)\|_{\Lambda_h^{-1}}. \quad (175)$$

We then proceed to upper bound the term (iii). Following steps similar to those in eq. (134) and Lemma B.2 in Jin et al. (2021), we have the following holds with probability at least $1 - \delta$

$$\begin{aligned} (iii) & \leq 2H \cdot \sqrt{\log(H \cdot \mathcal{N}_{\varepsilon,h}^v/\delta) + d \cdot \log(1 + N/\lambda_2) + 8\varepsilon^2 N/\lambda_2} \cdot \|\phi(x)\|_{\Lambda_h^{-1}}, \\ & \leq R_{\beta_2} \|\phi(x)\|_{\Lambda_h^{-1}} \end{aligned} \quad (176)$$

where R_{β_2} is an absolute constant satisfying

$$R_{\beta_2} \geq 2H \cdot \sqrt{\log(H \cdot \mathcal{N}_{\varepsilon,h}^v/\delta) + d \cdot \log(1 + N/\lambda_2) + 8\varepsilon^2 N^2/\lambda_2}, \quad (177)$$

and $\mathcal{N}_{\epsilon, h}^v$ is the cardinality of the following function class

$$\mathcal{V}_h(x, R_\theta, R_w, R_{\beta_1}, R_{\beta_2}, \lambda_1, \lambda_2) = \{\max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a)\} : \mathcal{S} \rightarrow [0, H]\}$$

$$\text{with } \|\Theta\|_2 \leq R_\theta, \|w\|_2 \leq R_w, \beta_1 \in [0, R_{\beta_1}], \beta_2 \in [0, R_{\beta_2}], \|\Sigma\|_2 \geq \lambda_1, \|\Lambda\|_2 \geq \lambda_2,$$

where $R_\theta = H\sqrt{dHN/\lambda_1}$, $R_w = H\sqrt{dN/\lambda_2}$ and

$$\begin{aligned} \bar{Q}_h(x) &= \min\{\langle \Phi_h(x), \Theta \rangle + \langle \phi(x), w \rangle \\ &\quad - \beta_1 \cdot \sqrt{\Phi_h(x)^\top \Sigma^{-1} \Phi_h(x)} - \beta_2 \cdot \sqrt{\phi(x)^\top \Lambda^{-1} \phi(x)}, H - h + 1\}^+. \end{aligned}$$

Then, following steps similar to those in Appendix B, we have

$$\begin{aligned} & \left| \max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a, \theta, w, \beta_1, \beta_2, \Sigma, \Lambda)\} - \max_{a \in \mathcal{A}} \{\bar{Q}_h(s, a, \theta', w', \beta'_1, \beta'_2, \Sigma', \Lambda')\} \right| \\ & \leq \max_{a \in \mathcal{A}} |\langle \Phi_h(x), \Theta - \Theta' \rangle| + \max_{a \in \mathcal{A}} |\langle \phi(x), w - w' \rangle| + \frac{1}{\sqrt{\lambda_1}} |\beta_1 - \beta'_1| + \frac{1}{\sqrt{\lambda_2}} |\beta_2 - \beta'_2| \\ & \quad + R_{\beta_1} \max_{a \in \mathcal{A}} \|\Phi_h(x)\|_{\Sigma^{-1}} - \|\Phi_h(x)\|_{\Sigma'^{-1}} + R_{\beta_2} \max_{a \in \mathcal{A}} \|\phi(x)\|_{\Lambda^{-1}} - \|\phi(x)\|_{\Lambda'^{-1}} \\ & \stackrel{(i)}{\leq} \|\Theta - \Theta'\|_2 + \|w - w'\|_2 + |\beta_1 - \beta'_1| + |\beta_2 - \beta'_2| \\ & \quad + R_{\beta_1} \sqrt{\|\Sigma^{-1} - \Sigma'^{-1}\|_F} + R_{\beta_2} \sqrt{\|\Lambda^{-1} - \Lambda'^{-1}\|_F}, \end{aligned} \quad (178)$$

where (i) follows from the fact that $\|\phi(x)\|_2 \leq 1$ and $\lambda_1, \lambda_2 \geq 1$. Following arguments similar to those used to obtain eq. (34) and applying Lemma 8.6 in Agarwal et al. (2019), we have

$$\begin{aligned} \log \mathcal{N}_{\epsilon, h}^v & \stackrel{(i)}{\leq} \mathcal{N}(\epsilon/6, \mathbb{R}^{dH}, R_\theta) + \mathcal{N}(\epsilon/6, \mathbb{R}^d, R_w) + \mathcal{N}(\epsilon/6, R_{\beta_1}) + \mathcal{N}(\epsilon/6, R_{\beta_2}) \\ & \quad + \mathcal{N}(\epsilon^2/(36R_{\beta_1}^2), \mathcal{F}, \sqrt{dH}/\lambda_1) + \mathcal{N}(\epsilon^2/(36R_{\beta_2}^2), \mathcal{F}, \sqrt{d}/\lambda_2) \\ & \stackrel{(ii)}{\leq} dH \log(1 + 12R_\theta/\epsilon) + d \log(1 + 12R_w/\epsilon) + \log(1 + 12R_{\beta_1}/\epsilon) + \log(1 + 12R_{\beta_2}/\epsilon) \\ & \quad + d^2 H^2 \log(1 + 36R_{\beta_1}^2 \sqrt{dH}/\epsilon^2) + d^2 \log(1 + 36R_{\beta_2}^2 \sqrt{d}/\epsilon^2) \\ & \stackrel{(iii)}{\leq} dH \log(1 + 12H\sqrt{dHN}/\epsilon) + d \log(1 + 12H\sqrt{dN}/\epsilon) \\ & \quad + \log(1 + 12C_{\beta_1} H \sqrt{dH} \log(N/\delta)/\epsilon) + \log(1 + 12R_{\beta_2}/\epsilon) \\ & \quad + d^2 H^2 \log(1 + 36C_{\beta_1}^2 dH^3 \sqrt{dH} \log(N/\delta)/\epsilon^2) + d^2 \log(1 + 36R_{\beta_2}^2 \sqrt{d}/\epsilon^2) \\ & \stackrel{(iv)}{\lesssim} C_1 d^2 H^2 \log(d^{3/2} H^{7/2} N^{1/2}/\epsilon^2) + C_2 d^2 \log(R_{\beta_2}^2 \sqrt{d}/\epsilon^2), \end{aligned} \quad (179)$$

where in (i) we use $\mathcal{N}(\epsilon, \mathbb{R}^d, B)$ to denote the ϵ -covering of ball with radius B in the space \mathbb{R}^d , $\mathcal{N}(\epsilon, B)$ to denote the ϵ -covering of interval $[0, B]$, and $\mathcal{N}(\epsilon, \mathcal{F}, B)$ to denote the ϵ -covering of the function class $\mathcal{F} = \{M : \|M\|_F \leq B\}$, (ii) follows from Lemma. 8.6 in Agarwal et al. (2019), (iii) follows from the definition of R_θ , R_w and R_{β_1} , and in (iv) we let C_1 and C_2 be sufficiently large and waive the $\log(\log(\cdot))$ term.

Substituting eq. (179) into eq. (176), we can obtain

$$\begin{aligned} & 2H \cdot \sqrt{\log(H \cdot \mathcal{N}_{\epsilon, h}^v/\delta) + d \cdot \log(1 + N/\lambda_2) + 8\epsilon^2 N/\lambda_2} \\ & \leq 2H \cdot \left(\sqrt{\log(H/\delta)} + \sqrt{\log \mathcal{N}_{\epsilon, h}^v} + \sqrt{d \cdot \log(1 + N)} + \sqrt{8\epsilon^2 N^2} \right) \\ & \leq 2H \cdot \left(\sqrt{\log(H/\delta)} + \sqrt{C_1 d^2 H^2 \log(d^{3/2} H^{7/2} N^{1/2}/\epsilon^2)} + \sqrt{C_2 d^2 \log(R_{\beta_2}^2 \sqrt{d}/\epsilon^2)} \right. \\ & \quad \left. + \sqrt{d \cdot \log(1 + N)} + \sqrt{8\epsilon^2 N^2} \right). \end{aligned} \quad (180)$$

Letting $\epsilon = (dH)^{1/4}/N$, we can see that when $R_{\beta_2} = C_{\beta_2} dH^2 \sqrt{\log(dH^3 N^{5/2}/\delta)}$, where C_{β_2} is a sufficiently large constant, we have

$$R_{\beta_2} \geq \text{R.H.S of eq. (180)},$$

which satisfies the inequality in eq. (177).

F.3 UPPER AND LOWER BOUNDS ON EVALUATION ERROR $\delta_h(\cdot)$

Using the properties that we obtained from Appendix F.1 & F.2 and following steps similar to those in Appendix D.3, we can obtain

$$0 \leq \delta_h(x) \leq 2[\beta_1 \cdot b_{r,h}(x) + \beta_2 \cdot b_{v,h}(x)], \quad \forall x \in \mathcal{X}, \quad \forall h \in [H],$$

where $\beta_1 = R_{\beta_1} = C_{\beta_1} H \sqrt{dH \log(N/\delta)}$ and $R_{\beta_2} = C_{\beta_2} dH^2 \sqrt{\log(dH^3 N^{5/2}/\delta)}$.

G SUPPORTING LEMMAS FOR OVERPARAMETERIZED NEURAL NETWORKS

The following lemma shows that an infinite-width neural network can be well-approximated by a finite-width neural network.

Lemma 4 (Approximation by Finite Sum). *Let $g(x) = \int_{\mathbb{R}^d} \sigma'(w^\top x) x^\top \ell(w) dp(w) \in \mathcal{F}_{g_1, g_2}$. Then for any $\epsilon > 0$, with probability at least $1 - \epsilon$ over w_1, \dots, w_m drawn i.i.d. from $N(0, I_d/d)$, there exist ℓ_1, \dots, ℓ_m where $\ell_i \in \mathbb{R}^d$ and $\|\ell_i\|_2 \leq g_2/\sqrt{dm}$ for all $i \in [m]$ such that the function $\hat{g}(x) = (1/\sqrt{m}) \sum_{i=1}^m \sigma'(w_i^\top x) x^\top \ell_i$ satisfies*

$$\sup_x |g(x) - \hat{g}(x)| \leq \frac{2L_\sigma g_2}{\sqrt{m}} + \frac{\sqrt{2} C_\sigma^2 g_2^2}{\sqrt{m}} \sqrt{\log\left(\frac{1}{\delta}\right)}$$

with probability at least $1 - \delta$.

Proof. The proof of Lemma 4 follows from the proof of Proposition C.1 in Gao et al. (2019) with some modifications. In Lemma 4 we consider a different distribution of w_i and upper bound on $\|\ell_i\|_2$ from those in Gao et al. (2019). First, we define the following random variable

$$a(w_1, \dots, w_m) = \sup_x |g(x) - \hat{g}(x)|.$$

Then, we proceed to show that $a(\cdot)$ is robust to the perturbation of one of its arguments. Let $\ell_i = \ell(w_i)/(\sqrt{dmp}(w_i))$. For w_1, \dots, w_m and \tilde{w}_i ($1 \leq i \leq m$), we have

$$\begin{aligned} & |a(w_1, \dots, w_m) - a(w_1, \dots, \tilde{w}_i, \dots, w_m)| \\ &= \frac{1}{\sqrt{dm}} \left| \sigma'(w_i^\top x) x^\top \ell_i - \sigma'(\tilde{w}_i^\top x) x^\top \ell_i \right| \\ &= \frac{1}{\sqrt{dm}} \left| \frac{\sigma'(w_i^\top x) x^\top \ell(w_i)}{p(w_i)} - \frac{\sigma'(\tilde{w}_i^\top x) x^\top \ell(\tilde{w}_i)}{p(\tilde{w}_i)} \right| \\ &\leq \frac{1}{\sqrt{dm}} \sup_{x \in \mathcal{X}} \left| \frac{\sigma'(w_i^\top x) x^\top \ell(w_i)}{p(w_i)} - \frac{\sigma'(\tilde{w}_i^\top x) x^\top \ell(\tilde{w}_i)}{p(\tilde{w}_i)} \right| \\ &\leq \frac{1}{\sqrt{dm}} \sup_{x \in \mathcal{X}} \left(\left| \frac{\sigma'(w_i^\top x) x^\top \ell(w_i)}{p(w_i)} \right| + \left| \frac{\sigma'(\tilde{w}_i^\top x) x^\top \ell(\tilde{w}_i)}{p(\tilde{w}_i)} \right| \right) \\ &\leq \frac{1}{\sqrt{dm}} \sup_{x \in \mathcal{X}} \left(\|\sigma'(w_i^\top x) x\|_2 \left\| \frac{\ell(w_i)}{p(w_i)} \right\|_2 + \|\sigma'(\tilde{w}_i^\top x) x\|_2 \left\| \frac{\ell(\tilde{w}_i)}{p(\tilde{w}_i)} \right\|_2 \right) \\ &\leq \frac{2C_\sigma g_2}{\sqrt{dm}} = \zeta, \end{aligned}$$

where the last inequality follows from the facts that $\|x\|_2 = 1$, $|\sigma'(\cdot)| \leq C_\sigma$ and $\sup_x \|\ell(w)/p(w)\|_2 \leq g_2$. Then, we proceed to bound the expectation of $a(\cdot)$. Note that our choice of ℓ_i ensures that $\sqrt{d} \cdot \mathbb{E}_{w_1, \dots, w_m} \hat{g}_h(\cdot) = g(\cdot)$. By symmetrization, we have

$$\begin{aligned} \mathbb{E}a &= \sqrt{d} \cdot \mathbb{E} \sup_{x \in \mathcal{X}} |\hat{g}(x) - \mathbb{E}\hat{g}(x)| \\ &\leq \frac{2\sqrt{d}}{\sqrt{m}} \cdot \mathbb{E}_{w, \varepsilon} \sup_{x \in \mathcal{X}} \left| \sum_{i=1}^m \varepsilon_i \sigma'(w_i^\top x) x^\top \ell_i \right|, \end{aligned}$$

where $\{\varepsilon_i\}_{i \in [m]}$ are a sequence of Rademacher random variables. Since $|x^\top \ell_i| \leq \|\ell_i\|_2 \leq g_2/\sqrt{m}$ and $\sigma'(\cdot)$ is L_σ -Lipschitz, we have that the function $b(\cdot) = \sigma'(\cdot)x^\top \ell_i$ is $(L_\sigma g_2/\sqrt{m})$ -Lipschitz. We then proceed as follows

$$\begin{aligned} \mathbb{E}a &\leq \frac{2\sqrt{d}}{\sqrt{m}} \cdot \mathbb{E}_{w, \varepsilon} \sup_{x \in \mathcal{X}} \left| \sum_{i=1}^m \varepsilon_i \sigma'(w_i^\top x) x^\top \ell_i \right| \\ &\stackrel{(i)}{\leq} \frac{2\sqrt{d}L_\sigma g_2}{m} \cdot \mathbb{E}_{w, \varepsilon} \sup_{x \in \mathcal{X}} \left| \left(\sum_{i=1}^m \varepsilon_i w_i \right)^\top x \right| \\ &\stackrel{(ii)}{\leq} \frac{2\sqrt{d}L_\sigma g_2}{m} \cdot \mathbb{E}_w \left\| \sum_{i=1}^m \varepsilon_i w_i \right\|_2 \\ &\stackrel{(iii)}{\leq} \frac{2\sqrt{d}L_\sigma g_2}{\sqrt{m}} \cdot \sqrt{\mathbb{E}_{w \sim N(0, I_d/d)} \|w\|_2^2} \\ &= \frac{2L_\sigma g_2}{\sqrt{m}}, \end{aligned}$$

where (i) follows from Talagrand's Lemma (Lemma 5.7) in [Mohri et al. \(2018\)](#), (ii) follows from the fact that $\|x\|_2 = 1$ for all $x \in \mathcal{X}$ and Cauchy-Schwartz inequality and (iii) follows from Jensen's inequality. Then, applying McDiarmid's inequality, we can obtain

$$\mathbb{P} \left(a \geq \frac{2L_\sigma g_2}{\sqrt{m}} + \epsilon \right) \leq \mathbb{P}(a \geq \mathbb{E}a + \epsilon) \leq \exp \left(-\frac{2\epsilon^2}{m\zeta^2} \right) = \exp \left(-\frac{m\epsilon^2}{2C_\sigma^2 g_2^2} \right).$$

Letting $\epsilon = \frac{\sqrt{2}C_\sigma^2 g_2^2}{\sqrt{m}} \sqrt{\log \left(\frac{1}{\delta} \right)}$, we have

$$\mathbb{P} \left(a \geq \frac{2L_\sigma g_2}{\sqrt{m}} + \frac{\sqrt{2}C_\sigma^2 g_2^2}{\sqrt{m}} \sqrt{\log \left(\frac{1}{\delta} \right)} \right) \leq \delta,$$

which completes the proof. \square

The following lemma bounds the perturbed gradient and value of local linearization of overparameterized neural networks around the initialization, which is provided as Lemma C.2 in [Yang et al. \(2020\)](#).

Lemma 5. *Consider the overparameterized neural network defined in Section 2.3. Consider any fixed input $x \in \mathcal{X}$. Let $R \leq c\sqrt{m}/(\log m)^3$ for some sufficiently small constant c . Then, with probability at least $1 - m^{-2}$ over the random initialization, we have for any $w \in \mathcal{B}(w_0, R)$, where $\mathcal{B}(w_0, R)$ denotes the Euclidean ball centred at w_0 with radius R , the followings hold*

$$\|\phi(x, w)\|_2 \leq C_\phi, \quad (181)$$

$$\|\phi(x, w) - \phi(x, w_0)\|_2 \leq \mathcal{O} \left(C_\phi \left(\frac{R}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right), \quad (182)$$

$$|f(x, w) - \langle \phi(x, w_0)^\top, (w - w_0) \rangle| \leq \mathcal{O} \left(C_\phi \left(\frac{R^4}{\sqrt{m}} \right)^{1/3} \sqrt{\log m} \right), \quad (183)$$

where $C_\phi = \mathcal{O}(1)$ is a constant independent from m and d .

Proof. Please see Lemma C.2 in [Yang et al. \(2020\)](#) for a detailed proof, which is based on Lemma F.1, F.2 in [Cai et al. \(2019\)](#), Lemma A.5, A.6 in [Gao et al. \(2019\)](#) and Theorem 1 in [Allen-Zhu et al. \(2019\)](#). \square

H SUPPORTING LEMMAS FOR RKHS

In this section, we provide some useful lemmas for general RKHS. Consider a variable space \mathcal{X} . Given a mapping $\phi(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^d$, we can assign a feature vector $\phi(x) \in \mathbb{R}^d$ for each $x \in \mathcal{X}$. We

further define a kernel function $K(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ as $K(x, x') = \phi(x)^\top \phi(x')$ for any $x, x' \in \mathcal{X}$. Let \mathcal{H} be a RKHS defined on \mathcal{X} with the kernel function $K(\cdot, \cdot)$. Let $\langle \cdot, \cdot \rangle_{\mathcal{H}} : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ and $\|\cdot\|_{\mathcal{H}} : \mathcal{H} \rightarrow \mathbb{R}$ denote the inner product and RKHS norm on \mathcal{H} , respectively. Since \mathcal{H} is a RKHS, there exists a feature mapping $\psi(\cdot) : \mathcal{X} \rightarrow \mathcal{H}$, such that $f(x) = \langle f(\cdot), \psi(x) \rangle_{\mathcal{H}}$ for all $f \in \mathcal{H}$ and all $x \in \mathcal{X}$. Moreover, for any $x, x' \in \mathcal{X}$ we have $K(x, x') = \langle \psi(x), \psi(x') \rangle_{\mathcal{H}}$. Without loss of generality, we further assume $\|\phi(x)\|_2 \leq C_\phi$ and $\|\psi(x)\|_{\mathcal{H}} \leq C_\psi$ for all $x \in \mathcal{X}$.

Let $\mathcal{L}^2(\mathcal{X})$ be the space of square-integrable functions on \mathcal{X} with respect to the Lebesgue measure and let $\langle \cdot, \cdot \rangle_{\mathcal{L}^2}$ be the inner product on $\mathcal{L}^2(\mathcal{X})$. The kernel function $K(\cdot, \cdot)$ induces an integral operator $T_K : \mathcal{L}^2(\mathcal{X}) \rightarrow \mathcal{L}^2(\mathcal{X})$ defined as

$$T_K f(z) = \int_{\mathcal{X}} K(x, x') \cdot f(x') dx', \quad \forall f \in \mathcal{L}^2(\mathcal{X}). \quad (184)$$

Consider the kernel function $K(\cdot, \cdot)$ of the RKHS \mathcal{H} . Let $\{x_i\}_{i=1}^\infty \subset \mathcal{X}$ be a discrete time stochastic process that is adapted to a filtration $\{\mathcal{F}_i\}_{i=0}^\infty$, i.e., x_i is \mathcal{F}_{i-1} measurable for all $i \geq 1$. We define the Gram matrix $K_N \in \mathbb{R}^{N \times N}$ and function $k_N(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^N$ as

$$K_N = [K(x_i, x_j)]_{i,j \in [N]} \in \mathbb{R}^{N \times N}, \quad k_N(x) = [K(x_1, x), \dots, K(x_N, x)]^\top \in \mathbb{R}^N. \quad (185)$$

Note that K_N and $k_N(x)$ can also be expressed as

$$K_N = \Phi \Phi^\top = \Psi \Psi^\top \in \mathbb{R}^{N \times N}, \quad \text{and} \quad k_N(x) = \Phi \phi(x) = \Psi \psi(x) \in \mathbb{R}^{N \times 1},$$

where $\Phi = [\phi(x_1), \dots, \phi(x_N)]^\top \in \mathbb{R}^{N \times d}$ and $\Psi = [\psi(x_1), \dots, \psi(x_N)]^\top \in \mathbb{R}^{N \times \infty}$. Given a regularization parameter $\lambda > 1$, we define the matrix Ω_N based on Φ and an operator Υ_N in RKHS \mathcal{H} based on Ψ as

$$\Omega_N = \Phi^\top \Phi + \lambda \cdot I_d, \quad \text{and} \quad \Upsilon_N = \Psi^\top \Psi + \lambda \cdot I_{\mathcal{H}}. \quad (186)$$

We next provide some fundamental properties for the RKHS \mathcal{H} .

Lemma 6. *For any $x \in \mathcal{X}$, considering K_N , $k_N(\cdot)$, Ω_N and Υ_N defined in eq. (185) and eq. (186), we have the followings hold*

$$\Phi^\top (K_N + I_N)^{-1} = \Omega_N^{-1} \Phi^\top, \quad (187)$$

$$\Psi^\top (K_N + I_N)^{-1} = \Upsilon_N^{-1} \Psi^\top, \quad (188)$$

$$\phi(x)^\top \Omega_N^{-1} \phi(x) \stackrel{(i)}{=} \frac{1}{\lambda} [K(x, x) - k_N(x)^\top (K_N + \lambda \cdot I_N)^{-1} k_N(x)] \stackrel{(ii)}{=} \psi(x)^\top \Upsilon_N^{-1} \psi(x). \quad (189)$$

Proof. The result in Lemma 6 can be obtained from steps spread out in Yang et al. (2020). We provide a detailed proof here for completeness.

We first proceed to prove eq. (187) and (i) in eq. (189). According to the definition of Ω_N , we have

$$\Omega_N \Phi^\top = \Phi^\top \Phi \Phi^\top + \lambda \Phi^\top = \Phi^\top (\Phi \Phi^\top + \lambda I_N) = \Phi^\top (K_N + I_N).$$

Multiplying Ω_N^{-1} on both sides of the above equality yields

$$\Phi^\top = \Omega_N^{-1} \Phi^\top (K_N + I_N),$$

which implies eq. (187) as follows

$$\Phi^\top (K_N + I_N)^{-1} = \Omega_N^{-1} \Phi^\top. \quad (190)$$

We next proceed as follows

$$\begin{aligned} \phi(x) &= \Omega_N^{-1} \Omega_N \phi(x) \\ &= \Omega_N^{-1} (\Phi^\top \Phi + \lambda \cdot I_d) \phi(x) \\ &= (\Omega_N^{-1} \Phi^\top) \Phi \phi(x) + \lambda \Omega_N^{-1} \phi(x) \\ &\stackrel{(i)}{=} \Phi^\top (K_N + I_N)^{-1} \Phi \phi(x) + \lambda \Omega_N^{-1} \phi(x), \end{aligned} \quad (191)$$

where (i) follows from eq. (190). Taking inter product with $\phi(x)$ on both sides of eq. (191) yields

$$\begin{aligned} K(x, x) &= \phi(x)^\top \phi(x) = \phi(x)^\top \Phi^\top (K_N + I_N)^{-1} \Phi \phi(x) + \lambda \phi(x)^\top \Omega_N^{-1} \phi(x) \\ &= k_N(x)^\top (K_N + I_N)^{-1} k_N(x) + \lambda \phi(x)^\top \Omega_N^{-1} \phi(x), \end{aligned}$$

which implies

$$\phi(x)^\top \Omega_N^{-1} \phi(x) = \frac{1}{\lambda} [K(x, x) - k_N(x)^\top (K_N + I_N)^{-1} k_N(x)]. \quad (192)$$

We next proceed to prove eq. (188) and (ii) in eq. (189). According to the definition of Υ_N , we have

$$\Upsilon_N \Psi^\top = \Psi^\top \Psi \Psi^\top + \lambda \Psi^\top = \Psi^\top (\Psi \Psi^\top + I_N) = \Psi^\top (K_N + I_N). \quad (193)$$

Multiplying Υ_N^{-1} on both sides of the above equality yields

$$\Psi^\top = \Upsilon_N^{-1} \Psi^\top (K_N + I_N),$$

which further implies eq. (188) as follows

$$\Psi^\top (K_N + I_N)^{-1} = \Upsilon_N^{-1} \Psi^\top. \quad (194)$$

We next proceed as follows

$$\begin{aligned} \psi(x) &= \Upsilon_N^{-1} \Upsilon_N \psi(x) \\ &= \Upsilon_N^{-1} (\Psi^\top \Psi + \lambda \cdot I_{\mathcal{H}}) \psi(x) \\ &= (\Upsilon_N^{-1} \Psi^\top) \Psi \psi(x) + \lambda \Upsilon_N^{-1} \psi(x) \\ &\stackrel{(i)}{=} \Psi^\top (K_N + I_N)^{-1} \Psi \psi(x) + \lambda \Upsilon_N^{-1} \psi(x), \end{aligned} \quad (195)$$

where (i) follows from eq. (194). Taking inter product with $\psi(x)$ on both sides of eq. (195) yields

$$\begin{aligned} K(x, x) &= \langle \psi(x), \psi(x) \rangle_{\mathcal{H}} = \psi(x)^\top \Psi^\top (K_N + I_N)^{-1} \Psi \psi(x) + \lambda \psi(x)^\top \Upsilon_N^{-1} \psi(x) \\ &= k_N(x)^\top (K_N + I_N)^{-1} k_N(x) + \lambda \psi(x)^\top \Upsilon_N^{-1} \psi(x), \end{aligned}$$

which implies

$$\psi(x)^\top \Upsilon_N^{-1} \psi(x) = \frac{1}{\lambda} [K(x, x) - k_N(x)^\top (K_N + I_N)^{-1} k_N(x)]. \quad (196)$$

Combining eq. (193) and eq. (196) completes the proof. \square

The following two lemmas characterize the concentration property of self-normalized processes.

Lemma 7 (Concentration of Self-Normalized Process in RKHS [Chowdhury & Gopalan \(2017\)](#)). *Let $\{\varepsilon_i\}_{i=1}^\infty$ be a real-valued stochastic process such that (i) $\varepsilon_i \in \mathcal{F}_t$ and (ii) ε_i is zero-mean and σ -sub-Gaussian conditioned on \mathcal{F}_{i-1} satisfying*

$$\mathbb{E}[\varepsilon_i | \mathcal{F}_{i-1}] = 0, \quad \mathbb{E}\left[e^{\kappa \varepsilon_i} \leq e^{\kappa^2 \sigma^2 / 2} | \mathcal{F}_{i-1}\right], \quad \forall \kappa \in \mathbb{R}. \quad (197)$$

Moreover, for any $t \geq 2$, let $E_N = [\varepsilon_1, \dots, \varepsilon_{N-1}]^\top \in \mathbb{R}^{N-1}$. For any $\eta > 0$ and any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have the following holds simultaneously for all $N \geq 1$:

$$E_N^\top [(K_N + \eta \cdot I_{N-1})^{-1} + I_{N-1}]^{-1} E_N \leq \sigma^2 \cdot \log \det[(1 + \eta) \cdot I_{N+1} + K_N] + 2\sigma^2 \cdot \log(1/\delta).$$

Moreover, if K_N is positive definite for all $N \geq 2$ with probability one, then the above inequality also holds with $\eta = 0$.

Lemma 8. *Let $\mathcal{G} \subset \{G : \mathcal{X} \rightarrow [0, C_g]\}$ be a class of bounded functions on \mathcal{X} . Let $\mathcal{G}_\epsilon \subset \mathcal{G}$ be the minimal ϵ -cover of \mathcal{G} such that $N_\epsilon = |\mathcal{G}_\epsilon|$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have*

$$\sup_{G \in \mathcal{G}} \left\| \sum_{i=1}^N \phi(x_i) (G(x_i) - \mathbb{E}[G(x_i) | \mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2$$

$$\leq 2C_g^2 \log \det(I + K_N/\lambda) + 2C_g^2 N(\lambda - 1) + 4C_g^2 \log(\mathcal{N}_\epsilon/\delta) + 8N^2 C_\phi^2 \epsilon^2/\lambda. \quad (198)$$

Moreover, if $G(\cdot)$ does not depend on $\{x_i\}_{i \in [N]}$, we have

$$\begin{aligned} & \left\| \sum_{i=1}^N \phi(x_i) (G(x_i) - \mathbb{E}[G(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 \\ & \leq C_g^2 \log \det(I + K_N/\lambda) + C_g^2 N(\lambda - 1) + 2C_g^2 \log(1/\delta). \end{aligned} \quad (199)$$

Proof. The proof is adapted but different from the proof of Lemma E.2 in Yang et al. (2020). We first proceed to prove eq. (198) and will show that eq. (199) can be obtained as a by-product of proving eq. (198). For any $G \in \mathcal{G}$, there exists a function G' in \mathcal{G}_ϵ such that $\sup_{x \in \mathcal{X}} |G(x) - G'(x)| \leq \epsilon$. Denote $\Delta_G(x) = G(x) - G'(x)$. We have the following holds

$$\begin{aligned} & \left\| \sum_{i=1}^N \phi(x_i) (G(x_i) - \mathbb{E}[G(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 \\ & \leq 2 \left\| \sum_{i=1}^N \phi(x_i) (G'(x_i) - \mathbb{E}[G'(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 + 2 \left\| \sum_{i=1}^N \phi(x_i) (\Delta_G(x_i) - \mathbb{E}[\Delta_G(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2. \end{aligned} \quad (200)$$

For the second term on the right hand side of eq. (200), we have

$$\left\| \sum_{i=1}^N \phi(x_i) (\Delta_G(x_i) - \mathbb{E}[\Delta_G(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 \leq N^2 C_\phi^2 \cdot (2\epsilon)^2/\lambda = 4N^2 C_\phi^2 \epsilon^2/\lambda. \quad (201)$$

To bound the first term on the right hand side of eq. (200), we apply Lemma 5 to $G'(x_i) - \mathbb{E}[G'(x_i)|\mathcal{F}_{i-1}]$. We fix $G' \in \mathcal{G}$ and let $\varepsilon_i = G'(x_i) - \mathbb{E}[G'(x_i)|\mathcal{F}_{i-1}]$ and $E_N = [\varepsilon_1, \dots, \varepsilon_{N-1}]^\top \in \mathbb{R}^{N-1}$. Using this notation, we have

$$\begin{aligned} & \left\| \sum_{i=1}^N \phi(x_i) (G'(x_i) - \mathbb{E}[G'(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 = \left\| \sum_{i=1}^N \phi(x_i) \varepsilon_i \right\|_{\Omega_N^{-1}}^2 = \|\Phi^\top E_N\|_{\Omega_N^{-1}}^2 \\ & = E_N^\top \Phi \Omega_N^{-1} \Phi^\top E_N \stackrel{(i)}{=} E_N^\top \Phi \Phi^\top (K_N + \lambda I_N)^{-1} E_N = E_N^\top K_N (K_N + \lambda I_N)^{-1} E_N \\ & \stackrel{(ii)}{\leq} E_N^\top (K_N + (\lambda - 1)I_N) (K_N + \lambda I_N)^{-1} E_N \\ & = E_N^\top (K_N + (\lambda - 1)I_N) [I_N + (K_N + (\lambda - 1)I_N)]^{-1} E_N \\ & = E_N^\top [(K_N + (\lambda - 1)I_N)^{-1} + I_N] E_N \end{aligned} \quad (202)$$

where (i) follows from eq. (187) in Lemma 6 and (ii) follows from the fact that $\lambda > 1$ and $K_N + \lambda I_N$ is positive definite. Note that each entry of E_N is bounded by C_g in absolute value. Applying Lemma 7 to eq. (202) and taking a union bound over \mathcal{G}_ϵ , for any $0 < \delta < 1$, we have the following holds with probability at least $1 - \delta$

$$\begin{aligned} & \sup_{G' \in \mathcal{G}_\epsilon} \left\| \sum_{i=1}^N \phi(x_i) (G'(x_i) - \mathbb{E}[G'(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 \\ & \leq C_g^2 \log \det[(1 + \eta)I + K_N] + 2C_g^2 \log(\mathcal{N}_\epsilon/\delta). \end{aligned} \quad (203)$$

Moreover, note that $(1 + \eta)I + K_N = [I + (1 + \eta)^{-1}K_N][(1 + \eta)I]$, which implies

$$\begin{aligned} \log \det[(1 + \eta)I + K_N] & = \log \det[I + (1 + \eta)^{-1}K_N] + N \log(1 + \eta) \\ & \leq \log \det[I + (1 + \eta)^{-1}K_N] + N\eta. \end{aligned} \quad (204)$$

Combining eq. (200), eq. (201), eq. (202), eq. (203) and eq. (204) and letting $\eta = \lambda - 1$, we have the following holds with probability $1 - \delta$

$$\left\| \sum_{i=1}^N \phi(x_i) (G(x_i) - \mathbb{E}[G(x_i)|\mathcal{F}_{i-1}]) \right\|_{\Omega_N^{-1}}^2 \leq 2C_g^2 \log \det(I + K_N/\lambda) + 2C_g^2 N(\lambda - 1) + 4C_g^2 \log(\mathcal{N}_\epsilon/\delta) + 8N^2 C_\phi^2 \epsilon^2/\lambda,$$

which completes the proof of eq. (198). To prove eq. (199) we do not need to go through the " ϵ -cover" argument since $G(\cdot)$ is independent from $\{x_i\}_{i \in N}$. We can directly apply Lemma 7 and then follow steps similar to those in eq. (204) to obtain eq. (199). \square

For any integer N and $\lambda > 0$, we define the maximal information gain associated with the RKHS \mathcal{H} as

$$\Gamma_K(N, \lambda) = \sup_{\mathcal{D} \subset \mathcal{X}} \{1/2 \cdot \log \det(I_d + \lambda^{-1} \cdot K_N)\},$$

where the supremum is taken over all discrete subset \mathcal{D} of \mathcal{X} with the cardinality no more than N .

Lemma 9 (Finite Spectrum/Effective Dimension Property). *Let $\{\sigma_j\}_{j \geq 1}$ be the eigenvalues of T_K defined in eq. (184) in the descending order. Let $\lambda \in [c_1, c_2]$ with c_1 and c_2 being absolute constants. If $\sigma_j = 0$ for all $j \geq D + 1$, where D is a positive integer. Then, we have $\Gamma_K(N, \lambda) = C_K \cdot D \cdot \log N$, where C_K is an absolute constant that depends on C_1, C_2, c_1, c_2 and C_ϕ .*

Proof. See the proof of Lemma D.5 in Yang et al. (2020) for a detailed proof. \square

I OTHER USEFUL LEMMAS

Lemma 10 (Matrix Bernstein Inequality Tropp (2015)). *Suppose that $\{A_i\}_{i=1}^N$ are independent and centered random matrices in $\mathbb{R}^{d_1 \times d_2}$, that is, $\mathbb{E}[A_i] = 0$ for all $i \in [N]$. Also, suppose $\|A_i\|_2 \leq C_A$ for all $i \in [n]$. Let $Z = \sum_{i=1}^N A_i$ and*

$$v(Z) = \max \{ \|\mathbb{E}[ZZ^\top]\|_2, \|\mathbb{E}[Z^\top Z]\|_2 \}.$$

For all $\xi \geq 0$, we have

$$\mathbb{P}(\|Z\|_2 \geq \xi) \leq (d_1 + d_2) \cdot \exp\left(-\frac{\xi^2/2}{v(Z) + C_A/3 \cdot \xi}\right).$$

Proof. See Theorem 1.6.2 in Tropp (2015) for a detailed proof. \square