

NEURAL POLICY GRADIENT METHODS: GLOBAL OPTIMALITY AND RATES OF CON- VERGENCE

Anonymous authors

Paper under double-blind review

ABSTRACT

Policy gradient methods with actor-critic schemes demonstrate tremendous empirical successes, especially when the actors and critics are parameterized by neural networks. However, it remains less clear whether such “neural” policy gradient methods converge to globally optimal policies and whether they even converge at all. We answer both the questions affirmatively in the overparameterized regime. In detail, we prove that neural natural policy gradient converges to a globally optimal policy at a sublinear rate. Also, we show that neural vanilla policy gradient converges sublinearly to a stationary point. Meanwhile, by relating the suboptimality of the stationary points to the representation power of neural actor and critic classes, we prove the global optimality of all stationary points under mild regularity conditions. Particularly, we show that a key to the global optimality and convergence is the “compatibility” between the actor and critic, which is ensured by sharing neural architectures and random initializations across the actor and critic. To the best of our knowledge, our analysis establishes the first global optimality and convergence guarantees for neural policy gradient methods.

1 INTRODUCTION

In reinforcement learning (Sutton and Barto, 2018), an agent aims to maximize its expected total reward by taking a sequence of actions according to a policy in a stochastic environment, which is modeled as a Markov decision process (MDP) (Puterman, 2014). To obtain the optimal policy, policy gradient methods (Williams, 1992; Baxter and Bartlett, 2000; Sutton et al., 2000) directly maximize the expected total reward via gradient-based optimization. As policy gradient methods are easily implementable and readily integrable with advanced optimization techniques such as variance reduction (Johnson and Zhang, 2013; Papini et al., 2018) and distributed optimization (Mnih et al., 2016; Espeholt et al., 2018), they enjoy wide popularity among practitioners. In particular, when the policy (actor) and action-value function (critic) are parameterized by neural networks, policy gradient methods achieve significant empirical successes in challenging applications, such as playing Go (Silver et al., 2016; 2017), real-time strategy gaming (Vinyals et al., 2019), robot manipulation (Peters and Schaal, 2006; Duan et al., 2016), and natural language processing (Wang et al., 2018). See Li (2017) for a detailed survey.

In stark contrast to the tremendous empirical successes, policy gradient methods remain much less well understood in terms of theory, especially when they involve neural networks. More specifically, most existing work analyzes the REINFORCE algorithm (Williams, 1992; Sutton et al., 2000), which estimates the policy gradient via Monte Carlo sampling. Based on the recent progress in non-convex optimization, Papini et al. (2018); Shen et al. (2019); Xu et al. (2019a); Karimi et al. (2019); Zhang et al. (2019) establish the rate of convergence of REINFORCE to a first- or second-order stationary point. However, the global optimality of the attained stationary point remains unclear. A more commonly used class of policy gradient methods is equipped with the actor-critic scheme

(Konda and Tsitsiklis, 2000), which alternately estimates the action-value function in the policy gradient via a policy evaluation step (critic update), and performs a policy improvement step using the estimated policy gradient (actor update). The global optimality and rate of convergence of such a class are even more challenging to analyze than that of REINFORCE. In particular, the policy evaluation step itself may converge to an undesirable stationary point or even diverge (Tsitsiklis and Van Roy, 1997), especially when it involves both nonlinear action-value function approximator, such as neural network, and temporal-difference update (Sutton, 1988). As a result, the estimated policy gradient may be biased, which possibly leads to divergence. Even if the algorithm converges to a stationary point, due to the nonconvexity of the expected total reward with respect to the policy as well as its parameter, the global optimality of such a stationary point remains unclear. The only exception is the linear-quadratic regulator (LQR) setting (Fazel et al., 2018; Malik et al., 2018; Tu and Recht, 2018; Yang et al., 2019a; Bu et al., 2019), which is, however, more restrictive than the general MDP setting that possibly involves neural networks.

To bridge the gap between practice and theory, we analyze neural policy gradient methods equipped with actor-critic schemes, where the actors and critics are represented by overparameterized two-layer neural networks. In detail, we study two settings, where the policy improvement steps are based on vanilla policy gradient and natural policy gradient, respectively. In both settings, the policy evaluation steps are based on the TD(0) algorithm (Sutton, 1988). In the first setting, we prove that neural vanilla policy gradient converges to a stationary point of the expected total reward at a $1/\sqrt{T}$ -rate in the expected squared norm of the policy gradient, where T is the number of policy improvement steps. Meanwhile, through a geometric characterization that relates the suboptimality of the stationary points to the representation power of the neural networks parameterizing the actor and critic, we establish the global optimality of all stationary points under mild regularity conditions. In the second setting, through the lens of Kullback-Leibler (KL) divergence regularization, we prove that neural natural policy gradient converges to a globally optimal policy at a $1/\sqrt{T}$ -rate in the expected total reward. In particular, a key to such global optimality and convergence guarantees is a notion of compatibility between the actor and critic, which connects the accuracy of policy evaluation steps with the efficacy of policy improvement steps. We show that such a notion of compatibility is ensured by using shared neural architectures and random initializations for both the actor and critic, which is often used as a practical heuristic (Mnih et al., 2016). To our best knowledge, our analysis gives the first global optimality and convergence guarantees for neural policy gradient methods, which corroborate their significant empirical successes.

Related Work. In contrast to the huge body of empirical literature on policy gradient methods, theoretical results on their convergence remain relatively scarce. In particular, Sutton et al. (2000) and Kakade (2002) analyze vanilla policy gradient (REINFORCE) and natural policy gradient with compatible action-value function approximators, respectively, which are further extended by Konda and Tsitsiklis (2000); Peters and Schaal (2008); Castro and Meir (2010) to incorporate actor-critic schemes. Most of this line of work only establishes the asymptotic convergence based on stochastic approximation techniques (Kushner and Yin, 2003; Borkar, 2009) and requires the actor and critic to be parameterized by linear functions. Another line of work (Papini et al., 2018; Xu et al., 2019a;b; Shen et al., 2019; Karimi et al., 2019; Zhang et al., 2019) builds on the recent progress in nonconvex optimization to establish the nonasymptotic rates of convergence of REINFORCE (Williams, 1992; Baxter and Bartlett, 2000; Sutton et al., 2000) and its variants, but only to first- or second-order stationary points, which, however, lacks global optimality guarantees. Moreover, when actor-critic schemes are involved, due to the error of policy evaluation steps and its impact on policy improvement steps, the nonasymptotic rates of convergence of policy gradient methods, even to first- or second-order stationary points, remain rather open.

Compared with the convergence of policy gradient methods, their global optimality is even less explored in terms of theory. Fazel et al. (2018); Malik et al. (2018); Tu and Recht (2018); Yang et al. (2019a); Bu et al. (2019) prove that policy gradient methods converge to globally optimal policies in the LQR setting, which is more restrictive. In very recent work, Bhandari and Russo (2019) establish the global optimality of vanilla policy gradient (REINFORCE) in the general MDP setting.

However, they require the policy class to be convex, which restricts its applicability to the tabular and LQR settings. In independent work, Agarwal et al. (2019) prove that vanilla policy gradient and natural policy gradient converge to globally optimal policies at $1/\sqrt{T}$ -rates in the tabular and linear settings. In the tabular setting, their rate of convergence of vanilla policy gradient depends on the size of the state space. In contrast, we focus on the nonlinear setting with the actor-critic scheme, where the actor and critic are parameterized by neural networks. It is worth mentioning that when such neural networks have linear activation functions, our analysis also covers the linear setting, which is, however, not our focus. In addition, Liu et al. (2019) analyze the proximal policy optimization (PPO) and trust region policy optimization (TRPO) algorithms (Schulman et al., 2015; 2017), where the actors and critics are parameterized by neural networks, and establish their $1/\sqrt{T}$ -rates of convergence to globally optimal policies. However, they require solving a subproblem of policy improvement in the functional space using multiple stochastic gradient steps in the parameter space, whereas vanilla policy gradient and natural policy gradient only require a single stochastic (natural) gradient step in the parameter space, which makes the analysis even more challenging.

There is also an emerging body of literature that analyzes the training and generalization error of deep supervised learning with overparameterized neural networks (Daniely, 2017; Jacot et al., 2018; Wu et al., 2018; Allen-Zhu et al., 2018a;b; Du et al., 2018a;b; Zou et al., 2018; Chizat and Bach, 2018; Jacot et al., 2018; Li and Liang, 2018; Cao and Gu, 2019a;b; Arora et al., 2019; Lee et al., 2019), especially when they are trained using stochastic gradient. See Fan et al. (2019) for a detailed survey. In comparison, our focus is on deep reinforcement learning with policy gradient methods. In particular, the policy evaluation steps are based on the TD(0) algorithm, which uses stochastic semigradient (Sutton, 1988) rather than stochastic gradient. Moreover, the interplay between the actor and critic makes our analysis even more challenging than that of deep supervised learning.

Notation. For distribution μ on Ω and $p > 0$, we define $\|f(\cdot)\|_{\mu,p} = (\int_{\Omega} |f|^p d\mu)^{1/p}$ as the $L_p(\mu)$ norm of f . We define $\|f(\cdot)\|_{\mu,\infty} = \inf\{C \geq 0 : |f(x)| \leq C \text{ for } \mu\text{-almost every } x\}$ as the $L_{\infty}(\mu)$ -norm of f . We write $\|f\|_{\mu,p}$ for notational simplicity when the variable of f is clear from the context. We further denote by $\|\cdot\|_{\mu}$ the $L_2(\mu)$ -norm for notational simplicity. For a vector $\phi \in \mathbb{R}^n$ and $p > 0$, we denote by $\|\phi\|_p$ the ℓ_p -norm of ϕ . We denote by $x = ([x]_1^{\top}, \dots, [x]_m^{\top})^{\top}$ a vector in \mathbb{R}^{md} , where $[x]_i \in \mathbb{R}^d$ is the i -th block of x for $i \in [m]$.

2 BACKGROUND

In this section, we introduce the background of reinforcement learning and policy gradient methods.

Reinforcement Learning. A discounted Markov decision process (MDP) is defined by tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \zeta, r, \gamma)$. Here \mathcal{S} and \mathcal{A} are the sets of all possible states and actions, respectively. Meanwhile, \mathcal{P} is the Markov transition kernel and r is the reward function, which is possibly stochastic. Specifically, when taking action $a \in \mathcal{A}$ at state $s \in \mathcal{S}$, the agent receives reward $r(s, a)$ and the environment transits into a new state according to transition probability $\mathcal{P}(\cdot | s, a)$. Meanwhile, ζ is the distribution of initial state $S_0 \in \mathcal{S}$ and $\gamma \in (0, 1)$ is the discount factor. In addition, policy $\pi(a | s)$ gives the probability of taking action a at state s . We denote the state- and action-value functions associated with π by $V^{\pi} : \mathcal{S} \rightarrow \mathbb{R}$ and $Q^{\pi} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which are defined respectively as

$$V^{\pi}(s) = (1 - \gamma) \cdot \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(S_t, A_t) \mid S_0 = s \right], \quad \forall s \in \mathcal{S}, \quad (2.1)$$

$$Q^{\pi}(s, a) = (1 - \gamma) \cdot \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(S_t, A_t) \mid S_0 = s, A_0 = a \right], \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad (2.2)$$

where $S_0 \sim \zeta(\cdot)$, $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \mathcal{P}(\cdot | S_t, A_t)$ for all $t \geq 0$. Also, we define the advantage function of policy π as the difference between Q^{π} and V^{π} , i.e., $A^{\pi}(s, a) = Q^{\pi}(s, a) -$

$V^\pi(s)$. By the definitions in (2.1) and (2.2), V^π and Q^π are related via

$$V^\pi(s) = \mathbb{E}_\pi[Q^\pi(s, a)] = \langle Q^\pi(s, \cdot), \pi(\cdot | s) \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the inner product in $\mathbb{R}^{|\mathcal{A}|}$. Here we write $\mathbb{E}_{a \sim \pi(\cdot | s)}[Q^\pi(s, a)]$ as $\mathbb{E}_\pi[Q^\pi(s, a)]$ for notational simplicity. Note that policy π together with transition probability \mathcal{P} induces a Markov chain over state space \mathcal{S} . We denote by ϱ_π the stationary state distribution of the Markov chain induced by π . We further define $\varsigma_\pi(s, a) = \pi(a | s) \cdot \varrho_\pi(s)$ as the stationary state-action distribution over $\mathcal{S} \times \mathcal{A}$. Meanwhile, policy π induces a state visitation measure over \mathcal{S} and a state-action visitation measure over $\mathcal{S} \times \mathcal{A}$, which are denoted by ν_π and σ_π , respectively. Specifically, for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we define

$$\nu_\pi(s) = (1 - \gamma) \cdot \sum_{t=0}^{\infty} \gamma^t \cdot \mathbb{P}(S_t = s), \quad \sigma_\pi(s, a) = (1 - \gamma) \cdot \sum_{t=0}^{\infty} \gamma^t \cdot \mathbb{P}(S_t = s, A_t = a), \quad (2.3)$$

where $S_0 \sim \zeta(\cdot)$, $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \mathcal{P}(\cdot | S_t, A_t)$ for all $t \geq 0$. By definition, we have $\sigma_\pi(s, a) = \pi(a | s) \cdot \nu_\pi(s)$. We define the expected total reward function $J(\pi)$ by

$$J(\pi) = (1 - \gamma) \cdot \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(S_t, A_t) \right] = \mathbb{E}_\zeta[V^\pi(s)] = \mathbb{E}_{\sigma_\pi}[r(s, a)], \quad \forall \pi, \quad (2.4)$$

where we write $\mathbb{E}_{\sigma_\pi}[r(s, a)] = \mathbb{E}_{(s, a) \sim \sigma_\pi(\cdot, \cdot)}[r(s, a)]$ for notational simplicity. The goal of reinforcement learning is to find the optimal policy that maximizes $J(\pi)$, which is denoted by π^* . When state space \mathcal{S} is large, a popular approach is to find the maximizer of $J(\pi)$ over a class of parameterized policies $\{\pi_\theta : \theta \in \mathcal{B}\}$, where $\theta \in \mathcal{B}$ is the parameter and \mathcal{B} is the parameter space. In this case, we obtain the optimization problem $\max_{\theta \in \mathcal{B}} J(\pi_\theta)$.

Policy Gradient Methods. Policy gradient methods maximize $J(\pi_\theta)$ using $\nabla_\theta J(\pi_\theta)$. These methods are based on the policy gradient theorem (Sutton and Barto, 2018), which states that

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\sigma_{\pi_\theta}}[Q^{\pi_\theta}(s, a) \cdot \nabla_\theta \log \pi_\theta(a | s)], \quad (2.5)$$

where σ_{π_θ} is the state-action visitation measure defined in (2.3). Based on (2.5), (vanilla) policy gradient maximizes the expected total reward via gradient ascent. Specifically, we generate a sequence of policy parameters $\{\theta_i\}_{i \in [T]}$ via

$$\theta_{i+1} \leftarrow \theta_i + \eta \cdot \nabla_\theta J(\pi_{\theta_i}), \quad (2.6)$$

where $\eta > 0$ is the learning rate. Meanwhile, natural policy gradient (Kakade, 2002) utilizes natural gradient ascent (Amari, 1998), which is invariant to the parameterization of policies. Specifically, let $F(\theta)$ be the Fisher information matrix corresponding to policy π_θ , which is given by

$$F(\theta) = \mathbb{E}_{\sigma_{\pi_\theta}} \left[\nabla_\theta \log \pi_\theta(a | s) [\nabla_\theta \log \pi_\theta(a | s)]^\top \right]. \quad (2.7)$$

At each iteration, natural policy gradient performs

$$\theta_{i+1} \leftarrow \theta_i + \eta \cdot F^{-1}(\theta_i) \cdot \nabla_\theta J(\pi_{\theta_i}), \quad (2.8)$$

where $F^{-1}(\theta_i)$ is the inverse of $F(\theta_i)$ and η is the learning rate. In practice, both Q^{π_θ} in (2.5) and $F(\theta)$ in (2.7) remain to be estimated, which yields approximations of the policy improvement steps in (2.6) and (2.8).

3 NEURAL POLICY GRADIENT METHODS

In this section, we represent π_θ by a two-layer neural network and study neural policy gradient methods, which estimate the policy gradient and natural policy gradient using the actor-critic scheme (Konda and Tsitsiklis, 2000).

3.1 OVERPARAMETERIZED NEURAL POLICY

We now introduce the parameterization of policies. For notational simplicity, we assume that $\mathcal{S} \times \mathcal{A} \subseteq \mathbb{R}^d$ with $d \geq 2$ and define $x = (s, a) \in \mathcal{S} \times \mathcal{A}$. Without loss of generality, we further assume

that $\|x\|_2 = \|(s, a)\|_2 \leq 1$ for all $x \in \mathcal{S} \times \mathcal{A}$. A two-layer neural network $f(x; W, b)$ with input x and width m takes the form of

$$f(x; W, b) = \frac{1}{\sqrt{m}} \sum_{r=1}^m b_r \cdot \text{ReLU}(x^\top [W]_r), \quad \forall x \in \mathcal{S} \times \mathcal{A}. \quad (3.1)$$

Here $\text{ReLU}: \mathbb{R} \rightarrow \mathbb{R}$ is the rectified linear unit (ReLU) activation function, which is defined as $\text{ReLU}(u) = \mathbb{1}\{u > 0\} \cdot u$. Also, $\{b_r\}_{r \in [m]}$ and $W = ([W]_1^\top, \dots, [W]_m^\top)^\top \in \mathbb{R}^{md}$ in (3.1) are the parameters. When training the neural network, we initialize the parameters via $[W_{\text{init}}]_r \sim N(0, I_d/d)$ and $b_r \sim \text{Unif}(\{-1, 1\})$ for all $r \in [m]$. Note that the ReLU activation function satisfies $\text{ReLU}(c \cdot u) = c \cdot \text{ReLU}(u)$ for all $c > 0$ and $u \in \mathbb{R}$. Hence, without loss of generality, we keep b_r fixed at the initial parameter throughout training and only update W in the sequel. For notational simplicity, we write $f(x; W, b)$ as $f(x; W)$ hereafter.

Using the two-layer neural network in (3.1), we define

$$\pi_\theta(a | s) = \frac{\exp[\tau \cdot f((s, a); \theta)]}{\sum_{a' \in \mathcal{A}} \exp[\tau \cdot f((s, a'); \theta)]}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad (3.2)$$

where $f(x; \theta)$ is defined in (3.1) with $\theta \in \mathbb{R}^{md}$ playing the role of W . Note that π_θ defined in (3.2) takes the form of an energy-based policy (Haarnoja et al., 2017), where τ is the temperature parameter and $f(x; \theta)$ is the energy function.

In the sequel, we investigate policy gradient methods for the class of neural policies defined in (3.2). We define the feature mapping $\phi_\theta = ([\phi_\theta]_1^\top, \dots, [\phi_\theta]_m^\top)^\top: \mathbb{R}^d \rightarrow \mathbb{R}^{md}$ of a two-layer neural network $f(x; \theta)$ as

$$[\phi_\theta]_r(x) = \frac{b_r}{\sqrt{m}} \cdot \mathbb{1}\{x^\top [\theta]_r > 0\} \cdot x, \quad \forall x \in \mathcal{S} \times \mathcal{A}, \forall r \in [m]. \quad (3.3)$$

By (3.1), it holds that $f(x; \theta) = \phi_\theta(x)^\top \theta$. Meanwhile, $f(x; \theta)$ is almost everywhere differentiable with respect to θ , and it holds that $\nabla_\theta f(x; \theta) = \phi_\theta(x)$. In the following proposition, we calculate the closed forms of the policy gradient $\nabla_\theta J(\pi_\theta)$ and the Fisher information matrix $F(\theta)$ for π_θ defined in (3.2).

Proposition 3.1 (Policy Gradient and Fisher Information Matrix). For π_θ defined in (3.2), we have

$$\nabla_\theta J(\pi_\theta) = \tau \cdot \mathbb{E}_{\sigma_{\pi_\theta}} \left[Q^{\pi_\theta}(s, a) \cdot \left(\phi_\theta(s, a) - \mathbb{E}_{\pi_\theta}[\phi_\theta(s, a')] \right) \right], \quad (3.4)$$

$$F(\theta) = \tau^2 \cdot \mathbb{E}_{\sigma_{\pi_\theta}} \left[\left(\phi_\theta(s, a) - \mathbb{E}_{\pi_\theta}[\phi_\theta(s, a')] \right) \left(\phi_\theta(s, a) - \mathbb{E}_{\pi_\theta}[\phi_\theta(s, a')] \right)^\top \right], \quad (3.5)$$

where $\phi_\theta(s, a)$ is defined in (3.3) with $x = (s, a)$, τ is the temperature parameter, and σ_{π_θ} is the state-action visitation measure defined in (2.3). Here we write $\mathbb{E}_{\pi_\theta}[\phi_\theta(s, a')] = \mathbb{E}_{a' \sim \pi_\theta(\cdot | s)}[\phi_\theta(s, a')]$ for notational simplicity.

Proof. See §H.1 for a detailed proof. □

Since the action-value function Q^{π_θ} in (3.4) is unknown, to obtain the policy gradient, we use another neural network to track the action-value function of policy π_θ . Specifically, we use a two-layer neural network $Q_\omega(s, a) = f((s, a); \omega)$ defined in (3.1) to represent the action-value function Q^{π_θ} , where ω plays the same role as W in (3.1). Such an approach is known as the actor-critic scheme (Konda and Tsitsiklis, 2000). We call π_θ and Q_ω the actor and critic, respectively. We highlight that in the overparameterized regime where the width of two-layer neural networks m is large, a shared architecture and random initialization between the actor Q_ω and the energy function of critic π_θ ensures approximate compatible function approximations. See §B for details.

3.2 NEURAL POLICY GRADIENT METHODS

Now we present neural policy gradient and neural natural policy gradient. Following the actor-critic scheme, they generate a sequence of policies and action-value functions $\{(\pi_{\theta_i}, Q_{\omega_i})\}_{i \in [T]}$.

3.2.1 ACTOR UPDATE

As introduced in §2, we aim to solve the optimization problem $\max_{\theta \in \mathcal{B}} J(\pi_\theta)$ iteratively via gradient-based methods, where \mathcal{B} is the parameter space. We set $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$, where $R > 1$ and W_{init} is the initial parameter defined in §3.1. For all $i \in [T]$, let θ_i be the policy parameter at the i -th iteration. For notational simplicity, in the sequel, we denote by σ_i and ζ_i the state-action visitation measure $\sigma_{\pi_{\theta_i}}$ and the stationary state-action distribution $\zeta_{\pi_{\theta_i}}$, respectively, which are defined in §2. Similarly, we write $\nu_i = \nu_{\pi_{\theta_i}}$ and $\varrho_i = \varrho_{\pi_{\theta_i}}$. To update θ_i , we set

$$\theta_{i+1} \leftarrow \Pi_{\mathcal{B}}(\theta_i + \eta \cdot G(\theta_i) \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i})), \quad (3.6)$$

where we define $\Pi_{\mathcal{B}}: \mathbb{R}^{md} \rightarrow \mathcal{B}$ as the projection operator onto the parameter space $\mathcal{B} \subseteq \mathbb{R}^{md}$. Here $G(\theta_i) \in \mathbb{R}^{md \times md}$ is a matrix specific to each algorithm. Specifically, $G(\theta_i)$ is the identity matrix for policy gradient and $G(\theta_i) = F^{-1}(\theta_i)$ for natural policy gradient, where $F(\theta_i)$ is the Fisher information matrix in (3.5). Meanwhile, η is the learning rate and $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ is an estimator of $\nabla_\theta J(\pi_{\theta_i})$, which takes the form of

$$\widehat{\nabla}_\theta J(\pi_{\theta_i}) = \frac{\tau_i}{B} \cdot \sum_{\ell=1}^B Q_{\omega_i}(s_\ell, a_\ell) \cdot \nabla_\theta \log \pi_{\theta_i}(a_\ell | s_\ell). \quad (3.7)$$

Here τ_i is the temperature parameter of π_{θ_i} , $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$ is sampled from the state-action visitation measure σ_i corresponding to the current policy π_{θ_i} , and $B > 0$ is the batch size. Also, Q_{ω_i} is the critic obtained by Algorithm 2. See §C for the sampling from visitation measures.

Inverting Fisher Information Matrix. Recall that $G(\theta_i)$ is the inverse of the Fisher information matrix used in natural policy gradient. In the overparameterized regime, inverting an estimator $\widehat{F}(\theta_i)$ of $F(\theta_i)$ can be infeasible as $\widehat{F}(\theta_i)$ is a high-dimensional matrix, which is possibly not invertible. To resolve this issue, we estimate the natural policy gradient $G(\theta_i) \cdot \nabla_\theta J(\pi_{\theta_i})$ by solving

$$\min_{\alpha \in \mathcal{B}} \|\widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i})\|_2, \quad (3.8)$$

where $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ is defined in (3.7), τ_i is the temperature parameter in π_{θ_i} , and \mathcal{B} is the parameter space. Meanwhile, $\widehat{F}(\theta)$ is an unbiased estimator of $F(\theta_i)$ based on $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$ sampled from σ_i , which is defined as

$$\widehat{F}(\theta_i) = \frac{\tau_i^2}{B-1} \cdot \sum_{\ell=1}^B \left(\phi_{\theta_i}(s_\ell, a_\ell) - \frac{1}{B} \cdot \sum_{j=1}^B \phi_{\theta_i}(s_j, a_j) \right) \left(\phi_{\theta_i}(s_\ell, a_\ell) - \frac{1}{B} \cdot \sum_{j=1}^B \phi_{\theta_i}(s_j, a_j) \right)^\top, \quad (3.9)$$

where ϕ_{θ_i} is defined in (3.3). The actor update of neural natural policy gradient takes the form of

$$\tau_{i+1} \leftarrow \tau_i + \eta, \quad \theta_{i+1} \leftarrow \tau_i \cdot \theta_i + \eta \cdot \operatorname{argmin}_{\alpha \in \mathcal{B}} \|\widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i})\|_2, \quad (3.10)$$

where we use an arbitrary minimizer of (3.8) if it is not unique. Note that we also update the temperature parameter by $\tau_{i+1} \leftarrow \tau_i + \eta$, which ensures $\theta_{i+1} \in \mathcal{B}$. It is worth mentioning that up to minor modifications, our analysis allows for approximately solving (3.8), which is the common practice of approximate second-order optimization (Martens and Grosse, 2015; Wu et al., 2017).

To summarize, at the i -th iteration, neural policy gradient obtains θ_{i+1} via projected gradient ascent using $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ defined in (3.7). Meanwhile, neural natural policy gradient solves (3.8) and obtains θ_{i+1} according to (3.10).

3.2.2 CRITIC UPDATE

To obtain $\widehat{\nabla}_\theta J(\pi_\theta)$, it remains to obtain the critic Q_{ω_i} in (3.7). For any policy π , Q^π is the unique solution to the Bellman equation $Q = \mathcal{T}^\pi Q$ (Sutton and Barto, 2018). Here \mathcal{T}^π is the Bellman operator that takes the form of

$$\mathcal{T}^\pi Q(s, a) = \mathbb{E}[(1 - \gamma) \cdot r(s, a) + Q(s', a')], \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where $s' \sim \mathcal{P}(\cdot | s, a)$ and $a' \sim \pi(\cdot | s')$. Correspondingly, we aim to solve the following optimization problem

$$\omega_i \leftarrow \operatorname{argmin}_{\omega \in \mathcal{B}} \mathbb{E}_{\varsigma_i} \left[\left(Q_\omega(s, a) - \mathcal{T}^{\pi_{\theta_i}} Q_\omega(s, a) \right)^2 \right], \quad (3.11)$$

where ς_i and $\mathcal{T}^{\pi_{\theta_i}}$ are the stationary state-action distribution and the Bellman operator associated with π_{θ_i} , respectively, and \mathcal{B} is the parameter space. We adopt the neural temporal-difference learning (TD) algorithm studied in Cai et al. (2019), which solves the optimization problem in (3.11) via stochastic semigradient descent (Sutton, 1988). Specifically, an iteration of neural TD takes the form of

$$\omega(t+1/2) \leftarrow \omega(t) - \eta_{\text{TD}} \cdot \left(Q_{\omega(t)}(s, a) - r(s, a) - \gamma Q_{\omega(t)}(s', a') \right) \cdot \nabla_\omega Q_{\omega(t)}(s, a), \quad (3.12)$$

$$\omega(t+1) \leftarrow \operatorname{argmin}_{\alpha \in \mathcal{B}} \|\alpha - \omega(t+1/2)\|_2, \quad (3.13)$$

where $(s, a) \sim \varsigma_i(\cdot)$, $s' \sim \mathcal{P}(\cdot | s, a)$, and $a' \sim \pi(\cdot | s')$. Here (3.12) is the stochastic semigradient step, and (3.13) projects the parameter obtained by (3.12) back to the parameter space \mathcal{B} . Meanwhile, the state-action pairs in (3.12) are sampled from the stationary state-action distribution ς_i , which is achieved by sampling from the Markov chain induced by π_{θ_i} until it mixes. See Algorithm 2 in §F for details. Finally, combining the actor updates and the critic update described in (3.6), (3.10), and (3.11), respectively, we obtain neural policy gradient and natural policy gradient, which are described in Algorithm 1.

Algorithm 1 Neural Policy Gradient Methods

Require: Number of iterations T , number of TD iterations T_{TD} , learning rate η , learning rate of TD iterations η_{TD} , temperature parameters τ_1, \dots, τ_T , batch size B .

- 1: **Initialization:** Initialize $b_r \sim \text{Unif}(\{-1, 1\})$ and $[W_{\text{init}}]_r \sim N(0, I_d/d)$ for all $r \in [m]$. Set $\mathcal{B} \leftarrow \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ and $\theta_1 \leftarrow W_{\text{init}}$.
- 2: **for** $i \in [T]$ **do**
- 3: Update ω_i using Algorithm 2 with π_{θ_i} as the input, $\omega(0) \leftarrow W_{\text{init}}$ and $\{b_r\}_{r \in [m]}$ as the initialization, T_{TD} as the number of iterations, and η_{TD} as the learning rate.
- 4: Sample $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$ from the visitation measure σ_i , and estimate $\widehat{\nabla}_\theta J(\pi_\theta)$ and $\widehat{F}(\theta_i)$ using (3.7) and (3.9), respectively.
- 5: If using policy gradient, update θ_{i+1} by

$$\theta_{i+1} \leftarrow \Pi_{\mathcal{B}}(\theta_i + \eta \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i})).$$

If using natural policy gradient, update θ_{i+1} and τ_{i+1} by

$$\tau_{i+1} \leftarrow \tau_i + \eta, \quad \tau_{i+1} \cdot \theta_{i+1} \leftarrow \tau_i \cdot \theta_i + \eta \cdot \operatorname{argmin}_{\alpha \in \mathcal{B}} \|\widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i})\|_2.$$

6: **end for**

7: **Output:** $\{\pi_{\theta_i}\}_{i \in [T+1]}$.

4 MAIN RESULTS

In this section, we establish the global optimality and convergence for neural policy gradient methods. Hereafter, we assume that the absolute value of the reward function r is upper bounded by an absolute constant $Q_{\max} > 0$. As a result, we obtain from (2.1) and (2.2) that $|V^\pi(s, a)| \leq Q_{\max}$, $|Q^\pi(s, a)| \leq Q_{\max}$, and $|A^\pi(s, a)| \leq 2Q_{\max}$ for all π and $(s, a) \in \mathcal{S} \times \mathcal{A}$. In what follows, we show that neural policy gradient converges to a stationary point of $J(\pi_\theta)$ with respect to θ at a sub-linear rate. We further characterize the geometry of $J(\pi_\theta)$ and establish the global optimality of the obtained stationary point. We defer the global optimality and convergence of neural natural policy gradient to §A.

In the sequel, we study the convergence of neural policy gradient, i.e., Algorithm 1 with (3.6) as actor update, where $G(\theta) = I_{md}$. In what follows, we lay out an assumption on the action-value function Q^π . Recall that we write $x = (s, a)$.

Assumption 4.1 (Action-Value Function Class). We define

$$\mathcal{F}_{R,\infty} = \left\{ f(x) = f_0(x) + \int \mathbb{1}\{w^\top x > 0\} \cdot x^\top \iota(w) d\mu(w) : \|\iota(w)\|_\infty \leq R/\sqrt{d} \right\},$$

where $\mu: \mathbb{R}^d \rightarrow \mathbb{R}$ is the density function of the Gaussian distribution $N(0, I_d/d)$ and $f_0(\cdot) = f(\cdot; W_{\text{init}})$ is the neural network corresponding to the initial parameter W_{init} . We assume that $Q^\pi \in \mathcal{F}_{R,\infty}$ for all π .

Assumption 4.1 is a mild regularity condition on Q^π , as $\mathcal{F}_{R,\infty}$ captures a sufficiently general family of functions, which constitute a subset of the reproducing kernel Hilbert space (RKHS) induced by the random feature $\mathbb{1}\{w^\top x > 0\} \cdot x$ with $w \sim N(0, I_d/d)$ (Rahimi and Recht, 2008; 2009) up to the shift of f_0 . Similar assumptions are imposed in the analysis of batch reinforcement learning in RKHS (Farahmand et al., 2016).

In what follows, we lay out a regularity condition on the state visitation measure ν_π and the stationary state distribution ϱ_π .

Assumption 4.2 (Regularity Condition on ν_π and ϱ_π). Let π and $\tilde{\pi}$ be two arbitrary policies. We assume that there exists an absolute constant $c > 0$ such that

$$\mathbb{E}_{\tilde{\pi}, \nu_\pi} \left[\mathbb{1}\{|y^\top(s, a)| \leq u\} \right] \leq c \cdot u / \|y\|_2, \quad \mathbb{E}_{\tilde{\pi}, \varrho_\pi} \left[\mathbb{1}\{|y^\top(s, a)| \leq u\} \right] \leq c \cdot u / \|y\|_2, \quad \forall y \in \mathbb{R}^d, \forall u > 0.$$

Here the expectations are taken over the joint distributions $\tilde{\pi}(\cdot | \cdot) \cdot \nu_\pi(\cdot)$ and $\tilde{\pi}(\cdot | \cdot) \cdot \varrho_\pi(\cdot)$ over $\mathcal{S} \times \mathcal{A}$, respectively.

Assumption 4.2 essentially imposes a regularity condition on the Markov transition kernel \mathcal{P} of the MDP as \mathcal{P} determines ν_π and ϱ_π for all π . Such a regularity condition holds if both ν_π and ϱ_π have upper-bounded density functions for all π .

After introducing these regularity conditions, we present the following proposition adapted from Cai et al. (2019), which characterizes the global convergence of neural TD for critic update.

Proposition 4.3 (Convergence of Critic Update). We set $\eta_{\text{TD}} = \min\{(1 - \gamma)/8, 1/\sqrt{T_{\text{TD}}}\}$ in Algorithm 1. Let Q_{ω_i} be the output of the i -th critic update in Line 3 of Algorithm 1, which is an estimator of $Q^{\pi_{\theta_i}}$ obtained by Algorithm 2 with T_{TD} iterations. Under Assumptions 4.1 and 4.2, it holds for $T_{\text{TD}} = \Omega(m)$ that

$$\mathbb{E}_{\text{init}} [\|Q_{\omega_i} - Q^{\pi_{\theta_i}}\|_{\zeta_i}^2] = \mathcal{O}(R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}), \quad (4.1)$$

where ζ_i is the stationary state-action distribution corresponding to π_{θ_i} . Here the expectation is taken over the random initialization.

Proof. See §F.1 for a detailed proof. \square

Cai et al. (2019) show that the error of critic update consists of two parts, namely the approximation error of neural networks and the algorithmic error of neural TD. The former decays as the width m grows, while the latter decays as the number of neural TD iterations T_{TD} in Algorithm 2 grows. By setting $T_{\text{TD}} = \Omega(m)$, the algorithmic error in (4.1) of Proposition 4.3 is dominated by the approximation error. In contrast with Cai et al. (2019), we obtain a more refined convergence characterization under the more restrictive assumption that $Q^\pi \in \mathcal{F}_{R,\infty}$. Specifically, such a restriction allows us to obtain the upper bound of the mean squared error in (4.1) of Proposition 4.3.

It now remains to establish the global convergence of actor update, which involves the estimator $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ of the policy gradient $\nabla_\theta J(\pi_{\theta_i})$ based on $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$. We introduce the following regularity condition on the variance of $\widehat{\nabla}_\theta J(\pi_{\theta_i})$.

Assumption 4.4 (Variance Upper Bound). Recall that σ_i is the state-action visitation measure corresponding to π_{θ_i} for all $i \in [T]$. Let $\xi_i = \widehat{\nabla}_\theta J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_\theta J(\pi_{\theta_i})]$, where $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ is defined in (3.7). We assume that there exists an absolute constant $\sigma_\xi > 0$ such that $\mathbb{E}[\|\xi_i\|_2^2] \leq \tau_i^2 \cdot \sigma_\xi^2 / B$ for all $i \in [T]$. Here the expectations are taken over σ_i given θ_i and ω_i .

Assumption 4.4 is a mild regularity condition. Such a regularity condition holds if the Markov chain that generates $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$ mixes sufficiently fast and $Q_{\omega_i}(s, a)$ with $(s, a) \sim \sigma_i$ have upper bounded second moments for all $i \in [T]$. Zhang et al. (2019) verify that under certain regularity conditions, similar unbiased policy gradient estimators have almost surely upper bounded norms, which implies Assumption 4.4. Similar regularity conditions are also imposed in the analysis of policy gradient methods by Xu et al. (2019a;b).

In what follows, we impose a regularity condition on the discrepancy between the state-action visitation measure and the stationary state-action distribution corresponding to the same policy.

Assumption 4.5 (Regularity Condition on σ_i and ς_i). We assume that there exists an absolute constant $\kappa > 0$ such that

$$\left\{ \mathbb{E}_{\varsigma_i} \left[\left(\frac{d\sigma_i}{d\varsigma_i}(s, a) \right)^2 \right] \right\}^{1/2} \leq \kappa, \quad \forall i \in [T]. \quad (4.2)$$

Here $d\sigma_i/d\varsigma_i$ is the Radon-Nikodym derivative of σ_i with respect to ς_i .

We highlight that if the MDP is initialized at the stationary distribution ς_i , the state-action visitation measure σ_i is the same as ς_i . Meanwhile, if the induced Markov state-action chain mixes sufficiently fast, such an assumption also holds. A similar regularity condition is imposed by Scherrer (2013), which assumes that the L_∞ -norm of $d\sigma_i/d\varsigma_i$ is upper bounded, whereas we only assume that its L_2 -norm is upper bounded.

Meanwhile, we impose the following regularity condition on the smoothness of the expected total reward $J(\pi_\theta)$ with respect to θ .

Assumption 4.6 (Lipschitz Continuous Policy Gradient). We assume that $\nabla_\theta J(\pi_\theta)$ is L -Lipschitz continuous with respect to θ , where $L > 0$ is an absolute constant.

Such an assumption holds when the transition probability $\mathcal{P}(\cdot | s, a)$ and the reward function r are both Lipschitz continuous with respect to their inputs (Piotto et al., 2015). Also, Karimi et al. (2019); Zhang et al. (2019); Xu et al. (2019b) verify the Lipschitz continuity of the policy gradient under certain regularity conditions.

Note that we restrict θ to the parameter space \mathcal{B} . Here we call $\hat{\theta} \in \mathcal{B}$ a stationary point of $J(\pi_\theta)$ if it holds for all $\theta \in \mathcal{B}$ that $\nabla_\theta J(\pi_{\hat{\theta}})^\top (\theta - \hat{\theta}) \leq 0$. We now show that the sequence $\{\theta_i\}_{i \in [T+1]}$ generated by neural policy gradient converges to a stationary point at a sublinear rate.

Theorem 4.7 (Convergence to Stationary Point). We set $\tau_i = 1$, $\eta = 1/\sqrt{T}$, $\eta_{TD} = \min\{(1 - \gamma)/8, 1/\sqrt{T_{TD}}\}$, $T_{TD} = \Omega(m)$, and $\mathcal{B} = \{\alpha : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ in Algorithm 1, where the actor update is given in (3.6) with $G(\theta) = I_{md}$. For all $i \in [T]$, we define

$$\rho_i = \eta^{-1} \cdot \left[\Pi_{\mathcal{B}}(\theta_i + \eta \cdot \nabla_\theta J(\pi_{\theta_i})) - \theta_i \right] \in \mathbb{R}^{md}, \quad (4.3)$$

where $\Pi_{\mathcal{B}}: \mathbb{R}^{md} \rightarrow \mathcal{B}$ is the projection operator onto $\mathcal{B} \subseteq \mathbb{R}^{md}$. Under the assumptions of Proposition 4.3 and Assumptions 4.4-4.6, for $T \geq 4L^2$ we have

$$\min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2^2] \leq 8/\sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{T+1}}) - J(\pi_{\theta_1})] + 8\sigma_\xi^2/B + \varepsilon_Q(T),$$

where κ is defined in (4.2) of Assumption 4.2 and

$$\varepsilon_Q(T) = \kappa \cdot \mathcal{O}(R^{5/2} \cdot m^{-1/4} \cdot T^{1/2} + R^{9/4} \cdot m^{-1/8} \cdot T^{1/2}).$$

Here the expectations are taken over all the randomness.

Proof. See §D.1 for a detailed proof. \square

By Theorem 4.7 with $m = \Omega(T^8 \cdot R^{18})$ and $B = \Omega(\sqrt{T})$, we obtain $\min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2^2] = \mathcal{O}(1/\sqrt{T})$. Therefore, when the neural networks are sufficiently wide and the batch size B is sufficiently large, neural policy gradient achieves a $1/\sqrt{T}$ -rate of convergence. Moreover, ρ_i defined in (4.3) is known as the gradient mapping at θ_i (Nesterov, 2018). It is known that $\hat{\theta} \in \mathcal{B}$ is a stationary point if and only if the gradient mapping at $\hat{\theta}$ is a zero vector. Therefore, (a subsequence

of $\{\theta_i\}_{i \in [T]}$ converges to a stationary point $\hat{\theta} \in \mathcal{B}$ as $\min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2^2]$ converges to zero. In other words, neural policy gradient converges to a stationary point at a $1/\sqrt{T}$ -rate. Also, we remark that the projection operator in actor update is adopted only for the purpose of simplicity, which can be removed with more refined analysis. Moreover, the projection-free version of neural policy gradient converges to a stationary point at a similar sublinear rate. See §G for details.

We now characterize the global optimality of the obtained stationary point $\hat{\theta}$. To this end, we compare the expected total reward of $\pi_{\hat{\theta}}$ with that of the global optimum π^* of $J(\pi)$.

Theorem 4.8 (Global Optimality of Stationary Point). Let $\hat{\theta} \in \mathcal{B}$ be a stationary point of $J(\pi_{\hat{\theta}})$. It holds that

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \inf_{\theta \in \mathcal{B}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}},$$

where Q_{\max} is the upper bound of $|r|$ and $u_{\hat{\theta}}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined as

$$u_{\hat{\theta}}(s, a) = \frac{d\sigma_{\pi^*}}{d\sigma_{\pi_{\hat{\theta}}}}(s, a) - \frac{d\nu_{\pi^*}}{d\nu_{\pi_{\hat{\theta}}}}(s) + \phi_{\hat{\theta}}(s, a)^\top \hat{\theta}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (4.4)$$

Here $d\sigma_{\pi^*}/d\sigma_{\pi_{\hat{\theta}}}$ and $d\nu_{\pi^*}/d\nu_{\pi_{\hat{\theta}}}$ are the Radon-Nikodym derivatives, and $\|\cdot\|_{\sigma_{\pi_{\hat{\theta}}}}$ is the $L_2(\sigma_{\pi_{\hat{\theta}}})$ -norm.

Proof. See §D.2 for a detailed proof. \square

To understand Theorem 4.8, we highlight that for $\theta, \hat{\theta} \in \mathcal{B}$, the function $\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta$ is well approximated by the overparameterized neural network $f((\cdot, \cdot); \theta)$. See Corollary E.4 for details. Therefore, the global optimality of $\pi_{\hat{\theta}}$ depends on the error of approximating $u_{\hat{\theta}}$ with an overparameterized two-layer neural network. Specifically, if $u_{\hat{\theta}}$ is well approximated by an overparameterized two-layer neural network, then $\pi_{\hat{\theta}}$ is nearly as optimal as π^* . In the following corollary, we formally establish a sufficient condition for any stationary point $\hat{\theta}$ to be globally optimal.

Theorem 4.9 (Global Optimality of Stationary Point). Let $\hat{\theta} \in \mathcal{B}$ be a stationary point of $J(\pi_{\hat{\theta}})$. We assume that $u_{\hat{\theta}} \in \mathcal{F}_{R, \infty}$ in Theorem 4.8. Under Assumption 4.2, it holds that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}}[J(\pi^*) - J(\pi_{\hat{\theta}})] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}).$$

More generally, without assuming $u_{\hat{\theta}} \in \mathcal{F}_{R, \infty}$ in Theorem 4.8, under Assumption 4.2, it holds that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}}[J(\pi^*) - J(\pi_{\hat{\theta}})] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}) + \mathbb{E}_{\text{init}}[\|\Pi_{\mathcal{F}_{R, \infty}} u_{\hat{\theta}} - u_{\hat{\theta}}\|_{\sigma_{\pi_{\hat{\theta}}}}].$$

Here the expectation is taken over the random initialization, and $\Pi_{\mathcal{F}_{R, \infty}}$ is the projection operator onto $\mathcal{F}_{R, \infty}$ with respect to the $L_2(\sigma_{\pi_{\hat{\theta}}})$ -norm.

Proof. See §H.2 for a detailed proof. \square

By Theorem 4.9, a stationary point $\hat{\theta}$ is globally optimal if $u_{\hat{\theta}} \in \mathcal{F}_{R, \infty}$ and $m \rightarrow \infty$. Moreover, following from the definition of ρ_i in (4.3) of Theorem 4.7, we have that

$$\nabla_{\theta} J(\pi_{\theta_i})^\top (\theta - \theta_i) \leq (2R + 2\eta \cdot Q_{\max}) \cdot \|\rho_i\|_2, \quad \forall \theta \in \mathcal{B}. \quad (4.5)$$

Since $\|\rho_i\|_2 = 0$ implies that θ_i is a stationary point, the right-hand side of (4.5) quantifies the deviation of θ_i from a stationary point $\hat{\theta}$. See §H.3 for details. Following similar analysis to §D.2 and §H.2, if $u_{\theta_i} \in \mathcal{F}_{R, \infty}$ for all $i \in [T]$, we obtain that

$$(1 - \gamma) \cdot \min_{i \in [T]} \mathbb{E}[J(\pi^*) - J(\pi_{\theta_i})] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}) + (2R + 2\eta \cdot Q_{\max}) \cdot \min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2].$$

Thus, by invoking Theorem 4.7, it holds for sufficiently large m and B that the expected total reward $J(\pi_{\theta_i})$ converges to the global optimum $J(\pi^*)$ at a $1/T^{1/4}$ -rate.

REFERENCES

- Agarwal, A., Kakade, S. M., Lee, J. D. and Mahajan, G. (2019). Optimality and approximation with policy gradient methods in Markov decision processes. *arXiv preprint arXiv:1908.00261*.
- Allen-Zhu, Z., Li, Y. and Liang, Y. (2018a). Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*.
- Allen-Zhu, Z., Li, Y. and Song, Z. (2018b). A convergence theory for deep learning via overparameterization. *arXiv preprint arXiv:1811.03962*.
- Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Computation*, **10** 251–276.
- Antos, A., Szepesvári, C. and Munos, R. (2008). Learning near-optimal policies with Bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, **71** 89–129.
- Arora, S., Du, S. S., Hu, W., Li, Z. and Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*.
- Barron, A. R. (1994). Approximation and estimation bounds for artificial neural networks. *Machine Learning*, **14** 115–133.
- Baxter, J. and Bartlett, P. L. (2000). Direct gradient-based reinforcement learning. In *International Symposium on Circuits and Systems*.
- Bhandari, J. and Russo, D. (2019). Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*.
- Borkar, V. S. (2009). *Stochastic Approximation: A Dynamical Systems Viewpoint*, vol. 48. Springer.
- Bu, J., Mesbahi, A., Fazel, M. and Mesbahi, M. (2019). LQR through the lens of first order methods: Discrete-time case. *arXiv preprint arXiv:1907.08921*.
- Cai, Q., Yang, Z., D. Lee, J. and Wang, Z. (2019). Neural temporal-difference learning converges to global optima. *arXiv preprint arXiv:1905.10027*.
- Cao, Y. and Gu, Q. (2019a). Generalization bounds of stochastic gradient descent for wide and deep neural networks. *arXiv preprint arXiv:1905.13210*.
- Cao, Y. and Gu, Q. (2019b). A generalization theory of gradient descent for learning overparameterized deep ReLU networks. *arXiv preprint arXiv:1902.01384*.
- Castro, D. D. and Meir, R. (2010). A convergent online single time scale actor critic algorithm. *Journal of Machine Learning Research*, **11** 367–410.
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. *arXiv preprint arXiv:1905.00360*.
- Chizat, L. and Bach, F. (2018). A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*.
- Daniely, A. (2017). SGD learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*.
- Du, S. S., Lee, J. D., Li, H., Wang, L. and Zhai, X. (2018a). Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*.

- Du, S. S., Zhai, X., Poczos, B. and Singh, A. (2018b). Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J. and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I. et al. (2018). IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*.
- Fan, J., Ma, C. and Zhong, Y. (2019). A selective overview of deep learning. *arXiv preprint arXiv:1904.05526*.
- Farahmand, A.-m., Ghavamzadeh, M., Szepesvári, C. and Mannor, S. (2016). Regularized policy iteration with nonparametric function spaces. *Journal of Machine Learning Research*, **17** 4809–4874.
- Farahmand, A.-m., Szepesvári, C. and Munos, R. (2010). Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*.
- Fazel, M., Ge, R., Kakade, S. M. and Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*.
- Funahashi, K.-I. (1989). On the approximate realization of continuous mappings by neural networks. *Neural Networks*, **2** 183–192.
- Haarnoja, T., Tang, H., Abbeel, P. and Levine, S. (2017). Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*.
- Jacot, A., Gabriel, F. and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*.
- Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*.
- Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*.
- Kakade, S. M. (2002). A natural policy gradient. In *Advances in Neural Information Processing Systems*.
- Karimi, B., Miasojedow, B., Moulines, E. and Wai, H.-T. (2019). Non-asymptotic analysis of biased stochastic approximation scheme. *arXiv preprint arXiv:1902.00629*.
- Klusowski, J. M. and Barron, A. R. (2016). Risk bounds for high-dimensional ridge function combinations including neural networks. *arXiv preprint arXiv:1607.01434*.
- Konda, V. (2002). *Actor-Critic Algorithms*. Ph.D. thesis, Massachusetts Institute of Technology.
- Konda, V. R. and Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *Advances in Neural Information Processing Systems*.
- Kushner, H. and Yin, G. G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications*, vol. 35. Springer Science & Business Media.
- Lazaric, A., Ghavamzadeh, M. and Munos, R. (2016). Analysis of classification-based policy iteration algorithms. *Journal of Machine Learning Research*, **17** 583–612.

- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J. and Sohl-Dickstein, J. (2018). Deep neural networks as Gaussian processes. In *International Conference on Learning Representations*.
- Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Sohl-Dickstein, J. and Pennington, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*.
- Li, Y. (2017). Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
- Li, Y. and Liang, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*.
- Liu, B., Cai, Q., Yang, Z. and Wang, Z. (2019). Neural proximal/trust region policy optimization attains globally optimal policy. *arXiv preprint arXiv:1906.10306*.
- Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P. L. and Wainwright, M. J. (2018). Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. *arXiv preprint arXiv:1812.08305*.
- Martens, J. and Grosse, R. (2015). Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D. and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*.
- Munos, R. and Szepesvári, C. (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, **9** 815–857.
- Nesterov, Y. (2018). *Lectures on Convex Optimization*. Springer.
- Pan, X. and Srikumar, V. (2016). Expressiveness of rectifier networks. In *International Conference on Machine Learning*.
- Papini, M., Binaghi, D., Canonaco, G., Pirodda, M. and Restelli, M. (2018). Stochastic variance-reduced policy gradient. *arXiv preprint arXiv:1806.05618*.
- Peters, J. and Schaal, S. (2006). Policy gradient methods for robotics. In *International Conference on Intelligent Robots and Systems*.
- Peters, J. and Schaal, S. (2008). Natural actor-critic. *Neurocomputing*, **71** 1180–1190.
- Pirodda, M., Restelli, M. and Bascetta, L. (2015). Policy gradient in Lipschitz Markov decision processes. *Machine Learning*, **100** 255–283.
- Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Rahimi, A. and Recht, B. (2008). Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*.
- Rahimi, A. and Recht, B. (2009). Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems*.
- Scherrer, B. (2013). On the performance bounds of some policy search dynamic programming algorithms. *arXiv preprint arXiv:1306.0539*.

- Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B. and Geist, M. (2015). Approximate modified policy iteration and its application to the game of Tetris. *Journal of Machine Learning Research*, **16** 1629–1676.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M. and Moritz, P. (2015). Trust region policy optimization. In *International Conference on Machine Learning*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shen, Z., Ribeiro, A., Hassani, H., Qian, H. and Mi, C. (2019). Hessian aided policy gradient. In *International Conference on Machine Learning*.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, **529** 484.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A. et al. (2017). Mastering the game of Go without human knowledge. *Nature*, **550** 354.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, **3** 9–44.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT press.
- Sutton, R. S., McAllester, D. A., Singh, S. P. and Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*.
- Szepesvári, C. and Munos, R. (2005). Finite time bounds for sampling based fitted value iteration. In *International Conference on Machine Learning*.
- Tsitsiklis, J. N. and Van Roy, B. (1997). Analysis of temporal-difference learning with function approximation. In *Advances in Neural Information Processing Systems*.
- Tu, S. and Recht, B. (2018). The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. *arXiv preprint arXiv:1812.03565*.
- van Handel, R. (2014). *Probability in High Dimension*. Princeton University.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., Dudzik, A., Huang, A., Georgiev, P., Powell, R., Ewalds, T., Horgan, D., Kroiss, M., Danihelka, I., Agapiou, J., Oh, J., Dalibard, V., Choi, D., Sifre, L., Sulsky, Y., Vezhnevets, S., Molloy, J., Cai, T., Budden, D., Paine, T., Gulcehre, C., Wang, Z., Pfaff, T., Pohlen, T., Wu, Y., Yogatama, D., Cohen, J., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Apps, C., Kavukcuoglu, K., Hassabis, D. and Silver, D. (2019). AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>.
- Wagner, P. (2011). A reinterpretation of the policy oscillation phenomenon in approximate policy iteration. In *Advances in Neural Information Processing Systems*.
- Wagner, P. (2013). Optimistic policy iteration and natural actor-critic: A unifying view and a non-optimality result. In *Advances in Neural Information Processing Systems*.
- Wang, W. Y., Li, J. and He, X. (2018). Deep reinforcement learning for NLP. In *Association for Computational Linguistics*.

- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, **8** 229–256.
- Wu, L., Ma, C. and Weinan, E. (2018). How SGD selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Advances in Neural Information Processing Systems*.
- Wu, Y., Mansimov, E., Grosse, R. B., Liao, S. and Ba, J. (2017). Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In *Advances in Neural Information Processing Systems*.
- Xu, P., Gao, F. and Gu, Q. (2019a). An improved convergence analysis of stochastic variance-reduced policy gradient. *arXiv preprint arXiv:1905.12615*.
- Xu, P., Gao, F. and Gu, Q. (2019b). Sample efficient policy gradient methods with recursive variance reduction. *arXiv preprint arXiv:1909.08610*.
- Yang, Z., Chen, Y., Hong, M. and Wang, Z. (2019a). On the global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. *arXiv preprint arXiv:1907.06246*.
- Yang, Z., Xie, Y. and Wang, Z. (2019b). A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*.
- Zhang, K., Koppel, A., Zhu, H. and Başar, T. (2019). Global convergence of policy gradient methods to (almost) locally optimal policies. *arXiv preprint arXiv:1906.08383*.
- Zou, D., Cao, Y., Zhou, D. and Gu, Q. (2018). Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *arXiv preprint arXiv:1811.08888*.

A NEURAL NATURAL POLICY GRADIENT

In the sequel, we study the convergence of neural natural policy gradient. As shown in Algorithm 1, neural natural policy gradient uses neural TD for policy evaluation and updates the actor using (3.10), where θ_i and τ_i in (3.2) are both updated. To analyze critic update, we impose Assumptions 4.1 and 4.2, which guarantee that Theorem 4.3 holds. Meanwhile, to analyze actor update, we impose the following regularity conditions.

In parallel to Assumption 4.4, we lay out the following regularity condition on the variance of the estimators of the policy gradient and Fisher information matrix.

Assumption A.1 (Variance Upper Bound). Let $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$, where W_{init} is the initial parameter. We define

$$\delta_i = (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i) / \eta = \underset{\alpha \in \mathcal{B}}{\operatorname{argmin}} \|\widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_{\theta} J(\pi_{\theta_i})\|_2,$$

where $\widehat{\nabla}_{\theta} J(\pi_{\theta_i})$ and $\widehat{F}(\theta_i)$ are defined in (3.7) and (3.9), respectively. With slight abuse of notation, for all $i \in [T]$, we define the function $\xi_i : \mathbb{R}^{md} \rightarrow \mathbb{R}^{md}$ as

$$\xi_i(\alpha) = \widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_{\theta} J(\pi_{\theta_i}) - \mathbb{E}[\widehat{F}(\theta_i) \cdot \alpha - \tau_i \cdot \widehat{\nabla}_{\theta} J(\pi_{\theta_i})].$$

We assume that there exists an absolute constant $\sigma_{\xi} > 0$ such that

$$\mathbb{E}[\|\xi_i(\delta_i)\|_2^2] \leq \tau_i^4 \cdot \sigma_{\xi}^2 / B, \quad \mathbb{E}[\|\xi_i(\omega_i)\|_2^2] \leq \tau_i^4 \cdot \sigma_{\xi}^2 / B, \quad \forall i \in [T].$$

Here the expectations are taken over σ_i given θ_i and ω_i .

Next, we lay out a regularity condition on the visitation measures σ_i, ν_i and the stationary distributions ς_i, ϱ_i , respectively.

Assumption A.2 (Upper Bounded Concentrability Coefficient). We denote by ν_* and σ_* the state and state-action visitation measures corresponding to the global optimum π^* . For all $i \in [T]$, we define the concentrability coefficients $\varphi_i, \psi_i, \varphi'_i$, and ψ'_i as

$$\begin{aligned} \varphi_i &= \left\{ \mathbb{E}_{\sigma_i} [(d\sigma_*/d\sigma_i)^2] \right\}^{1/2}, & \psi_i &= \left\{ \mathbb{E}_{\nu_i} [(d\nu_*/d\nu_i)^2] \right\}^{1/2}, \\ \varphi'_i &= \left\{ \mathbb{E}_{\varsigma_i} [(d\sigma_*/d\varsigma_i)^2] \right\}^{1/2}, & \psi'_i &= \left\{ \mathbb{E}_{\varrho_i} [(d\nu_*/d\varrho_i)^2] \right\}^{1/2}, \end{aligned} \quad (\text{A.1})$$

where $d\sigma_*/d\sigma_i, d\nu_*/d\nu_i, d\sigma_*/d\varsigma_i$, and $d\nu_*/d\varrho_i$ are the Radon-Nikodym derivatives. We assume that the concentrability coefficients defined in (A.1) are uniformly upper bounded by an absolute constant $c_0 > 0$.

The regularity condition on upper bounded concentrability coefficients is commonly imposed in the reinforcement learning literature and is standard for theoretical analysis (Szepesvári and Munos, 2005; Munos and Szepesvári, 2008; Antos et al., 2008; Lazaric et al., 2016; Farahmand et al., 2010; 2016; Scherrer, 2013; Scherrer et al., 2015; Yang et al., 2019b; Chen and Jiang, 2019).

Finally, we introduce the following regularity condition on the initial parameter W_{init} in Algorithm 1.

Assumption A.3 (Upper Bounded Moment at Random Initialization). Let $\phi_0(s, a) \in \mathbb{R}^{md}$ be the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. We assume that there exists an absolute constant $M > 0$ such that

$$\mathbb{E}_{\text{init}} \left[\sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} |f((s, a); W_{\text{init}})|^2 \right] = \mathbb{E}_{\text{init}} \left[\sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} |\phi_0(s, a)^\top W_{\text{init}}|^2 \right] \leq M^2.$$

Here the expectations are taken over the random initialization.

Note that as $m \rightarrow \infty$, the initial neural network $\phi_0(s, a)^\top W_{\text{init}}$ converges to a Gaussian process indexed by (s, a) (Lee et al., 2018), which lies in a compact subset of \mathbb{R}^d . It is known that under certain regularity conditions, the maximum of a Gaussian process over a compact index set is a sub-Gaussian random variable (van Handel, 2014). Therefore, the regularity condition that $\max_{(s,a)} |\phi_0(s, a)^\top W_{\text{init}}|$ has a finite second moment is mild.

We now establish the global optimality and rate of convergence of neural natural policy gradient.

Theorem A.4 (Global Optimality and Convergence). We set $\eta = 1/\sqrt{T}$, $\eta_{\text{TD}} = \min\{(1-\gamma)/8, 1/\sqrt{T_{\text{TD}}}\}$, $T_{\text{TD}} = \Omega(m)$, $\tau_i = (i-1) \cdot \eta$, and $\mathcal{B} = \{\alpha : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ in Algorithm 1, where the actor update is given in (3.10). Under the assumptions of Theorem 4.3 and Assumptions A.1-A.3, we have

$$\min_{i \in [T]} \mathbb{E}[J(\pi^*) - J(\pi_{\theta_i})] \leq \frac{\log |\mathcal{A}| + 9R^2 + M}{(1-\gamma) \cdot \sqrt{T}} + \frac{1}{(1-\gamma) \cdot T} \cdot \sum_{i=1}^T \bar{\epsilon}_i(T). \quad (\text{A.2})$$

Here M is defined in Assumption A.3 and $\bar{\epsilon}_i(T)$ satisfies

$$\bar{\epsilon}_i(T) = \underbrace{\sqrt{8c_0} \cdot R^{1/2} \cdot (\sigma_\xi^2/B)^{1/4}}_{(a)} + \underbrace{\mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4} \cdot T^{1/2} + R^{5/4} \cdot m^{-1/8})}_{(b)} + \underbrace{\varepsilon_{Q,i}}_{(c)}, \quad (\text{A.3})$$

where c_0 is defined in Assumption A.2 and $\varepsilon_{Q,i} = c_0 \cdot \mathcal{O}(R^{3/2} \cdot m^{-1/4} + R^{5/4} \cdot m^{-1/8})$. Here the expectation is taken over all the randomness.

Proof. See §D.3 for a detailed proof. \square

As shown in (A.2) of Theorem A.4, the optimality gap $\min_{i \in [T]} \mathbb{E}[J(\pi^*) - J(\pi_{\theta_i})]$ is upper bounded by two terms. Intuitively, the first $\mathcal{O}(1/\sqrt{T})$ term characterizes the convergence of neural natural policy gradient as $m, B \rightarrow \infty$. Meanwhile, the second term aggregates the errors incurred by both actor update and critic update due to finite m and B . Specifically, in (A.3) of Theorem A.4, term (a) corresponds to the estimation error of $\hat{F}(\theta)$ and $\hat{V}_\theta J(\pi_\theta)$ due to the finite batch size B , which vanishes as $B \rightarrow \infty$. Term (b) corresponds to the incompatibility between the parameterizations of the actor and critic. As introduced in §3.1, we use shared architecture and random initialization to ensure approximately compatible function approximations. In particular, term (b) vanishes as $m \rightarrow \infty$. Meanwhile, term (c) corresponds to the policy evaluation error, i.e., the error of approximating $Q^{\pi_{\theta_i}}$ using Q_{ω_i} . As shown in Theorem 4.3, such an error is sufficiently small when both m and T_{TD} are sufficiently large. To conclude, when m , B , and T_{TD} are sufficiently large, the expected total reward of (a subsequence of) $\{\pi_{\theta_i}\}_{i \in [T+1]}$ obtained from the neural natural policy gradient converges to the global optimum $J(\pi^*)$ at a $1/\sqrt{T}$ -rate. Formally, we have the following corollary.

Corollary A.5 (Global Optimality and Convergence). Under the same assumptions of Theorem A.4, it holds for $m = \Omega(R^{10} \cdot T^6)$ and $B = \Omega(R^2 \cdot T^2 \cdot \sigma_\xi^2)$ that

$$\min_{i \in [T]} \mathbb{E}[J(\pi^*) - J(\pi_{\theta_i})] = \mathcal{O}\left(\frac{\log |\mathcal{A}|}{(1-\gamma) \cdot \sqrt{T}}\right).$$

Here the expectation is taken over all the randomness.

Proof. See §H.4 for a detailed proof. \square

Corollary A.5 establishes both the global optimality and rate of convergence of neural natural policy gradient. Combining Theorem 4.7 and Corollary A.5, we conclude that when we use overparameterized two-layer neural networks, both neural policy gradient and neural natural policy gradient converge at $1/\sqrt{T}$ -rates. In comparison, when m and B are sufficiently large, neural policy gradient is only shown to converge to a stationary point under the additional regularity condition that $\nabla_\theta J(\pi_\theta)$ is Lipschitz continuous (Assumption 4.6). Moreover, by Theorem 4.8, the global optimality of such a stationary point hinges on the representation power of the overparameterized two-layer neural network. In contrast, neural natural policy gradient is shown to attain the global optimum when both m and B are sufficiently large without additional regularity conditions such as Assumption 4.6, which reveals the benefit of incorporating more sophisticated optimization techniques to reinforcement learning. A similar phenomenon is observed in the LQR setting (Fazel et al., 2018; Malik et al., 2018; Tu and Recht, 2018), where natural policy gradient enjoys an improved rate of convergence.

In recent work, Liu et al. (2019) study the global optimality and rates of convergence of neural proximal policy optimization (PPO) and trust region policy optimization (TRPO) (Schulman et al., 2015; 2017). Although Liu et al. (2019) establish a similar $1/\sqrt{T}$ -rate of convergence to the global optimum, neural PPO is different from neural natural policy gradient, as it requires solving a sub-problem of policy improvement in the functional space by fitting a neural network using multiple stochastic gradient steps in the parameter space. In contrast, neural natural policy gradient only requires a single stochastic natural gradient step in the parameter space, which makes the analysis even more challenging.

B SHARED INITIALIZATION AND COMPATIBLE FUNCTION APPROXIMATION

Sutton et al. (2000) introduce the notion of compatible function approximations. Specifically, the action-value function Q_ω is compatible with π_θ if we have $\nabla_\omega A_\omega(s, a) = \nabla_\theta \log \pi_\theta(a | s)$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, where $A_\omega(s, a) = Q_\omega(s, a) - \langle Q_\omega(s, \cdot), \pi_\theta(\cdot | s) \rangle$ is the advantage function corresponding to Q_ω . Compatible function approximations enable us to construct unbiased estimators of the policy gradient, which are essential for the optimality and convergence of policy gradient methods (Konda and Tsitsiklis, 2000; Sutton et al., 2000; Kakade, 2002; Peters and Schaal, 2008; Wagner, 2011; 2013).

To approximately obtain compatible function approximations when both the actor and critic are represented by neural networks, we use a shared architecture between the action-value function Q_ω and the energy function of π_θ , and initialize Q_ω and π_θ with the same parameter W_{init} , where $[W_{\text{init}}]_r \sim N(0, I_d/d)$ for all $r \in [m]$. We show that in the overparameterized regime where m is large, the shared architecture and random initialization ensures Q_ω to be approximately compatible with π_θ in the following sense. We define $\bar{\phi}_0 = ([\bar{\phi}_0]_1^\top, \dots, [\bar{\phi}_0]_m^\top)^\top : \mathbb{R}^d \rightarrow \mathbb{R}^{md}$ as the centered feature mapping corresponding to the initialization, which takes the form of

$$\begin{aligned} [\bar{\phi}_0]_r(s, a) &= \frac{b_r}{\sqrt{m}} \cdot \mathbf{1}\{(s, a)^\top [W_{\text{init}}]_r > 0\} \cdot (s, a) \\ &\quad - \mathbb{E}_{\pi_\theta} \left[\frac{b_r}{\sqrt{m}} \cdot \mathbf{1}\{(s, a')^\top [W_{\text{init}}]_r > 0\} \cdot (s, a') \right], \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \end{aligned} \quad (\text{B.1})$$

where W_{init} is the initialization shared by both the actor and critic, and we omit the dependency on θ for notational simplicity. Similarly, we define for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ the following centered feature mappings,

$$\bar{\phi}_\theta(s, a) = \phi_\theta(s, a) - \mathbb{E}_{\pi_\theta} [\phi_\theta(s, a')], \quad \bar{\phi}_\omega(s, a) = \phi_\omega(s, a) - \mathbb{E}_{\pi_\theta} [\phi_\omega(s, a')]. \quad (\text{B.2})$$

Here $\phi_\theta(s, a)$ and $\phi_\omega(s, a)$ are the feature mappings defined in (3.3), which correspond to θ and ω , respectively. By (3.1), we have

$$A_\omega(s, a) = Q_\omega(s, a) - \mathbb{E}_{\pi_\theta} [Q_\omega(s, a')] = \bar{\phi}_\omega(s, a)^\top \omega, \quad \nabla_\theta \log \pi_\theta(a | s) = \bar{\phi}_\theta(s, a), \quad (\text{B.3})$$

which holds almost everywhere for $\theta \in \mathbb{R}^{md}$. As shown in Corollary E.3 in §E, when the width m is sufficiently large, in policy gradient methods, both $\bar{\phi}_\theta$ and $\bar{\phi}_\omega$ are well approximated by $\bar{\phi}_0$ defined in (B.1). Therefore, by (B.3), we conclude that in the overparameterized regime with shared architecture and random initialization, Q_ω is approximately compatible with π_θ .

C SAMPLING FROM VISITATION MEASURE

Recall that the policy gradient $\nabla_\theta J(\pi_\theta)$ in (3.4) involves an expectation with respect to the state-action visitation measure σ_{π_θ} . Thus, to obtain an unbiased estimator of the policy gradient, we need to sample from the visitation measure σ_{π_θ} . To achieve such a goal, we introduce an artificial MDP $(\mathcal{S}, \mathcal{A}, \tilde{\mathcal{P}}, \zeta, r, \gamma)$. Such an MDP only differs from the original MDP in the Markov transition kernel $\tilde{\mathcal{P}}$, which is defined as

$$\tilde{\mathcal{P}}(s' | s, a) = \gamma \cdot \mathcal{P}(s' | s, a) + (1 - \gamma) \cdot \zeta(s'), \quad \forall (s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}. \quad (\text{C.1})$$

Here \mathcal{P} is the Markov transition kernel of the original MDP. That is, at each state transition of the artificial MDP, the next state is sampled from the initial state distribution $\zeta(\cdot)$ with probability $1 - \gamma$. In other words, at each state transition, we restart the original MDP with probability $1 - \gamma$. As shown in [Konda \(2002\)](#), the stationary state distribution of the induced Markov chain is exactly the state visitation measure ν_{π_θ} . Therefore, when we sample a trajectory $\{(S_t, A_t)\}_{t \geq 0}$, where $S_0 \sim \zeta(\cdot)$, $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \tilde{\mathcal{P}}(\cdot | S_t, A_t)$ for all $t \geq 0$, the marginal distribution of (S_t, A_t) converges to the state-action visitation measure σ_{π_θ} .

D PROOF OF MAIN RESULTS

In this section, we present the proofs of [Theorems 4.7, 4.8, and A.4](#). Our proofs utilize the following lemma, which establishes the one-point convexity of $J(\pi)$ at the global optimum π^* . Such a lemma is adapted from [Kakade and Langford \(2002\)](#).

Lemma D.1 (Performance Difference ([Kakade and Langford, 2002](#))). It holds for all π that

$$J(\pi^*) - J(\pi) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{\nu_*} [\langle Q^\pi(s, \cdot), \pi^*(\cdot | s) - \pi(\cdot | s) \rangle], \quad (\text{D.1})$$

where ν_* is the state visitation measure corresponding to π^* .

Proof. Following from [Lemma J.1](#), which is the Lemma 6.1 in [Kakade and Langford \(2002\)](#), it holds for all π that

$$J(\pi^*) - J(\pi) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{\sigma_*} [A^\pi(s, a)], \quad (\text{D.2})$$

where σ_* is the state-action visitation measure corresponding to π^* , and A^π is the advantage function associated with π . By definition, we have $\sigma_*(\cdot, \cdot) = \pi^*(\cdot | \cdot) \cdot \nu_*(\cdot)$. Meanwhile, it holds for all $s \in \mathcal{S}$ that

$$\begin{aligned} \mathbb{E}_{\pi^*} [A^\pi(s, a)] &= \mathbb{E}_{\pi^*} [Q^\pi(s, a)] - V^\pi(s) = \langle Q^\pi(s, \cdot), \pi^*(\cdot | s) \rangle - \langle Q^\pi(s, \cdot), \pi(\cdot | s) \rangle \\ &= \langle Q^\pi(s, \cdot), \pi^*(\cdot | s) - \pi(\cdot | s) \rangle. \end{aligned} \quad (\text{D.3})$$

Combining [\(D.2\)](#) and [\(D.3\)](#), we conclude that

$$J(\pi^*) - J(\pi) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{\nu_*} [\langle Q^\pi(s, \cdot), \pi^*(\cdot | s) - \pi(\cdot | s) \rangle],$$

which concludes the proof of [Lemma D.1](#). \square

D.1 PROOF OF THEOREM 4.7

Proof. We first lower bound the difference between the expected total rewards of $\pi_{\theta_{i+1}}$ and π_{θ_i} . By [Assumption 4.6](#), $\nabla_\theta J(\pi_\theta)$ is L -Lipschitz continuous. Thus, it holds that

$$J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i}) \geq \eta \cdot \nabla_\theta J(\pi_{\theta_i})^\top \delta_i - L/2 \cdot \|\theta_{i+1} - \theta_i\|_2^2, \quad (\text{D.4})$$

where $\delta_i = (\theta_{i+1} - \theta_i)/\eta$. Recall that $\xi_i = \widehat{\nabla}_\theta J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_\theta J(\pi_{\theta_i})]$, where the expectation is taken over σ_i given θ_i and ω_i . It holds that

$$\nabla_\theta J(\pi_{\theta_i})^\top \delta_i = \left(\nabla_\theta J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_\theta J(\pi_{\theta_i})] \right)^\top \delta_i - \xi_i^\top \delta_i + \widehat{\nabla}_\theta J(\pi_{\theta_i})^\top \delta_i. \quad (\text{D.5})$$

On the right-hand side of [\(D.5\)](#), the first term represents the error of estimating $\nabla_\theta J(\pi_{\theta_i})$ using $\mathbb{E}[\widehat{\nabla}_\theta J(\pi_{\theta_i})] = \mathbb{E}_{\sigma_i} [\nabla_\theta \log \pi_{\theta_i}(a | s) \cdot Q_{\omega_i}(s, a)]$, the second term is related to the variance of the estimator $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ of the policy gradient $\nabla_\theta J(\pi_{\theta_i})$, and the last term relates the increment δ_i of actor update to $\widehat{\nabla}_\theta J(\pi_{\theta_i})$. In the following lemma, we establish a lower bound of the first term.

Lemma D.2. It holds that

$$\left| \left(\nabla_\theta J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_\theta J(\pi_{\theta_i})] \right)^\top \delta_i \right| \leq 2\kappa \cdot R/\eta \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i},$$

where $\widehat{\nabla}_\theta J(\pi_{\theta_i})$ is defined in [\(3.7\)](#), ς_i is the stationary state-action distribution, and κ is the absolute constant defined in [Assumption 4.5](#). Here the expectation is taken over σ_i given θ_i and ω_i .

Proof. See [§H.5](#) for a detailed proof. \square

For the second term on the right-hand side of (D.5), we have

$$-\xi_i^\top \delta_i \geq -\|\xi_i\|_2^2/2 - \|\delta_i\|_2^2/2. \quad (\text{D.6})$$

Now it remains to lower bound the third term on the right-hand side of (D.5). For notational simplicity, we define

$$e_i = \theta_{i+1} - (\theta_i + \eta \cdot \widehat{\nabla} J(\pi_{\theta_i})) = \Pi_{\mathcal{B}}(\theta_i + \eta \cdot \widehat{\nabla} J(\pi_{\theta_i})) - (\theta_i + \eta \cdot \widehat{\nabla} J(\pi_{\theta_i})),$$

where $\Pi_{\mathcal{B}}$ is the projection operator onto \mathcal{B} . It then holds that

$$e_i^\top \left[\Pi_{\mathcal{B}}(\theta_i + \eta \cdot \widehat{\nabla} J(\pi_{\theta_i})) - x \right] = e_i^\top (\theta_{i+1} - x) \leq 0, \quad \forall x \in \mathcal{B}. \quad (\text{D.7})$$

Specifically, setting $x = \theta_i$ in (D.7), we obtain that $e_i^\top \delta_i \leq 0$, which implies

$$\widehat{\nabla} J_{\theta}(\pi_{\theta_i})^\top \delta_i = (\delta_i - e_i/\eta)^\top \delta_i \geq \|\delta_i\|_2^2. \quad (\text{D.8})$$

By plugging Lemma D.2, (D.6), and (D.8) into (D.5), we obtain that

$$\nabla_{\theta} J(\pi_{\theta_i})^\top \delta_i \geq -2\kappa \cdot R/\eta \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i} + \|\delta_i\|_2^2/2 - \|\xi_i\|_2^2/2. \quad (\text{D.9})$$

Thus, by plugging (D.9) and the definition that $\delta_i = (\theta_{i+1} - \theta_i)/\eta$ into (D.4), we obtain for all $i \in [T]$ that

$$\begin{aligned} & (1 - L \cdot \eta) \cdot \mathbb{E}[\|\delta_i\|_2^2/2] \\ & \leq \eta^{-1} \cdot \mathbb{E}[J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i})] + 2\kappa \cdot R/\eta \cdot \mathbb{E}_{\text{init}}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}] + \mathbb{E}[\|\xi_i\|_2^2/2]. \end{aligned} \quad (\text{D.10})$$

Now we turn to characterize $\|\rho_i - \delta_i\|_2$. By the definition of ρ_i in (4.3), we have

$$\begin{aligned} \|\rho_i - \delta_i\|_2 &= \eta^{-1} \cdot \left\| \Pi_{\mathcal{B}}(\theta_i + \eta \cdot \nabla_{\theta} J(\pi_{\theta_i})) - \theta_i - \left(\Pi_{\mathcal{B}}(\theta_i + \eta \cdot \widehat{\nabla}_{\theta} J(\pi_{\theta_i})) - \theta_i \right) \right\|_2 \\ &= \eta^{-1} \cdot \left\| \Pi_{\mathcal{B}}(\theta_i + \eta \cdot \nabla_{\theta} J(\pi_{\theta_i})) - \Pi_{\mathcal{B}}(\theta_i + \eta \cdot \widehat{\nabla}_{\theta} J(\pi_{\theta_i})) \right\|_2 \\ &\leq \|\nabla_{\theta} J(\pi_{\theta_i}) - \widehat{\nabla}_{\theta} J(\pi_{\theta_i})\|_2. \end{aligned} \quad (\text{D.11})$$

The following lemma further upper bounds the right-hand side of (D.11).

Lemma D.3. It holds for all $i \in [T]$ that

$$\mathbb{E}[\|\nabla_{\theta} J(\pi_{\theta_i}) - \widehat{\nabla}_{\theta} J(\pi_{\theta_i})\|_2^2] \leq 2\mathbb{E}[\|\xi_i\|_2^2] + 8\kappa^2 \cdot \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}^2].$$

Here the expectations are taken over all the randomness.

Proof. See §H.6 for a detailed proof. □

Recall that we set $\eta = 1/\sqrt{T}$. Upon telescoping (D.10), it holds for $T \geq 4L^2$ that

$$\begin{aligned} \min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2^2] &\leq 1/T \cdot \sum_{i=1}^T \mathbb{E}[\|\rho_i\|_2^2] \\ &\leq 1/T \cdot \sum_{i=1}^T \left(2\mathbb{E}[\|\delta_i\|_2^2] + 2\mathbb{E}[\|\rho_i - \delta_i\|_2^2] \right) \\ &\leq 1/T \cdot \sum_{i=1}^T 4(1 - L \cdot \eta) \cdot \mathbb{E}[\|\delta_i\|_2^2] + 2\mathbb{E}[\|\rho_i - \delta_i\|_2^2] \\ &\leq 8/\sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{T+1}}) - J(\pi_{\theta_1})] + 8/T \cdot \sum_{i=1}^T \mathbb{E}[\|\xi_i\|_2^2] + \varepsilon_Q(T), \end{aligned} \quad (\text{D.12})$$

where the third inequality follows from the fact that $1 - L \cdot \eta \geq 1/2$, while the fourth inequality follows from (D.10), (D.11), and Lemma D.3. Here the expectations are taken over all the randomness, and $\varepsilon_Q(T)$ is defined as

$$\varepsilon_Q(T) = 16\kappa \cdot R/\sqrt{T} \cdot \sum_{i=1}^T \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}] + 16\kappa^2/\sqrt{T} \cdot \sum_{i=1}^T \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}^2].$$

By Proposition 4.3, we have

$$\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2] = O(R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}).$$

Meanwhile, by Assumption 4.4, it holds for all $i \in [T]$ that $\mathbb{E}[\|\xi_i\|_2^2/2] \leq \sigma_\xi^2/B$. By plugging the upper bounds of $\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2]$ and $\mathbb{E}[\|\xi_i\|_2^2/2]$ into (D.12), we conclude that

$$\min_{i \in [T]} \mathbb{E}[\|\rho_i\|_2^2] \leq 8/\sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{T+1}}) - J(\pi_{\theta_1})] + 8\sigma_\xi^2/B + \varepsilon_Q(T),$$

where

$$\varepsilon_Q(T) = \kappa \cdot \mathcal{O}(R^{5/2} \cdot m^{-1/4} \cdot T + R^{9/4} \cdot m^{-1/8} \cdot T).$$

Thus, we complete the proof of Theorem 4.7. \square

D.2 PROOF OF THEOREM 4.8

Proof. Since $\hat{\theta}$ is a stationary point of $J(\pi_\theta)$, it holds that

$$\nabla_\theta J(\pi_{\hat{\theta}})^\top (\theta - \hat{\theta}) \leq 0, \quad \forall \theta \in \mathcal{B}. \quad (\text{D.13})$$

Therefore, by Proposition 3.1, we obtain from (D.13) that

$$\nabla_\theta J(\pi_{\hat{\theta}})^\top (\theta - \hat{\theta}) = \mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\bar{\phi}_{\hat{\theta}}(s, a)^\top (\theta - \hat{\theta}) \cdot Q^{\pi_{\hat{\theta}}}(s, a)] \leq 0, \quad \forall \theta \in \mathcal{B}. \quad (\text{D.14})$$

Here $\phi_{\hat{\theta}}$ and $\bar{\phi}_{\hat{\theta}}$ are defined in (3.3) and (B.2) with $\theta = \hat{\theta}$, respectively. Since it holds that

$$\begin{aligned} \mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\bar{\phi}_{\hat{\theta}}(s, a)^\top (\theta - \hat{\theta}) \cdot V^{\pi_{\hat{\theta}}}(s)] &= 0, \\ \mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\mathbb{E}_{\pi_{\hat{\theta}}} [\phi_{\hat{\theta}}(s, a')^\top (\theta - \hat{\theta})] \cdot A^{\pi_{\hat{\theta}}}(s, a)] &= 0, \end{aligned}$$

where $a' \sim \pi_{\hat{\theta}}(\cdot | s)$ and $(s, a) \sim \sigma_{\pi_{\hat{\theta}}}(\cdot, \cdot)$, we obtain from (D.14) that

$$\mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\bar{\phi}_{\hat{\theta}}(s, a)^\top (\theta - \hat{\theta}) \cdot Q^{\pi_{\hat{\theta}}}(s, a)] = \mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\phi_{\hat{\theta}}(s, a)^\top (\theta - \hat{\theta}) \cdot A^{\pi_{\hat{\theta}}}(s, a)] \leq 0, \quad \forall \theta \in \mathcal{B}. \quad (\text{D.15})$$

Meanwhile, by Lemma D.1 we have

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) = \mathbb{E}_{\nu_*} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \pi^*(\cdot | s) - \pi_{\hat{\theta}}(\cdot | s) \rangle]. \quad (\text{D.16})$$

In what follows, we write $\Delta_\theta = \theta - \hat{\theta}$. Combining (D.14) and (D.16), we obtain that

$$\begin{aligned} (1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) & \leq \mathbb{E}_{\nu_*} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \pi^*(\cdot | s) - \pi_{\hat{\theta}}(\cdot | s) \rangle] - \mathbb{E}_{\sigma_{\pi_{\hat{\theta}}}} [\phi_{\hat{\theta}}(s, a)^\top \Delta_\theta \cdot A^{\pi_{\hat{\theta}}}(s, a)] \\ & = \mathbb{E}_{\nu_*} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \pi^*(\cdot | s) - \pi_{\hat{\theta}}(\cdot | s) \rangle] - \mathbb{E}_{\nu_{\pi_{\hat{\theta}}}} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \phi_{\hat{\theta}}(s, \cdot)^\top \Delta_\theta \cdot \pi_{\hat{\theta}}(\cdot | s) \rangle], \end{aligned} \quad (\text{D.17})$$

where we use the fact that $\sigma_{\pi_{\hat{\theta}}}(\cdot, \cdot) = \pi_{\hat{\theta}}(\cdot | \cdot) \cdot \nu_{\pi_{\hat{\theta}}}(\cdot)$. It remains to upper bound the right-hand side of (D.17). By calculation, it holds for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned} & (\pi^*(a | s) - \pi_{\hat{\theta}}(a | s)) d\nu_*(s) - \phi_{\hat{\theta}}(s, a)^\top \Delta_\theta \cdot \pi_{\hat{\theta}}(a | s) d\nu_{\hat{\theta}}(s) \\ & = \left(\frac{\pi^*(a | s) - \pi_{\hat{\theta}}(a | s)}{\pi_{\hat{\theta}}(a | s)} \cdot \frac{d\nu_*(s)}{d\nu_{\hat{\theta}}(s)} - \phi_{\hat{\theta}}(s, a)^\top \Delta_\theta \right) \cdot \pi_{\hat{\theta}}(a | s) d\nu_{\pi_{\hat{\theta}}}(s) \\ & = (u_{\hat{\theta}}(s, a) - \phi_{\hat{\theta}}(s, a)^\top \theta) d\sigma_{\pi_{\hat{\theta}}}(s, a), \end{aligned} \quad (\text{D.18})$$

where $u_{\hat{\theta}}$ is defined as

$$u_{\hat{\theta}}(s, a) = \frac{d\sigma_{\pi^*}}{d\sigma_{\pi_{\hat{\theta}}}}(s, a) - \frac{d\nu_{\pi^*}}{d\nu_{\pi_{\hat{\theta}}}}(s) + \phi_{\hat{\theta}}(s, a)^\top \hat{\theta}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

Here $d\sigma_{\pi^*}/d\sigma_{\pi_{\hat{\theta}}}$ and $d\nu_{\pi^*}/d\nu_{\pi_{\hat{\theta}}}$ are the Radon-Nikodym derivatives. By plugging (D.18) into (D.17), we obtain that

$$\begin{aligned}
& (1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) & \text{(D.19)} \\
& \leq \mathbb{E}_{\nu_*} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \pi^*(\cdot | s) - \pi_{\hat{\theta}}(\cdot | s) \rangle] - \mathbb{E}_{\nu_{\pi_{\hat{\theta}}}} [\langle A^{\pi_{\hat{\theta}}}(s, \cdot), \phi_{\hat{\theta}}(s, \cdot)^\top \Delta_\theta \cdot \pi_{\hat{\theta}}(\cdot | s) \rangle] \\
& = \int_{\mathcal{S}} \sum_{a \in \mathcal{A}} A^{\pi_{\hat{\theta}}}(s, a) \cdot \left((\pi^*(a | s) - \pi_{\hat{\theta}}(a | s)) d\nu_*(s) - \phi_{\hat{\theta}}(s, a)^\top \Delta_\theta \cdot \pi_{\hat{\theta}}(a | s) d\nu_{\hat{\theta}}(s) \right) \\
& = \int_{\mathcal{S} \times \mathcal{A}} A^{\pi_{\hat{\theta}}}(s, a) \cdot (u_{\hat{\theta}}(s, a) - \phi_{\hat{\theta}}(s, a)^\top \Delta_\theta) d\sigma_{\pi_{\hat{\theta}}}(s, a) \\
& \leq \|A^{\pi_{\hat{\theta}}}(\cdot, \cdot)\|_{\sigma_{\pi_{\hat{\theta}}}} \cdot \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}},
\end{aligned}$$

where the second equality follows from (D.18) and the last inequality is from the Cauchy-Schwartz inequality. Note that $|A^{\pi_{\hat{\theta}}}(s, a)| \leq 2Q_{\max}$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Therefore, it follows from (D.19) that

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}}, \quad \forall \theta \in \mathcal{B}. \quad \text{(D.20)}$$

Finally, by taking the infimum of the right-hand side of (D.20) with respect to $\theta \in \mathcal{B}$, we obtain that

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \inf_{\theta \in \mathcal{B}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}},$$

which concludes the proof of Theorem 4.8. \square

D.3 PROOF OF THEOREM A.4

Proof. For notational simplicity, we write $\pi_i = \pi_{\theta_i}$ hereafter. In the following lemma, we characterize the performance difference $J(\pi^*) - J(\pi_i)$ based on Lemma D.1.

Lemma D.4. It holds that

$$\begin{aligned}
(1 - \gamma) \cdot \eta \cdot (J(\pi^*) - J(\pi_i)) &= \mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_i(\cdot | s)) - D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_{i+1}(\cdot | s)) \right. \\
&\quad \left. - D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] - H_i,
\end{aligned}$$

where H_i is defined as

$$\begin{aligned}
H_i &= \mathbb{E}_{\nu_*} \left[\underbrace{\langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle}_{\text{(i)}} \right] \\
&\quad + \eta \cdot \mathbb{E}_{\nu_*} \left[\underbrace{\langle Q_{\omega_i}(s, \cdot) - Q^{\pi_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle}_{\text{(ii)}} \right] \\
&\quad + \mathbb{E}_{\nu_*} \left[\underbrace{\langle \log(\pi_i(\cdot | s) / \pi_{i+1}(\cdot | s)), \pi_{i+1}(\cdot | s) - \pi_i(\cdot | s) \rangle}_{\text{(iii)}} \right]. \quad \text{(D.21)}
\end{aligned}$$

Proof. See §H.7 for a detailed proof. \square

Here H_i defined in (D.21) of Lemma D.4 consists of three terms. Specifically, term (i) is related to the error of estimating natural policy gradient using (3.8). Term (ii) is related to the error of estimating Q^{π_i} using Q_{ω_i} . Term (iii) is the remainder term. We upper bound these three terms in §H.8. Combining these upper bounds, we obtain the following lemma.

Lemma D.5. Under Assumptions 4.2 and A.3, we have

$$\mathbb{E} \left[|H_i| - \mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] \right] \leq \eta^2 \cdot (9R^2 + M^2) + \eta \cdot (\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i} + \varepsilon_i.$$

Here the expectation is taken over all the randomness. Meanwhile, φ'_i and ψ'_i are the concentration coefficients defined in (A.1) of Assumption A.2, $\varepsilon_{Q,i} = \mathbb{E}[\|Q^{\pi_i} - Q_{\omega_i}\|_{\varsigma_i}]$, M is defined in

Assumption A.3, and ε_i is defined as

$$\begin{aligned} \varepsilon_i &= \sqrt{2} \cdot R^{1/2} \cdot \eta \cdot (\varphi_i + \psi_i) \cdot \tau_i^{-1} \cdot \left(\mathbb{E}[\|\xi_i(\delta_i)\|_2] + \mathbb{E}[\|\xi_i(\omega_i)\|_2] \right)^{1/2} \\ &\quad + \mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4} + \eta \cdot R^{5/4} \cdot m^{-1/8}). \end{aligned} \quad (\text{D.22})$$

Here $\xi_i(\delta_i)$ and $\xi_i(\omega_i)$ are defined in Assumption A.1, where $\delta_i = \eta^{-1} \cdot (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i)$, while φ_i and ψ_i are the concentrability coefficients defined in (A.1) of Assumption A.2.

Proof. See §H.8 for a detailed proof. \square

By Lemmas D.4 and D.5, we obtain that

$$\begin{aligned} (1 - \gamma) \cdot \mathbb{E}[J(\pi^*) - J(\pi_i)] &\leq \eta^{-1} \cdot \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_*(\cdot | s) \| \pi_i(\cdot | s)) - D_{\text{KL}}(\pi_*(\cdot | s) \| \pi_{i+1}(\cdot | s)) \right] \right] \\ &\quad + \eta \cdot (9R^2 + M^2) + \eta^{-1} \cdot \varepsilon_i + (\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i}, \end{aligned} \quad (\text{D.23})$$

where $\varepsilon_{Q,i} = \mathbb{E}[\|Q^{\pi_i} - Q_{\omega_i}\|_{\mathcal{S}_i}]$, M is defined in Assumption A.3, ε_i is defined in (D.22) of Lemma D.5, and the expectations are taken over all the randomness. Recall that we set $\eta = 1/\sqrt{T}$. Upon telescoping (D.23), we obtain that

$$\begin{aligned} (1 - \gamma) \cdot \min_{i \in [T]} \mathbb{E}[J(\pi^*) - J(\pi_i)] &\leq \frac{1 - \gamma}{T} \cdot \sum_{i=1}^T \mathbb{E}[J(\pi^*) - J(\pi_i)] \\ &\leq \frac{1}{\sqrt{T}} \cdot \left(\mathbb{E} \left[\mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_*(\cdot | s) \| \pi_1(\cdot | s)) \right] \right] + 9R^2 + M^2 \right) \\ &\quad + \frac{1}{T} \cdot \sum_{i=1}^T (\sqrt{T} \cdot \varepsilon_i + (\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i}), \end{aligned} \quad (\text{D.24})$$

where the expectations are taken over all the randomness and the last inequality follows from the fact that

$$D_{\text{KL}}(\pi_*(\cdot | s) \| \pi_{T+1}(\cdot | s)) \geq 0, \quad \forall s \in \mathcal{S}, \forall \theta_{T+1} \in \mathbb{R}^{md}.$$

In what follows, we upper bound the right-hand side of (D.24). Note that we set $\tau_1 = 0$. Following from (3.2), it then holds that $\pi_1(\cdot | s)$ is uniform over \mathcal{A} for all $s \in \mathcal{S}$ and $\theta_1 \in \mathbb{R}^{md}$. Therefore, we obtain that

$$D_{\text{KL}}(\pi_*(\cdot | s) \| \pi_1(\cdot | s)) \leq \log |\mathcal{A}|, \quad \forall s \in \mathcal{S}, \forall \theta_1 \in \mathbb{R}^{md}. \quad (\text{D.25})$$

Meanwhile, by Assumption A.1, we have

$$\mathbb{E}[\|\xi_i(\delta_i)\|_2] \leq \left\{ \mathbb{E} \left[\mathbb{E}_{\sigma_i} [\|\xi_i(\delta_i)\|_2^2] \right] \right\}^{1/2} \leq \tau_i^2 \cdot \sigma_\xi \cdot B^{-1/2},$$

where the expectation $\mathbb{E}_{\sigma_i}[\|\xi_i(\delta_i)\|_2^2]$ is taken over σ_i given θ_i and ω_i , while the other expectations are taken over all the randomness. A similar upper bound holds for $\mathbb{E}[\|\xi_i(\omega_i)\|_2]$. Therefore, by plugging the upper bounds of $\mathbb{E}[\|\xi_i(\sigma_i)\|_2]$ and $\mathbb{E}[\|\xi_i(\omega_i)\|_2]$ into ε_i defined in (D.22) of Lemma D.5, we obtain from Assumption A.2 that

$$\sqrt{T} \cdot \varepsilon_i \leq \sqrt{8}c_0 \cdot R^{1/2} \cdot \sigma_\xi^{1/2} \cdot B^{-1/4} + \mathcal{O}(\tau_{i+1} \cdot T^{1/2} \cdot R^{3/2} \cdot m^{-1/4} + R^{5/4} \cdot m^{-1/8}). \quad (\text{D.26})$$

Also, combining Assumption A.2 and Theorem 4.3, it holds that

$$(\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i} \leq 2c_0 \cdot \mathbb{E}[\|Q^{\pi_i} - Q_{\omega_i}\|_{\mathcal{S}_i}] = c_0 \cdot \mathcal{O}(R^{3/2} \cdot m^{-1/4} + R^{5/4} \cdot m^{-1/8}). \quad (\text{D.27})$$

Finally, by plugging (D.25), (D.26), and (D.27) into (D.24) and setting

$$\bar{\varepsilon}_i(T) = \sqrt{T} \cdot \varepsilon_i + (\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i},$$

we complete the proof of Theorem A.4. \square

E LINEARIZATION ERROR

In this section, we lay out a fundamental lemma that characterizes the distance between a two-layer neural network $\phi_\theta^\top \theta$ and its linearization $\phi_0^\top \theta$, where ϕ_θ is the feature mapping of the two-layer

neural network defined in (3.3) and ϕ_0 is the feature mapping corresponding to the initial parameter W_{init} .

We first introduce a function class that consists of linearizations of $f(x; W)$ defined in (3.1). Recall that we write $x = (s, a)$.

Definition E.1 (Function Class). Let $R > 0$ be an absolute constant. For all $m \in \mathbb{N}$, we define

$$\tilde{\mathcal{F}}_{R,m} = \left\{ \hat{f}(x; W) = \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m b_r \cdot \mathbb{1}\{[W_{\text{init}}]_r^\top x > 0\} \cdot [W]_r^\top x : \|W - W_{\text{init}}\|_2 \leq R \right\}, \quad (\text{E.1})$$

where $[W_{\text{init}}]_r \sim N(0, I_d/d)$ and $b_r \sim \text{Unif}(\{-1, 1\})$ are the initial parameters of the two-layer neural network defined in (3.1).

Note that $\tilde{\mathcal{F}}_{R,m}$ in (E.1) is a class of functions that are linear in W but nonlinear in x . Meanwhile, it holds that $\nabla_W \hat{f}(x; W) = \nabla_W f(x; W)|_{W=W_{\text{init}}}$ for all $x \in \mathcal{S} \times \mathcal{A}$, where $f(\cdot; W)$ is defined in (3.1). Thus, $\hat{f}(\cdot; W)$ can be viewed as the linearization of $f(\cdot; W)$ at the initial parameter W_{init} . Moreover, for a fixed R , the linearization error of $\hat{f}(\cdot; W)$ decays to zero as the width $m \rightarrow \infty$. Intuitively, since the norm $\|W - W_{\text{init}}\|_2$ is upper bounded by R , the differences between blocks $\|[W]_r - [W_{\text{init}}]_r\|_2$ are sufficiently small for a sufficiently large m and all $r \in [m]$. As a result, for a sufficiently large m , we have $\mathbb{1}\{[W_{\text{init}}]_r^\top x > 0\} = \mathbb{1}\{[W]_r^\top x > 0\}$ with high probability for all $r \in [m]$, and thus $f(x; W)$ is well approximated by its linearization $\hat{f}(x; W)$.

The following lemma formally characterizes the corresponding linearization error.

Lemma E.2 (Linearization Error (Cai et al., 2019)). Let W_{init} be the initial parameter of the two-layer neural network defined in (3.1). Let $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$. Under Assumption 4.2, it holds for all $\theta, \theta' \in \mathcal{B}$ that

$$\mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_\sigma^2] = \mathcal{O}(R^3 \cdot m^{-1/2}),$$

where the expectation is taken over the random initialization. Here $\phi_\theta, \phi_0 : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{md}$ are the feature mappings defined in (3.3) with parameter θ and W_{init} , respectively, and $\sigma(\cdot, \cdot) = \pi(\cdot | \cdot) \cdot \nu(\cdot)$ is a distribution over $\mathcal{S} \times \mathcal{A}$ such that Assumption 4.2 holds.

Proof. Following from (3.3), it holds that

$$\begin{aligned} & \phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta' \\ &= \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m \left(\mathbb{1}\{(s, a)^\top [\theta]_r > 0\} - \mathbb{1}\{(s, a)^\top [W_{\text{init}}]_r > 0\} \right) \cdot (s, a)^\top [\theta']_r. \end{aligned} \quad (\text{E.2})$$

Meanwhile, for $\mathbb{1}\{(s, a)^\top [\theta]_r > 0\} \neq \mathbb{1}\{(s, a)^\top [W_{\text{init}}]_r > 0\}$, we have

$$|(s, a)^\top [W_{\text{init}}]_r| \leq |(s, a)^\top [\theta]_r - (s, a)^\top [W_{\text{init}}]_r| \leq \|(s, a)\|_2 \cdot \|[\theta]_r - [W_{\text{init}}]_r\|_2, \quad (\text{E.3})$$

where the last inequality follows from the Cauchy-Schwartz inequality. Recall that $\|(s, a)\|_2 \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Thus, it follows from (E.3) that

$$|\mathbb{1}\{(s, a)^\top [\theta]_r > 0\} - \mathbb{1}\{(s, a)^\top [W_{\text{init}}]_r > 0\}| \leq \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\}. \quad (\text{E.4})$$

By plugging (E.4) into (E.2), we obtain that

$$\begin{aligned} & |\phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta'| \\ & \leq \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\} \cdot |(s, a)^\top [\theta']_r| \\ & \leq \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\} \cdot (|(s, a)^\top [W_{\text{init}}]_r| + \|[\theta]_r - [W_{\text{init}}]_r\|_2), \end{aligned} \quad (\text{E.5})$$

where the last inequality follows from the Cauchy-Schwartz inequality and the fact that $\|(s, a)\|_2 \leq 1$. Following from the fact that $\mathbb{1}\{|x| \leq y\} \cdot |x| \leq \mathbb{1}\{|x| \leq y\} \cdot y$, we obtain from (E.5) that

$$\begin{aligned} & |\phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta'| \\ & \leq \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\} \cdot (\|[\theta]_r - [W_{\text{init}}]_r\|_2 + \|[\theta']_r - [W_{\text{init}}]_r\|_2). \end{aligned} \quad (\text{E.6})$$

Therefore, following from the Cauchy-Schwartz inequality, we obtain from (E.6) that

$$\begin{aligned} & |\phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta'|^2 \\ & \leq \frac{1}{m} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\} \cdot \sum_{r=1}^m (2\|[\theta]_r - [W_{\text{init}}]_r\|_2^2 + 2\|[\theta']_r - [W_{\text{init}}]_r\|_2^2) \\ & \leq \frac{1}{m} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\} \cdot 2(\|\theta - W_{\text{init}}\|_2^2 + \|\theta' - W_{\text{init}}\|_2^2), \end{aligned} \quad (\text{E.7})$$

where the first inequality follows from the fact that $(x + y)^2 \leq 2x^2 + 2y^2$. Recall that $\theta, \theta' \in \mathcal{B}$, where $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$. Thus, following from (E.7), we have

$$|\phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta'|^2 \leq \frac{4R^2}{m} \cdot \sum_{r=1}^m \mathbb{1}\{|(s, a)^\top [W_{\text{init}}]_r| \leq \|[\theta]_r - [W_{\text{init}}]_r\|_2\}. \quad (\text{E.8})$$

Following from Assumption 4.2, we obtain from (E.8) that

$$\begin{aligned} \|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_\sigma^2 & = \mathbb{E}_\sigma [\|\phi_\theta(s, a)^\top \theta' - \phi_0(s, a)^\top \theta'\|^2] \\ & \leq \frac{4cR^2}{m} \cdot \sum_{r=1}^m \frac{\|[\theta]_r - [W_{\text{init}}]_r\|_2}{\|[W_{\text{init}}]_r\|_2}, \end{aligned} \quad (\text{E.9})$$

where c is the absolute constant defined by Assumption 4.2. It now suffices to take the expectation of the right-hand side of (E.9) over the random initialization. Following from the Cauchy-Schwartz inequality, we obtain that

$$\begin{aligned} \left(\sum_{r=1}^m \frac{\|[\theta]_r - [W_{\text{init}}]_r\|_2}{\|[W_{\text{init}}]_r\|_2} \right)^2 & \leq \left(\sum_{r=1}^m \|[\theta]_r - [W_{\text{init}}]_r\|_2^2 \right) \cdot \left(\sum_{r=1}^m 1/\|[W_{\text{init}}]_r\|_2^2 \right) \\ & = \|\theta - W_{\text{init}}\|_2^2 \cdot \sum_{r=1}^m 1/\|[W_{\text{init}}]_r\|_2^2 \\ & \leq R^2 \cdot \sum_{r=1}^m 1/\|[W_{\text{init}}]_r\|_2^2, \end{aligned} \quad (\text{E.10})$$

where the last inequality follows from the fact that $\theta \in \mathcal{B}$. Therefore, combining (E.9) and (E.10), we conclude that

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_\sigma^2] & \leq \frac{4cR^3}{m} \cdot \mathbb{E}_{\text{init}} \left[\left(\sum_{r=1}^m 1/\|[W_{\text{init}}]_r\|_2^2 \right)^{1/2} \right] \\ & \leq \frac{4cR^3}{m} \cdot \left(\sum_{r=1}^m \mathbb{E}_{\text{init}} [1/\|[W_{\text{init}}]_r\|_2^2] \right)^{1/2} \\ & = 4c_1 R^3 \cdot m^{-1/2}, \end{aligned}$$

where the second inequality follows from the Jensen's inequality and $c_1 = c \cdot \mathbb{E}_{x \sim N(0, I_d/d)} [1/\|x\|_2^2]$. Thus, we complete the proof of Lemma E.2. \square

Following from Lemma E.2, the linearization $\phi_0^\top \theta$ converges to the two-layer neural network $\phi_\theta^\top \theta$ as the width $m \rightarrow \infty$. Based on Lemma E.2, the following corollary characterizes a similar convergence where the feature mappings ϕ_θ and ϕ_0 are replaced by the centered feature mappings $\bar{\phi}_0$ and $\bar{\phi}_\theta$ defined in (B.1) and (B.2), respectively.

Corollary E.3. Let W_{init} be the initial parameter and $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ be the parameter space. Under Assumption 4.2, it holds for all $\theta, \theta' \in \mathcal{B}$ that

$$\mathbb{E}_{\text{init}} [\|\bar{\phi}_\theta(\cdot, \cdot)^\top \theta' - \bar{\phi}_0(\cdot, \cdot)^\top \theta'\|_\sigma^2] = \mathcal{O}(R^3 \cdot m^{-1/2}),$$

where the expectation is taken over the random initialization. Here $\bar{\phi}_0$ and $\bar{\phi}_\theta$ are the centered feature mappings defined in (B.1) and (B.2), respectively, and $\sigma(\cdot, \cdot) = \pi(\cdot | \cdot) \cdot \nu(\cdot)$ is a distribution over $\mathcal{S} \times \mathcal{A}$ such that Assumption 4.2 holds.

Proof. Following from the definitions of $\bar{\phi}_0$ and $\bar{\phi}_\theta$ in (B.1) and (B.2), respectively, we obtain that

$$\begin{aligned} \|\bar{\phi}_\theta(\cdot, \cdot)^\top \theta' - \bar{\phi}_0(\cdot, \cdot)^\top \theta'\|_\sigma^2 &= \|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta' - \mathbb{E}_{\pi_\theta} [\phi_\theta(\cdot, a')^\top \theta' - \phi_0(\cdot, a')^\top \theta']\|_\sigma^2 \\ &\leq 2\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_\sigma^2 + 2\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_{\pi_\theta, \nu}^2, \end{aligned}$$

where the second inequality follows from the Jensen's inequality and the fact that $\|x + y\|_2^2 \leq 2\|x\|_2^2 + 2\|y\|_2^2$. Therefore, following from Assumption 4.2 and Lemma E.2, we obtain that

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|\bar{\phi}_\theta(\cdot, \cdot)^\top \theta' - \bar{\phi}_0(\cdot, \cdot)^\top \theta'\|_\sigma^2] &\leq 2\mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_\sigma^2] + 2\mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta' - \phi_0(\cdot, \cdot)^\top \theta'\|_{\pi_\theta, \nu}^2] = \mathcal{O}(R^3 \cdot m^{-1/2}), \end{aligned}$$

which concludes the proof of Corollary E.3. \square

In what follows, we present a corollary that quantifies the difference between the function $\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta$ and the two-layer neural network $f((\cdot, \cdot); \theta) = \phi_\theta(\cdot, \cdot)^\top \theta$ by the $L_2(\sigma)$ -norm, where $\sigma(\cdot, \cdot) = \pi(\cdot | \cdot) \cdot \nu(\cdot)$ is a distribution over $\mathcal{S} \times \mathcal{A}$ such that Assumption 4.2 holds.

Corollary E.4. Let $\mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$. Under Assumption 4.2, it holds for all $\theta, \hat{\theta} \in \mathcal{B}$ that

$$\mathbb{E}_{\text{init}} [\|\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta - \phi_\theta(\cdot, \cdot)^\top \theta\|_\sigma] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}),$$

where the expectation is taken over the random initialization. Here $\phi_\theta: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{md}$ is the feature mapping with parameter θ , which is defined in (3.3), and σ is a distribution over $\mathcal{S} \times \mathcal{A}$ such that Assumption 4.2 holds.

Proof. By the triangle inequality, we have

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta - \phi_\theta(\cdot, \cdot)^\top \theta\|_\sigma] &\leq \mathbb{E}_{\text{init}} [\|\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta - \phi_0(\cdot, \cdot)^\top \theta\|_\sigma] + \mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta - \phi_0(\cdot, \cdot)^\top \theta\|_\sigma], \end{aligned} \quad (\text{E.11})$$

where ϕ_0 is the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. Meanwhile, for all $\theta, \hat{\theta} \in \mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$, it follows from Lemma E.2 that

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta - \phi_0(\cdot, \cdot)^\top \theta\|_\sigma] &= \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \\ \mathbb{E}_{\text{init}} [\|\phi_\theta(\cdot, \cdot)^\top \theta - \phi_0(\cdot, \cdot)^\top \theta\|_\sigma] &= \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{E.12})$$

where the expectations are taken over the random initialization. Combining (E.11) and (E.12), we obtain that

$$\mathbb{E}_{\text{init}} [\|\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta - \phi_\theta(\cdot, \cdot)^\top \theta\|_\sigma] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}),$$

which concludes the proof of Corollary E.4. \square

Corollary E.4 implies that when the width m is sufficiently large, $\phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta$ is well approximated by the two-layer neural network $f((\cdot, \cdot); \theta)$ in $L_2(\sigma)$ -norm, where $\sigma(\cdot, \cdot)$ is the distribution such that Assumption 4.2 holds.

F NEURAL TD

In this section, we introduce the details of neural TD (Cai et al., 2019) for critic update in Algorithm 1. Neural TD solves the optimization problem in (3.11) using the TD iterations defined in (3.12) and (3.13), which is summarized in Algorithm 2.

Algorithm 2 Neural TD (Cai et al., 2019)

Require: MDP $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, the policy π , number of TD iterations T_{TD} , learning rate of TD iterations η_{TD} .

- 1: **Initialization:** Initialize $b_r \sim \text{Unif}(\{-1, 1\})$ and $[W_{\text{init}}]_r \sim N(0, I_d/d)$. Set $\mathcal{B} \leftarrow \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ and $\omega(0) \leftarrow W_{\text{init}}$.
- 2: **for** $t = 0, \dots, T_{\text{TD}} - 1$ **do**
- 3: Sample a tuple (s, a, r, s', a') , where $(s, a) \sim \zeta_i$, $s' \sim \mathcal{P}(\cdot | s, a)$, $r = r(s, a)$, and $a' \sim \pi(\cdot | s')$. Set $x \leftarrow (s, a)$ and $x' \leftarrow (s', a')$.
- 4: Compute the Bellman residue $\delta \leftarrow Q_\omega(x) - (1 - \gamma) \cdot r - \gamma \cdot Q(x')$.
- 5: Perform a TD update step: $\omega(t + 1/2) \leftarrow \omega(t) - \eta \cdot \delta \cdot \nabla_\omega Q_{\omega(t)}(x)$.
- 6: Perform a projection step: $\omega(t + 1) \leftarrow \Pi_{\mathcal{B}}(\omega(t + 1/2))$.
- 7: Perform an averaging step: $\bar{\omega} \leftarrow \frac{t+1}{t+2} \cdot \bar{\omega} + \frac{1}{t+2} \cdot \omega(t + 1)$.
- 8: **end for**
- 9: **Output:** $Q_{\text{out}}(\cdot) \leftarrow Q_{\bar{\omega}}(\cdot)$

The following theorem by Cai et al. (2019) characterizes the rate of convergence of Algorithm 2.

Theorem F.1 (Convergence of Neural TD (Cai et al., 2019)). We set $\eta_{\text{TD}} = \min\{(1 - \gamma)/8, 1/\sqrt{T_{\text{TD}}}\}$ in Algorithm 2. Under Assumption 4.2, it holds that

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|Q_{\text{out}} - Q^\pi\|_{\zeta_\pi}^2] &\leq 2\mathbb{E}_{\text{init}} [\|\Pi_{\tilde{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\zeta_\pi}^2] \\ &\quad + \mathcal{O}(R^2 \cdot T_{\text{TD}}^{-1/2} + R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{F.1})$$

where $\Pi_{\tilde{\mathcal{F}}_{R,m}}$ is the projection operator onto $\tilde{\mathcal{F}}_{R,m}$, and $\zeta_\pi(\cdot)$ is the stationary state-action distribution corresponding to π .

Proof. See Proposition 4.7 in Cai et al. (2019) for a detailed proof. \square

F.1 PROOF OF PROPOSITION 4.3

Proof. By Theorem F.1, to establish the rate of convergence of neural TD, it suffices to characterize the approximation error $\mathbb{E}_{\text{init}} [\|\Pi_{\tilde{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\zeta_\pi}^2]$ in (F.1). To this end, we first define a new function class

$$\bar{\mathcal{F}}_{R,m} = \left\{ \hat{f}(x; W) = \frac{1}{\sqrt{m}} \cdot \sum_{r=1}^m b_r \cdot \mathbf{1}\{[W_{\text{init}}]_r^\top x > 0\} \cdot W_r^\top x : \|[W]_r - [W_{\text{init}}]_r\|_\infty \leq R/\sqrt{md} \right\},$$

where $[W_{\text{init}}]_r \sim N(0, I_d/d)$ and $b_r \sim \text{Unif}(\{-1, 1\})$ are the initial parameters. By definition, $\bar{\mathcal{F}}_{R,m}$ is a subset of $\tilde{\mathcal{F}}_{R,m}$ defined in Definition E.1. The following lemma obtained from Rahimi and Recht (2009) characterizes the deviation of $\bar{\mathcal{F}}_{R,m}$ from $\mathcal{F}_{R,\infty}$ given in Assumption 4.1.

Lemma F.2 (Projection Error of $\bar{\mathcal{F}}_{R,m}$ (Rahimi and Recht, 2009)). Let $f(\cdot) \in \mathcal{F}_{R,\infty}$, where $\mathcal{F}_{R,\infty}$ is defined in Assumption 4.1. For any $\delta > 0$, it holds with probability at least $1 - \delta$ that

$$\|\Pi_{\bar{\mathcal{F}}_{R,m}} f - f\|_\zeta \leq R \cdot m^{-1/2} \cdot [1 + \sqrt{2 \log(1/\delta)}], \quad (\text{F.2})$$

where ζ is a distribution over $\mathcal{S} \times \mathcal{A}$.

Proof. See Rahimi and Recht (2009) for a detailed proof. \square

Following from (F.2) in Lemma F.2, for all $f \in \mathcal{F}_{R,\infty}$ and $t > 0$, we have

$$\mathbb{P}(\|\Pi_{\bar{\mathcal{F}}_{R,m}} f - f\|_\zeta \geq t) \leq \exp(-1/2 \cdot (t \cdot \sqrt{m}/R - 1)^2). \quad (\text{F.3})$$

Meanwhile, following from Assumption 4.1, we have $Q^\pi \in \mathcal{F}_{R,\infty}$. Therefore, by setting $f = Q^\pi$ and $\varsigma = \varsigma_\pi$ in (F.3), we obtain that

$$\begin{aligned} \mathbb{E}_{\text{init}} [\|\Pi_{\bar{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\varsigma_\pi}^2] &= \int_0^\infty t \cdot \mathbb{P}(\|\Pi_{\bar{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\varsigma_\pi}^2 \geq t) dt \\ &\leq \int_0^\infty t \cdot \exp(-1/2 \cdot (t \cdot \sqrt{m}/R - 1)^2) dt = \mathcal{O}(R^2 \cdot m^{-1}). \end{aligned} \quad (\text{F.4})$$

Also, note that $\bar{\mathcal{F}}_{R,m} \subseteq \tilde{\mathcal{F}}_{R,m}$, where $\tilde{\mathcal{F}}_{R,m}$ is defined in Definition E.1. Therefore, it follows from (F.4) that

$$\mathbb{E}_{\text{init}} [\|\Pi_{\tilde{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\varsigma_\pi}^2] \leq \mathbb{E}_{\text{init}} [\|\Pi_{\bar{\mathcal{F}}_{R,m}} Q^\pi - Q^\pi\|_{\varsigma_\pi}^2] = \mathcal{O}(R^2 \cdot m^{-1}). \quad (\text{F.5})$$

Combining (F.5) and Theorem F.1, we obtain for $\eta_{\text{TD}} = \min\{(1 - \gamma)/8, 1/\sqrt{T_{\text{TD}}}\}$ that

$$\mathbb{E}_{\text{init}} [\|Q_{\text{out}} - Q^\pi\|_{\sigma_\pi}^2] = \mathcal{O}(R^2 \cdot m^{-1} + R^2 \cdot T_{\text{TD}}^{-1/2} + R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}). \quad (\text{F.6})$$

Finally, by setting $T_{\text{TD}} = \Omega(m)$ in (F.6), we obtain

$$\mathbb{E}_{\text{init}} [\|Q_{\omega_i} - Q^{\pi_{\theta_i}}\|_{\varsigma_i}^2] = \mathcal{O}(R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}),$$

which concludes the proof of Proposition 4.3. \square

G PROJECTION-FREE NEURAL POLICY GRADIENT

In this section, we study the global convergence of neural policy gradient where we do not impose projection in actor update. In specific, the actor update of projection-free neural natural gradient takes the form of

$$\theta_{i+1} \leftarrow \theta_i + \eta \cdot \hat{\nabla}_\theta J(\pi_{\theta_i}).$$

Here $\hat{\nabla}_\theta J(\pi_{\theta_i})$ is the estimator defined in (3.7). We summarize projection-free neural policy gradient in Algorithm 3.

Algorithm 3 Projection-Free Neural Policy Gradient

Require: Number of iterations T , number of TD iterations T_{TD} , learning rate η , learning rate of TD iterations η_{TD} , temperature parameters τ_1, \dots, τ_T , batch size B .

- 1: **Initialization:** Initialize $b_r \sim \text{Unif}\{-1, 1\}$ and $[W_{\text{init}}]_r \sim N(0, I_d/d)$ for all $r \in [m]$. Set $\mathcal{B} \leftarrow \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$ and $\theta_1 \leftarrow W_{\text{init}}$.
 - 2: **for** $i \in [T]$ **do**
 - 3: Update ω_i using Algorithm 2 with π_{θ_i} as the input, $\omega(0) \leftarrow W_{\text{init}}$ and $\{b_r\}_{r \in [m]}$ as the initialization, T_{TD} as the number of iterations, and η_{TD} as the learning rate.
 - 4: Sample $\{(s_\ell, a_\ell)\}_{\ell \in [B]}$ from the visitation measure σ_i , and estimate $\hat{\nabla}_\theta J(\pi_\theta)$ using (3.7).
 - 5: Update θ_{i+1} by $\theta_{i+1} \leftarrow \theta_i + \eta \cdot \hat{\nabla}_\theta J(\pi_{\theta_i})$.
 - 6: **end for**
 - 7: **Output:** $\{\pi_{\theta_i}\}_{i \in [T+1]}$.
-

Similar to Theorem 4.7, in the following theorem, we show that the sequence $\{\theta_i\}_{i \in [T+1]}$ generated by Algorithm 3 converges to a stationary point $\hat{\theta}$ with $\nabla_\theta J(\pi_{\hat{\theta}}) = 0$ at a sublinear rate.

Theorem G.1 (Convergence to Stationary Point). Let $\eta = 1/\sqrt{T}$, $\tau_i = 1$, $\eta_{\text{TD}} = \min\{(1 - \gamma)/8, 1/\sqrt{T_{\text{TD}}}\}$, and $T_{\text{TD}} = \Omega(m)$ in Algorithm 3. Under the assumptions of Theorem 4.3, Assumptions 4.4-4.6, and Assumption A.3, it holds for $T \geq 4L^2$ and $B = \Omega(\sigma_\varepsilon \cdot T^{1/2})$ that

$$\min_{i \in [T]} \mathbb{E} [\|\nabla_\theta J(\pi_{\theta_i})\|_2^2] \leq 8/\sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{T+1}}) - J(\pi_{\theta_1})] + \epsilon,$$

where

$$\epsilon = \mathcal{O}(T^{-1/2} + T \cdot R^{5/2} \cdot m^{-1/4} + T \cdot R^{9/4} \cdot m^{-1/8}).$$

Here the expectations are taken over all the randomness.

Proof. Our proof aligns closely to that of Theorem 4.7. We first lower bound the difference $J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i})$. By Assumption 4.6, we have

$$J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i}) \geq \eta \cdot \nabla_{\theta} J(\pi_{\theta_i})^{\top} \delta_i - L/2 \cdot \|\theta_{i+1} - \theta_i\|_2^2. \quad (\text{G.1})$$

In what follows, we define

$$\delta_i = (\theta_{i+1} - \theta_i)/\eta = \widehat{\nabla}_{\theta} J(\pi_{\theta_i}), \quad \forall i \in [T].$$

Following from the actor update in Algorithm 3, we have $\delta_i = \widehat{\nabla}_{\theta} J(\pi_{\theta_i})$. Meanwhile, following the proof of Lemma D.2 in §H.5, we obtain that

$$\left| \left(\nabla_{\theta} J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_i})] \right)^{\top} \delta_i \right| \leq \kappa/\eta \cdot 2\|\theta_{i+1} - \theta_i\|_2 \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}, \quad (\text{G.2})$$

where the expectation $\mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_i})]$ is taken over σ_i given θ_i and ω_i . Recall that $\xi_i = \widehat{\nabla}_{\theta} J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_i})]$, where the expectation is taken over σ_i given θ_i and ω_i . Following from (G.2), we obtain

$$\begin{aligned} \nabla_{\theta} J(\pi_{\theta_i})^{\top} \delta_i &= \left(\nabla_{\theta} J(\pi_{\theta_i}) - \mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_i})] \right)^{\top} \delta_i - \xi_i^{\top} \delta_i + \widehat{\nabla}_{\theta} J(\pi_{\theta_i})^{\top} \delta_i \\ &\geq -2\kappa \cdot \|\theta_{i+1} - \theta_i\|_2/\eta \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i} - \|\xi_i\|_2^2/2 - \|\delta_i\|_2^2/2, \end{aligned} \quad (\text{G.3})$$

where the second inequality follows from a similar argument to §D.1. Hence, by plugging (G.3) into (G.1), we have

$$\begin{aligned} J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i}) &\geq (\eta - L \cdot \eta^2)/2 \cdot \|\delta_i\|_2^2 - \eta \cdot \|\xi_i\|_2^2/2 - 2\kappa \cdot \|\theta_{i+1} - \theta_i\|_2/\eta \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}. \end{aligned} \quad (\text{G.4})$$

It remains to upper bound the distance $\|\theta_{i+1} - \theta_i\|_2$. To this end, we use the fact that $\|\theta_{i+1} - \theta_i\|_2 \leq \|\theta_i - W_{\text{init}}\|_2 + \|\theta_{i+1} - W_{\text{init}}\|_2$ and upper bound the norms $\|\theta_i - W_{\text{init}}\|_2$ and $\|\theta_{i+1} - W_{\text{init}}\|_2$. Following from the actor update in Algorithm 3, we obtain that

$$\|\theta_i - W_{\text{init}}\|_2 \leq \sum_{j=1}^{i-1} \eta \cdot \|\widehat{\nabla}_{\theta} J(\pi_{\theta_j})\|_2 \leq \sum_{j=1}^{i-1} \eta \cdot \left(\|\mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_j})]\|_2 + \|\xi_j\|_2 \right), \quad (\text{G.5})$$

where the expectation $\mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_j})]$ is taken over σ_j given θ_j and ω_j . Meanwhile, it holds that

$$\|\mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_j})]\|_2 = \|\mathbb{E}_{\sigma_j}[\bar{\phi}_{\theta_j}(s, a) \cdot Q_{\omega_j}(s, a)]\|_2 \leq \mathbb{E}_{\sigma_j}[\|\bar{\phi}_{\theta_j}(s, a)\|_2 \cdot |Q_{\omega_j}(s, a)|], \quad (\text{G.6})$$

where $\bar{\phi}_{\theta_j}$ is the centered feature mapping defined in (B.2), and the last inequality follows from the Jensen's inequality. We now upper bound the right-hand side of (G.6). Note that $\|\bar{\phi}_{\theta_j}(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Meanwhile, for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned} |Q_{\omega_j}(s, a)| &= |\phi_{\omega_j}(s, a)^{\top} \omega_j| \leq |\phi_{\omega_j}(s, a)^{\top} \omega_j - \phi_0(s, a)^{\top} W_{\text{init}}| + |\phi_0(s, a)^{\top} W_{\text{init}}|, \\ &\leq \|\omega_j - W_{\text{init}}\|_2 + M_0 \leq R + M_0, \end{aligned} \quad (\text{G.7})$$

where ϕ_{ω_j} and ϕ_0 are the feature mappings defined in (3.3) with $\theta = \omega_j$ and $\theta = W_{\text{init}}$, respectively. Here the second inequality follows from the 1-Lipschitz continuity of the two-layer neural network $f(x; W)$ with respect to the parameter W , and the last inequality follows from the fact that $\omega_j \in \mathcal{B}$. Meanwhile, W_{init} is the initial parameter and M_0 is defined as

$$M_0 = \sup_{(s, a) \in \mathcal{S} \times \mathcal{A}} |\phi_0(s, a)^{\top} W_{\text{init}}|.$$

By plugging (G.7) into (G.6), we obtain that

$$\|\mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_j})]\|_2 \leq 2R + 2M_0, \quad (\text{G.8})$$

By further plugging (G.8) into (G.5), we obtain for all $i \in [T + 1]$ that

$$\|\theta_i - W_{\text{init}}\|_2 \leq 2\eta \cdot T \cdot (R + M_0) + \sum_{j=1}^{i-1} \eta \cdot \|\xi_j\|_2. \quad (\text{G.9})$$

We now lower bound the right-hand side of (G.4) based on (G.9). Following from the Cauchy-Schwartz inequality, Assumption 4.4, and Assumption A.3, we obtain that

$$\begin{aligned} & \mathbb{E}[\|\theta_i - W_{\text{init}}\|_2 \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}] \tag{G.10} \\ & \leq 2\eta \cdot T \cdot \left(R + \{\mathbb{E}[M_0^2]\}^{1/2}\right) \cdot \left\{\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2]\right\}^{1/2} + \sum_{j=1}^{i-1} \eta \cdot \left\{\mathbb{E}[\|\xi_j\|_2^2]\right\}^{1/2} \cdot \left\{\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2]\right\}^{1/2}, \\ & \leq (2\eta \cdot T \cdot (R + M) + \sigma_\xi \cdot \eta \cdot T \cdot B^{-1/2}) \cdot \left\{\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2]\right\}^{1/2}, \end{aligned}$$

where the expectation is taken over all the randomness, and σ_ξ , M are the absolute constants defined in Assumptions 4.4 and A.3, respectively. By plugging (G.10) into (G.4), we obtain that

$$\begin{aligned} & (\eta - L \cdot \eta^2)/2 \cdot \mathbb{E}[\|\delta_i\|_2^2] \tag{G.11} \\ & \leq \mathbb{E}[J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i})] + \eta \cdot \sigma_\xi^2/(2B) + R_0 \cdot \left(\mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2]\right)^{1/2}, \end{aligned}$$

where we use the fact that $\|\theta_{i+1} - \theta_i\|_2 \leq \|\theta_{i+1} - W_{\text{init}}\|_2 + \|\theta_i - W_{\text{init}}\|_2$. Here the expectations are taken over all the randomness, and R_0 is defined as

$$R_0 = 4\eta \cdot T \cdot (R + M) + 2\sigma_\xi \cdot \eta \cdot T \cdot B^{-1/2},$$

Following from Theorem 4.3 and Assumption 4.2, it holds for $\eta = 1/\sqrt{T}$, $B = \Omega(\sigma_\xi \cdot T^{1/2})$, and $T_{\text{TD}} = \Omega(m)$ that

$$R_0 = \mathcal{O}(R \cdot \sqrt{T}), \quad \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2] = \mathcal{O}(R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}). \tag{G.12}$$

By plugging (G.12) into (G.11) with $\eta = 1/\sqrt{T}$ and $B = \Omega(\sigma_\xi \cdot T^{1/2})$, we obtain that

$$(1 - L/\sqrt{T})/2 \cdot \mathbb{E}[\|\delta_i\|_2^2] \leq \sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i})] + \epsilon, \tag{G.13}$$

where

$$\epsilon = \mathcal{O}(T^{-1/2} + R^{5/2} \cdot m^{-1/4} \cdot T + R^{9/4} \cdot m^{-1/8} \cdot T). \tag{G.14}$$

It remains to upper bound $\|\delta_i - \nabla_\theta J(\pi_{\theta_i})\|_2$, where $\delta_i = \widehat{\nabla} J(\pi_{\theta_i})$. Following from Lemma D.3, we obtain that

$$\mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i}) - \widehat{\nabla} J(\pi_{\theta_i})\|_2^2] \leq 2\mathbb{E}[\|\xi_i\|_2^2] + 4\kappa^2 \cdot \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2],$$

where the expectations are taken over all the randomness. Therefore, following from Proposition 4.3 and Assumption 4.4, it holds for $\eta = 1/\sqrt{T}$, $B = \Omega(T^{1/2})$, and $T_{\text{TD}} = \Omega(m)$ that

$$\mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i}) - \widehat{\nabla} J(\pi_{\theta_i})\|_2^2] = \mathcal{O}(\sigma_\xi^2 \cdot T^{-1/2} + R^3 \cdot m^{-1/2} + R^{5/2} \cdot m^{-1/4}). \tag{G.15}$$

Thus, combining (G.13) and (G.15), we obtain that

$$\begin{aligned} \mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i})\|_2^2] & \leq 2\mathbb{E}[\|\delta_i\|_2^2] + 2\mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i}) - \widehat{\nabla} J(\pi_{\theta_i})\|_2^2] \\ & \leq 4(1 - L/\sqrt{T}) \cdot \mathbb{E}[\|\delta_i\|_2^2] + 2\mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i}) - \widehat{\nabla} J(\pi_{\theta_i})\|_2^2] \\ & \leq 8\sqrt{T} \cdot \mathbb{E}[J(\pi_{\theta_{i+1}}) - J(\pi_{\theta_i})] + \epsilon, \end{aligned} \tag{G.16}$$

where we use the fact that $T \geq 4L^2$ and we define ϵ in (G.14). By telescoping (G.16), we obtain that

$$\min_{i \in [T]} \mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i})\|_2^2] \leq \frac{1}{T} \cdot \sum_{i=1}^T \mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i})\|_2^2] \leq 8\mathbb{E}[J(\pi_{\theta_{T+1}}) - J(\pi_{\theta_1})]/\sqrt{T} + \epsilon,$$

where

$$\epsilon = \mathcal{O}(T^{-1/2} + R^{5/2} \cdot m^{-1/4} \cdot T + R^{9/4} \cdot m^{-1/8} \cdot T).$$

Here the expectations are taken over all the randomness. Thus, we complete the proof of Theorem G.1. \square

Following from Theorem G.1, it holds for $m = \Omega(R^{18} \cdot T^{12})$ that

$$\min_{i \in [T]} \mathbb{E}[\|\nabla_\theta J(\pi_{\theta_i})\|_2^2] = \mathcal{O}(1/\sqrt{T}).$$

Therefore, θ_i converges to a stationary point at a $1/\sqrt{T}$ -rate if the width m of the two-layer neural network and the batch size B are sufficiently large. We highlight that comparing with neural policy gradient with projection in actor update, Algorithm 3 needs a larger width m for the $1/\sqrt{T}$ -rate of convergence. Such an extra requirement on m corresponds to an extra price paid for having the projection-free actor update.

G.1 GLOBAL OPTIMALITY OF STATIONARY POINT

In this section, we characterize the global optimality of a stationary point $\hat{\theta}$ such that

$$\nabla_{\theta} J(\pi_{\hat{\theta}})^{\top} (\theta - \hat{\theta}) \leq 0, \quad \forall \theta \in \mathbb{R}^{md}.$$

Following from Theorem 4.8, we obtain that

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \inf_{\theta \in \mathbb{R}^{md}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}}. \quad (\text{G.17})$$

We define $\tilde{R} = 2 \cdot \|\hat{\theta} - W_{\text{init}}\|_2$ and $\tilde{\mathcal{B}} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq \tilde{R}\}$, where W_{init} is the initialization of parameters. It then holds that $\hat{\theta} \in \tilde{\mathcal{B}}$. Moreover, following from (G.17) and the fact that $\tilde{\mathcal{B}} \subseteq \mathbb{R}^{md}$, we obtain that

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \inf_{\theta \in \tilde{\mathcal{B}}} \|u(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \quad (\text{G.18})$$

Meanwhile, following from Corollary E.4, it holds that $\phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta$ is well approximated by $f((\cdot, \cdot); \theta)$ for $\theta, \hat{\theta} \in \tilde{\mathcal{B}}$. Thus, following from (G.18), the suboptimality of $\hat{\theta}$ is characterized by the approximation error $\inf_{\theta \in \tilde{\mathcal{B}}} \|u_{\hat{\theta}}(\cdot, \cdot) - f((\cdot, \cdot); \theta)\|_{\sigma_{\pi_{\hat{\theta}}}}$, which quantifies the representation power of neural networks. In the following theorem, we present a sufficient condition for a stationary point to be globally optimal.

Theorem G.2 (Global Optimality of Stationary Point). Let $\hat{\theta}$ be a stationary point of $J(\pi_{\theta})$. Let $\tilde{R} = 2 \cdot \|\hat{\theta} - W_{\text{init}}\|_2$. It then holds for $u_{\hat{\theta}}(s, a) \in \mathcal{F}_{\tilde{R}, \infty}$ that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}} [J(\pi^*) - J(\pi_{\hat{\theta}})] = \mathcal{O}(\tilde{R}^{3/2} \cdot m^{-1/4}),$$

where $u_{\hat{\theta}}(s, a)$ is defined in (4.4) and the expectation is taken over the random initialization.

Proof. The proof aligns closely to that of Theorem 4.9. Recall that we define $\tilde{R} = 2 \cdot \|\hat{\theta} - W_{\text{init}}\|_2$ and $\tilde{\mathcal{B}} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq \tilde{R}\}$, where W_{init} is the initial parameter. To prove Theorem G.2, it suffices to upper bound the right-hand side of (G.18) under the expectation over the random initialization. Following from the triangle inequality, we obtain that

$$\begin{aligned} & \inf_{\theta \in \tilde{\mathcal{B}}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \\ & \leq \inf_{\theta \in \tilde{\mathcal{B}}} \left\{ \|u_{\hat{\theta}}(\cdot, \cdot) - \Pi_{\tilde{\mathcal{F}}_{\tilde{R}, m}} u_{\hat{\theta}}(\cdot, \cdot)\|_{\sigma_{\pi_{\hat{\theta}}}} + \|\Pi_{\tilde{\mathcal{F}}_{\tilde{R}, m}} u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \right\} \\ & = \|u_{\hat{\theta}}(\cdot, \cdot) - \Pi_{\tilde{\mathcal{F}}_{\tilde{R}, m}} u_{\hat{\theta}}(\cdot, \cdot)\|_{\sigma_{\pi_{\hat{\theta}}}} + \inf_{\theta \in \tilde{\mathcal{B}}} \|\Pi_{\tilde{\mathcal{F}}_{\tilde{R}, m}} u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}}, \end{aligned} \quad (\text{G.19})$$

where $\tilde{\mathcal{F}}_{\tilde{R}, m}$ is defined in Definition E.1. It remains to upper bound the right-hand side of (G.19).

In what follows, we define $\tilde{\theta} \in \tilde{\mathcal{B}}$ as

$$\phi_0(\cdot, \cdot)^{\top} \tilde{\theta} = \Pi_{\tilde{\mathcal{F}}_{\tilde{R}, m}} u_{\hat{\theta}}(\cdot, \cdot) \in \tilde{\mathcal{F}}_{\tilde{R}, m},$$

where ϕ_0 is the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. Following from (G.19) and the fact that $\tilde{\theta} \in \tilde{\mathcal{B}}$, we have

$$\inf_{\theta \in \tilde{\mathcal{B}}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \leq \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_0(\cdot, \cdot)^{\top} \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}} + \|\phi_0(\cdot, \cdot)^{\top} \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^{\top} \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}. \quad (\text{G.20})$$

Since $u_{\hat{\theta}} \in \tilde{\mathcal{F}}_{\tilde{R}, \infty}$, following from the proof of Proposition 4.3 in §F.1, we obtain that

$$\begin{aligned} & \mathbb{E}_{\text{init}} [\|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}] \\ & \leq \left\{ \mathbb{E}_{\text{init}} [\|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}^2] \right\}^{1/2} = \mathcal{O}(\tilde{R} \cdot m^{-1/2}), \end{aligned} \quad (\text{G.21})$$

where the first inequality follows from the Jensen's inequality, and the expectations are taken over the random initialization. Meanwhile, following from Lemma E.2, we obtain for all $\hat{\theta}, \tilde{\theta} \in \mathcal{B}$ that

$$\begin{aligned} & \mathbb{E}_{\text{init}} [\|\phi_0(\cdot, \cdot)^\top \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}] \\ & \leq \left\{ \mathbb{E}_{\text{init}} [\|\phi_0(\cdot, \cdot)^\top \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}^2] \right\}^{1/2} = \mathcal{O}(\tilde{R}^{3/2} \cdot m^{-1/4}). \end{aligned} \quad (\text{G.22})$$

Finally, upon taking expectation in (G.20) over the random initialization and further plugging (G.21) and (G.22) into (G.20), we conclude that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}} [J(\pi^*) - J(\pi_{\hat{\theta}})] \leq 2Q_{\max} \cdot \mathbb{E}_{\text{init}} \left[\inf_{\theta \in \tilde{\mathcal{B}}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \right] = \mathcal{O}(\tilde{R}^{3/2} \cdot m^{-1/4}),$$

where the first inequality follows from (G.18). Thus, we complete the proof of Theorem G.2. \square

Following from Theorem G.2, any stationary point $\hat{\theta}$ is near-optimal if it holds that $u_{\hat{\theta}} \in \mathcal{F}_{\tilde{R}, \infty}$ and the width m is sufficiently large.

H PROOF OF AUXILIARY RESULTS

In this section, we lay out the proofs of the auxiliary results.

H.1 PROOF OF PROPOSITION 3.1

Proof. The proof is based on the policy gradient theorem (Sutton and Barto, 2018) in (2.5) and the definition of Fisher information matrix in (2.7). It suffices to compute the gradient $\nabla_{\theta} \log \pi_{\theta}(\cdot | \cdot)$. By the parameterization of $\pi_{\theta}(\cdot | \cdot)$ in (3.2), it holds that

$$\begin{aligned} \nabla_{\theta} \log \pi_{\theta}(a | s) &= \tau \cdot \nabla_{\theta} f((s, a); \theta) - \tau \cdot \frac{\sum_{a' \in \mathcal{A}} \nabla_{\theta} f((s, a'); \theta) \cdot \exp[\tau \cdot f((s, a'); \theta)]}{\sum_{a' \in \mathcal{A}} \exp[\tau \cdot f((s, a'); \theta)]} \\ &= \tau \cdot \nabla_{\theta} f((s, a); \theta) - \tau \cdot \mathbb{E}_{\pi_{\theta}} [\nabla_{\theta} f((s, a'); \theta)], \end{aligned} \quad (\text{H.1})$$

where we write $\mathbb{E}_{\pi_{\theta}} [\nabla_{\theta} f((s, a'); \theta)] = \mathbb{E}_{a' \sim \pi_{\theta}(\cdot | s)} [\nabla_{\theta} f((s, a'); \theta)]$ for notational simplicity. Meanwhile, recall that $\nabla_{\theta} f((s, a); \theta) = \phi_{\theta}(s, a)$, where ϕ_{θ} is the feature mapping defined in (3.3). Thus, (H.1) implies that

$$\nabla_{\theta} \log \pi_{\theta}(a | s) = \tau \cdot \phi_{\theta}(s, a) - \tau \cdot \mathbb{E}_{\pi_{\theta}} [\phi_{\theta}(s, a')]. \quad (\text{H.2})$$

Finally, combining (2.5), (2.7), and (H.2), we have

$$\begin{aligned} \nabla_{\theta} J(\pi_{\theta}) &= \tau \cdot \mathbb{E}_{\sigma_{\pi_{\theta}}} \left[Q^{\pi_{\theta}}(s, a) \cdot \left(\phi_{\theta}(s, a) - \mathbb{E}_{\pi_{\theta}} [\phi_{\theta}(s, a')] \right) \right], \\ F(\theta) &= \tau^2 \cdot \mathbb{E}_{\sigma_{\pi_{\theta}}} \left[\left(\phi_{\theta}(s, a) - \mathbb{E}_{\pi_{\theta}} [\phi_{\theta}(s, a')] \right) \left(\phi_{\theta}(s, a) - \mathbb{E}_{\pi_{\theta}} [\phi_{\theta}(s, a')] \right)^\top \right], \end{aligned}$$

which concludes the proof of Proposition 3.1. \square

H.2 PROOF OF THEOREM 4.9

Proof. By Theorem 4.8, we have

$$(1 - \gamma) \cdot (J(\pi^*) - J(\pi_{\hat{\theta}})) \leq 2Q_{\max} \cdot \inf_{\theta \in \tilde{\mathcal{B}}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}}, \quad (\text{H.3})$$

where $u_{\hat{\theta}}$ is defined in (4.4). It suffices to upper bound the right-hand side of (H.3) under the expectation over the random initialization. Following from the triangle inequality, we obtain that

$$\begin{aligned} & \inf_{\theta \in \mathcal{B}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \\ & \leq \inf_{\theta \in \mathcal{B}} \left\{ \|u_{\hat{\theta}}(\cdot, \cdot) - \Pi_{\tilde{\mathcal{F}}_{R,m}} u_{\hat{\theta}}(\cdot, \cdot)\|_{\sigma_{\pi_{\hat{\theta}}}} + \|\Pi_{\tilde{\mathcal{F}}_{R,m}} u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \right\} \\ & = \|u_{\hat{\theta}}(\cdot, \cdot) - \Pi_{\tilde{\mathcal{F}}_{R,m}} u_{\hat{\theta}}(\cdot, \cdot)\|_{\sigma_{\pi_{\hat{\theta}}}} + \inf_{\theta \in \mathcal{B}} \|\Pi_{\tilde{\mathcal{F}}_{R,m}} u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}}, \end{aligned} \quad (\text{H.4})$$

where $\tilde{\mathcal{F}}_{R,m}$ is defined in Definition E.1. It remains to upper bound the right-hand side of (H.4). In what follows, we define $\tilde{\theta}$ as

$$\phi_0(\cdot, \cdot)^\top \tilde{\theta} = \Pi_{\tilde{\mathcal{F}}_{R,m}} u_{\hat{\theta}}(\cdot, \cdot) \in \tilde{\mathcal{F}}_{R,m},$$

where ϕ_0 is the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. By the definition of $\tilde{\mathcal{F}}_{R,m}$ in Definition E.1, it holds that $\tilde{\theta} \in \mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$. Thus, by (H.4) and the fact that $\tilde{\theta} \in \mathcal{B}$, we have

$$\begin{aligned} & \inf_{\theta \in \mathcal{B}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \\ & \leq \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_0(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}} + \|\phi_0(\cdot, \cdot)^\top \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}. \end{aligned} \quad (\text{H.5})$$

Following from the proof of Proposition 4.3 in §F.1, it holds for $u_{\hat{\theta}} \in \mathcal{F}_{R,\infty}$ that

$$\begin{aligned} & \mathbb{E}_{\text{init}} [\|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}] \\ & \leq \left\{ \mathbb{E}_{\text{init}} [\|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}^2] \right\}^{1/2} = \mathcal{O}(R \cdot m^{-1/2}), \end{aligned} \quad (\text{H.6})$$

where the first inequality follows from the Jensen's inequality, and the expectations are taken over the random initialization. Meanwhile, following from Lemma E.2, we obtain for all $\hat{\theta}, \tilde{\theta} \in \mathcal{B}$ that

$$\begin{aligned} & \mathbb{E}_{\text{init}} [\|\phi_0(\cdot, \cdot)^\top \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}] \\ & \leq \left\{ \mathbb{E}_{\text{init}} [\|\phi_0(\cdot, \cdot)^\top \tilde{\theta} - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \tilde{\theta}\|_{\sigma_{\pi_{\hat{\theta}}}}^2] \right\}^{1/2} = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{H.7})$$

where the expectations are taken over the random initialization. Finally, by plugging (H.6) and (H.7) into (H.5), we obtain that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}} [J(\pi^*) - J(\pi_{\hat{\theta}})] \leq 2Q_{\max} \cdot \mathbb{E}_{\text{init}} \left[\inf_{\theta \in \mathcal{B}} \|u_{\hat{\theta}}(\cdot, \cdot) - \phi_{\hat{\theta}}(\cdot, \cdot)^\top \theta\|_{\sigma_{\pi_{\hat{\theta}}}} \right] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}),$$

where the first inequality follows from (H.3). Meanwhile, if $u_{\hat{\theta}} \in \mathcal{F}_{R,\infty}$ is not imposed, similarly, we conclude that

$$(1 - \gamma) \cdot \mathbb{E}_{\text{init}} [J(\pi^*) - J(\pi_{\hat{\theta}})] \leq \mathcal{O}(R^{3/2} \cdot m^{-1/4}) + \mathbb{E}_{\text{init}} [\|\Pi_{\mathcal{F}_{R,\infty}} u_{\hat{\theta}} - u_{\hat{\theta}}\|_{\sigma_{\pi_{\hat{\theta}}}}],$$

which completes the proof of Theorem 4.9. \square

H.3 PROOF OF INEQUALITY (4.5)

Proof. Recall that we define ρ_i as

$$\rho_i = \eta^{-1} \cdot \left(\Pi_{\mathcal{B}}(\theta_i + \eta \cdot \nabla_{\theta} J(\pi_{\theta_i})) - \theta_i \right), \quad (\text{H.8})$$

where $\Pi_{\mathcal{B}}$ is the projection operator onto \mathcal{B} . Following from (H.8) and the fact that $(\Pi_{\mathcal{B}}y - y)^\top (\Pi_{\mathcal{B}}y - x) \leq 0$ for all $x \in \mathcal{B}$, we have

$$(\eta \cdot \rho_i - \eta \cdot \nabla_{\theta} J(\pi_{\theta_i}))^\top (\eta \cdot \rho_i + \theta_i - \theta) \leq 0, \quad \forall \theta \in \mathcal{B}. \quad (\text{H.9})$$

Thus, following from (H.9), we obtain that

$$\begin{aligned} \nabla_{\theta} J(\pi_{\theta_i})^\top (\theta - \theta_i) & \leq \rho_i^\top (\theta - \theta_i) - \eta \cdot \|\rho_i\|_2^2 + \eta \cdot \rho_i^\top \nabla_{\theta} J(\pi_{\theta_i}) \\ & \leq \|\rho_i\|_2 \cdot (\|\theta - \theta_i\|_2 + \eta \cdot \|\nabla_{\theta} J(\pi_{\theta_i})\|_2), \quad \forall \theta \in \mathcal{B}, \end{aligned} \quad (\text{H.10})$$

where the last inequality follows from the Cauchy-Schwartz inequality and the fact that $-\eta \cdot \|\rho_i\|_2^2 \leq 0$. It remains to upper bound the right-hand side of (H.10). For all $\theta, \theta_i \in \mathcal{B} = \{\alpha \in \mathbb{R}^{md} : \|\alpha - W_{\text{init}}\|_2 \leq R\}$, we obtain that $\|\theta - \theta_i\|_2 \leq 2R$. Meanwhile, recall that we set $\tau_i = 1$.

Therefore, following from Proposition 3.1, we obtain that

$$\|\nabla_{\theta} J(\pi_{\theta_i})\|_2 \leq \mathbb{E}_{\sigma_i} [|Q^{\pi_{\theta_i}}(s, a)| \cdot \|\bar{\phi}_{\theta_i}(s, a)\|_2] \leq 2Q_{\max}, \quad (\text{H.11})$$

where the first inequality follows from the Jensen's inequality, and the second inequality follows from the facts that $|Q^{\pi_{\theta_i}}(s, a)| \leq Q_{\max}$ and $\|\bar{\phi}_{\theta_i}(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. By plugging (H.11) and the upper bound $\|\theta - \theta_i\|_2 \leq 2R$ into (H.10), we conclude that

$$\nabla_{\theta} J(\pi_{\theta_i})^{\top} (\theta - \theta_i) \leq (2R + 2\eta \cdot Q_{\max}) \cdot \|\rho_i\|_2, \quad \forall \theta \in \mathcal{B},$$

which concludes the proof of (4.5). \square

H.4 PROOF OF COROLLARY A.5

Proof. It suffices to calculate $\bar{\epsilon}_i(T)$ defined in (A.3) in Theorem A.4. Note that $\tau_i = (i-1)/\sqrt{T}$ leads to the fact that $\tau_i = \mathcal{O}(\sqrt{T})$ for all $i \in [T]$. Thus, it holds for $m = \Omega(R^{10} \cdot T^6)$ that

$$\begin{aligned} \mathcal{O}(\tau_{i+1} \cdot T^{1/2} \cdot R^{3/2} \cdot m^{-1/4}) &= \mathcal{O}(R^{-1} \cdot T^{-1/2}), \quad \forall i \in [T], \\ \mathcal{O}(R^{5/4} \cdot m^{-1/8}) &= \mathcal{O}(T^{-1/2}). \end{aligned} \quad (\text{H.12})$$

Meanwhile, it holds for $B = \Omega(R^2 \cdot T^2 \cdot \sigma_{\xi}^2)$ that

$$R^{1/2} \cdot (\sigma_{\xi}^2/B)^{1/4} = \mathcal{O}(T^{-1/2}). \quad (\text{H.13})$$

Therefore, combining (H.12) and (H.13), we obtain that

$$\bar{\epsilon}_i(T) = \sqrt{8}c_0 \cdot R^{1/2} \cdot (\sigma_{\xi}^2/B)^{1/4} + \mathcal{O}((1 + \tau_{i+1} \cdot T^{1/2}) \cdot R^{3/2} \cdot m^{-1/4} + R^{5/4} \cdot m^{-1/8}) = \mathcal{O}(T^{-1/2}).$$

Following from Theorem A.4, we obtain that

$$\begin{aligned} \min_{i \in [T]} \mathbb{E} [J(\pi^*) - J(\pi_{\theta_i})] &= \frac{\log |\mathcal{A}| + 9R^2 + M}{(1 - \gamma) \cdot \sqrt{T}} + \mathcal{O}((1 - \gamma)^{-1} \cdot T^{1/2}) \\ &= \mathcal{O}\left(\frac{\log |\mathcal{A}|}{(1 - \gamma) \cdot \sqrt{T}}\right), \end{aligned} \quad (\text{H.14})$$

which concludes the proof of Corollary A.5. \square

H.5 PROOF OF LEMMA D.2

Proof. In the sequel, we write $g_i = \mathbb{E}[\widehat{\nabla}_{\theta} J(\pi_{\theta_i})]$ for notational simplicity, where $\widehat{\nabla}_{\theta} J(\pi_{\theta_i})$ is defined in (3.7), and the expectation is taken over σ_i given θ_i and ω_i . Recall that we set $\tau_i = 1$. By Proposition 3.1, we obtain that

$$\begin{aligned} |(\nabla_{\theta} J(\pi_{\theta_i}) - g_i)^{\top} \delta_i| &= \left| \mathbb{E}_{\sigma_i} \left[\bar{\phi}_{\theta_i}(s, a) \cdot (Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)) \right]^{\top} \delta_i \right| \\ &\leq \|\delta_i\|_2 \cdot \mathbb{E}_{\sigma_i} [\|\bar{\phi}_{\theta_i}(s, a)\|_2 \cdot |Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)|], \end{aligned} \quad (\text{H.15})$$

where $\bar{\phi}_{\theta_i}(s, a)$ is defined in (B.2) with $\theta = \theta_i$ and the inequality follows from the Jensen's inequality. Since both θ_i and θ_{i+1} are in set \mathcal{B} , we have $\|\delta_i\|_2 \leq 2R/\eta$. Meanwhile, note that $\|\bar{\phi}_{\theta_i}(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Therefore, it follows from Assumption 4.5 and (H.15) that

$$\begin{aligned} |(\nabla_{\theta} J(\pi_{\theta_i}) - g_i)^{\top} \delta_i| &\leq 2R/\eta \cdot \mathbb{E}_{\sigma_i} [|Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)|] \\ &\leq 2R/\eta \cdot \left(\mathbb{E}_{\sigma_i} [(d\sigma_i/d\varsigma_i)^2] \right)^{1/2} \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i} \\ &\leq 2\kappa \cdot R/\eta \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\varsigma_i}, \end{aligned}$$

where the second inequality follows from the Cauchy-Schwartz inequality, $d\sigma_i/d\varsigma_i$ is the Radon-Nikodym derivative, and κ is defined in Assumption 4.5. Thus, we complete the proof of Lemma D.2. \square

H.6 PROOF OF LEMMA D.3

Proof. In what follows, we write $g_i = \mathbb{E}[\widehat{\nabla} J(\pi_{\theta_i})]$ for notational simplicity, where the expectation is taken over σ_i given θ_i and ω_i . Note that

$$\mathbb{E}[\|\nabla_{\theta} J(\pi_{\theta_i}) - \widehat{\nabla}_{\theta} J(\pi_{\theta_i})\|_2^2] \leq 2\mathbb{E}[\|\xi_i\|_2^2] + 2\mathbb{E}[\|\nabla_{\theta} J(\pi_{\theta_i}) - g_i\|_2^2], \quad (\text{H.16})$$

where we use the fact that $\|x + y\|_2^2 \leq 2\|x\|_2^2 + 2\|y\|_2^2$, while the expectations are taken over all the randomness. Following from Proposition 3.1, we obtain that

$$\begin{aligned} \|\nabla_{\theta} J(\pi_{\theta_i}) - g_i\|_2 &= \left\| \mathbb{E}_{\sigma_i} \left[\bar{\phi}_{\theta_i}(s, a) \cdot (Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)) \right] \right\|_2 \\ &\leq \mathbb{E}_{\sigma_i} [\|\bar{\phi}_{\theta_i}(s, a)\|_2 \cdot |Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)|], \end{aligned} \quad (\text{H.17})$$

where $\bar{\phi}_{\theta_i}(s, a)$ is defined in (B.2) with $\theta = \theta_i$ and the second inequality follows from the Jensen's inequality. Since $\|\bar{\phi}_{\theta_i}(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we obtain from (H.17) that

$$\begin{aligned} \|\nabla_{\theta} J(\pi_{\theta_i}) - g_i\|_2^2 &\leq \left(\mathbb{E}_{\sigma_i} [\|\bar{\phi}_{\theta_i}(s, a)\|_2 \cdot |Q^{\pi_{\theta_i}}(s, a) - Q_{\omega_i}(s, a)|] \right)^2 \\ &\leq 4\kappa^2 \cdot \|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2, \end{aligned} \quad (\text{H.18})$$

where κ is defined in Assumption 4.5 and the inequality follows from the Cauchy-Schwartz inequality. Combining (H.16) and (H.18), we obtain that

$$\mathbb{E}[\|\nabla_{\theta} J(\pi_{\theta_i}) - \widehat{\nabla}_{\theta} J(\pi_{\theta_i})\|_2^2] \leq 2\mathbb{E}[\|\xi_i\|_2^2] + 8\kappa^2 \cdot \mathbb{E}[\|Q^{\pi_{\theta_i}} - Q_{\omega_i}\|_{\zeta_i}^2],$$

which concludes the proof of Lemma D.3. \square

H.7 PROOF OF LEMMA D.4

Proof. Following from the definition of the KL divergence, we have

$$D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_i(\cdot | s)) - D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_{i+1}(\cdot | s)) = \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi^*(\cdot | s) \rangle. \quad (\text{H.19})$$

Meanwhile, the right-hand side of (H.19) can be expanded as follows,

$$\begin{aligned} &\langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi^*(\cdot | s) \rangle \\ &= \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi^*(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle + \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_{i+1}(\cdot | s) \rangle \\ &= \underbrace{\langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi^*(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle}_{L_i} + D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)). \end{aligned} \quad (\text{H.20})$$

Combining (H.19) and (H.20), we obtain that

$$L_i = D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_i(\cdot | s)) - D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_{i+1}(\cdot | s)) - D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)). \quad (\text{H.21})$$

In what follows, we compute the difference

$$\mathbb{E}_{\nu_*}[L_i] - (1 - \gamma) \cdot \eta \cdot (J(\pi^*) - J(\pi_i)).$$

Following from Lemma D.1, it holds that

$$J(\pi^*) - J(\pi_i) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{\nu_*} [\langle Q^{\pi_i}(s, \cdot), \pi^*(s, \cdot) - \pi_i(s, \cdot) \rangle]. \quad (\text{H.22})$$

Meanwhile, for L_i defined in (H.20), we obtain that

$$\begin{aligned} &L_i - \eta \cdot \langle Q^{\pi_i}(s, \cdot), \pi^*(s, \cdot) - \pi_i(s, \cdot) \rangle \\ &= \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_*(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle - \eta \cdot \langle Q^{\pi_i}(s, \cdot), \pi^*(s, \cdot) - \pi_i(s, \cdot) \rangle \\ &= \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle \\ &\quad + \eta \cdot \langle Q_{\omega_i}(s, \cdot) - Q^{\pi_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle + \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle. \end{aligned} \quad (\text{H.23})$$

Note that upon taking the expectation over $s \sim \nu_*(\cdot)$ in (H.23), the right-hand side of (H.23) equals to H_i defined in (D.21). Thus, combining (H.22) and (H.23), we obtain that

$$\mathbb{E}_{\nu_*}[L_i] - (1 - \gamma) \cdot \eta \cdot (J(\pi^*) - J(\pi_i)) = H_i, \quad (\text{H.24})$$

where H_i is defined in (D.21). By plugging (H.21) into (H.24), we conclude that

$$(1 - \gamma) \cdot \eta \cdot (J(\pi^*) - J(\pi_i)) = \mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_i(\cdot | s)) - D_{\text{KL}}(\pi^*(\cdot | s) \| \pi_{i+1}(\cdot | s)) \right. \\ \left. - D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] - H_i,$$

which concludes the proof of Lemma D.4. \square

H.8 PROOF OF LEMMA D.5

Proof. By (D.21), we have

$$\mathbb{E}[|H_i|] \leq \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle \right| \right] \right. \\ \left. + \eta \cdot \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle Q_{\omega_i}(s, \cdot) - Q^{\pi_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle \right| \right] \right] \right. \\ \left. + \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle \log(\pi_i(\cdot | s) / \pi_{i+1}(\cdot | s)), \pi_{i+1}(\cdot | s) - \pi_i(\cdot | s) \rangle \right| \right] \right], \quad (\text{H.25})$$

where the inequality follows from the Jensen's inequality, and the expectations are taken over all the randomness. To prove Lemma D.5, we establish upper bounds of the three terms in the left-hand side of (H.25) separately in the following lemmas.

Lemma H.1. It holds that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle Q_{\omega_i}(s, \cdot) - Q^{\pi_i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle \right| \right] \right] \leq (\phi'_i + \psi'_i) \cdot \mathbb{E}[\|Q_{\omega_i} - Q^{\pi_i}\|_{\mathfrak{C}_i}],$$

where ϕ'_i and ψ'_i are the concentrability coefficients defined in (A.1) of Assumption A.2. Here the expectations are taken over all the randomness.

Proof. See §I.1 for a detailed proof. \square

Lemma H.2. Under Assumptions 4.2 and A.3, it holds that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle \right| \right] \right] \\ \leq \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] \right] + \eta^2 \cdot (9R^2 + M^2) + \mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4}),$$

where M is the absolute constant defined in Assumption A.3, and ϕ_0 is the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. Here the expectations are taken over all the randomness.

Proof. See §I.2 for a detailed proof. \square

Lemma H.3. Under Assumption 4.2, it holds that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle \right| \right] \right] \\ \leq \sqrt{2}(\varphi_i + \psi_i) \cdot \eta \cdot R^{1/2} \cdot \tau_i^{-1} \cdot \left(\mathbb{E}[\|\xi_i(\delta_i)\|_2] + \mathbb{E}[\|\xi_i(\omega_i)\|_2] \right)^{1/2} \\ + \mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4} + \eta \cdot R^{5/4} \cdot m^{-1/8}),$$

where φ_i and ψ_i are the concentrability coefficients defined in (A.1) of Assumption A.2 and $\xi_i(\delta_i)$, $\xi_i(\omega_i)$ are defined in Assumption A.1. Here the expectations are taken over all the randomness.

Proof. See §I.3 for a detailed proof. \square

Finally, applying Lemmas H.1, H.2, and H.3 to (H.25), it holds under Assumptions 4.2 and A.3 that

$$\mathbb{E} \left[|H_i| - \mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] \right] \leq \eta^2 \cdot (6R^2 + M^2) + \eta \cdot (\varphi'_i + \psi'_i) \cdot \varepsilon_{Q,i} + \varepsilon_i, \quad (\text{H.26})$$

where

$$\varepsilon_i = \left\{ \mathbb{E} [\|\xi_i(\delta_i)\|_2 + \|\xi_i(\omega_i)\|_2] \right\}^{1/2} + \mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4} + \eta \cdot R^{5/4} \cdot m^{-1/8}), \quad \varepsilon_{Q,i} = \mathbb{E} [\|Q^{\pi_i} - Q_{\omega_i}\|_{\zeta_i}].$$

Here the expectations are taken over all the randomness. Therefore, we complete the proof of Lemma D.5. \square

I PROOF OF SUPPORTING LEMMAS

In this section, we provide the proofs of the lemmas in §H.

I.1 PROOF OF LEMMA H.1

Proof. We define $\Delta_{Q,i}(s, a) = Q_{\omega_i}(s, a) - Q^{\pi_i}(s, a)$. It holds that

$$\mathbb{E}_{\nu^*} [\langle \Delta_{Q,i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle] = \int_{\mathcal{S}} \left| \sum_{a \in \mathcal{A}} \Delta_{Q,i}(s, a) \cdot [\pi^*(a | s) - \pi_i(a | s)] \right| d\nu_*(s). \quad (\text{I.1})$$

Meanwhile, it holds for any $s \in \mathcal{S}$ that

$$\begin{aligned} & \left| \sum_{a \in \mathcal{A}} \Delta_{Q,i}(s, a) \cdot [\pi^*(a | s) - \pi_i(a | s)] \right| \\ &= \left| \int_{a \in \mathcal{A}} \Delta_{Q,i}(s, a) \cdot [\pi^*(a | s) - \pi_i(a | s)] / \pi_i(a | s) d\pi_i(a | s) \right| \\ &\leq \int_{a \in \mathcal{A}} |\Delta_{Q,i}(s, a) \cdot [\pi^*(a | s) - \pi_i(a | s)] / \pi_i(a | s)| d\pi_i(a | s). \end{aligned} \quad (\text{I.2})$$

Combining (I.1) and (I.2), we obtain that

$$\begin{aligned} & \mathbb{E}_{\nu^*} [\langle \Delta_{Q,i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle] \\ &\leq \int_{\mathcal{S} \times \mathcal{A}} |\Delta_{Q,i}(s, a) \cdot [\pi^*(a | s) - \pi_i(a | s)] / \pi_i(a | s)| d\tilde{\sigma}(s, a), \end{aligned} \quad (\text{I.3})$$

where we define $\tilde{\sigma}(\cdot, \cdot) = \pi_i(\cdot | \cdot) \cdot \nu_*(\cdot)$. Recall that $\zeta_i(\cdot, \cdot) = \pi_i(\cdot | \cdot) \cdot \varrho_i(\cdot)$ and $\sigma_*(\cdot, \cdot) = \pi^*(\cdot | \cdot) \cdot \nu_*(\cdot)$. Therefore, following from (I.3), it holds that

$$\begin{aligned} & \mathbb{E}_{\nu^*} [\langle \Delta_{Q,i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle] \\ &\leq \int_{\mathcal{S} \times \mathcal{A}} |\Delta_{Q,i}(s, a)| d\sigma_* + \int_{\mathcal{S} \times \mathcal{A}} |\Delta_{Q,i}(s, a)| \cdot \frac{d\nu^*}{d\varrho_i}(s) d\zeta_i(s, a). \end{aligned} \quad (\text{I.4})$$

Finally, applying the Cauchy-Schwartz inequality to (I.4) yields that

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E}_{\nu^*} [\langle \Delta_{Q,i}(s, \cdot), \pi^*(\cdot | s) - \pi_i(\cdot | s) \rangle] \right] \\ &\leq \left(\left\{ \mathbb{E}_{\zeta_i} [(d\sigma_*/d\zeta_i)^2] \right\}^{1/2} + \left\{ \mathbb{E}_{\varrho_i} [(d\nu^*/d\varrho_i)^2] \right\}^{1/2} \right) \cdot \mathbb{E} \left[\left\{ \mathbb{E}_{\zeta_i} [|\Delta_{Q,i}(s, a)|^2] \right\}^{1/2} \right] \\ &= (\varphi'_i + \psi'_i) \cdot \mathbb{E} \left[\left\{ \mathbb{E}_{\zeta_i} [|\Delta_{Q,i}(s, a)|^2] \right\}^{1/2} \right] = (\varphi'_i + \psi'_i) \cdot \mathbb{E} [\|\Delta_{Q,i}\|_{\zeta_i}], \end{aligned}$$

where $d\sigma_*/d\zeta_i$ and $d\nu^*/d\varrho_i$ are the Radon-Nikodym derivatives, φ'_i and ψ'_i are the concentrability coefficients defined in Assumption A.2, and the expectations are taken over all the randomness. Thus, we complete the proof of Lemma H.1. \square

I.2 PROOF OF LEMMA H.2

Proof. In what follows, we denote the two-layer neural network $f((s, a); \theta)$ defined in (3.1) by $f_\theta(s, a)$ for notational simplicity. Following from the parameterization of π_θ in (3.2), we obtain that

$$\begin{aligned} & \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle \\ &= \langle \tau_{i+1} \cdot f_{\theta_{i+1}}(s, \cdot) - \tau_i \cdot f_{\theta_i}(s, \cdot), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle - \langle C_i(s), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle, \end{aligned} \quad (\text{I.5})$$

where we define $C_i(s)$ by

$$C_i(s) = \log \left(\sum_{a \in \mathcal{A}} \exp(\tau_i \cdot f_{\theta_i}(s, a)) \right) - \log \left(\sum_{a \in \mathcal{A}} \exp(\tau_{i+1} \cdot f_{\theta_{i+1}}(s, a)) \right).$$

Note that both $\pi_i(\cdot | s)$ and $\pi_{i+1}(\cdot | s)$ are distributions over \mathcal{A} , which implies that

$$\langle C_i(s), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle = C_i(s) - C_i(s) = 0, \quad \forall s \in \mathcal{S}. \quad (\text{I.6})$$

Meanwhile, recall that we define the feature mapping $\phi_\theta(s, a)$ in (3.3). For the two-layer neural network $f_\theta(s, a) = f((s, a); \theta)$, it holds that

$$f_\theta(s, a) = \phi_\theta(s, a)^\top \theta, \quad (\text{I.7})$$

In what follows, we write $\phi_i(s, a) = \phi_{\theta_i}(s, a)$ and $\Delta_i(a | s) = \pi_i(a | s) - \pi_{i+1}(a | s)$ for notational simplicity. By plugging (I.7) and (I.6) into (I.5), we obtain that

$$\begin{aligned} & \left| \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \Delta_i(\cdot | s) \rangle \right| = \left| \langle \tau_{i+1} \cdot \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \tau_i \cdot \phi_i(s, \cdot)^\top \theta_i, \Delta_i(\cdot | s) \rangle \right| \\ & \hspace{20em} (\text{I.8}) \\ & \leq \left| \langle \phi_i(s, \cdot)^\top (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i), \Delta_i(\cdot | s) \rangle \right| + \tau_{i+1} \cdot \left| \langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_i, \Delta_i(\cdot | s) \rangle \right| \\ & \leq \underbrace{\|\phi_i(s, \cdot)^\top (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i)\|_{\infty, \mathcal{A}} \cdot \|\Delta_i(\cdot | s)\|_{1, \mathcal{A}}}_{(\text{i})} + \underbrace{\tau_{i+1} \cdot \left| \langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_i, \Delta_i(\cdot | s) \rangle \right|}_{(\text{ii})}, \end{aligned}$$

where the last inequality follows from the Hölder's inequality. Here we denote by $\|\cdot\|_{\infty, \mathcal{A}}$ and $\|\cdot\|_{1, \mathcal{A}}$ the vector ℓ_∞ - and ℓ_1 -norms defined on $\mathbb{R}^{|\mathcal{A}|}$, respectively. In what follows, we upper bound terms (i) and (ii) on the right-hand side of (I.8) separately.

Upper Bounding (i) in (I.8). Recall that we define

$$\delta_i = \eta^{-1} \cdot (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i) = \underset{\omega \in \mathcal{B}}{\operatorname{argmin}} \|\widehat{F}(\theta_i) \cdot \omega - \tau_i \cdot \widehat{\nabla} J(\pi_{\theta_i})\|_2.$$

Thus, it holds that $\delta_i \in \mathcal{B}$ and $\|\delta_i - W_{\text{init}}\|_2 \leq R$, where W_{init} is the initial parameter. In what follows, we denote by $\phi_0(s, a)$ the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$. Then for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned} & |\phi_i(s, a)^\top (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i)| = \eta \cdot |\phi_i(s, a)^\top \delta_i| \\ & \leq \eta \cdot (|\phi_0(s, a)^\top W_{\text{init}}| + |\phi_i(s, a)^\top \delta_i - \phi_i(s, a)^\top \theta_i| + |\phi_i(s, a)^\top \theta_i - \phi_0(s, a)^\top W_{\text{init}}|) \\ & \leq \eta \cdot (M_0 + \|\phi_i(s, a)\|_2 \cdot \|\delta_i - \theta_i\|_2 + |\phi_i(s, a)^\top \theta_i - \phi_0(s, a)^\top W_{\text{init}}|), \quad (\text{I.9}) \end{aligned}$$

where the first inequality follows from the triangle inequality, the second inequality follows from the Cauchy-Schwartz inequality, and M_0 is defined as

$$M_0 = \sup_{(s, a) \in \mathcal{S} \times \mathcal{A}} |\phi_0(s, a)^\top W_{\text{init}}|. \quad (\text{I.10})$$

In what follows, we upper bound the right-hand side of (I.9). Note that $\tau_{i-1} + \eta = \tau_i$. Therefore, we obtain that

$$\|\theta_i - W_{\text{init}}\|_2 \leq \tau_{i-1} / \tau_i \cdot \|\theta_{i-1} - W_{\text{init}}\|_2 + \eta / \tau_i \cdot \|\delta_{i-1} - W_{\text{init}}\|_2, \quad (\text{I.11})$$

which holds for all $i > 1$. Recursively, since $\theta_1 = W_{\text{init}} \in \mathcal{B}$ and $\delta_i \in \mathcal{B}$ for all $i \in [T]$, it then follows from (I.11) that $\theta_i \in \mathcal{B}$ for all $i \in [T]$. Thus it holds that $\|\delta_i - \theta_i\|_2 \leq 2R$. Meanwhile, following from (3.3), it holds for all $\theta \in \mathbb{R}^{md}$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$ that $\|\phi_\theta(s, a)\|_2 \leq 1$. Therefore, we obtain that

$$\|\phi_i(s, a)\|_2 \cdot \|\delta_i - \theta_i\|_2 \leq 2R, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (\text{I.12})$$

It remains to upper bound $|\phi_i(s, a)^\top \theta_i - \phi_0(s, a)^\top W_{\text{init}}|$, which is equal to $|f_\theta(s, a) - f_{W_{\text{init}}}(s, a)|$ by (I.7). Recall that $f_\theta(s, a)$ is differentiable with respect to $\theta \in \mathbb{R}^{md}$ almost everywhere, and the gradient $\nabla_\theta f_\theta = ([\nabla_\theta f_\theta]_1^\top, \dots, [\nabla_\theta f_\theta]_m^\top)^\top$ is given by

$$[\nabla_\theta f_\theta]_r(s, a) = \frac{b_r}{\sqrt{m}} \cdot \mathbf{1}\{(s, a)^\top [\theta]_r > 0\} \cdot (s, a) = [\phi_\theta]_r(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where $\phi_\theta(s, a)$ is defined in (3.3). Since $\|\phi_\theta(s, a)\|_2 \leq 1$ for all $\theta \in \mathbb{R}^{md}$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, we obtain for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned} |\phi_i(s, a)^\top \theta_i - \phi_0(s, a)^\top W_{\text{init}}| &= |f_{\theta_i}(s, a) - f_{W_{\text{init}}}(s, a)| \leq \sup_{\theta \in \mathbb{R}^{md}} \|\nabla_\theta f_\theta(s, a)\|_2 \cdot \|\theta_i - W_{\text{init}}\|_2 \\ &\leq \sup_{\theta \in \mathbb{R}^{md}} \|\phi_\theta(s, a)\|_2 \cdot \|\theta_i - W_{\text{init}}\|_2 \leq R, \end{aligned} \quad (\text{I.13})$$

where the last inequality holds since $\theta_i \in \mathcal{B}$.

By plugging (I.12) and (I.13) into (I.9), we have

$$|\tau_{i+1} \cdot \phi_i(s, a)^\top \theta_{i+1} - \tau_i \cdot \phi_i(s, a)^\top \theta_i| \leq \eta \cdot (M_0 + 3R), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where M_0 is defined in (I.10). Therefore, it holds for all $s \in \mathcal{S}$ that

$$\begin{aligned} \|\tau_{i+1} \cdot \phi_i(s, \cdot)^\top \theta_{i+1} - \tau_i \cdot \phi_i(s, \cdot)^\top \theta_i\|_{\infty, \mathcal{A}} &= \sup_{a \in \mathcal{A}} |\tau_{i+1} \cdot \phi_i(s, a)^\top \theta_{i+1} - \tau_i \cdot \phi_i(s, a)^\top \theta_i| \\ &\leq \eta \cdot (M_0 + 3R). \end{aligned} \quad (\text{I.14})$$

Finally, by Pinsker's inequality, it follows from (I.14) that

$$\begin{aligned} \|\phi_i(s, \cdot)^\top (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i)\|_{\infty, \mathcal{A}} \cdot \|\Delta_i(\cdot | s)\|_{1, \mathcal{A}} - D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \\ \leq \eta \cdot (M_0 + 3R) \cdot \|\pi_{i+1}(\cdot | s) - \pi_i(\cdot | s)\|_{1, \mathcal{A}} - 1/2 \cdot \|\pi_{i+1}(\cdot | s) - \pi_i(\cdot | s)\|_{1, \mathcal{A}}^2. \end{aligned} \quad (\text{I.15})$$

By completing the squares, we further bound (I.15) by

$$\begin{aligned} \|\phi_i(s, \cdot)^\top (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i)\|_{\infty, \mathcal{A}} \cdot \|\Delta_i(\cdot | s)\|_{1, \mathcal{A}} - D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \\ = -1/2 \cdot [\|\pi_{i+1}(\cdot | s) - \pi_i(\cdot | s)\|_{1, \mathcal{A}} - \eta \cdot (M_0 + 3R)]^2 + 1/2 \cdot \eta^2 \cdot (M_0 + 3R)^2 \\ \leq 1/2 \cdot \eta^2 \cdot (M_0 + 3R)^2 \leq \eta^2 \cdot (M_0^2 + 9R^2), \end{aligned} \quad (\text{I.16})$$

which holds for all $s \in \mathcal{S}$. Here the last inequality follows from the fact that $(x + y)^2 \leq 2x^2 + 2y^2$.

Upper Bounding (ii) in (I.8). It holds for all $s \in \mathcal{S}$ that

$$\begin{aligned} &|\langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}, \Delta_i(\cdot | s) \rangle| \quad (\text{I.17}) \\ &\leq |\langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}, \pi_i(\cdot | s) \rangle| + |\langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}, \pi_{i+1}(\cdot | s) \rangle| \\ &\leq \|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1} + \|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}\|_{\pi_{i+1}, 1}, \end{aligned}$$

where for any distribution $\pi \in \mathcal{P}(\mathcal{A})$, we denote by $\|\cdot\|_{\pi, p}$ the $L_p(\pi)$ -norm, i.e. $\|v\|_{\pi, p} = [\sum_{a \in \mathcal{A}} \pi(a) \cdot |v(a)|^p]^{1/p}$. Following from Assumption 4.2 and Lemma E.2, it holds that

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1}] \right] \\ &\leq \mathbb{E} [\|\phi_{i+1}(\cdot, \cdot)^\top \theta_{i+1} - \phi_0(\cdot, \cdot)^\top \theta_{i+1}\|_{\pi_i \cdot \nu_*}] = O(R^{3/2} \cdot m^{-1/4}), \\ &\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_i(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1}] \right] \\ &\leq \mathbb{E} [\|\phi_i(\cdot, \cdot)^\top \theta_{i+1} - \phi_0(\cdot, \cdot)^\top \theta_{i+1}\|_{\pi_i \cdot \nu_*}] = O(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{I.18})$$

where the inequalities follow from the Cauchy-Schwartz inequality, and the expectations are taken over all the randomness. Meanwhile, it holds that

$$\begin{aligned} &\|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1} \quad (\text{I.19}) \\ &\leq \|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1} + \|\phi_i(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1}. \end{aligned}$$

Combining (I.18) and (I.19), we obtain that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}\|_{1, \pi_i}] \right] = O(R^{3/2} \cdot m^{-1/4}). \quad (\text{I.20})$$

Similarly, it holds that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}\|_{1, \pi_{i+1}}] \right] = O(R^{3/2} \cdot m^{-1/4}), \quad (\text{I.21})$$

where the expectation is taken over all the randomness. By plugging (I.20) and (I.21) into (I.17), we obtain that

$$\tau_{i+1} \cdot \mathbb{E} \left[\mathbb{E}_{\nu_*} [\langle \phi_{i+1}(s, \cdot)^\top \theta_{i+1} - \phi_i(s, \cdot)^\top \theta_{i+1}, \Delta_i(\cdot | s) \rangle] \right] = O(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4}). \quad (\text{I.22})$$

Finally, by plugging (I.16) and (I.22) into (I.8), it holds under Assumptions 4.2 and A.3 that

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[\left| \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)), \pi_i(\cdot | s) - \pi_{i+1}(\cdot | s) \rangle \right| \right] \right] \\ & \leq \mathbb{E} \left[\mathbb{E}_{\nu_*} \left[D_{\text{KL}}(\pi_{i+1}(\cdot | s) \| \pi_i(\cdot | s)) \right] \right] + \eta^2 \cdot (9R^2 + M^2) + O(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4}), \end{aligned}$$

where M is defined in Assumption A.3. Thus, we complete the proof of Lemma H.2. \square

I.3 PROOF OF LEMMA H.3

Proof. Note that $\mathbb{E}_{\pi_{\theta_i}}[\phi_{\theta_i}(s, a')]$ and $\mathbb{E}_{\pi_{\omega_i}}[\phi_{\omega_i}(s, a')]$ depend solely on $s \in \mathcal{S}$, where we write $\mathbb{E}_{\pi_{\theta_i}}[\phi_{\theta_i}(s, a')] = \mathbb{E}_{a' \sim \pi_{\theta_i}(\cdot | s)}[\phi_{\theta_i}(s, a')]$ for notational simplicity. Thus, we have

$$\langle \mathbb{E}_{\pi_{\theta_i}}[\phi_{\theta_i}(s, a')^\top \delta_i - \phi_{\omega_i}(s, a')^\top \omega_i], \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle = 0, \quad \forall s \in \mathcal{S}. \quad (\text{I.23})$$

Meanwhile, following from the parameterization of π_θ in (3.2) and (I.6) in §I.2, we obtain that

$$\begin{aligned} & \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle \\ & = \langle \tau_{i+1} \cdot \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \tau_i \cdot \phi_{\theta_i}(s, \cdot)^\top \theta_i - \eta \cdot \phi_{\omega_i}(s, \cdot)^\top \omega_i, \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle. \end{aligned} \quad (\text{I.24})$$

In what follows, we define $\Delta_i^*(a | s) = \pi_*(a | s) - \pi_i(a | s)$ for notational simplicity. Then, combining (I.23), (I.24), and (I.26), we obtain that

$$\begin{aligned} & \langle \log(\pi_{i+1}(\cdot | s) / \pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \Delta_i^*(\cdot | s) \rangle \\ & = \eta \cdot \langle \phi_{\theta_i}(s, \cdot)^\top \delta_i - \phi_{\omega_i}(s, \cdot)^\top \omega_i, \Delta_i^*(\cdot | s) \rangle + \tau_{i+1} \cdot \langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_i, \Delta_i^*(\cdot | s) \rangle \\ & = \underbrace{\eta \cdot \langle \bar{\phi}_{\theta_i}(s, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(s, \cdot)^\top \omega_i, \Delta_i^*(\cdot | s) \rangle}_{(\text{iii})} + \underbrace{\tau_{i+1} \cdot \langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_i, \Delta_i^*(\cdot | s) \rangle}_{(\text{iv})}, \end{aligned} \quad (\text{I.25})$$

where $\bar{\phi}_{\theta_i}$ and $\bar{\phi}_{\omega_i}$ are the centered feature mappings defined in (B.2) and δ_i is defined as

$$\delta_i = \eta^{-1} \cdot (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i) = \underset{\omega \in \mathcal{B}}{\text{argmin}} \|\widehat{F}(\theta_i)\omega - \tau_i \cdot \widehat{\nabla} J(\pi_{\theta_i})\|_2. \quad (\text{I.26})$$

In what follows, we upper bound the expectation of (iii) and (iv) over the visitation measure ν_* separately.

Upper Bounding (iii) in (I.25). It holds that

$$\begin{aligned} & \mathbb{E}_{\nu_*} \left[\left| \langle \bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i, \pi_*(a | s) \rangle \right| \right] \\ & \leq \int_{\mathcal{S} \times \mathcal{A}} |\bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i| d\sigma_*(s, a) \\ & = \int_{\mathcal{S} \times \mathcal{A}} |\bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i| \cdot \frac{d\sigma_*(s, a)}{d\sigma_i(s, a)} d\sigma_i(s, a) \\ & \leq \varphi_i \cdot \|\bar{\phi}_{\theta_i}(\cdot, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(\cdot, \cdot)^\top \omega_i\|_{\sigma_i}, \end{aligned} \quad (\text{I.27})$$

where $d\sigma_*/d\sigma_i$ is the Radon-Nikodym derivative, φ_i is defined in (A.1) of Assumption A.2, and the last inequality follows from the Cauchy-Schwartz inequality. Similarly, it holds that

$$\begin{aligned} & \mathbb{E}_{\nu_*} \left[\left| \langle \bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i, \pi_i(a | s) \rangle \right| \right] \\ & \leq \int_{\mathcal{S} \times \mathcal{A}} |\bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i| d\pi_i(a | s) \cdot \nu_*(s) \\ & = \int_{\mathcal{S} \times \mathcal{A}} |\bar{\phi}_{\theta_i}(s, a)^\top \delta_i - \bar{\phi}_{\omega_i}(s, a)^\top \omega_i| \cdot \frac{d\nu_*(s)}{d\nu_i(s)} d\sigma_i(s, a) \\ & \leq \psi_i \cdot \|\bar{\phi}_{\theta_i}(\cdot, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(\cdot, \cdot)^\top \omega_i\|_{\sigma_i}, \end{aligned} \quad (\text{I.28})$$

where $d\nu_*/d\nu_i$ is the Radon-Nikodym derivative, ψ_i is defined in (A.1) of Assumption A.2, and the last inequality follows from the Cauchy-Schwartz inequality. Combining (I.27) and (I.28), we

obtain that

$$\mathbb{E}_{\nu_*} [| \langle \bar{\phi}_{\theta_i}(s, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(s, \cdot)^\top \omega_i, \Delta_i^*(\cdot | s) \rangle |] \leq (\varphi_i + \psi_i) \cdot \| \bar{\phi}_{\theta_i}(\cdot, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(\cdot, \cdot)^\top \omega_i \|_{\sigma_i}. \quad (\text{I.29})$$

It now suffices to upper bound the norm $\| \bar{\phi}_{\theta_i}(\cdot, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(\cdot, \cdot)^\top \omega_i \|_{\sigma_i}$. With a slight abuse of notation, we write $\bar{\phi}_{\theta_i} = \bar{\phi}_{\theta_i}(\cdot, \cdot)$ and $\bar{\phi}_{\omega_i} = \bar{\phi}_{\omega_i}(\cdot, \cdot)$ hereafter for notational simplicity. Note that

$$\begin{aligned} \| \bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i \|_{\sigma_i} &= \sqrt{ \mathbb{E}_{\sigma_i} [(\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i) (\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)] } \\ &\leq \underbrace{ \sqrt{ | (\delta_i - \omega_i)^\top \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} \cdot (\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)] | } }_{(\text{iii.a})} + \underbrace{ \sqrt{ \mathbb{E}_{\sigma_i} [(\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}) (\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i})] } }_{(\text{iii.b})}. \end{aligned} \quad (\text{I.30})$$

We now upper bound the expectations of the right-hand side of (I.30) over all the randomness.

Upper Bounding (iii.a) in (I.30). Note that $\omega_i \in \mathcal{B}$ and $\delta_i \in \mathcal{B}$, where δ_i is defined in (I.26) and $\mathcal{B} = \{ \alpha \in \mathbb{R}^{md} : \| \alpha - W_{\text{init}} \|_2 \leq R \}$. Therefore, we obtain that

$$\| \omega_i - \delta_i \|_2 \leq 2R. \quad (\text{I.31})$$

Meanwhile, following from Proposition 3.1 and (3.7), it holds that

$$\mathbb{E}_{\sigma_i} [\widehat{F}(\theta_i)] = F(\theta_i) = \tau_i^2 \cdot \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} \bar{\phi}_{\theta_i}^\top], \quad \mathbb{E}_{\sigma_i} [\widehat{\nabla}_\theta J(\pi_{\theta_i})] = \tau_i \cdot \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} \bar{\phi}_{\omega_i}^\top \omega_i], \quad (\text{I.32})$$

where \mathbb{E}_{σ_i} is taken over σ_i given θ_i . In what follows, we write $g_i = \mathbb{E}_{\sigma_i} [\widehat{\nabla}_\theta J(\pi_{\theta_i})]$ for notational simplicity, where \mathbb{E}_{σ_i} is taken over σ_i given θ_i . By plugging (I.31) and (I.32) into (iii.a) in (I.30), we obtain that

$$\begin{aligned} | (\delta_i - \omega_i)^\top \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} \cdot (\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)] | &= \tau_i^{-2} \cdot | (\delta_i - \omega_i)^\top (F(\theta_i) \cdot \delta_i - \tau_i \cdot g_i) | \\ &\leq 2R \cdot \tau_i^{-2} \cdot \| F(\theta_i) \cdot \delta_i - \tau_i \cdot g_i \|_2, \end{aligned} \quad (\text{I.33})$$

where the last inequality follows from the Cauchy-Schwartz inequality and (I.31). Following from (I.33), we obtain that

$$\begin{aligned} \mathbb{E} \left[\left(| (\delta_i - \omega_i)^\top \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} (\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)] | \right)^{1/2} \right] &\leq C_i \cdot \mathbb{E} \left[\left(\| F(\theta_i) \cdot \delta_i - \tau_i \cdot g_i \|_2 \right)^{1/2} \right] \\ &\leq C_i \cdot \mathbb{E} \left[\left(\| \widehat{F}(\theta_i) \cdot \delta_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2 + \| \xi_i(\delta_i) \|_2 \right)^{1/2} \right] \\ &\leq C_i \cdot \left\{ \mathbb{E} [\| \widehat{F}(\theta_i) \cdot \delta_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2] + \mathbb{E} [\| \xi_i(\delta_i) \|_2] \right\}^{1/2}, \end{aligned} \quad (\text{I.34})$$

where the expectations are taken over all the randomness. Here the last inequality follows from the Jensen's inequality, $C_i = \sqrt{2R} \cdot \tau_i^{-1}$, and $\xi_i(\delta_i)$ is defined as

$$\xi_i(\delta_i) = \widehat{F}(\theta_i) \cdot \delta_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) - (F(\theta_i) \cdot \delta_i - \tau_i \cdot g_i). \quad (\text{I.35})$$

In what follows, we upper bound the term $\| \widehat{F}(\theta_i) \cdot \delta_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2$ in the right-hand side of (I.34). Recall that we define δ_i as

$$\delta_i = \eta^{-1} \cdot (\tau_{i+1} \cdot \theta_{i+1} - \tau_i \cdot \theta_i) = \underset{\omega \in \mathcal{B}}{\operatorname{argmin}} \| \widehat{F}(\theta_i) \cdot \omega - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2. \quad (\text{I.36})$$

Therefore, since $\omega_i \in \mathcal{B}$, we obtain from (I.36) that

$$\begin{aligned} \| \widehat{F}(\theta_i) \cdot \delta_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2 &\leq \| \widehat{F}(\theta_i) \cdot \omega_i - \tau_i \cdot \widehat{\nabla}_\theta J(\pi_{\theta_i}) \|_2 \\ &\leq \| F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i \|_2 + \| \xi_i(\omega_i) \|_2, \end{aligned} \quad (\text{I.37})$$

where recall that, similar to (I.35), we define the error term $\xi_i(\omega_i)$ as follows,

$$\xi_i(\omega_i) = \widehat{F}(\theta_i) \cdot \omega_i - \tau_i \cdot \widehat{J}(\pi_{\theta_i}) - (F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i). \quad (\text{I.38})$$

By plugging (I.37) into (I.34), we obtain that

$$\begin{aligned} \mathbb{E} \left[\left(| (\delta_i - \omega_i)^\top \mathbb{E}_{\sigma_i} [\bar{\phi}_{\theta_i} (\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)] | \right)^{1/2} \right] \\ \leq C_i \cdot \left(\mathbb{E} [\| F(\theta_i) \cdot \omega_i - \tau_i \cdot \nabla_\theta J(\pi_{\theta_i}) \|_2] + \mathbb{E} [\| \xi_i(\delta_i) \|_2] + \mathbb{E} [\| \xi_i(\omega_i) \|_2] \right)^{1/2}, \end{aligned} \quad (\text{I.39})$$

where $C_i = \sqrt{2R} \cdot \tau_i^{-1}$ and $\xi_i(\delta_i)$, $\xi_i(\omega_i)$ are defined in (I.35) and (I.38), respectively. To upper bound the right-hand side of (I.39), it now suffices to upper bound the expectation $\mathbb{E}[\|F(\theta_i) \cdot \delta_i - \tau_i \cdot \nabla_{\theta} J(\pi_{\theta_i})\|_2]$. By (I.32), we obtain that

$$\begin{aligned} \|F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i\|_2 &= \tau_i^2 \cdot \|\mathbb{E}_{\sigma_i}[\bar{\phi}_{\theta_i} \cdot (\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i}^c)^\top \omega_i]\|_2 \\ &\leq \tau_i^2 \cdot \mathbb{E}_{\sigma_i}[\|\bar{\phi}_{\theta_i} \cdot (\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i})^\top \omega_i\|_2] = \tau_i^2 \cdot \mathbb{E}_{\sigma_i}[\|\bar{\phi}_{\theta_i}\|_2 \cdot |(\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i})^\top \omega_i|], \end{aligned} \quad (\text{I.40})$$

where the inequality follows from the Jensen's inequality. In what follows, we upper bound the right-hand side of (I.40). Note that $\|\bar{\phi}_{\theta_i}(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. By further plugging into (I.40), we obtain that

$$\|F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i\|_2 \leq 2\tau_i^2 \cdot \mathbb{E}_{\sigma_i}[|(\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i})^\top \omega_i|] \leq 2\tau_i^2 \cdot \|(\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i})^\top \omega_i\|_{\sigma_i}, \quad (\text{I.41})$$

where the last inequality follows from the Jensen's inequality. Recall that $\omega_i, \theta_i \in \mathcal{B}$. Therefore, following from Assumption 4.2 and Corollary E.3, it further holds that

$$\begin{aligned} &\mathbb{E}[\|(\bar{\phi}_{\theta_i} - \bar{\phi}_{\omega_i})^\top \omega_i\|_{\sigma_i}] \\ &\leq \mathbb{E}[\|(\bar{\phi}_{\theta_i} - \bar{\phi}_0)^\top \omega_i\|_{\sigma_i}] + \mathbb{E}[\|(\bar{\phi}_0 - \bar{\phi}_{\omega_i})^\top \omega_i\|_{\sigma_i}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{I.42})$$

where the expectations are taken over all the randomness. Combining (I.41) and (I.42), we obtain that

$$\mathbb{E}[\|F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i\|_2] = \mathcal{O}(2\tau_i^2 \cdot R^{3/2} \cdot m^{-1/4}), \quad (\text{I.43})$$

where the expectation is taken over all the randomness. Finally, by plugging (I.43) into (I.39), we conclude that

$$\begin{aligned} &\mathbb{E}\left[\left(|(\delta_i - \omega_i)^\top \mathbb{E}_{\sigma_i}[\bar{\phi}_{\theta_i}(\bar{\phi}_{\theta_i}^\top \delta_i - \bar{\phi}_{\omega_i}^\top \omega_i)]|\right)^{1/2}\right] \\ &\leq C_i \cdot \left(\mathbb{E}[\|F(\theta_i) \cdot \omega_i - \tau_i \cdot g_i\|_2] + \mathbb{E}[\|\xi_i(\omega_i)\|_2] + \mathbb{E}[\|\xi_i(\delta_i)\|_2]\right)^{1/2} \\ &= \mathcal{O}(R^{5/4} \cdot m^{-1/8}) + \sqrt{2R} \cdot \tau_i^{-1} \cdot \left(\mathbb{E}[\|\xi_i(\delta_i)\|_2] + \mathbb{E}[\|\xi_i(\omega_i)\|_2]\right)^{1/2}, \end{aligned} \quad (\text{I.44})$$

where $C_i = \sqrt{2R} \cdot \tau_i^{-1}$ and $\xi_i(\delta_i)$, $\xi_i(\omega_i)$ are as defined in Assumption A.1.

Upper Bounding (iii.b) in (I.30). Following from the Cauchy-Schwartz inequality, it holds that

$$\sqrt{\mathbb{E}_{\sigma_i}[(\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i})(\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i})]} \leq (\|\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i} \cdot \|\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i})^{1/2}. \quad (\text{I.45})$$

To upper bound the right-hand side of (I.45), we first upper bound the term $\|\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i}$. Recall that $\omega_i, \theta_i \in \mathcal{B}$. Following from Assumption 4.2 and Corollary E.3, it holds that

$$\mathbb{E}[\|\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \quad \mathbb{E}[\|\omega_i^\top \bar{\phi}_{\omega_i} - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \quad (\text{I.46})$$

where $\bar{\phi}_0$ is defined in (B.1) and the expectations are taken over all the randomness. Therefore, following from (I.46), we obtain that

$$\begin{aligned} &\mathbb{E}[\|\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i}] \\ &\leq \mathbb{E}[\|\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] + \mathbb{E}[\|\omega_i^\top \bar{\phi}_{\omega_i} - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}). \end{aligned} \quad (\text{I.47})$$

It remains to upper bound the term $\|\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i}$ in the right-hand side of (I.45). Since $\delta_i \in \mathcal{B}$, by Assumption 4.2 and Corollary E.3, we obtain that

$$\mathbb{E}[\|\delta_i^\top \bar{\phi}_{\theta_i} - \delta_i^\top \bar{\phi}_0\|_{\sigma_i}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \quad (\text{I.48})$$

where the expectation is taken over all the randomness. Meanwhile, following from the fact that $\|\bar{\phi}_0(s, a)\|_2 \leq 2$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we obtain that

$$|\delta_i^\top \bar{\phi}_0(s, a) - \omega_i^\top \bar{\phi}_0(s, a)| \leq \|\bar{\phi}_0(s, a)\|_2 \cdot \|\delta_i - \omega_i\|_2 \leq 4R, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, W_{\text{init}} \in \mathbb{R}^{md}, \quad (\text{I.49})$$

where the first inequality follows from the Cauchy-Schwartz inequality and the second inequality follows from the fact that $\delta_i, \omega_i \in \mathcal{B}$. Combining (I.46), (I.48), and (I.49), we obtain that

$$\begin{aligned} \mathbb{E}[\|\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i}\|_{\sigma_i}] &\leq \mathbb{E}[\|\delta_i^\top \bar{\phi}_{\theta_i} - \delta_i^\top \bar{\phi}_0\|_{\sigma_i}] + \mathbb{E}[\|\delta_i^\top \bar{\phi}_0 - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] \\ &\quad + \mathbb{E}[\|\omega_i^\top \bar{\phi}_{\omega_i} - \omega_i^\top \bar{\phi}_0\|_{\sigma_i}] = \mathcal{O}(R), \end{aligned} \quad (\text{I.50})$$

where the expectation is taken over all the randomness. Finally, plugging (I.47) and (I.50) into (I.45), we obtain that

$$\left\{ \mathbb{E}_{\sigma_i} [(\omega_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i})(\delta_i^\top \bar{\phi}_{\theta_i} - \omega_i^\top \bar{\phi}_{\omega_i})] \right\}^{1/2} = \mathcal{O}(R^{5/4} \cdot m^{-1/8}), \quad (\text{I.51})$$

which concludes the upper bound of (iii.b) in (I.30).

Finally, by plugging (I.44), (I.51), and (I.30) into (I.29), we obtain that

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E}_{\nu_*} [|\langle \bar{\phi}_{\theta_i}(s, \cdot)^\top \delta_i - \bar{\phi}_{\omega_i}(s, \cdot)^\top \omega_i, \Delta_i^*(\cdot | s) \rangle|] \right] \\ &= \eta \cdot (\varphi_i + \psi_i) \cdot \left[\mathcal{O}(R^{5/4} \cdot m^{-1/8}) + \sqrt{2R} \cdot \tau_i^{-1} \cdot \left(\mathbb{E}[\|\xi_i(\delta_i)\|_2 + \|\xi_i(\omega_i)\|_2] \right)^{1/2} \right], \end{aligned} \quad (\text{I.52})$$

where $\xi_i(\delta_i)$ and $\xi_i(\omega_i)$ are defined in Assumption A.1. Here the expectations are taken over all the randomness.

Upper Bounding (iv) in (I.25). The bounding of the term (iv) is similar to that of the term (ii) in §H.8. It holds that

$$\begin{aligned} &|\langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}, \Delta_i^*(\cdot | s) \rangle| \\ &\leq |\langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}, \pi^*(\cdot | s) \rangle| + |\langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}, \pi_i(\cdot | s) \rangle| \\ &\leq \|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1} + \|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1}. \end{aligned} \quad (\text{I.53})$$

Note that $\theta_i, \theta_{i+1} \in \mathcal{B}$. Following from Assumption 4.2 and Lemma E.2, it holds that

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1}] \right] \\ &\leq \mathbb{E} [\|\phi_{\theta_{i+1}}(\cdot, \cdot)^\top \theta_{i+1} - \phi_0(\cdot, \cdot)^\top \theta_{i+1}\|_{\sigma_*}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \\ &\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_i}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1}] \right] \\ &\leq \mathbb{E} [\|\phi_{\theta_i}(\cdot, \cdot)^\top \theta_{i+1} - \phi_0(\cdot, \cdot)^\top \theta_{i+1}\|_{\sigma_*}] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{I.54})$$

where the inequalities follow from the Jensen's inequality, ϕ_0 is the feature mapping defined in (3.3) with $\theta = W_{\text{init}}$, and the expectations are taken over all the randomness. Following from (I.54), we obtain that

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1}] \right] \\ &\leq \mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1}] \right] + \mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_i}(s, \cdot)^\top \theta_{i+1} - \phi_0(s, \cdot)^\top \theta_{i+1}\|_{\pi^*, 1}] \right] \\ &= \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \end{aligned} \quad (\text{I.55})$$

where the expectations are taken over all the randomness. Similarly, it holds that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} [\|\phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}\|_{\pi_i, 1}] \right] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}), \quad (\text{I.56})$$

By plugging (I.55) and (I.56) into (I.53), we obtain that

$$\mathbb{E} \left[\mathbb{E}_{\nu_*} [|\langle \phi_{\theta_{i+1}}(s, \cdot)^\top \theta_{i+1} - \phi_{\theta_i}(s, \cdot)^\top \theta_{i+1}, \Delta_i^*(\cdot | s) \rangle|] \right] = \mathcal{O}(R^{3/2} \cdot m^{-1/4}). \quad (\text{I.57})$$

which concludes the upper bound of (iv) in (I.25).

Finally, by plugging (I.52) and (I.57) into (I.25), we obtain that

$$\begin{aligned} &\mathbb{E} \left[\mathbb{E}_{\nu_*} [|\langle \log(\pi_{i+1}(\cdot | s)/\pi_i(\cdot | s)) - \eta \cdot Q_{\omega_i}(s, \cdot), \pi_*(\cdot | s) - \pi_i(\cdot | s) \rangle|] \right] \\ &\leq \sqrt{2}(\varphi_i + \psi_i) \cdot \eta \cdot R^{1/2} \cdot \tau_i^{-1} \cdot \left(\mathbb{E}[\|\xi_i(\delta_i)\|_2] + \mathbb{E}[\|\xi_i(\omega_i)\|_2] \right)^{1/2} \\ &\quad + \mathcal{O}(\tau_{i+1} \cdot R^{3/2} \cdot m^{-1/4} + \eta \cdot R^{5/4} \cdot m^{-1/8}), \end{aligned} \quad (\text{I.58})$$

where φ_i, ψ_i are defined in Assumption A.2 and $\xi_i(\delta_i), \xi_i(\omega_i)$ are defined in Assumption A.1. Thus, we complete the proof of Lemma H.3. \square

J AUXILLIARY LEMMA

Lemma J.1 (Performance Difference (Kakade and Langford, 2002)). It holds for any π and $\tilde{\pi}$ that

$$J(\tilde{\pi}) - J(\pi) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{\tilde{\pi} \cdot \nu_{\tilde{\pi}}} [A^\pi(s, a)].$$

Here $\nu_{\tilde{\pi}}(\cdot)$ is the state visitation measure corresponding to $\tilde{\pi}$, which is defined in (2.3).

Proof. See Kakade and Langford (2002) for a detailed proof. \square