
Wasserstein Flow Meets Replicator Dynamics: A Mean-Field Analysis of Representation Learning in Actor-Critic

Yufeng Zhang[†]

Northwestern University
yufengzhang2023@u.northwestern.edu

Siyu Chen[†]

Tsinghua University
chensy18@mails.tsinghua.edu.cn

Zhuoran Yang

Princeton University
zy6@princeton.edu

Michael I. Jordan

UC Berkeley
jordan@cs.berkeley.edu

Zhaoran Wang

Northwestern University
zhaoranwang@gmail.com

Abstract

Actor-critic (AC) algorithms, empowered by neural networks, have had significant empirical success in recent years. However, most of the existing theoretical support for AC algorithms focuses on the case of linear function approximations, or linearized neural networks, where the feature representation is fixed throughout training. Such a limitation fails to capture the key aspect of representation learning in neural AC, which is pivotal in practical problems. In this work, we take a mean-field perspective on the evolution and convergence of feature-based neural AC. Specifically, we consider a version of AC where the actor and critic are represented by overparameterized two-layer neural networks and are updated with two-timescale learning rates. The critic is updated by temporal-difference (TD) learning with a larger stepsize while the actor is updated via proximal policy optimization (PPO) with a smaller stepsize. In the continuous-time and infinite-width limiting regime, when the timescales are properly separated, we prove that neural AC finds the globally optimal policy at a sublinear rate. Additionally, we prove that the feature representation induced by the critic network is allowed to evolve within a neighborhood of the initial one.

1 Introduction

In reinforcement learning (RL) [56], an agent aims to learn the optimal policy that maximizes the expected total reward by interacting with the environment. Policy-based RL algorithms achieve such a goal by directly optimizing the expected total reward as a function of the policy, which often involves two components: policy evaluation and policy improvement. Specifically, policy evaluation refers to estimating the value function of the current policy, which characterizes the performance of the current policy and reveals the updating direction for finding a better policy, which is known as policy improvement. Algorithms with these two ingredients are also called actor-critic (AC) methods [36], where the actor and the critic refer to the policy and its corresponding value function, respectively.

[†] Equal contribution.

Recently, in RL applications with large state spaces, actor-critic empowered by expressive function approximators such as neural networks have achieved striking empirical successes [3, 4, 9, 20, 51, 52, 60]. These successes benefit from the data-dependent representations learned by neural networks. Unfortunately, however, the theoretical understanding of this data-dependent benefit is very limited. The classical theory of AC focuses on the case of linear function approximation, where the actor and critic are represented using linear functions with the feature mapping fixed throughout learning [10, 11, 36]. Meanwhile, a few recent works establish convergence and optimality of AC with overparameterized neural networks [26, 39, 61], where the neural network training is captured by the Neural Tangent Kernel (NTK) [30]. Specifically, with properly designed parameter initialization and stepsizes, and sufficiently large network widths, the neural networks employed by both actor and critic can be assumed to be well approximated by linear functions of a random feature determined by initial parameters. In other words, concerning representation learning, the features induced by these algorithms are by assumption infinitesimally close to the initial featural representation, which is data-independent.

In this work, we make initial steps towards understanding how representation learning comes into play in neural AC. Specifically, we address the following questions:

Going beyond the NTK regime, does neural AC provably find the globally optimal policy? How does the feature representation associated with the neural network evolve along with neural AC?

We focus on a version of AC where the critic performs temporal-difference (TD) learning [55] for policy evaluation and the actor improves its policy via proximal policy optimization (PPO) [49], which corresponds to a Kullback-Leibler (KL) divergence regularized optimization problem, with the critic providing the update direction. Moreover, we utilize two-timescale updates where both the actor and critic are updated at each iteration but with the critic having a much larger stepsize. In other words, the critic is updated at a faster timescale. Meanwhile, we represent the critic explicitly as a two-layer overparameterized neural network and parameterize the actor implicitly via the critic and PPO updates. To examine convergence, we study the evolution of the actor and critic in the continuous-time limiting regime with the network width going to infinity. In such a regime, the actor update is closely connected to replicator dynamics [12, 28, 50] and the critic update is captured by a semigradient flow in the Wasserstein space [59]. Moreover, the semigradient flow runs at a faster timescale according to the two-timescale mechanism.

It turns out that the separation of timescales plays an important role in the convergence analysis. In particular, in the continuous-time limit, it enables us to first separately analyze the evolution of actor and critic and then combine these results to get final theoretical guarantees. Specifically, focusing solely on the actor, we prove that the time-averaged suboptimality of the actor converges sublinearly to zero up to the time-averaged policy evaluation error associated with critic updates. Moreover, for the critic, under proper regularity conditions, we connect the Bellman error to the Wasserstein distance and show that the time-averaged policy evaluation error also converges sublinearly to zero. Therefore, we show that neural AC provably achieves global optimality at a sublinear rate. Furthermore, regarding representation learning, we show that the critic induces a data-dependent feature representation within an $O(1/\alpha)$ neighborhood of the initial representation in terms of the Wasserstein distance, where α is a sufficiently large scaling parameter.

The key to our technical analysis reposes on three ingredients: (i) infinite-dimensional variational inequalities with a one-point monotonicity [27], (ii) a mean-field perspective on neural networks [19, 41, 42, 53, 54], and (iii) the two-timescale stochastic approximation [13, 37]. In particular, in the infinite-width limit, the neural network and its induced feature representation are identified with a distribution over the parameter space. The mean-field perspective enables us to characterize the evolution of such a distribution within the Wasserstein space via a partial differential equation (PDE) [5, 6, 58, 59]. For policy evaluation, such a PDE is given by the semigradient flow induced by TD learning. We characterize the error of policy evaluation by showing that mean-field Bellman error satisfies a version of one-point monotonicity tailored to the Wasserstein space. Moreover, our actor analysis utilizes the geometry of policy optimization, which shows that the expected total reward,

as a function of the policy, also enjoys the property of one-point monotonicity in the policy space. Finally, the actor and critic errors are connected via two-timescale stochastic approximation. To the best of our knowledge, this is the first time that convergence and global optimality guarantees have been obtained for neural AC.

Related Work. AC with linear function approximation has been studied extensively in the literature. In particular, using a two-timescale stochastic approximation via ordinary differential equations, [10, 11, 36] establish asymptotic convergence guarantees in the continuous-time limiting regime. More recently, using more sophisticated optimization techniques, various works [29, 35, 64–66] have established discrete-time convergence guarantees that show that linear AC converges sublinearly to either a stationary point or the globally optimal policy. Furthermore, when overparameterized neural networks are employed, [26, 39, 61] prove that neural AC converges to the global optimum at a sublinear rate. In these works, the initial value of the network parameters and the learning rates are chosen such that both actor and critic updates are captured by the NTK. In other words, when the network width is sufficiently large, such a version of neural AC is well approximated by its linear counterpart via the neural tangent feature. In comparison, we establish a mean-field analysis that has a different scaling than the NTK regime. We also establish finite-time convergence to global optimality, and more importantly, the feature representation induced by the critic is data-dependent and allowed to evolve within a much larger neighborhood around the initial one.

Furthermore, our work is also related to the recent line of research on understanding stochastic gradient descent (SGD) for supervised learning problems involving an overparameterized two-layer neural network under the mean-field regime. See, e.g., [16, 19, 21, 22, 31, 40–42, 53, 54, 63] and the references therein. In the continuous-time and infinite-width limit, these works show that SGD for neural network training is captured by a Wasserstein gradient flow [5, 6, 59] of an energy function that corresponds to the objective function in supervised learning. In contrast, our analysis combines such a mean-field analysis with TD learning and two-timescale stochastic approximation, which are tailored specifically to AC. Moreover, our critic is updated via TD learning, which is a semigradient algorithm and there is no objective functional making TD learning a gradient-based algorithm. Thus, in the mean-field regime, our critic is given by a Wasserstein semigradient flow, which also differs from these existing works.

Additionally, our work is closely related to [1, 69], who provide mean-field analyses for neural TD-learning and Q-learning [62]. In comparison, we focus on AC, which is a two-timescale policy optimization algorithm. Finally, [2] studies softmax policy gradient with neural network policies in the mean-field regime, where policy gradient is cast as a Wasserstein gradient flow with respect to the expected total reward. The algorithm assumes that the critic directly gets the desired value function and thus the algorithm is single-timescale. Moreover, the convergence guarantee in [2] is asymptotic. In comparison, our AC is two-timescale and we establish non-asymptotic sublinear convergence guarantees to global optimality.

Notation. We denote by $\mathcal{P}(\mathcal{X})$ the set of probability measures over the measurable space \mathcal{X} . Given a curve $\rho : \mathbb{R} \rightarrow \mathcal{P}(\mathcal{X})$, we denote by $\dot{\rho}_s = \partial_t \rho_t |_{t=s}$ its derivative with respect to the time. For an operator $F : \mathcal{X} \rightarrow \mathcal{X}$ and a measure $\mu \in \mathcal{P}(\mathcal{X})$, we denote by $F_{\#}\mu = \mu \circ F^{-1}$ the push forward of μ through F . We denote by $\chi^2(\rho \| \mu)$ the chi-squared divergence between probability measures ρ and μ , which is defined as $\chi^2(\rho \| \mu) = \int (\rho/\mu - 1)^2 d\mu$. Given two probability measures ρ and μ , we denote the Kullback-Leibler divergence or the relative entropy from μ to ρ by $\text{KL}(\rho \| \mu) = \int \log(\rho/\mu) d\rho$. For $\nu_1, \nu_2, \mu \in \mathcal{P}(\mathcal{X})$, we define the $\dot{H}^{-1}(\mu)$ weighted homogeneous Sobolev norm as $\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} = \sup \{ |\langle f, \nu_1 - \nu_2 \rangle| \mid \|f\|_{\dot{H}^1(\mu)} \leq 1 \}$. We denote by $\|f(x)\|_{p,\mu} = (\int |f(x)|^p \mu(dx))^{1/p}$ the ℓ_p -norm with respect to probability measure μ . We denote by \otimes the semidirect product, i.e., $\mu \otimes K = K(y|x)\mu(x)$ for $\mu \in \mathcal{P}(\mathcal{X})$ and transition kernel $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$. For a function $f : \mathcal{X} \rightarrow \mathbb{R}$, we denote by $\text{Lip}(f) = \sup_{x,y \in \mathcal{X}, x \neq y} |f(x) - f(y)| / \|x - y\|$ its Lipschitz constant. We denote a normal distribution on \mathbb{R}^D by $\mathcal{N}(\mu, \Sigma)$, where μ is the mean value and Σ is the covariance matrix.

2 Background

In this section, we first introduce the policy optimization problem and the actor-critic method. We then present the definition of the Wasserstein space.

2.1 Policy Optimization and Actor-Critic

We consider a Markov decision process (MDP) given by $(\mathcal{S}, \mathcal{A}, \gamma, P, r, \mathcal{D}_0)$, where $\mathcal{S} \subseteq \mathbb{R}^{d_1}$ is the state, $\mathcal{A} \subseteq \mathbb{R}^{d_2}$ is the action space, $\gamma \in (0, 1)$ is the discount factor, $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$ is the transition kernel, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_+$ is the reward function, and $\mathcal{D}_0 \in \mathcal{P}(\mathcal{S})$ is the initial state distribution. Without loss of generality, we assume that $\mathcal{S} \times \mathcal{A} \subseteq \mathbb{R}^d$ and $\|(s, a)\|_2 \leq 1$, where $d = d_1 + d_2$. We remark that as long as the state-action space is bounded, we can normalize the space to be within the unit sphere. Given a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$, at the m th step, the agent takes an action a_m at state s_m according to $\pi(\cdot | s_m)$ and observes a reward $r_m = r(s_m, a_m)$. The environment then transits to the next state s_{m+1} according to the transition kernel $P(\cdot | s_m, a_m)$. Note that the policy π induces Markov chains on both \mathcal{S} and $\mathcal{S} \times \mathcal{A}$. Considering the Markov chain on \mathcal{S} , we denote the induced Markov transition kernel by $P^\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{S})$, which is defined as $P^\pi(s' | s) = \int_{\mathcal{A}} P(s' | s, a) \pi(da | s)$. Likewise, we denote the Markov transition kernel on $\mathcal{S} \times \mathcal{A}$ by $\tilde{P}^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S} \times \mathcal{A})$, which is defined as $\tilde{P}^\pi(s', a' | s, a) = \pi(a' | s') P(s' | s, a)$. Let $\tilde{\mathcal{D}}$ be a probability measure on $\mathcal{S} \times \mathcal{A}$. We then define the visitation measure induced by policy π and starting from $\tilde{\mathcal{D}}$ as

$$\tilde{\mathcal{E}}_{\tilde{\mathcal{D}}}^\pi(d(s, a)) = (1 - \gamma) \cdot \sum_{m \geq 0} \gamma^m \cdot \mathbb{P}((s_m, a_m) \in d(s, a) | (s_0, a_0) \sim \tilde{\mathcal{D}}), \quad (2.1)$$

where (s_m, a_m) is the trajectory generated by starting from $(s_0, a_0) \sim \tilde{\mathcal{D}}$ and following policy π thereafter. If $\tilde{\mathcal{D}} = \mathcal{D} \otimes \pi$ holds, we then denote such a visitation measure by $\tilde{\mathcal{E}}_{\tilde{\mathcal{D}}}^\pi$. Furthermore, we denote by $\mathcal{E}(ds) = \int_{\mathcal{A}} \tilde{\mathcal{E}}(ds, da)$ the marginal distribution of visitation measure $\tilde{\mathcal{E}}$ with respect to \mathcal{S} . In particular, when $(s_0, a_0) \sim \mathcal{D} \otimes \pi$ holds in (2.1), it follows that $\tilde{\mathcal{E}}_{\tilde{\mathcal{D}}}^\pi = \mathcal{E}_{\mathcal{D}}^\pi \otimes \pi$. In policy optimization, we aim to maximize the expected total rewards $J(\pi)$ defined as follows,

$$J(\pi) = \mathbb{E}^\pi \left[\sum_{m \geq 0} \gamma^m \cdot r(s_m, a_m) \mid s_0 \sim \mathcal{D}_0 \right],$$

where we denote by \mathbb{E}^π the expectation with respect to $a_m \sim \pi(\cdot | s_m)$ and $s_{m+1} \sim P(\cdot | s_m, a_m)$ for $m \geq 0$. We define the action value function $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and the state value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ as follows,

$$Q^\pi(s, a) = \mathbb{E}^\pi \left[\sum_{m \geq 0} \gamma^m \cdot r(s_m, a_m) \mid s_0 = s, a_0 = a \right], \quad V^\pi(s) = \langle Q^\pi(s, \cdot), \pi(\cdot | s) \rangle_{\mathcal{A}}, \quad (2.2)$$

where we denote by $\langle \cdot, \cdot \rangle_{\mathcal{A}}$ the inner product on the action space \mathcal{A} . Correspondingly, the advantage function $A^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined as

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s).$$

It is known that the action value function Q^π is the unique global minimizer to the following mean-squared Bellman error (MSBE),

$$\text{MSBE}(Q; \pi) = \frac{1}{2} \mathbb{E}_{(s, a) \sim \tilde{\Phi}^\pi} \left[(Q(s, a) - r(s, a) - \gamma \mathbb{E}_{(s', a') \sim \tilde{P}^\pi(\cdot | s, a)} [Q(s', a')])^2 \right], \quad (2.3)$$

where $\tilde{\Phi}^\pi$ is a weighting distribution depending on policy π and is with full support, i.e., $\text{supp}(\tilde{\Phi}^\pi) = \mathcal{S} \times \mathcal{A}$. Therefore, the policy optimization problem can be written as the following bilevel optimization problem,

$$\max_{\pi} J(\pi) = \mathbb{E}_{s \sim \mathcal{D}_0} \left[\langle Q^\pi(s, \cdot), \pi(\cdot | s) \rangle_{\mathcal{A}} \right], \quad \text{subject to } Q^\pi = \underset{Q}{\text{argmin}} \text{MSBE}(Q; \pi). \quad (2.4)$$

The inner problem in (2.4) is known as a policy evaluation subproblem, while the outer problem is the policy improvement subproblem. One of the most popular way to solve the policy optimization problem is actor-critic (AC) methods [56], where the job of the critic is to evaluate current policy and then the actor updates its policy according to the critic's evaluation.

2.2 Wasserstein Space

Let $\Theta \subseteq \mathbb{R}^D$ be a Polish space. We denote by $\mathcal{P}_2(\Theta) \subseteq \mathcal{P}(\Theta)$ the set of probability measures with finite second moments. Then, the Wasserstein-2 (W_2) distance between $\mu, \nu \in \mathcal{P}_2(\Theta)$ is defined as follows,

$$W_2(\mu, \nu) = \inf \left\{ \mathbb{E}[\|X - Y\|^2]^{1/2} \mid \text{law}(X) = \mu, \text{law}(Y) = \nu \right\},$$

where the infimum is taken over the random variables X and Y on Θ and we denote by $\text{law}(X)$ the distribution of a random variable X . We call $\mathcal{M} = (\mathcal{P}_2(\Theta), W_2)$ the Wasserstein (W_2) space, which is an infinite-dimensional manifold [59]. See §A.1 for more details.

3 Algorithm

Two-timescale Actor-critic. We consider a two-timescale Actor-critic (AC) algorithm [34, 45] for the policy optimization problem in (2.4). For policy evaluation, we parameter the critic Q with a neural network and update the parameter via temporal-difference (TD) learning [55]. For policy improvement, we update the actor policy π via proximal policy optimization (PPO) [49]. Our algorithm is two-timescale since both the actor and critic are updated at each iteration with different stepsizes. Specifically, we parameterize the critic Q by the following neural network with width M and parameter $\hat{\theta} = (\hat{\theta}^{(1)}, \dots, \hat{\theta}^{(M)}) \in \mathbb{R}^{D \times M}$,

$$Q_{\hat{\theta}}(s, a) = \frac{\alpha}{M} \sum_{i=1}^M \sigma(s, a; \hat{\theta}^{(i)}). \quad (3.1)$$

Here $\sigma(s, a; \theta) : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^D \rightarrow \mathbb{R}$ is the activation function and $\alpha > 0$ is the scaling parameter. Such a structure also appears in [17, 18, 41]. In a discrete-time finite-width (DF) scenario, at the k th iteration, the critic and actor are updated as follows,

$$\text{DF-TD: } \hat{\theta}_{k+1}^{(i)} = \hat{\theta}_k^{(i)} - \frac{\varepsilon'}{\alpha} (Q_{\hat{\theta}_k}(s_k, a_k) - r(s_k, a_k) - \gamma Q_{\hat{\theta}_k}(s'_k, a'_k)) \nabla_{\theta} \sigma(s, a; \hat{\theta}_k^{(i)}), \quad (3.2)$$

$$\text{DF-PPO: } \hat{\pi}_{k+1}(\cdot | s) = \underset{\pi}{\operatorname{argmax}} \left\{ \langle Q_{\hat{\theta}_k}(s, \cdot), \pi(\cdot | s) \rangle_{\mathcal{A}} - \varepsilon^{-1} \cdot \text{KL}(\pi(\cdot | s) \| \hat{\pi}_k(\cdot | s)) \right\}, \quad (3.3)$$

where $(s_k, a_k) \sim \tilde{\Phi}^{\hat{\pi}_k}$ and $(s'_k, a'_k) \sim \tilde{P}^{\hat{\pi}_k}(\cdot | s_k, a_k)$. Here $\hat{\pi}_k$ is the policy for the actor at the k th iteration, $\tilde{\Phi}^{\hat{\pi}_k}$ is the corresponding weighting distribution, ε and ε' are the stepsizes for the DF-PPO update and the DF-TD update, respectively. In (3.2), the scaling of α^{-1} arises since our update falls into the lazy-training regime [18]. In the sequel, we denote by $\eta = \varepsilon'/\varepsilon$ the relative TD timescale. Note that in a double-loop AC algorithm, the critic can usually be solved with high precision. In the two-timescale AC however, even with the KL-divergence term in (3.3) which regularizes the policy update and helps to improve the local estimation quality of the TD update, the critic $Q_{\hat{\theta}_k}$ for updating the actor's policy $\hat{\pi}_k$ can still be far from the true action value function $Q^{\hat{\pi}_k}$. Since the policy evaluation problem is not fully solved at each iteration, the two-timescale AC can be more efficient in computation while more challenging to establish a theoretical guarantee.

Mean-field (MF) Limit. To analyze the convergence of the two-timescale AC with neural networks, we employ the analysis that studies the mean-field limit regime [41, 42]. Here, by saying the mean-field limit, we refer to the infinite-width limit, i.e., $M \rightarrow \infty$ for the neural network width M in (3.1), and the continuous-time limit, i.e., $t = k\varepsilon$ where $\varepsilon \rightarrow 0$ for the stepsize in (3.2) and (3.3). For $\hat{\theta} = \{\hat{\theta}^{(i)}\}_{i=1}^M$ independently sampled from a distribution ρ , we can write the infinite-width limit of (3.1) as

$$Q(s, a; \rho) = \alpha \int \sigma(s, a; \theta) \rho(d\theta). \quad (3.4)$$

In the sequel, we denote by $\hat{\rho}_k$ the distribution of $\hat{\theta}_k^{(i)}$ for the infinite-width limit of the neural network at the k th iteration. We further let ρ_t and π_t be the continuous-time limits of $\hat{\rho}_k$ and $\hat{\pi}_k$, respectively.

As studied in [69], the mean-field limit of the DF-TD update in (3.2) is

$$\text{MF-TD: } \partial_t \rho_t = -\eta \operatorname{div}(\rho_t \cdot g(\cdot; \rho_t, \pi_t)), \quad (3.5)$$

where η is the relative TD timescale and

$$g(\theta; \rho, \pi) = -\mathbb{E}_{\tilde{\Phi}^\pi} \left\{ [Q(s, a; \rho) - r(s, a) - \gamma \cdot Q(s', a'; \rho)] \cdot \alpha^{-1} \nabla_{\theta} \sigma(s, a; \theta) \right\} \quad (3.6)$$

is a vector field. Here $\mathbb{E}_{\tilde{\Phi}^\pi}$ is taken with respect to $(s, a) \sim \tilde{\Phi}^\pi$ and $(s', a') \sim \tilde{P}^\pi(\cdot | s, a)$. It remains to characterize the mean-field limit of the DF-PPO update in (3.3). By solving the maximization problem in (3.3), the infinite-width limit of DF-PPO update can be written in closed form as

$$\varepsilon^{-1} \cdot \left\{ \log [\hat{\pi}_{k+1}(a | s)] - \log [\hat{\pi}_k(a | s)] \right\} = Q(s, a; \hat{\rho}_k) - \hat{Z}_k(s),$$

where $\hat{Z}_k(s)$ is the normalizing factor such that $\int \hat{\pi}_k(da | s) = 1$ for any $s \in \mathcal{S}$. By letting $t = k\varepsilon$ and $\varepsilon \rightarrow 0$, we have $\partial_t \log \pi_t = Q_t - Z_t$, which can be further written as $\partial_t \pi_t = \pi_t \cdot (Q_t - Z_t)$. Here we have $Q_t(a, s) = Q(a, s; \rho_t)$ and Z_t is the continuous-time limit of \hat{Z}_k . Furthermore, noting that $\partial_t \int \pi_t(da | s) = 0$, the mean-field limit of the DF-PPO update in (3.3) is

$$\text{MF-PPO: } \frac{d}{dt} \pi_t = \pi_t \cdot A_t, \quad \text{where } A_t(s, a) = Q_t(s, a) - \int Q_t(s, a) \pi_t(da | s). \quad (3.7)$$

The two updates (3.5) and (3.7) correspond to the mean-field limits of (3.2) and (3.3), respectively, and together serve as the mean-field limit of the two-timescale AC. In particular, we remark that the MF-TD update in (3.5) for the critic is captured by a semigradient flow in the Wasserstein space [59] while the MF-PPO update in (3.7) for the actor resembles the replicator dynamics [12, 28, 50]. Note that such a framework is applicable to continuous state and action space. In this paper, we aim to provide a theoretical analysis of the mean-field limit of the two-timescale AC.

4 Main Result

In this section, we first establish the convergence of the MF-PPO update in §4.1. Then, with additional assumptions, we establish the optimality and convergence of the mean-field two-timescale AC in §4.2.

4.1 Convergence of Mean-field PPO

For the MF-PPO update in (3.7), we establish the following theorem on its global optimality and convergence rate.

Theorem 4.1 (Convergence of MF-PPO). Let $\pi^* = \operatorname{argmax}_{\pi} J(\pi)$ be the optimal policy and π_0 be the initial policy. Then, it holds that

$$\frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt \leq \frac{\zeta}{T} + 4\kappa \cdot \underbrace{\frac{1}{T} \int_0^T \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}} dt}_{\text{policy evaluation error}}, \quad (4.1)$$

where $\tilde{\phi}^{\pi_t} \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$ is an evaluation distribution for the policy evaluation error and $\zeta = \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\operatorname{KL}(\pi^*(\cdot | s) \| \pi_0(\cdot | s)) \right]$ is the expected KL-divergence between π^* and π_0 . Furthermore, letting $\tilde{\phi}^{\pi_t} = \frac{1}{2} \tilde{\phi}_0 + \frac{1}{2} \phi_0 \otimes \pi_t$, where $\tilde{\phi}_0 \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$ is a base distribution and $\phi_0 = \int_{\mathcal{A}} \tilde{\phi}_0(\cdot, da)$, the concentrability coefficient κ is then given by

$$\kappa = \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^{\pi^*}}{\tilde{\phi}_0} \right\|_{\infty}.$$

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

The concentrability coefficient commonly appears in the reinforcement learning literature [7, 23, 24, 38, 39, 43, 48, 57, 61]. In contrast to a more standard concentrability coefficient form, note that κ is irrelevant to the update of the algorithm. To show the convergence of the MF-PPO, our condition here is much weaker since we only need a given base distribution $\tilde{\phi}_0$ such that $\kappa < \infty$.

Theorem 4.1 shows that the MF-PPO converges to the globally optimal policy at a rate of $\mathcal{O}(T^{-1})$ up to the policy evaluation error. Such a theorem implies the global optimality and convergence of a double-loop AC algorithm, where the critic Q_t is solved to high precision and the policy evaluation error is sufficiently small. In the sequel, we consider a more challenging setting, where the critic Q_t is updated simultaneously along with the update of the actor’s policy π_t .

4.2 Global Optimality and Convergence of Two-timescale AC

In what follows, we aim to characterize the upper bound of the policy evaluation error when the critic and the actor are updated simultaneously. Specifically, the actor is updated via MF-PPO in (3.7) and the critic $Q_t = Q(\cdot; \rho_t)$ is updated via the MF-TD in (3.5). For the smooth function σ in the parameterization of the Q function in (3.4), we consider it to be the following two-layer neural network,

$$\sigma(s, a; \theta) = B_\beta \cdot \beta(b) \cdot \tilde{\sigma}(w^\top(s, a, 1)), \quad (4.2)$$

where $\tilde{\sigma} : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function, $\theta = (b, w)$ is the parameter, and $\beta : \mathbb{R} \rightarrow (-1, 1)$ is an odd and invertible function with scaling hyper-parameter $B_\beta > 0$. It then holds that $D = d + 2$, where d and D are the dimensions of (s, a) and θ , respectively. It is worth noting that the function class of $\int \sigma(s, a; \theta) \rho(d\theta)$ for $\rho \in \mathcal{P}_2(\mathbb{R}^D)$ is the same as

$$\mathcal{F} = \left\{ \int \beta' \cdot \tilde{\sigma}(w^\top(s, a, 1)) \nu(dw) \mid \nu \in \mathcal{P}_2((-B_\beta, B_\beta) \times \mathbb{R}^{d+1}) \right\}, \quad (4.3)$$

which captures a vast function class because of the universal function approximation theorem [8, 47]. We remark that we introduce the rescaling function β in (4.2) to avoid the study of the space of probability measures over $(-B_\beta, B_\beta) \times \mathbb{R}^{d+1}$ in (4.3), which has boundary and thus lacks the regularity in the study of optimal transport. Furthermore, note that we introduce a hyper-parameter $\alpha > 1$ in the Q function in (3.4). Thus, we are using $\alpha \cdot \mathcal{F}$ to represent \mathcal{F} , which causes an “over-representation” when $\alpha > 1$. Such over-representation appears to be essential for our analysis. For a brief peek, we remark that α actually controls the gap in the average total reward over time when the relative time-scale η is properly selected according to Theorem 4.6. Furthermore, such an influence is imposed through Lemma 4.4, which shows that the Wasserstein distance between ρ_0 and ρ_{π_t} is upper bounded by $O(1/\alpha)$. In what follows, we consider the initialization of the TD update to be $\rho_0 = \mathcal{N}(0, I_D)$, which implies that $Q(s, a; \rho_0) = 0$. We next impose the following assumption on the two-layer neural network σ .

Assumption 4.2 (Regularity of the Neural Network). For the two-layer neural network σ defined in (4.2), we assume that the following properties hold.

- (i) The rescaling function $\beta : \mathbb{R} \rightarrow (-1, 1)$ is odd, $L_{0,\beta}$ -Lipschitz continuous, $L_{1,\beta}$ -smooth, and invertible. Meanwhile, the inverse β^{-1} is locally Lipschitz continuous. In particular, we assume that β^{-1} is ℓ_β -Lipschitz continuous in $[-2/3, 2/3]$.
- (ii) The activation function $\tilde{\sigma} : \mathbb{R} \rightarrow \mathbb{R}$ is odd, $B_{\tilde{\sigma}}$ -bounded, $L_{0,\tilde{\sigma}}$ -Lipschitz continuous, and $L_{1,\tilde{\sigma}}$ -smooth.

We remark that Assumption 4.2 is not restrictive and can be satisfied by a large family of neural networks, e.g., $\tilde{\sigma}(x) = \tanh(x)$ and $\beta(b) = \tanh(b)$. Noting that $\|(s, a)\|_2 \leq 1$, Assumption 4.2 implies that the function $\sigma(s, a; \theta)$ in (4.2) is odd with respect to w and b and is also bounded, Lipschitz continuous, and smooth in the parameter domain, that is,

$$|\nabla_\theta \sigma(s, a; \theta)| < B_1, \quad |\nabla_{\theta\theta}^2 \sigma(s, a; \theta)| < B_2. \quad (4.4)$$

We then impose the following assumption on the MDP.

Assumption 4.3 (Regularity of the MDP). For the MDP $(\mathcal{S}, \mathcal{A}, \gamma, P, r, \mathcal{D}_0)$, we assume the following properties hold.

- (i) The reward function r and the transition kernel P admit the following representations with respect to the activation function $\tilde{\sigma}$,

$$r(s, a) = B_r \cdot \int \tilde{\sigma}((s, a, 1)^\top w) \mu(dw), \quad (4.5)$$

$$P(s' | s, a) = \int \tilde{\sigma}((s, a, 1)^\top w) \varphi(s') \psi(s'; dw), \quad (4.6)$$

where μ and $\psi(s'; \cdot)$ are probability measures in $\mathcal{P}_2(\mathbb{R}^{d+1})$ for any $s' \in \mathcal{S}$, B_r is a positive scaling parameter, and $\varphi(s') : \mathcal{S} \rightarrow \mathbb{R}_+$ is a nonnegative function.

- (ii) The reward function r satisfies that $r(s, a) \geq 0$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. For the representation of r in (4.5) and the representation of the transition kernel P in (4.6), we assume that

$$\begin{aligned} \chi^2(\mu \| \rho_{w,0}) &< M_\mu, & \chi^2(\psi(s; \cdot) \| \rho_{w,0}) &< M_\psi, & \forall s \in \mathcal{S}, \\ \int \varphi(s) ds &\leq M_{1,\varphi}, & \int \varphi(s)^2 ds &\leq M_{2,\varphi}, \end{aligned}$$

where $\rho_{w,0}$ is the marginal distribution of ρ_0 with respect to w , i.e., $\rho_{w,0} = \int \rho_0(db, \cdot)$, χ^2 is the chi-squared divergence, and $M_\mu, M_\psi, M_{1,\varphi}, M_{2,\varphi}$ are absolute constants.

- (iii) We assume that there exists an absolute constant \mathcal{G} such that

$$\begin{aligned} \|\psi(s; \cdot) - \psi(s'; \cdot)\|_{\dot{H}^{-1}(\mu)} &< \mathcal{G}, & \|\psi(s; \cdot) - \mu\|_{\dot{H}^{-1}(\mu)} &< \mathcal{G}, \\ \|\psi(s; \cdot) - \mu\|_{\dot{H}^{-1}(\psi(s'; \cdot))} &< \mathcal{G}, & \|\psi(s; \cdot) - \psi(s'; \cdot)\|_{\dot{H}^{-1}(\psi(s''; \cdot))} &< \mathcal{G}, \end{aligned} \quad \forall s, s', s'' \in \mathcal{S},$$

where $\|\cdot\|_{\dot{H}^{-1}(\cdot)}$ is the weighted homogeneous Sobolev norm.

We remark that by assuming ψ to be a probability measure and that $\varphi(s') \geq 0$ in (4.6), the representation of the transition kernel does not lose generality. Specifically, the function class of (4.6) is the same as

$$\mathcal{P} = \left\{ \int \tilde{\sigma}((s, a, 1)^\top w) \tilde{\psi}(s'; dw) \mid \tilde{\psi}(s'; \cdot) \text{ is a signed measure for any } s' \in \mathcal{S} \right\}.$$

See §C.1 for a detailed proof. Assumption 4.3 generalizes the linear MDP in [14, 32, 67, 68]. In contrast, our representation of the reward function and the transition kernel benefits from the universal function approximation theorem and is thus not as restrictive as the original linear MDP assumption. Note that the infinite-width neural network has a two-layer structure by (4.2). We establish the following lemma on the regularity of the representation of the action value function Q^π by such a neural network.

Lemma 4.4 (Regularity of Representation of Q^π). Suppose that Assumptions 4.2 and 4.3 hold. For any policy π , there exists a probability measure $\rho_\pi \in \mathcal{P}_2(\mathbb{R}^D)$ for the representation of Q^π with the following properties.

- (i) For function $Q(s, a; \rho_\pi)$ defined by (3.4) with $\rho = \rho_\pi$ and the action value function $Q^\pi(s, a)$ defined by (2.2), we have $Q(s, a; \rho_\pi) = Q^\pi(s, a)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.
- (ii) By letting $B_\beta \geq 2(B_r + \gamma(1 - \gamma)^{-1} B_r M_{1,\varphi})$ for the neural network defined in (4.2) and $\rho_0 \sim \mathcal{N}(0, I_D)$ for the initial distribution, we have $\widetilde{W}_2(\rho_\pi, \rho_0) \leq \bar{D}$ for any policy π , where we define $\widetilde{W}_2(\cdot, \cdot) = \alpha W_2(\cdot, \cdot)$ as the scaled W_2 metric. Here constant \bar{D} depends on the discount factor γ and the absolute constants $L_{0,\beta}, L_{1,\beta}, l_\beta, B_r, M_\mu, M_\psi, M_{1,\varphi}, M_{2,\varphi}$ defined in Assumptions 4.2 and 4.3.

Proof. See §B.2 for a detailed proof. □

Property (i) of Lemma 4.4 shows that the action value function Q^π can be parameterized with the infinite-width two-layer neural network $Q(\cdot; \rho_\pi)$ in (3.4). Note that a larger B_β captures a larger function class in (4.3). Without loss of generality, we consider that $B_\beta \geq 2(B_r + \gamma(1-\gamma)^{-1}B_r M_{1,\varphi})$ holds in the sequel. Hence, by Property (ii), it holds that $\widetilde{W}_2(\rho_\pi, \rho_0) \leq O(1)$ for any policy π . In particular, it holds by Property (i) of Lemma 4.4 that $\|Q_t - Q^{\pi_t}\|_{2, \widetilde{\phi}^{\pi_t}} = \|Q(\cdot; \rho_t) - Q(\cdot; \rho_{\pi_t})\|_{2, \widetilde{\phi}^{\pi_t}}$ and we have the following theorem to characterize such an error with regard to the W_2 space.

Theorem 4.5 (Upper Bound of Policy Evaluation Error). Suppose that Assumptions 4.2 and 4.3 hold and $\rho_0 \sim \mathcal{N}(0, I_D)$ is the initial distribution. We specify the weighting distribution $\widetilde{\Phi}^{\pi_t}$ in MF-TD (3.5) as $\widetilde{\Phi}^{\pi_t} = \widetilde{\mathcal{E}}_{\widetilde{\phi}^{\pi_t}}^{\pi_t}$, where $\widetilde{\phi}^{\pi_t} \in \mathcal{P}(\mathcal{S} \times \mathcal{A})$ is the evaluation distribution for the policy evaluation error in Theorem 4.1. Then, it holds that

$$(1 - \sqrt{\gamma}) \cdot \|Q_t - Q^{\pi_t}\|_{2, \widetilde{\phi}^{\pi_t}}^2 \leq -\frac{d}{dt} \frac{\widetilde{W}_2^2(\rho_t, \rho_{\pi_t})}{2\eta} + \Delta_t, \quad (4.7)$$

where

$$\begin{aligned} \Delta_t = & 2\alpha^{1/2}\eta^{-1}\mathcal{B}B_1 \cdot \widetilde{W}_2(\rho_t, \rho_0)\widetilde{W}_2(\rho_t, \rho_{\pi_t}) \\ & + \alpha^{-1}B_2 \cdot \left(4B_1 \max\{\widetilde{W}_2(\rho_{\pi_t}, \rho_0), \widetilde{W}_2(\rho_t, \rho_0)\} + B_r\right) \widetilde{W}_2(\rho_t, \rho_{\pi_t})^2. \end{aligned}$$

Here B_1 and B_2 are defined in (4.4) of Assumption 4.2, η is the relative TD timescale, α is the scaling parameter of the neural network, and $\widetilde{W}_2 = \alpha W_2$ is the scaled W_2 metric. Moreover, constant \mathcal{B} depends on the discount factor γ , the scaling parameter B_β in (4.2), and the absolute constants $l_\beta, B_r, M_{1,\varphi}, \mathcal{G}$ defined in Assumptions 4.2 and 4.3.

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

Here we give a nonrigorous discussion on how to upper bound Δ_t in (4.7). If $\widetilde{W}_2(\rho_t, \rho_0) \leq O(1)$ holds for any $t \in [0, T]$, by $\widetilde{W}_2(\rho_{\pi_t}, \rho_0) \leq O(1)$ in Lemma 4.4 and the triangle inequality of W_2 distance [59], it follows that $\widetilde{W}_2(\rho_t, \rho_{\pi_t}) \leq O(1)$ and $\Delta_t \leq O(\alpha^{1/2}\eta^{-1} + \alpha^{-1})$. Taking a time average of integration on both sides of (4.7), the policy evaluation error $\frac{1}{T} \int_0^T \|Q_t - Q^{\pi_t}\|_{2, \widetilde{\phi}^{\pi_t}} dt$ is then upper bounded by $O(\eta^{-1}T^{-1} + \alpha^{1/2}\eta^{-1} + \alpha^{-1})$. Inspired by such a fact, we introduce the following restarting mechanism to ensure $\widetilde{W}_2(\rho_t, \rho_0) \leq O(1)$.

Restarting Mechanism. Let $\widetilde{W}_0 = \lambda\bar{D}$ be a threshold, where \bar{D} is the upper bound for $\widetilde{W}_2(\rho_\pi, \rho_0)$ by Lemma 4.4, $\lambda \geq 3$ is a constant scaling parameter for the restarting threshold, ρ_t is the distribution of the parameters in the neural network at time t , and ρ_0 is the initial distribution. Whenever we detect that $\widetilde{W}_2(\rho_t, \rho_0)$ reaches \widetilde{W}_0 in the update, we pause and reset ρ_t to ρ_0 by resampling the parameters from ρ_0 . Then, we reset the critic with the newly sampled parameters while keeping the actor's policy π_t unchanged and continue the update.

The restarting mechanism guarantees $\widetilde{W}_2(\rho_t, \rho_0) \leq \lambda\bar{D}$ by restricting the distribution ρ_t of the parameters to be close to ρ_0 . Moreover, by letting $\lambda \geq 3$, we ensure that ρ_{π_t} is realizable by ρ_t since $\widetilde{W}_2(\rho_{\pi_t}, \rho_0) \leq \bar{D} \leq \lambda\bar{D}$, which means that the neural network is capable of capturing the representation of the action value function Q^{π_t} . We remark that by letting $\widetilde{W}_0 = O(1)$, we allow ρ_t to deviate from ρ_0 up to $W_2(\rho_t, \rho_0) \leq O(\alpha^{-1})$ in the restarting mechanism. In contrast, the NTK regime [15] which corresponds to letting $\alpha = \sqrt{M}$ in (3.1) only allows ρ_t to deviate from ρ_0 by the chi-squared divergence $\chi^2(\rho_t \| \rho_0) \leq O(M^{-1}) = o(1)$. That is, the NTK regime fails to induce a feature representation significantly different from the initial one. Before moving on, we summarize the construction of the weighting distribution $\widetilde{\Phi}^{\pi_t}$ in Theorem 4.1 and 4.5 as follows,

$$\widetilde{\Phi}^{\pi_t} = \widetilde{\mathcal{E}}_{\widetilde{\phi}^{\pi_t}}^{\pi_t}, \quad \widetilde{\phi}^{\pi_t} = \frac{1}{2}\widetilde{\phi}_0 + \frac{1}{2}\phi_0 \otimes \pi_t, \quad \phi_0 = \int_{\mathcal{A}} \widetilde{\phi}_0(\cdot, da), \quad (4.8)$$

where $\widetilde{\phi}_0$ is the base distribution. Now we have the following theorem that characterizes the global optimality and convergence of the two-timescale AC with restarting mechanism.

Theorem 4.6 (Global Optimality and Convergence Rate of Two-timescale AC with Restarting Mechanism). Suppose that (4.8) and Assumptions 4.2 and 4.3 hold. With the restarting mechanism, it holds that

$$\frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt \leq \underbrace{\frac{\zeta}{T}}_{(a)} + 4\kappa \underbrace{\sqrt{\alpha^{-1}S_1 + \alpha^{1/2}\eta^{-1}S_2 + \frac{\eta^{-1}\bar{D}^2}{2T(1-\sqrt{\gamma})}}}_{(b)}, \quad (4.9)$$

where we have

$$\zeta = \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^*} \left[\text{KL}(\pi^*(\cdot | s) \| \pi_0(\cdot | s)) \right], \quad \kappa = \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}_0} \right\|_{\infty},$$

$$S_1 = \frac{(1 + \lambda)^2 \bar{D}^2 B_2 (4B_1 \lambda \bar{D} + B_r)}{1 - \sqrt{\gamma}}, \quad S_2 = \frac{2\mathcal{B}B_1 \lambda (1 + \lambda) \bar{D}^2}{1 - \sqrt{\gamma}}.$$

Here B_r , B_1 and B_2 are defined in Assumption 4.2 and 4.3, \bar{D} is the upper bound for $\tilde{W}_2(\rho_\pi, \rho_0)$ in Lemma 4.4, \mathcal{B} depends on the discount factor γ and the absolute constants defined in Assumption 4.2 and 4.3, and λ is the scaling parameter for the restarting threshold. Besides, it holds for the total restarting number N that

$$N \leq (\lambda - 2)^{-1} ((\alpha^{-1}\eta S_1 + \alpha^{1/2}S_2)2T\bar{D}^{-2}(1 - \sqrt{\gamma}) + 1).$$

Proof. See §B.4 for a detailed proof. □

Note that for a given MDP with starting distribution \mathcal{D}_0 , the expected KL-divergence ζ and the concentrability coefficient κ are both independent of the two-timescale update. We remark that our condition for (4.9) to be bounded is not restrictive. Specifically, we only need a given π_0 and $\tilde{\phi}_0$ such that the KL-divergence $\zeta < \infty$ and the concentrability coefficient $\kappa < \infty$, which is weaker than the concentrability coefficient used in [7, 23, 24, 38, 39, 43, 48, 57, 61].

The first term (a) on the right-hand side of (4.9) diminishes as $T \rightarrow \infty$. The second term (b) corresponds to the policy evaluation error. We give an example to demonstrate the convergence of the two-time AC. We let the scaling parameter $\lambda = 3$ for the restarting threshold. By letting $\eta = \alpha^{3/2}$, it holds that (b) = $O(\alpha^{-1/2})$ as $\alpha \rightarrow \infty$. Thus, we have that $\frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt$ descends at a rate of $O(T^{-1} + O(\alpha^{-1/2}) + O(\alpha^{-3/4}T^{-1/2}))$. Note that $\eta = \alpha^{3/2}$ shows that the critic has a larger relative TD timescale in (3.5). As for the total number of restartings N , it holds that $N \leq O(\alpha^{1/2}T)$ as $\alpha \rightarrow \infty$, which induces a tradeoff, i.e., a larger α guarantees a smaller gap in $\frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt$ but yields in more restartings and a larger relative TD timescale.

5 Acknowledgement

Zhaoran Wang acknowledges National Science Foundation (Awards 2048075, 2008827, 2015568, 1934931), Simons Institute (Theory of Reinforcement Learning), Amazon, J.P. Morgan, and Two Sigma for their supports. Zhuoran Yang acknowledges Simons Institute (Theory of Reinforcement Learning). This work was also supported in part by the Vannevar Bush Faculty Fellowship program under grant number N00014-21-1-2941.

References

- [1] Agazzi, A. and Lu, J. (2019). Temporal-difference learning for nonlinear value function approximation in the lazy training regime. *arXiv preprint arXiv:1905.10917*.
- [2] Agazzi, A. and Lu, J. (2020). Global optimality of softmax policy gradient with single hidden layer neural networks in the mean-field regime. *arXiv preprint arXiv:2010.11858*.

- [3] Agostinelli, F., McAleer, S., Shmakov, A. and Baldi, P. (2019). Solving the rubik’s cube with deep reinforcement learning and search. *Nature Machine Intelligence*, **1** 356–363.
- [4] Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R. et al. (2019). Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*.
- [5] Ambrosio, L. and Gigli, N. (2013). A user’s guide to optimal transport. In *Modelling and Optimisation of Flows on Networks*. Springer, 1–155.
- [6] Ambrosio, L., Gigli, N. and Savaré, G. (2008). *Gradient flows: In metric spaces and in the space of probability measures*. Springer.
- [7] 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.
- [8] Barron, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, **39** 930–945.
- [9] Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C. et al. (2019). Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- [10] Bhatnagar, S., Ghavamzadeh, M., Lee, M. and Sutton, R. S. (2008). Incremental natural actor-critic algorithms. In *Advances in Neural Information Processing Systems*.
- [11] Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M. and Lee, M. (2009). Natural actor-critic algorithms. *Automatica*, **45** 2471–2482.
- [12] Börgers, T. and Sarin, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of economic theory*, **77** 1–14.
- [13] Borkar, V. S. (2009). *Stochastic approximation: A dynamical systems viewpoint*. Springer.
- [14] Cai, Q., Yang, Z., Jin, C. and Wang, Z. (2019). Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*.
- [15] Cai, Q., Yang, Z., Lee, J. D. and Wang, Z. (2019). Neural temporal-difference learning converges to global optima. In *Advances in Neural Information Processing Systems*.
- [16] Chen, M., Wang, Y., Liu, T., Yang, Z., Li, X., Wang, Z. and Zhao, T. (2020). On computation and generalization of generative adversarial imitation learning. *arXiv preprint arXiv:2001.02792*.
- [17] Chen, Z., Cao, Y., Gu, Q. and Zhang, T. (2020). Mean-field analysis of two-layer neural networks: Non-asymptotic rates and generalization bounds. *arXiv preprint arXiv:2002.04026*.
- [18] Chizat, L. and Bach, F. (2018). A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*.
- [19] Chizat, L. and Bach, F. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems*.
- [20] 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*.
- [21] Fang, C., Dong, H. and Zhang, T. (2019). Over parameterized two-level neural networks can learn near optimal feature representations. *arXiv preprint arXiv:1910.11508*.

- [22] Fang, C., Gu, Y., Zhang, W. and Zhang, T. (2019). Convex formulation of overparameterized deep neural networks. *arXiv preprint arXiv:1911.07626*.
- [23] Farahmand, A.-m., Ghavamzadeh, M., Szepesvári, C. and Mannor, S. (2016). Regularized policy iteration with nonparametric function spaces. *The Journal of Machine Learning Research*, **17** 4809–4874.
- [24] Farahmand, A.-m., Szepesvári, C. and Munos, R. (2010). Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*.
- [25] Friedrichs, K. O. (1944). The identity of weak and strong extensions of differential operators. *Transactions of the American Mathematical Society*, **55** 132–151.
- [26] Fu, Z., Yang, Z. and Wang, Z. (2020). Single-timescale actor-critic provably finds globally optimal policy. *arXiv preprint arXiv:2008.00483*.
- [27] Harker, P. T. and Pang, J.-S. (1990). Finite-dimensional variational inequality and nonlinear complementarity problems: A survey of theory, algorithms and applications. *Mathematical Programming*, **48** 161–220.
- [28] Hennes, D., Morrill, D., Omidshafiei, S., Munos, R., Perolat, J., Lanctot, M., Gruslys, A., Lespiau, J.-B., Parmas, P., Duéñez-Guzmán, E. et al. (2020). Neural replicator dynamics: Multi-agent learning via hedging policy gradients. In *International Conference on Autonomous Agents and MultiAgent Systems*.
- [29] Hong, M., Wai, H.-T., Wang, Z. and Yang, Z. (2020). A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic. *arXiv preprint arXiv:2007.05170*.
- [30] Jacot, A., Gabriel, F. and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*.
- [31] Javanmard, A., Mondelli, M. and Montanari, A. (2019). Analysis of a two-layer neural network via displacement convexity. *arXiv preprint arXiv:1901.01375*.
- [32] Jin, C., Yang, Z., Wang, Z. and Jordan, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*.
- [33] Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*, vol. 2.
- [34] Kakade, S. M. (2002). A natural policy gradient. In *Advances in Neural Information Processing Systems*.
- [35] Khodadadian, S., Doan, T. T., Maguluri, S. T. and Romberg, J. (2021). Finite sample analysis of two-time-scale natural actor-critic algorithm. *arXiv preprint arXiv:2101.10506*.
- [36] Konda, V. R. and Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *Advances in Neural Information Processing Systems*.
- [37] Kushner, H. and Yin, G. G. (2003). *Stochastic approximation and recursive algorithms and applications*. Springer.
- [38] Lazaric, A., Ghavamzadeh, M. and Munos, R. (2016). Analysis of classification-based policy iteration algorithms. *The Journal of Machine Learning Research*, **17** 583–612.
- [39] 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*.

- [40] Lu, Y., Ma, C., Lu, Y., Lu, J. and Ying, L. (2020). A mean field analysis of deep resnet and beyond: Towards provably optimization via overparameterization from depth. In *International Conference on Machine Learning*. PMLR.
- [41] Mei, S., Misiakiewicz, T. and Montanari, A. (2019). Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*. PMLR.
- [42] Mei, S., Montanari, A. and Nguyen, P.-M. (2018). A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, **115** E7665–E7671.
- [43] Munos, R. and Szepesvári, C. (2008). Finite-time bounds for fitted value iteration. *The Journal of Machine Learning Research*, **9** 815–857.
- [44] Otto, F. and Villani, C. (2000). Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, **173** 361–400.
- [45] Peters, J. and Schaal, S. (2008). Natural actor-critic. *Neurocomputing*, **71** 1180–1190.
- [46] Peyre, R. (2011). Comparison between w_2 distance and \dot{H}^{-1} norm, and localisation of wasserstein distance. *arXiv preprint arXiv:1104.4631*.
- [47] Pinkus, A. (1999). Approximation theory of the MLP model in neural networks. *Acta Numerica*, **8** 143–195.
- [48] Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B. and Geist, M. (2015). Approximate modified policy iteration and its application to the game of Tetris. *The Journal of Machine Learning Research*, **16** 1629–1676.
- [49] Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- [50] Schuster, P. and Sigmund, K. (1983). Replicator dynamics. *Journal of theoretical biology*, **100** 533–538.
- [51] 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** 484–489.
- [52] 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.
- [53] Sirignano, J. and Spiliopoulos, K. (2020). Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, **130** 1820–1852.
- [54] Sirignano, J. and Spiliopoulos, K. (2020). Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, **80** 725–752.
- [55] Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, **3** 9–44.
- [56] Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- [57] Szepesvári, C. and Munos, R. (2005). Finite time bounds for sampling based fitted value iteration. In *International Conference on Machine Learning*. ACM.
- [58] Villani, C. (2003). *Topics in optimal transportation*. American Mathematical Society.
- [59] Villani, C. (2008). *Optimal transport: Old and new*. Springer.

- [60] Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P. et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, **575** 350–354.
- [61] Wang, L., Cai, Q., Yang, Z. and Wang, Z. (2019). Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*.
- [62] Watkins, C. and Dayan, P. (1992). Q-learning. *Machine Learning*, **8** 279–292.
- [63] Wei, C., Lee, J. D., Liu, Q. and Ma, T. (2019). Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*.
- [64] Wu, Y., Zhang, W., Xu, P. and Gu, Q. (2020). A finite time analysis of two time-scale actor critic methods. *arXiv preprint arXiv:2005.01350*.
- [65] Xu, T., Wang, Z. and Liang, Y. (2020). Improving sample complexity bounds for actor-critic algorithms. *arXiv preprint arXiv:2004.12956*.
- [66] Xu, T., Wang, Z. and Liang, Y. (2020). Non-asymptotic convergence analysis of two time-scale (natural) actor-critic algorithms. *arXiv preprint arXiv:2005.03557*.
- [67] Yang, L. F. and Wang, M. (2019). Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. *arXiv preprint arXiv:1905.10389*.
- [68] Yang, L. F. and Wang, M. (2019). Sample-optimal parametric q-learning using linearly additive features. *arXiv preprint arXiv:1902.04779*.
- [69] Zhang, Y., Cai, Q., Yang, Z., Chen, Y. and Wang, Z. (2020). Can temporal-difference and q-learning learn representation? a mean-field theory. *arXiv preprint arXiv:2006.04761*.

A Supplement to the Background

In this section, we present some backgrounds on Wasserstein space and replicator dynamics.

A.1 Wasserstein Space

Let $\Theta \subseteq \mathbb{R}^D$ be a Polish space. We denote by $\mathcal{P}_2(\Theta) \subseteq \mathcal{P}(\Theta)$ the set of probability measures with finite second moments. Then, the Wasserstein-2 (W_2) distance between $\mu, \nu \in \mathcal{P}_2(\Theta)$ is defined as follows,

$$W_2(\mu, \nu) = \inf \left\{ \mathbb{E}[\|X - Y\|^2]^{1/2} \mid \text{law}(X) = \mu, \text{law}(Y) = \nu \right\}, \quad (\text{A.1})$$

where the infimum is taken over the random variables X and Y on Θ and we denote by $\text{law}(X)$ the distribution of a random variable X . We call $\mathcal{M} = (\mathcal{P}_2(\Theta), W_2)$ the Wasserstein space, which is an infinite-dimensional manifold [59]. In particular, we define the tangent vector at $\mu \in \mathcal{M}$ as $\dot{\rho}$ for the corresponding curve $\rho : [0, 1] \rightarrow \mathcal{P}_2(\Theta)$ with $\rho_0 = \mu$. Under certain regularity conditions, the continuity equation $\partial_t \rho_t = -\text{div}(\rho_t v_t)$ corresponds to a vector field $v : [0, 1] \times \Theta \rightarrow \mathbb{R}^D$, which endows the infinite-dimensional manifold $\mathcal{P}_2(\Theta)$ with a weak Riemannian structure in the following sense [59]. Given any tangent vectors u and \tilde{u} at $\mu \in \mathcal{M}$ and the corresponding vector fields v, \tilde{v} , which satisfy $u + \text{div}(\mu v) = 0$ and $\tilde{u} + \text{div}(\mu \tilde{v}) = 0$, respectively, we define the inner product of u and \tilde{u} as follows,

$$\langle u, \tilde{u} \rangle_{\mu, W_2} = \int v \cdot \tilde{v} d\mu = \langle v, \tilde{v} \rangle_{\mu}, \quad (\text{A.2})$$

which yields a Riemannian metric. Such a Riemannian metric further induces a norm $\|u\|_{\mu, W_2} = \langle u, u \rangle_{\mu, W_2}^{1/2}$ for any tangent vector $u \in T_{\mu} \mathcal{M}$ at any $\mu \in \mathcal{M}$, which allows us to write the Wasserstein-2 distance defined in (A.1) as follows,

$$W_2(\mu, \nu) = \inf \left\{ \left(\int_0^1 \|\dot{\rho}_t\|_{\rho_t, W_2}^2 dt \right)^{1/2} \mid \rho : [0, 1] \rightarrow \mathcal{M}, \rho_0 = \mu, \rho_1 = \nu \right\}. \quad (\text{A.3})$$

Here $\dot{\rho}_s$ denotes $\partial_t \rho_t \big|_{t=s}$ for any $s \in [0, 1]$. In particular, the infimum in (A.3) is attained by the geodesic $\tilde{\rho} : [0, 1] \rightarrow \mathcal{P}_2(\Theta)$ connecting $\mu, \nu \in \mathcal{M}$. Moreover, the geodesics on \mathcal{M} are constant-speed, that is,

$$\|\dot{\tilde{\rho}}_t\|_{\tilde{\rho}_t, W_2} = W_2(\mu, \nu), \quad \forall t \in [0, 1]. \quad (\text{A.4})$$

In Wasserstein space \mathcal{M} , a curve $\rho : [0, 1] \rightarrow \mathcal{P}_2(\Theta)$ is defined to be absolutely continuous if there exists $m \in L^1(a, b)$, i.e., $\int_a^b |\dot{m}(t)| dt < \infty$, such that

$$W_2(\rho_s, \rho_t) \leq \int_s^t m(r) dr, \quad \forall a < s \leq t < b.$$

Such an absolutely continuous curve ρ_t allows us to define the metric derivative in \mathcal{M} as follows,

$$|\dot{\rho}_t|_{W_2} = \lim_{s \rightarrow t} \frac{W_2(\rho_s, \rho_t)}{|s - t|}. \quad (\text{A.5})$$

By [6], the metric derivative $|\dot{\rho}_t|_{W_2}$ is connected to the norm of the tangent vector by

$$|\dot{\rho}_t|_{W_2} = \|\dot{\rho}_t\|_{\rho_t, W_2}. \quad (\text{A.6})$$

Furthermore, we introduce the Wasserstein-1 distance, which is defined as

$$W_1(\mu^1, \mu^2) = \inf \left\{ \mathbb{E}[\|X - Y\|] \mid \text{law}(X) = \mu^1, \text{law}(Y) = \mu^2 \right\}$$

for any $\mu^1, \mu^2 \in \mathcal{P}(\mathbb{R}^D)$ with finite first moments. The Wasserstein-1 distance has the following dual representation [6],

$$W_1(\mu^1, \mu^2) = \sup \left\{ \int f(x) d(\mu^1 - \mu^2)(x) \mid \text{continuous } f : \mathbb{R}^D \rightarrow \mathbb{R}, \text{Lip}(f) \leq 1 \right\}. \quad (\text{A.7})$$

A.2 Replicator Dynamics

The replicator dynamics originally arises in the study of evolutionary game theory [50]. For a function f , the replicator dynamics is given by the differential equation

$$\frac{d}{dt}x_t(a) = x_t(a)[f(a, x_t) - \phi(x)],$$

where $\phi(x) = \int x(a)f(a, x)$. As for the PPO update in (3.7), for a fixed s , let $x(a) = \pi(a | s)$ and $f(a, x) = Q^\pi(s, a)$, we see that (3.7) corresponds to a replicator dynamics if $Q_t = Q^{\pi_t}$. Note that in the simultaneous update of both the critic and actor, we do not have access to the true action value function Q^π . Thus, we use the estimator Q_t calculated by the critic step to guide the update of the actor in the PPO update, which takes the form of a replicator dynamics in the continuous-time limit.

B Proofs of Main Results

In this section, we give detailed proof of the theorems and present a detailed statement of Lemma 4.4.

B.1 Proof of Theorem 4.1

Proof. Following from the performance difference lemma [33], we have

$$J(\pi^*) - J(\pi_t) = (1 - \gamma)^{-1} \cdot \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} [\langle A^{\pi_t}(s, \cdot), \pi^*(\cdot | s) - \pi_t(\cdot | s) \rangle_{\mathcal{A}}],$$

where $\mathcal{E}_{\mathcal{D}_0}^{\pi^*}$ is the visitation measure induced by π^* from \mathcal{D}_0 and A^{π_t} is the advantage function. Note that the continuous PPO dynamics in (3.7) can be equivalently written as $\partial_t \log \pi_t = A_t$. Thus, we have

$$\begin{aligned} (1 - \gamma) \cdot (J(\pi^*) - J(\pi_t)) &= \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\left\langle \frac{d}{dt} \log \pi_t(\cdot | s) + A^{\pi_t}(s, \cdot) - A_t(s, \cdot), \pi^*(\cdot | s) - \pi_t(\cdot | s) \right\rangle_{\mathcal{A}} \right] \\ &= \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\left\langle \frac{d}{dt} \log \pi_t(\cdot | s), \pi^*(\cdot | s) - \pi_t(\cdot | s) \right\rangle_{\mathcal{A}} \right] \\ &\quad + \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\langle Q^{\pi_t}(s, \cdot) - Q_t(s, \cdot), \pi^*(\cdot | s) - \pi_t(\cdot | s) \rangle_{\mathcal{A}} \right]. \end{aligned} \quad (\text{B.1})$$

For the first term on the right-hand side of (B.1), it holds that,

$$\begin{aligned} &\left\langle \frac{d}{dt} \log \pi_t(\cdot | s), \pi^*(\cdot | s) - \pi_t(\cdot | s) \right\rangle_{\mathcal{A}} \\ &= \left\langle \frac{d}{dt} \log \pi_t(\cdot | s), \pi^*(\cdot | s) \right\rangle_{\mathcal{A}} - \left\langle \frac{d}{dt} \log \pi_t(\cdot | s), \pi_t(\cdot | s) \right\rangle_{\mathcal{A}} \\ &= -\frac{d}{dt} \text{KL}(\pi^*(\cdot | s) \| \pi_t(\cdot | s)), \end{aligned} \quad (\text{B.2})$$

where the last equality holds by noting that $\langle \partial_t \log \pi_t(\cdot | s), \pi_t(\cdot | s) \rangle_{\mathcal{A}} = \partial_t \int_{\mathcal{A}} \pi_t(da | s) = 0$. For the second term on the right-hand side of (B.1), by the Cauchy-Schwartz inequality, we have

$$\begin{aligned} \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\langle Q^{\pi_t}(s, \cdot) - Q_t(s, \cdot), \pi^*(\cdot | s) \rangle_{\mathcal{A}} \right] &= \mathbb{E}_{(s,a) \sim \tilde{\phi}^{\pi_t}} \left[(Q^{\pi_t}(s, a) - Q_t(s, a)) \pi^*(a | s) \frac{\mathcal{E}_{\mathcal{D}_0}^{\pi^*}(s)}{\tilde{\phi}^{\pi_t}(x)} \right] \\ &\leq \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^{\pi^*}}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} \cdot \|Q^{\pi_t} - Q_t\|_{2, \tilde{\phi}^{\pi_t}}, \end{aligned} \quad (\text{B.3})$$

$$\begin{aligned} \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^{\pi^*}} \left[\langle Q^{\pi_t}(s, \cdot) - Q_t(s, \cdot), \pi_t(\cdot | s) \rangle_{\mathcal{A}} \right] &= \mathbb{E}_{(x) \sim \tilde{\phi}^{\pi_t}} \left[(Q^{\pi_t}(s, \cdot) - Q_t(s, \cdot)) \pi_t(a | s) \frac{\mathcal{E}_{\mathcal{D}_0}^{\pi^*}(s)}{\tilde{\phi}^{\pi_t}(x)} \right] \\ &\leq \left\| \frac{\mathcal{E}_{\mathcal{D}_0}^{\pi^*} \otimes \pi_t}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} \cdot \|Q^{\pi_t} - Q_t\|_{2, \tilde{\phi}^{\pi_t}}. \end{aligned} \quad (\text{B.4})$$

Plugging (B.2), (B.3), and (B.4) into (B.1), we have

$$J(\pi^*) - J(\pi_t) \leq -\frac{d}{dt} \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^*} \left[\text{KL}(\pi^*(\cdot | s) \| \pi_t(\cdot | s)) \right] \\ + \left(\left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} + \left\| \frac{\mathcal{E}_{\mathcal{D}_0}^* \otimes \pi_t}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} \right) \cdot \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}. \quad (\text{B.5})$$

By further letting $\tilde{\phi}^{\pi_t} = \frac{1}{2}\tilde{\phi}_0 + \frac{1}{2}\phi_0 \otimes \pi_t$, it holds for the concentrability coefficient $\left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}}$ in (B.3) that

$$\left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} \leq \left\| \frac{2\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}_0} \right\|_{2, \tilde{\phi}^{\pi_t}} \leq 2 \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}_0} \right\|_{\infty}. \quad (\text{B.6})$$

By further letting $\phi_0(s) = \int_{\mathcal{A}} \tilde{\phi}_0(s, da)$, it holds for the concentrability coefficient $\left\| \frac{\mathcal{E}_{\mathcal{D}_0}^* \otimes \pi_t}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}}$ in (B.4) that

$$\left\| \frac{\mathcal{E}_{\mathcal{D}_0}^* \otimes \pi_t}{\tilde{\phi}^{\pi_t}} \right\|_{2, \tilde{\phi}^{\pi_t}} \leq \left\| \frac{2\mathcal{E}_{\mathcal{D}_0}^*}{\phi_0} \right\|_{2, \tilde{\phi}^{\pi_t}} \leq 2 \left\| \frac{\int \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*(x)}{\tilde{\phi}_0(x)} \tilde{\phi}_0(s, da)}{\phi_0(s)} \right\|_{\infty} \leq 2 \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}_0} \right\|_{\infty}. \quad (\text{B.7})$$

Plugging (B.6) and (B.7) into (B.5) and taking integration on both sides of (B.5), we have

$$\frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt \leq \frac{1}{T} \mathbb{E}_{s \sim \mathcal{E}_{\mathcal{D}_0}^*} \left[\text{KL}(\pi^*(\cdot | s) \| \pi_0(\cdot | s)) \right] + \frac{4\kappa}{T} \cdot \int_0^T \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}} dt,$$

where $\kappa = \left\| \frac{\tilde{\mathcal{E}}_{\mathcal{D}_0}^*}{\tilde{\phi}_0} \right\|_{\infty}$. Thus, we complete the proof of Theorem 4.1. \square

B.2 Detailed Statement of Lemma 4.4

We give a detailed version of Lemma 4.4 as follows.

Lemma B.1 (Regularity of Representation of Q^π). Suppose that Assumptions 4.2 and 4.3 hold. For any policy π , there exists a probability measure $\rho_\pi \in \mathcal{P}_2(\mathbb{R}^D)$ for the representation of Q^π with the following properties.

- (i) For function $Q(s, a; \rho_\pi)$ defined by (3.4) with $\rho = \rho_\pi$ and the action value function $Q^\pi(s, a)$ defined by (2.2), we have $Q(s, a; \rho_\pi) = Q^\pi(s, a)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.
- (ii) For g defined in (3.6), we have $g(\cdot; \rho_\pi) = 0$ for any policy π .
- (iii) By letting $B_\beta \geq 2(B_r + \gamma(1 - \gamma)^{-1}B_r M_{1, \varphi})$ for the neural network defined in (4.2) and $\rho_0 \sim \mathcal{N}(0, I_D)$ for the initial distribution, we have $\tilde{W}_2(\rho_\pi, \rho_0) < \bar{D}$ for any policy π , where we define $\tilde{W}(\cdot, \cdot) = \alpha W(\cdot, \cdot)$ as the scaled W_2 metric. Here constant \bar{D} depends on the discount factor γ and the absolute constants $L_{0, \beta}, L_{1, \beta}, l_\beta, B_r, M_\mu, M_\psi, M_{1, \varphi}, M_{2, \varphi}$ defined in Assumptions 4.2 and 4.3.
- (iv) For any two policies π_1 and π_2 , it holds that,

$$W_2(\rho_{\pi_1}, \rho_{\pi_2}) \leq \alpha^{-\frac{1}{2}} \mathcal{B} \cdot \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\|\pi_1(\cdot | s') - \pi_2(\cdot | s')\|_1 \right],$$

where \mathcal{B} depends on the discount factor γ , the scaling parameter B_β in (4.2), and the absolute constants $l_\beta, B_r, M_{1, \varphi}, \mathcal{G}$ defined in Assumptions 4.2 and 4.3.

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

B.3 Proof of Theorem 4.5

Proof. For notation simplicity, we let $x = (s, a)$. By Property (i) of Lemma B.1, it holds that $\|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}^2 = \|Q_t - Q(\cdot; \rho_{\pi_t})\|_{2, \tilde{\phi}^{\pi_t}}^2$, where $Q_t = Q(\cdot; \rho_t)$. Thus it prompts us to study the W_2 distance between ρ_t and ρ_{π_t} . By the first variation formula in Lemma E.2, it holds that

$$\frac{d}{dt} \frac{W_2^2(\rho_t, \rho_{\pi_t})}{2} = -\langle \dot{\rho}_t, \tilde{\alpha}_t^0 \rangle_{\rho_t, W_2} - \langle \dot{\rho}_{\pi_t}, \tilde{\beta}_t^0 \rangle_{\rho_{\pi_t}, W_2}. \quad (\text{B.8})$$

Here $\tilde{\alpha}_t^{[0,1]}$ is the geodesic connecting ρ_t and ρ_{π_t} , and $\tilde{\beta}_t^{[0,1]}$ is its time-inverse, i.e., $\tilde{\beta}_t^s = \tilde{\alpha}_t^{1-s}$. Besides, we denote by $\tilde{\alpha}_t^s = \partial_s \tilde{\alpha}_t^s$ and $\tilde{\beta}_t^s = \partial_s \tilde{\beta}_t^s$ the derivation of geodesics $\tilde{\alpha}_t^s$ and $\tilde{\beta}_t^s$ with respect to s , respectively. We denote by v_s the corresponding vector field at $\tilde{\alpha}_t^s$, which satisfies $\partial_s \tilde{\alpha}_t^s = -\text{div}(\tilde{\alpha}_t^s v_s)$. For the first term of (B.8), it holds that

$$\begin{aligned} -\langle \dot{\rho}_t, \tilde{\alpha}_t^0 \rangle_{\rho_t, W_2} &= -\eta \cdot \langle g(\cdot; \rho_t, \pi_t), v_0 \rangle_{\rho_t} \\ &= \eta \cdot \int_0^1 \partial_s \langle g(\cdot; \tilde{\alpha}_t^s, \pi_t), v_s \rangle_{\tilde{\alpha}_t^s} ds - \eta \cdot \langle g(\cdot; \rho_{\pi_t}, \pi_t), v_1 \rangle_{\rho_{\pi_t}} \\ &= \eta \cdot \int_0^1 \langle \partial_s g(\cdot; \tilde{\alpha}_t^s, \pi_t), v_s \rangle_{\tilde{\alpha}_t^s} ds + \eta \cdot \int_0^1 \int g(\theta; \tilde{\alpha}_t^s, \pi_t) \cdot \partial_s (v_s \tilde{\alpha}_t^s)(\theta) d\theta ds, \end{aligned} \quad (\text{B.9})$$

where the first equality follows from (A.2) and the third equation follows from $g(\cdot; \rho_{\pi_t}, \pi_t) = 0$ by Property (ii) of Lemma B.1. For the first term on the right-hand side of (B.9), we have

$$\begin{aligned} &\eta \cdot \int_0^1 \langle \partial_s g(\cdot; \tilde{\alpha}_t^s, \pi_t), v_s \rangle_{\tilde{\alpha}_t^s} ds \\ &= -\alpha^{-1} \eta \cdot \int_0^1 \int \left\langle \mathbb{E}_{\tilde{\Phi}^{\pi_t}}^{\pi_t} \left[\partial_s (Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \nabla \sigma(x; \theta) \right], (v_s \tilde{\alpha}_t^s)(\theta) \right\rangle d\theta ds \\ &= \alpha^{-1} \eta \cdot \int_0^1 \int \left\langle \mathbb{E}_{\tilde{\Phi}^{\pi_t}}^{\pi_t} \left[\partial_s (Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \sigma(x; \theta) \right], \text{div}(v_s \tilde{\alpha}_t^s)(\theta) \right\rangle d\theta ds, \end{aligned} \quad (\text{B.10})$$

where the first equality holds by definition of g in (3.6) and the last equality follows from Stokes' formula. Note that we have $\text{div}(v_s \tilde{\alpha}_t^s) = -\partial_s \tilde{\alpha}_t^s$ by the definition of vector field v_s . Thus, it holds for (B.10) that

$$\begin{aligned} &\eta \cdot \int_0^1 \langle \partial_s g(\cdot; \tilde{\alpha}_t^s, \pi_t), v_s \rangle_{\tilde{\alpha}_t^s} ds \\ &= -\alpha^{-1} \eta \cdot \int_0^1 \int \left\langle \mathbb{E}_{\tilde{\Phi}^{\pi_t}}^{\pi_t} \left[\partial_s (Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \sigma(x; \theta) \right], \partial_s \tilde{\alpha}_t^s(\theta) \right\rangle d\theta ds \\ &= -\alpha^{-2} \eta \cdot \int_0^1 \mathbb{E}_{\tilde{\Phi}^{\pi_t}}^{\pi_t} \left[\partial_s (Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \partial_s Q(x; \tilde{\alpha}_t^s) \right] ds, \end{aligned} \quad (\text{B.11})$$

where the last equality follows from the definition of Q in (3.4). We let $f(\tilde{D}) = (\mathbb{E}_{x \sim \tilde{\mathcal{E}}_{\tilde{D}}^{\pi_t}} [(\partial_s Q(x; \tilde{\alpha}_t^s))^2])^{1/2}$ with respect to a specific s and t . Recall that for the weighting distribution $\tilde{\Phi}^{\pi_t}$, we set $\tilde{\Phi}^{\pi_t} = \tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t}$. Hence, for the integrand of (B.11), we have

$$\begin{aligned} &-\mathbb{E}_{\tilde{\Phi}^{\pi_t}}^{\pi_t} \left[\partial_s (Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \partial_s Q(x; \tilde{\alpha}_t^s) \right] \\ &= -f(\tilde{\phi}^{\pi_t})^2 + \gamma \cdot \int \partial_s Q(x; \tilde{\alpha}_t^s) \partial_s Q(x'; \tilde{\alpha}_t^s) \tilde{P}^{\pi_t}(x' | x) \tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t}(x) dx \\ &\leq -f(\tilde{\phi}^{\pi_t})^2 + \gamma \cdot \sqrt{\int (\partial_s Q(x; \tilde{\alpha}_t^s))^2 \tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t}(dx)} \cdot \sqrt{\int (\partial_s Q(x'; \tilde{\alpha}_t^s))^2 \tilde{P}^{\pi_t}(dx' | x) \tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t}(dx)}, \end{aligned} \quad (\text{B.12})$$

where the equality follows from the definition of $\mathbb{E}_{\tilde{\phi}^{\pi_t}}$ in (3.6) and the inequality follows from the Cauchy-Schwarz inequality. We define $\mathcal{T}^\pi : \mathcal{P}(\mathcal{S} \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S} \times \mathcal{A})$ as a mapping operator such that $\mathcal{T}^\pi \tilde{\mathcal{D}}(x') = \int \tilde{\mathcal{D}}(x) \tilde{P}^\pi(x'|dx)$. We rewrite (B.12) as

$$-\mathbb{E}_{\tilde{\phi}^{\pi_t}}^{\pi_t} [\partial_s(Q(x; \tilde{\alpha}_t^s) - \gamma \cdot Q(x'; \tilde{\alpha}_t^s)) \partial_s Q(x; \tilde{\alpha}_t^s)] \leq -f(\tilde{\phi}^{\pi_t})^2 + \gamma \cdot f(\tilde{\phi}^{\pi_t}) \cdot f(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t}). \quad (\text{B.13})$$

By the definition of \mathcal{T}^{π_t} and the definition of visitation measure in (2.1), it holds that $\tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t} - \gamma \tilde{\mathcal{E}}_{\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t}}^{\pi_t} = (1 - \gamma) \tilde{\phi}^{\pi_t}$. Hence, we have $f(\tilde{\phi}^{\pi_t})^2 - \gamma f^2(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t}) = (1 - \gamma) \mathbb{E}_{\tilde{\phi}^{\pi_t}} [(\partial_s Q(x; \tilde{\alpha}_t^s))^2]$ and it holds for (B.13) that

$$\begin{aligned} -f(\tilde{\phi}^{\pi_t})^2 + \gamma f(\tilde{\phi}^{\pi_t}) \cdot f(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t}) &= -\frac{f(\tilde{\phi}^{\pi_t})}{f(\tilde{\phi}^{\pi_t}) + \gamma f(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t})} \cdot (f(\tilde{\phi}^{\pi_t})^2 - \gamma^2 f^2(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t})) \\ &\leq -\frac{f(\tilde{\phi}^{\pi_t})^2 - \gamma \left(f(\tilde{\phi}^{\pi_t})^2 - (1 - \gamma) \mathbb{E}_{\tilde{\phi}^{\pi_t}} [(Q(x; \tilde{\alpha}_t^s))^2] \right)}{1 + \sqrt{\gamma}} \\ &\leq -(1 - \sqrt{\gamma}) \cdot \mathbb{E}_{\tilde{\phi}^{\pi_t}} [(\partial_s Q(x; \tilde{\alpha}_t^s))^2], \end{aligned} \quad (\text{B.14})$$

where the first inequality holds by noting that $f(\tilde{\phi}^{\pi_t}) \geq \sqrt{\gamma} f(\mathcal{T}^{\pi_t} \tilde{\phi}^{\pi_t})$ and the last inequality holds by noting that $f(\tilde{\phi}^{\pi_t})^2 \geq (1 - \gamma) \mathbb{E}_{\tilde{\phi}^{\pi_t}} [(\partial_s Q(x; \tilde{\alpha}_t^s))^2]$. Combining (B.11), (B.13), and (B.14) together, it holds for the first term of (B.9) that

$$\begin{aligned} \eta \cdot \int_0^1 \langle \partial_s g(\cdot; \tilde{\alpha}_t^s, \pi_t), v_s \rangle_{\tilde{\alpha}_t^s} ds &\leq -\alpha^{-2} \eta (1 - \sqrt{\gamma}) \cdot \int_0^1 \mathbb{E}_{\tilde{\phi}^{\pi_t}} [(\partial_s Q(x; \tilde{\alpha}_t^s))^2] ds \\ &\leq -\alpha^{-2} \eta (1 - \sqrt{\gamma}) \cdot \|Q(x; \rho_{\pi_t}) - Q(x; \rho_t)\|_{2, \tilde{\phi}^{\pi_t}}^2, \end{aligned} \quad (\text{B.15})$$

where the last inequality holds by the Cauchy-Schwarz inequality. For the second term of (B.9), we have

$$\begin{aligned} \eta \cdot \int_0^1 \int g(\theta; \tilde{\alpha}_t^s, \pi_t) \cdot \partial_s (v_s \tilde{\alpha}_t^s)(\theta) d\theta ds &= \eta \cdot \int_0^1 \int \langle \nabla g(\theta; \tilde{\alpha}_t^s, \pi_t), \tilde{\alpha}_t^s(\theta) \cdot v_s(\theta) \otimes v_s(\theta) \rangle d\theta ds \\ &\leq \eta \cdot \int_0^1 \sup_{\theta} \|\nabla g(\theta; \tilde{\alpha}_t^s, \pi_t)\|_{\text{F}} \cdot W_2(\rho_t, \rho_{\pi_t})^2 ds \\ &= \eta \cdot \sup_{\theta, s} \|\nabla g(\theta; \tilde{\alpha}_t^s, \pi_t)\|_{\text{F}} \cdot W_2(\rho_t, \rho_{\pi_t})^2, \end{aligned} \quad (\text{B.16})$$

where the first equality holds by Eulerian representation of geodesic in Lemma E.3 that $\partial_s (v_s \cdot \tilde{\alpha}_t^s) = -\text{div}(\tilde{\alpha}_t^s \cdot v_s \otimes v_s)$ and Stokes's formula. Here, we denote by \otimes the outer product between two vectors. The inequality of (B.16) follows from $\|v_s(\theta) \otimes v_s(\theta)\|_{\text{F}} = \|v_s(\theta)\|^2$ and the property of geodesic in (A.4) that $\|v_s\|_{\tilde{\alpha}_t^s, W_2} = W_2(\rho_t, \rho_{\pi_t})$.

For the Second term on the right-hand side of (B.8), we denote by u_t the corresponding vector field at ρ_{π_t} such that $\partial_t \rho_{\pi_t} = -\text{div}(\rho_{\pi_t} u_t)$. Then, it holds that

$$-\langle \dot{\rho}_{\pi_t}, \tilde{\beta}_t^0 \rangle_{\rho_{\pi_t}, W_2} = \langle u_t, v_1 \rangle_{\rho_{\pi_t}} \leq \|\dot{\rho}_{\pi_t}\|_{\rho_{\pi_t}, W_2} \cdot W_2(\rho_t, \rho_{\pi_t}) = |\dot{\rho}_{\pi_t}|_{W_2} \cdot W_2(\rho_t, \rho_{\pi_t}), \quad (\text{B.17})$$

where the first equality follows from (A.2), the inequality follows from the Cauchy-Schwarz inequality and the facts that $\|u_t\|_{\rho_{\pi_t}} = \|\dot{\rho}_{\pi_t}\|_{\rho_{\pi_t}, W_2}$ by (A.2) and that $\|v_1\|_{\rho_{\pi_t}} = W_2(\rho_t, \rho_{\pi_t})$ by (A.4). The last equality of (B.17) holds by (A.6). Plugging the definition of metric derivative $|\dot{\rho}_t|_{W_2}$ in (A.5) into (B.17), we have

$$\begin{aligned} -\langle \dot{\rho}_{\pi_t}, \tilde{\beta}_t^0 \rangle_{\rho_{\pi_t}, W_2} &\leq \lim_{\Delta t \rightarrow 0} \frac{W_2(\rho_{\pi_t}, \rho_{\pi_t + \Delta t})}{|\Delta t|} \cdot W_2(\rho_t, \rho_{\pi_t}) \\ &\leq \alpha^{-1/2} \mathcal{B} \cdot \lim_{\Delta t \rightarrow 0} \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_t}} \left[\left\| \frac{\pi_{t+\Delta t}(\cdot | s') - \pi_t(\cdot | s')}{\Delta t} \right\|_1 \right] \cdot W_2(\rho_t, \rho_{\pi_t}) \\ &= \alpha^{-1/2} \mathcal{B} \cdot \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_t}} \left[\left\| A_t(s', \cdot) \pi_t(\cdot | s') \right\|_1 \right] \cdot W_2(\rho_t, \rho_{\pi_t}), \end{aligned} \quad (\text{B.18})$$

where the second inequality follows from Property (iv) of Lemma B.1 and the equality follows from the MF-PPO update in (3.7). For the approximation of the advantage function A_t , it holds that

$$\begin{aligned} \sup_{x \in \mathcal{S} \times \mathcal{A}} |A_t(x)| &= \sup_{x \in \mathcal{S} \times \mathcal{A}} \left| Q(x; \rho_t) - \int Q(s, a') \pi_t(\mathrm{d}a' | s) \right| \\ &\leq 2 \sup_{x \in \mathcal{S} \times \mathcal{A}} |Q(x; \rho_t)|. \end{aligned} \quad (\text{B.19})$$

Plugging (B.19) into (B.18), we have

$$\begin{aligned} -\langle \dot{\rho}_{\pi_t}, \tilde{\beta}_t^0 \rangle_{\rho_{\pi_t}} &\leq \alpha^{-1/2} \mathcal{B} \cdot \sup_{x \in \mathcal{X}} |A_t(x)| \cdot \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}^{\pi_t}} \left[\|\pi_t(\cdot | s')\|_1 \right] \cdot W_2(\rho_t, \rho_{\pi_t}) \\ &\leq 2\alpha^{-1/2} \mathcal{B} \cdot \sup_{x \in \mathcal{X}} |Q(x; \rho_t)| \cdot W_2(\rho_t, \rho_{\pi_t}). \end{aligned} \quad (\text{B.20})$$

where the last inequality follows from $\|\pi_t(\cdot | s')\|_1 = 1$. By first plugging (B.15) and (B.16) into (B.9), and then plugging (B.9) and (B.20) into (B.8), we have

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} \frac{W_2^2(\rho_t, \rho_{\pi_t})}{2} &\leq -\alpha^{-2} \eta (1 - \sqrt{\gamma}) \cdot \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}^2 + \eta \cdot \sup_{\theta, s} \|\nabla g(\theta; \tilde{\alpha}_t^s, \pi_t)\|_{\mathbb{F}} \cdot W_2(\rho_t, \rho_{\pi_t})^2 \\ &\quad + 2\alpha^{-1/2} \mathcal{B} \cdot \sup_{x \in \mathcal{X}} |Q(x; \rho_t)| \cdot W_2(\rho_t, \rho_{\pi_t}). \end{aligned} \quad (\text{B.21})$$

Plugging Lemma D.1 into (B.21), we have

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} \frac{W_2^2(\rho_t, \rho_{\pi_t})}{2} &\leq -\eta \alpha^{-2} (1 - \sqrt{\gamma}) \cdot \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}^2 \\ &\quad + \eta \alpha^{-1} B_2 \cdot \left(2\alpha B_1 \sup_{s \in [0,1]} W_2(\tilde{\alpha}_t^s, \rho_0) + B_r \right) \cdot W_2(\rho_t, \rho_{\pi_t})^2 \\ &\quad + 2\alpha^{1/2} \mathcal{B} B_1 \cdot W_2(\rho_t, \rho_0) W_2(\rho_t, \rho_{\pi_t}). \end{aligned} \quad (\text{B.22})$$

Note that $\tilde{\alpha}_t^s$ is the geodesic connecting ρ_t and ρ_{π_t} . By Lemma D.2, we have

$$\sup_{s \in [0,1]} W_2(\tilde{\alpha}_t^s, \rho_0) \leq 2 \max \{ W_2(\rho_{\pi_t}, \rho_0), W_2(\rho_t, \rho_0) \}. \quad (\text{B.23})$$

Plugging (B.23) into (B.22), it follows that

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} \frac{\widetilde{W}_2^2(\rho_t, \rho_{\pi_t})}{2\eta} &\leq -(1 - \sqrt{\gamma}) \cdot \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}^2 + 2\eta^{-1} \alpha^{1/2} \mathcal{B} B_1 \cdot \widetilde{W}_2(\rho_t, \rho_0) \widetilde{W}_2(\rho_t, \rho_{\pi_t}) \\ &\quad + \alpha^{-1} B_2 \cdot \left(4B_1 \max \{ \widetilde{W}_2(\rho_{\pi_t}, \rho_0), \widetilde{W}_2(\rho_t, \rho_0) \} + B_r \right) \widetilde{W}_2(\rho_t, \rho_{\pi_t})^2, \end{aligned}$$

Where $\widetilde{W}_2 = \alpha^{-1} W_2$ is the scaled W_2 metric. Thus, we complete the proof of Theorem 4.5. \square

B.4 Proof of Theorem 4.6

Proof. We remark that the restarting mechanism produces discontinuity in ρ_t while π_t remains continuous. Let T_0, T_1, \dots, T_N denote the restarting points in $[0, T]$, where $T_0 = 0$ and N is the total restarting number in $[0, T]$. Let T_n^- and T_n^+ denote the moments just before and after the restarting occurring at T_n , respectively. According to the restarting mechanism, we have $\widetilde{W}_2(\rho_t, \rho_0) \leq \lambda \bar{D}$, $\widetilde{W}_2(\rho_{T_n^-}, \rho_0) = \lambda \bar{D}$ and $\rho_{T_n^+} = \rho_0$. Recall that we set $\tilde{\Phi}^{\pi_t} = \tilde{\mathcal{E}}_{\tilde{\phi}^{\pi_t}}^{\pi_t}$. By (4.7) of Theorem 4.5, it holds that

$$\begin{aligned} (1 - \sqrt{\gamma}) \cdot \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}^2 &\leq -\frac{\mathrm{d}}{\mathrm{d}t} \frac{\widetilde{W}_2^2(\rho_t, \rho_{\pi_t})}{2\eta} + 2\eta^{-1} \alpha^{1/2} \mathcal{B} B_1 \lambda \bar{D} \cdot \widetilde{W}_2(\rho_t, \rho_{\pi_t}) \\ &\quad + \alpha^{-1} B_2 \cdot (4B_1 \lambda \bar{D} + B_r) \widetilde{W}_2(\rho_t, \rho_{\pi_t})^2. \end{aligned} \quad (\text{B.24})$$

For simplicity, we let $S_1 = \frac{(1+\lambda)^2 \bar{D}^2 B_2 (4B_1 \lambda \bar{D} + B_r)}{1 - \sqrt{\gamma}}$, $S_2 = \frac{2\mathcal{B} B_1 \lambda (1+\lambda) \bar{D}^2}{1 - \sqrt{\gamma}}$, $\xi = \frac{1}{2(1 - \sqrt{\gamma})}$, $Q(t) = \|Q_t - Q^{\pi_t}\|_{2, \tilde{\phi}^{\pi_t}}$, and $\widetilde{W}(t) = \widetilde{W}_2(\rho_t, \rho_{\pi_t})$. By sum of the integrals of (B.24) on

$[T_0, T_1], \dots, [T_{N-1}, T_N]$, and $[T_N, T]$, we have

$$\begin{aligned} \frac{1}{T} \int_0^T Q^2(t) dt &\leq \frac{1}{T} \int_0^T \left(\frac{\alpha^{-1} S_1}{(1+\lambda)^2 \bar{D}^2} \widetilde{W}(t)^2 + \frac{\alpha^{1/2} \eta^{-1} S_2}{(1+\lambda) \bar{D}} \widetilde{W}(t) \right) dt \\ &\quad - \frac{\xi}{T\eta} \left(\sum_{n=0}^{N-1} \int_{T_n}^{T_{n+1}} \frac{d}{dt} \widetilde{W}(t)^2 dt + \int_{T_N}^T \frac{d}{dt} \widetilde{W}(t)^2 dt \right). \end{aligned} \quad (\text{B.25})$$

Note that we have $\widetilde{W}(t) = \widetilde{W}_2(\rho_t, \rho_{\pi_t}) \leq \widetilde{W}_2(\rho_t, \rho_0) + \widetilde{W}_2(\rho_{\pi_t}, \rho_0) \leq (1+\lambda) \bar{D}$ by the triangle inequality of W_2 distance, the restarting mechanism and Property (iii) of Lemma B.1. It thus holds for (B.25) that

$$\begin{aligned} \frac{1}{T} \int_0^T Q^2(t) dt &\leq \frac{1}{T} \int_0^T (\alpha^{-1} S_1 + \alpha^{1/2} \eta^{-1} S_2) dt \\ &\quad - \frac{\xi}{T\eta} \left(\sum_{n=0}^{N-1} (\widetilde{W}^2(T_{n+1}^-) - \widetilde{W}^2(T_n^+)) + (\widetilde{W}^2(T) - \widetilde{W}^2(T_N^+)) \right). \end{aligned} \quad (\text{B.26})$$

Note that we have $\widetilde{W}(T_{n+1}^-) \geq \widetilde{W}_2(\rho_0, \rho_{T_{n+1}^-}) - \widetilde{W}_2(\rho_0, \rho_{\pi_{T_{n+1}^-}}) \geq (\lambda-1) \bar{D} \geq \bar{D} \geq \widetilde{W}(\rho_0, \rho_{\pi_{T_n^+}}) = \widetilde{W}(T_n^+)$ by the triangle inequality of W_2 distance, the restarting mechanism and Property (iii) of Lemma B.1. It thus holds for (B.26) that

$$\frac{1}{T} \int_0^T Q^2(t) dt \leq \alpha^{-1} S_1 + \alpha^{1/2} \eta^{-1} S_2 + \frac{\eta^{-1} \bar{D}^2}{2T(1-\sqrt{\gamma})}. \quad (\text{B.27})$$

By setting $\tilde{\phi}^{\pi_t} = \frac{1}{2} \tilde{\phi}_0 + \frac{1}{2} \phi_0 \otimes \pi_t$ and plugging (B.27) into (4.1) of Theorem 4.1, we have

$$\begin{aligned} \frac{1}{T} \int_0^T (J(\pi^*) - J(\pi_t)) dt &\leq \frac{\zeta}{T} + \frac{4\kappa}{T} \int_0^T Q(t) dt \\ &\leq \frac{\zeta}{T} + 4\kappa \sqrt{\frac{1}{T} \int_0^T Q^2(t) dt} \\ &\leq \frac{\zeta}{T} + 4\kappa \sqrt{\alpha^{-1} S_1 + \alpha^{1/2} \eta^{-1} S_2 + \frac{\eta^{-1} \bar{D}^2}{2T(1-\sqrt{\gamma})}}, \end{aligned} \quad (\text{B.28})$$

where $\kappa = \|\tilde{\mathcal{E}}_{\mathcal{D}_0}^* / \tilde{\phi}_0\|_\infty$ is the concentrability coefficient and $\zeta = \mathbb{E}_{s \sim \mathcal{E}_{\pi^*}} [\text{KL}(\pi^*(\cdot | s) \| \pi_0(\cdot | s))]$ is the KL-divergence between π^* and π_0 . Here the second inequality follows from the Cauchy-Schwarz inequality. Therefore, we complete the proof of (4.9) in Theorem 4.6. In what follows, we aim to upper bound the total restarting number N . Recall that we have $\widetilde{W}(T_n^+) \leq \bar{D}$ and $\widetilde{W}(T_{n+1}^-) \geq (\lambda-1) \bar{D}$ according to the restarting mechanism. Thus, it holds for (B.26) that

$$\frac{1}{T} \int_0^T Q^2(t) dt \leq \alpha^{-1} S_1 + \alpha^{1/2} \eta^{-1} S_2 - \frac{\xi}{T\eta} (N(\lambda-2) \bar{D}^2 - \bar{D}^2).$$

Since $Q^2(t) \geq 0$, it follows that

$$\begin{aligned} N &\leq ((\alpha^{-1} S_1 + \alpha^{1/2} \eta^{-1} S_2) 2T(1-\sqrt{\gamma})\eta + \bar{D}^2) \cdot \frac{\bar{D}^{-2}}{\lambda-2} \\ &= (\lambda-2)^{-1} ((\alpha^{-1} \eta S_1 + \alpha^{1/2} S_2) 2T \bar{D}^{-2} (1-\sqrt{\gamma}) + 1), \end{aligned}$$

which upper bounds the total restarting number N . Hence, we complete the proof of Theorem 4.6. \square

C Proofs of Supporting Lemmas

In this section, we give detailed proof of supporting lemmas.

C.1 Generality of Transition Kernel Representation

Recall that we have the following representation of the transition kernel in Assumption 4.3,

$$P(s' | s, a) = \int \tilde{\sigma}((s, a, 1)^\top w) \varphi(s') \psi(s'; w) dw, \quad (\text{C.1})$$

where $\varphi(s') \geq 0$ and $\tilde{\psi}(s; \cdot) \in \mathcal{P}(w)$. Such a representation has the same function class as

$$\tilde{P}(s' | s, a) = \int \tilde{\sigma}((s, a, 1)^\top w) \tilde{\psi}(s'; w) dw, \quad (\text{C.2})$$

where $\tilde{\psi}$ is a finite signed measure.

Proof. For simplicity, let \mathcal{P} and $\tilde{\mathcal{P}}$ denote the function class represented by (C.1) and (C.2), respectively. Note that for any transition kernel $P(s' | s, a)$ represented by (C.1), by letting $\tilde{\psi}(s'; w) = \varphi(s') \psi(s'; w)$, such a transition kernel $P(s' | s, a)$ can be equivalently represented by $\tilde{P}(s' | s, a)$ in (C.2). Thus, it holds that $\mathcal{P} \subseteq \tilde{\mathcal{P}}$. Therefore, we only need to prove $\tilde{\mathcal{P}} \subseteq \mathcal{P}$, which is equivalent to proving that for any $\tilde{P}(s' | s, a) \in \tilde{\mathcal{P}}$ given by (C.2), the signed measure $\tilde{\psi}$ can be non-negative. If that is the case, by letting $\varphi(s') = \int (\tilde{\psi}_+(s'; w) + \tilde{\psi}_-(s'; -w)) dw$ and $\psi(s'; w) = \varphi(s')^{-1} (\tilde{\psi}_+(s'; w) + \tilde{\psi}_-(s'; -w))$, we can have (C.2) equivalently represented by (C.1). Note that there always exist non-negative functions $\tilde{\psi}_+$ and $\tilde{\psi}_-$ such that $\tilde{\psi} = \tilde{\psi}_+ - \tilde{\psi}_-$. Since $\tilde{\sigma}$ is an odd function, it holds that

$$\begin{aligned} \tilde{P}(s' | s, a) &= \int \tilde{\sigma}(w^\top x) \tilde{\psi}(s'; w) dw \\ &= \int \tilde{\sigma}(w^\top x) \tilde{\psi}_+(s'; w) dw + \int \tilde{\sigma}(-w^\top x) \tilde{\psi}_-(s'; w) dw \\ &= \int \tilde{\sigma}(w^\top x) (\tilde{\psi}_+(s'; w) + \tilde{\psi}_-(s'; -w)) dw. \end{aligned}$$

Thus, by letting $\varphi(s') = \int (\tilde{\psi}_+(s'; w) + \tilde{\psi}_-(s'; -w)) dw$ and $\psi(s'; w) = \varphi(s')^{-1} (\tilde{\psi}_+(s'; w) + \tilde{\psi}_-(s'; -w))$, it holds that $\tilde{P}(s' | s, a) = \int \tilde{\sigma}(w^\top x) \varphi(s') \psi(s'; w) dw = P(s' | s, a)$, where $\psi \in \mathcal{P}(\mathcal{W})$ and $\varphi \geq 0$. Hence, any $\tilde{P}(s' | s, a) \in \tilde{\mathcal{P}}$ can be equivalently represented by (C.1) and it follows that $\tilde{\mathcal{P}} \subseteq \mathcal{P}$. Thus, (C.1) has the same function class as (C.2) and we complete the proof. \square

C.2 Proof of Detailed Version of Lemma 4.4

In this section, we prove a more detailed version of Lemma 4.4, i.e., Lemma B.1.

Proof. We begin by a sketch of the proof of Lemma B.1. We first construct functions Z_π and $\nu_\pi(w)$. With the use of mollifiers, we prove that there exists function $p_\pi(b)$ satisfying (C.6) and then formulate a construction of $\bar{\rho}_\pi(\theta)$, which gives way to obtain ρ_π . For the proof of Property (iii), using the technique of Talagrand's inequality and the chi-squared divergence, we establish a constant upper bound for $W_2(\rho_\pi, \rho_0)$. For the proof of Property (iv), by exploiting the inequality between W_2 distance and the weighted homogeneous Sobolev norm, we upper bound $W_2(\rho_{\pi_1}, \rho_{\pi_2})$ up to $O(\alpha^{-1/2})$.

Proof of Property (i) of Lemma B.1. We give a proof of Property (i) by a construction of ρ_π . For notational simplicity, we let $x = (s, a, 1)$. By definitions of the action value function Q^π and the state value function V^π in (2.2), we have

$$\begin{aligned} Q^\pi(s, a) &= r(s, a) + \gamma \cdot \int P(s' | s, a) V^\pi(s') ds' \\ &= \int \tilde{\sigma}(w^\top x) \left\{ B_r \cdot \mu(w) + \gamma \cdot \int \varphi(s') \psi(s'; w) V^\pi(s') ds' \right\} dw \\ &= \int Z_\pi \cdot \tilde{\sigma}(w^\top x) \nu_\pi(w) dw, \end{aligned} \quad (\text{C.3})$$

where

$$\nu_\pi(w) = Z_\pi^{-1} \cdot \left(B_r \mu(w) + \gamma \cdot \int \varphi(s') \psi(s'; w) V^\pi(s') ds' \right), \quad (\text{C.4})$$

$$Z_\pi = B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds'. \quad (\text{C.5})$$

Here, the second equality in (C.3) holds by (i) of Assumption 4.3. We construct $\bar{\rho}_\pi$ by $\bar{\rho}_\pi = \nu_\pi \times p_\pi$, i.e., $\bar{\rho}_\pi(w, b) = \nu_\pi(w) p_\pi(b)$, where $p_\pi(b)$ is defined to be a probability measure in $\mathcal{P}_2(\mathbb{R})$ such that

$$\int B_\beta \cdot \beta(b) p_\pi(db) = Z_\pi. \quad (\text{C.6})$$

We remark that such a p_π exists and we will provide a construction later. Since we have $V^\pi \geq 0$, $\varphi \geq 0$, and $\psi \geq 0$ by Assumption 4.3, it turns out that ν_π is a probability density function according to (C.4), which further suggests that $\bar{\rho}_\pi \in \mathcal{P}_2(\mathbb{R}^D)$. Plugging (C.6) into (C.3), it holds that

$$Q^\pi(x) = \int \sigma(x; \theta) \bar{\rho}_\pi(\theta) d\theta,$$

where the equality holds by noting that $\sigma(x; \theta) = B_\beta \cdot \beta(b) \cdot \tilde{\sigma}(w^\top x)$ in (4.2) and that $\bar{\rho}_\pi = \nu_\pi \times p_\pi$. Furthermore, by letting $\rho_\pi = \bar{\rho}_\pi + (1 - \alpha^{-1})(\rho_0 - \bar{\rho}_\pi)$, we have

$$Q^\pi(x) = \alpha \int \sigma(x; \theta) (\rho_\pi - (1 - \alpha^{-1})\rho_0) d\theta = \alpha \int \sigma(x; \theta) \rho_\pi d\theta = Q(x; \rho_\pi),$$

where the second equality holds by noting that σ is odd with respect to w and b and that $\rho_0 \sim \mathcal{N}(0, I_D)$ is an even function. Thus, we finish the construction of ρ_π and also complete the proof of Property (i) in Lemma B.1.

A Construction for $p_\pi(b)$. Recall that we have p_π defined in (C.6). Here, we provide a construction for p_π which has some properties that will facilitate our analysis. Ideally, we want p_π to have global support and concentrate to its mean, which motivates us to consider p_π to be Gaussian distribution with high variance. Recall that we assume that β^{-1} is ℓ_β -Lipschitz continuous on $[-2/3, 2/3]$ in Assumption 4.2. Let $q(b)$ be the probability density function of the standard Gaussian distribution, i.e., $q \sim \mathcal{N}(0, 1)$. Then, $q_\epsilon(b - z) = \epsilon^{-1} \cdot q((b - z)/\epsilon)$ is the probability density function such that $q_\epsilon \sim \mathcal{N}(z, \epsilon^2)$. We define function β_ϵ as follows,

$$\beta_\epsilon(z) = \int \beta(b) q_\epsilon(b - z) db = (\beta * q_\epsilon)(z), \quad (\text{C.7})$$

where $*$ denotes the convolution. Note that $\{q_\epsilon\}_{\epsilon>0}$ can be viewed as a class of mollifiers [25]. In particular, let

$$\bar{\epsilon} = \min \left\{ \sqrt{\frac{\pi}{2}} \cdot \frac{1}{6L_{0,\beta}}, \sqrt{\frac{\pi}{2}} \cdot \frac{1}{2\ell_\beta L_{1,\beta}}, 1 \right\}, \quad (\text{C.8})$$

where $L_{0,\beta}$ and $L_{1,\beta}$ characterize the Lipschitz continuity and smoothness of β respectively and β^{-1} is ℓ_β -Lipschitz continuous in $[-2/3, 2/3]$ by Assumption 4.2. For the approximation error of mollifier β_ϵ , it holds that

$$\begin{aligned} |\beta(z) - \beta_\epsilon(z)| &= \left| \int (\beta(z) - \beta(z')) q_\epsilon(z - z') dz' \right| \\ &\leq L_{0,\beta} \cdot \int |z - z'| q_\epsilon(z - z') dz' \\ &= L_{0,\beta} \cdot \bar{\epsilon} \cdot \sqrt{\frac{2}{\pi}}, \end{aligned}$$

where the last inequality follows from $\int |z| q_\epsilon(z) dz = \epsilon \cdot \sqrt{2/\pi}$. Similarly, we have $|\dot{\beta}(z) - \dot{\beta}_\epsilon(z)| \leq L_{1,\beta} \cdot \bar{\epsilon} \cdot \sqrt{2/\pi}$. By definition of $\bar{\epsilon}$ in (C.8), it further holds that

$$\sup_{b \in \beta^{-1}([-2/3, 2/3])} |\beta(b) - \beta_\epsilon(b)| \leq 1/6, \quad (\text{C.9})$$

$$\sup_{b \in \beta^{-1}([-2/3, 2/3])} |\dot{\beta}(b) - \dot{\beta}_\epsilon(b)| \leq \frac{1}{2\ell_\beta}. \quad (\text{C.10})$$

Note that $\beta(b)$ is a monotonic function with $|\dot{\beta}(b)| \geq 1/\ell_\beta$ in $\beta^{-1}([-2/3, 2/3])$ by Assumption 4.2. With regard to (C.10), it follows that $\beta_\epsilon(b)$ is also monotonic in $\beta^{-1}([-2/3, 2/3])$ and that

$$|\dot{\beta}_\epsilon(b)| \geq |\dot{\beta}(b)| - |\dot{\beta}(b) - \dot{\beta}_\epsilon(b)| \geq \frac{1}{2\ell_\beta}, \quad \forall b \in \beta^{-1}([-2/3, 2/3]). \quad (\text{C.11})$$

Furthermore, by (C.9) and the continuity of β_ϵ in $\beta^{-1}([-2/3, 2/3])$, we have

$$[-1/2, 1/2] \subseteq \beta_\epsilon(\beta^{-1}([-2/3, 2/3])). \quad (\text{C.12})$$

The monotonicity of $\beta_\epsilon(b)$, (C.11), and (C.12) together show that β_ϵ^{-1} exists and is $2\ell_\beta$ -Lipschitz continuous in $[-1/2, 1/2]$. Moreover, since β is an odd function, it holds by (C.7) that β_ϵ is also an odd function with $\beta_\epsilon(0) = 0$. Hence, it holds that $\beta_\epsilon^{-1}([-1/2, 1/2]) \subseteq [-\ell_\beta, \ell_\beta]$. Furthermore, by (C.5), it holds that

$$\begin{aligned} Z_\pi &= B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds' \\ &\leq B_r + \gamma \cdot (1 - \gamma)^{-1} \cdot B_r \cdot \int \varphi(s') ds' \\ &\leq B_r + \gamma(1 - \gamma)^{-1} B_r M_{1,\varphi}, \end{aligned}$$

where the first inequality follows from the fact that $V^\pi(s) \leq (1 - \gamma)^{-1} \cdot \sup_{s,a} r(s, a)$ and the last inequality follows from Assumption 4.3 that $\int \varphi(s') ds' \leq M_{1,\varphi}$. By setting $B_\beta \geq 2(B_r + \gamma(1 - \gamma)^{-1} B_r M_{1,\varphi})$, it holds that $|Z_\pi/B_\beta| \leq 1/2$, which indicates that $\beta_\epsilon^{-1}(Z_\pi/B_\beta)$ exists and allows $p_\pi(b) = q_\epsilon(b - \beta_\epsilon^{-1}(Z_\pi/B_\beta))$ to be the probability density function such that $p_\pi(b) \sim \mathcal{N}(\beta_\epsilon^{-1}(Z_\pi/B_\beta), \epsilon^2)$. For the mean value $\beta_\epsilon^{-1}(Z_\pi/B_\beta)$, recalling that $\beta_\epsilon^{-1}([-1/2, 1/2]) \subseteq [-\ell_\beta, \ell_\beta]$, it thus holds that

$$|\beta_\epsilon^{-1}(Z_\pi/B_\beta)| \leq \ell_\beta. \quad (\text{C.13})$$

Following from (C.7), we have

$$\int \beta(b) p_\pi(b) db = \int \beta(b) q_\epsilon(b - \beta_\epsilon^{-1}(Z_\pi/B_\beta)) db = \beta_\epsilon(\beta_\epsilon^{-1}(Z_\pi/B_\beta)) = Z_\pi/B_\beta.$$

Hence, our construction of p_π here is in line with the definition of p_π in (C.6). In the sequel, we consider $p_\pi(b) = q_\epsilon(b - \beta_\epsilon^{-1}(Z_\pi/B_\beta))$ to hold all along.

Proof of Property (ii) of Lemma B.1. Here we show that $g(\cdot; \pi, \rho_\pi) = 0$ is a direct result of $Q^\pi(x) = Q(x; \rho_\pi)$ in Property (i). Note that

$$Q^\pi(x) - r(x) - \gamma \mathbb{E}_{x' \sim \bar{P}^\pi(\cdot | x)} Q^\pi(x') = 0 \quad (\text{C.14})$$

holds by the definition of the action value function Q^π in (2.2) for any $x \in \mathcal{S} \times \mathcal{A}$. Since we have $Q^\pi(x) = Q(x; \rho_\pi)$ proved, by plugging (C.14) into the definition of g in (3.6), where $Q(\cdot; \rho_\pi)$ is substituted for Q^π , it follows that $g(x; \pi, \rho_\pi) = 0$. Thus, we complete the proof of Property (ii) of Lemma B.1.

Proof of Property (iii) of Lemma B.1.

In what follows, we aim to upper bound $W^2(\rho_\pi, \rho_0)$. We summarize our aforementioned constructions as follows,

$$Z_\pi = B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds',$$

$$\nu_\pi(w) = Z_\pi^{-1} \left(B_r \mu(w) + \gamma \cdot \int \varphi(s') \psi(s'; w) V^\pi(s') ds' \right), \quad (\text{C.15})$$

$$p_\pi(b) = q_\epsilon(b - \beta_\epsilon^{-1}(Z_\pi/B_\beta)), \quad (\text{C.16})$$

$$\bar{\rho}_\pi(\theta) = p_\pi(b) \nu_\pi(w), \quad (\text{C.17})$$

$$\rho_\pi = \bar{\rho}_\pi + (1 - \alpha^{-1})(\rho_0 - \bar{\rho}_\pi). \quad (\text{C.18})$$

Plugging (C.16) into the definition of Chi-squared divergence, we have

$$\chi^2(p_\pi \| \rho_{0,b}) = \int \frac{p_\pi^2}{\rho_{0,b}} db - 1 = \frac{1}{\bar{\epsilon}\sqrt{2-\bar{\epsilon}^2}} \cdot \exp\left\{\frac{(\beta_{\bar{\epsilon}}^{-1}(Z_\pi/B_\beta))^2}{2-\bar{\epsilon}^2}\right\} - 1, \quad (\text{C.19})$$

Note that we have $\bar{\epsilon} = \min\left\{\sqrt{\pi/2} \cdot (6L_{0,\beta})^{-1}, \sqrt{\pi/2} \cdot (2\ell_\beta L_{1,\beta})^{-1}, 1\right\}$ by (C.8) and $|\beta_{\bar{\epsilon}}^{-1}(Z_\pi/B_\beta)| \leq \ell_\beta$ by (C.13). Hence, we have $\chi^2(p_\pi \| \rho_{0,b})$ upper bounded. As for ν_π in (C.15), we have

$$\begin{aligned} \chi^2(\nu_\pi \| \rho_{0,w}) &= \chi^2\left(B_r Z_\pi^{-1} \mu + \int V^\pi(s') Z_\pi^{-1} \varphi(s') \psi(s'; w) ds' \middle\| \rho_{0,w}\right) \\ &\leq 3\left(\chi^2(B_r Z_\pi^{-1} \mu \| \rho_{0,w}) + \chi^2\left(\int V^\pi(s') Z_\pi^{-1} \varphi(s') \psi(s'; \cdot) ds' \middle\| \rho_{0,w}\right) + 1\right), \end{aligned} \quad (\text{C.20})$$

where the inequality holds by Property (iii) of Lemma D.3. For the first term on the right-hand side of (C.20), we have

$$\begin{aligned} \chi^2(B_r Z_\pi^{-1} \mu \| \rho_{0,w}) &= B_r^2 Z_\pi^{-2} \chi^2(\mu \| \rho_{0,w}) + (1 - B_r Z_\pi^{-1})^2 \\ &\leq \chi^2(\mu \| \rho_{0,w}) + 1 \\ &\leq M_\mu + 1, \end{aligned} \quad (\text{C.21})$$

where the equality holds by Property (i) of Lemma D.3, the first inequality holds by noting that $|B_r Z_\pi^{-1}| \leq 1$ and the last inequality follows from $\chi^2(\mu \| \rho_{0,w}) < M_\mu$ by Assumption 4.3. Hence, the first term on the right-hand side of (C.20) is upper bounded. As for the second term, by Property (iv) of Lemma D.3, we have

$$\begin{aligned} &\chi^2\left(\int V^\pi(s') Z_\pi^{-1} \varphi(s') \psi(s'; \cdot) ds' \middle\| \rho_{0,w}\right) \\ &\leq \int (V^\pi(s') Z_\pi^{-1} \varphi(s'))^2 ds' \cdot \int \chi^2(\psi(s'; \cdot) \| \rho_{0,w}) ds' + \left(\int V^\pi(s') Z_\pi^{-1} \varphi(s') ds' - 1\right)^2 \\ &\leq (1-\gamma)^{-2} M_{2,\varphi} \cdot M_\psi + (1-\gamma)^{-2} M_{1,\varphi}^2 + 1, \end{aligned} \quad (\text{C.22})$$

where the last inequality holds by noting that $|V^\pi| \leq (1-\gamma)^{-1} B_r$, $Z_\pi \geq B_r$, $\|\mathcal{S}\| \leq 1$, and that $\chi^2(\psi(s'; \cdot) \| \rho_{0,w}) < M_\psi$ by Assumption 4.3. Hence, it holds that the second term of (C.20) is also upper bounded. Plugging (C.21) and (C.22) into (C.20), we can establish the upper bound for $\chi^2(\nu_\pi \| \rho_{0,w})$. Furthermore, by Property (ii) of Lemma D.3 and noting that $\bar{\rho}_\pi = p_\pi \times \nu_\pi$, we have $\chi^2(\bar{\rho}_\pi \| \rho_0)$ upper bounded as well, that is,

$$\chi^2(\bar{\rho}_\pi \| \rho_0) < \frac{1}{2} \bar{D}^2,$$

where \bar{D} depends on absolute constants occurring in (C.19), (C.21), and (C.22), i.e., $\ell_\beta, L_{0,\beta}, L_{1,\beta}, B_r, M_{1,\varphi}, M_{2,\varphi}, M_\mu, M_\psi$ in Assumptions 4.2 and 4.3. Since $\rho_0 \sim \mathcal{N}(0, I_D)$, it holds for any ρ_π that,

$$\begin{aligned} \frac{1}{2} W_2(\rho_\pi, \rho_0)^2 &\leq \text{KL}(\rho_\pi \| \rho_0) \leq \int \left(\frac{\rho_\pi}{\rho_0} - 1\right) \frac{\rho_\pi}{\rho_0} \rho_0(d\theta) = \int \left(\frac{\rho_\pi}{\rho_0} - 1\right)^2 \rho_0(d\theta) \\ &= \int \left(\frac{(1-\alpha^{-1})\rho_0(\theta) + \alpha^{-1}\bar{\rho}_\pi(\theta)}{\rho_0(\theta)} - 1\right)^2 \rho_0(d\theta) = \alpha^{-2} \chi^2(\bar{\rho}_\pi \| \rho_0) < \frac{\alpha^{-2}}{2} \bar{D}^2, \end{aligned} \quad (\text{C.23})$$

where the first inequality follows from Talagrand's inequality in Lemma E.4. Plugging $\widetilde{W}_2 = \alpha W_2$ into (C.23), we complete the proof of Property (iii) of Lemma B.1.

Proof of Property (iv) of Lemma B.1.

Upper Bounding $W_2(p_{\pi_1}, p_{\pi_2})$. Following the property of W_2 distance with respect to Gaussian distribution, we have $W_2(p_{\pi_1}, p_{\pi_2}) = |\beta_{\bar{\epsilon}}^{-1}(Z_{\pi_1}/B_\beta) - \beta_{\bar{\epsilon}}^{-1}(Z_{\pi_2}/B_\beta)|$. Recall that we have $|Z_\pi/B_\beta| \leq 1/2$ and that $\beta_{\bar{\epsilon}}^{-1}$ is $2\ell_\beta$ -Lipschitz continuous on $[-1/2, 1/2]$. It then holds that

$$W_2(p_{\pi_1}, p_{\pi_2}) = |\beta_{\bar{\epsilon}}^{-1}(Z_{\pi_1}/B_\beta) - \beta_{\bar{\epsilon}}^{-1}(Z_{\pi_2}/B_\beta)| \leq 2\ell_\beta \cdot |Z_{\pi_1} - Z_{\pi_2}|/B_\beta. \quad (\text{C.24})$$

Meanwhile, we have

$$\begin{aligned} |Z_{\pi_1} - Z_{\pi_2}| &\leq \gamma \cdot \int \varphi(s') \cdot |V^{\pi_1}(s') - V^{\pi_2}(s')| ds' \\ &\leq \gamma \cdot \int \varphi(s') ds' \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')| \\ &\leq \gamma \cdot M_{1,\varphi} \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|, \end{aligned} \quad (\text{C.25})$$

where the last inequality holds by (ii) of Assumption 4.3. Plugging (C.25) into (C.24), it holds for $W_2(p_{\pi_1}, p_{\pi_2})$ that

$$W_2(p_{\pi_1}, p_{\pi_2}) \leq \frac{2\ell_\beta \cdot \gamma \cdot M_{1,\varphi} \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|}{B_\beta}. \quad (\text{C.26})$$

Upper Bounding $W_2(\nu_{\pi_1}, \nu_{\pi_2})$. By definition of ν_π in (C.15), we have

$$\nu_{\pi_1} - \nu_{\pi_2} = \frac{B_r(Z_{\pi_2} - Z_{\pi_1})}{Z_{\pi_1}Z_{\pi_2}}\mu + \gamma \int \varphi(s')\psi(s'; \cdot) \left(\frac{V^{\pi_1}}{Z_{\pi_1}} - \frac{V^{\pi_2}}{Z_{\pi_2}} \right) ds'. \quad (\text{C.27})$$

For $V^{\pi_1}/Z_{\pi_1} - V^{\pi_2}/Z_{\pi_2}$, it holds that

$$\begin{aligned} \left| \frac{V^{\pi_1}}{Z_{\pi_1}} - \frac{V^{\pi_2}}{Z_{\pi_2}} \right| &\leq \max\{V^{\pi_1}, V^{\pi_2}\} \cdot \frac{|Z_{\pi_1} - Z_{\pi_2}|}{Z_{\pi_1}Z_{\pi_2}} + \max\{Z_{\pi_1}^{-1}, Z_{\pi_2}^{-1}\} |V^{\pi_1} - V^{\pi_2}| \\ &\leq (B_r^{-1}(1 - \gamma)^{-1}\gamma M_{1,\varphi} + B_r^{-1}) \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|, \end{aligned} \quad (\text{C.28})$$

where the last inequality holds by noting that

$$\left| \frac{Z_{\pi_2} - Z_{\pi_1}}{Z_{\pi_1}Z_{\pi_2}} \right| \leq \gamma M_{1,\varphi} B_r^{-2} \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|. \quad (\text{C.29})$$

Here the inequality in (C.29) holds by (C.25) and the fact that $Z_\pi \geq B_r$. For $W_2(\nu_{\pi_1}, \nu_{\pi_2})$, by Lemma E.5, it holds that

$$W_2(\nu_{\pi_1}, \nu_{\pi_2}) \leq 2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\nu_{\pi_1})}. \quad (\text{C.30})$$

Recall that we have ν_π defined in (C.15) that

$$\begin{aligned} \nu_\pi(w) &= Z_\pi^{-1} \cdot \left(B_r \mu(w) + \gamma \cdot \int \varphi(s') \psi(s'; w) V^\pi(s') ds' \right), \\ Z_\pi &= B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds' \end{aligned}$$

where $Z_\pi^{-1}(B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds') = 1$, $B_r Z_\pi^{-1} \geq 0$, and $\gamma \varphi(s') V^\pi(s') Z_\pi^{-1} \geq 0$, which indicate that ν_π is in the convex hull of μ and $\psi(s'; \cdot)$. Hence, by Property (i) of Lemma D.4, it holds that

$$2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\nu_{\pi_1})} \leq 2 \max \left\{ \sup_s \|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\psi(s; \cdot))}, \|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\mu)} \right\}. \quad (\text{C.31})$$

Furthermore, following from (C.27) and $Z_\pi^{-1}(B_r + \gamma \cdot \int \varphi(s') V^\pi(s') ds') = 1$, by Property (ii) of Lemma D.4, it holds for $2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\mu)}$ that

$$\begin{aligned} 2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\mu)} &\leq \max \left\{ \sup_{s' \in \mathcal{S}} \|\mu - \psi(s'; \cdot)\|_{\dot{H}^{-1}(\mu)}, \sup_{(s', s'') \in \mathcal{S} \times \mathcal{S}} \|\psi(s'; \cdot) - \psi(s''; \cdot)\|_{\dot{H}^{-1}(\mu)} \right\} \\ &\quad \cdot \left(\left| \frac{B_r(Z_{\pi_2} - Z_{\pi_1})}{Z_{\pi_1}Z_{\pi_2}} \right| + \gamma \int \varphi(s') \left| \frac{V^{\pi_1}}{Z_{\pi_1}} - \frac{V^{\pi_2}}{Z_{\pi_2}} \right| ds' \right). \end{aligned} \quad (\text{C.32})$$

Plugging (iii) of Assumption 4.3, (C.28), and (C.29) into (C.32), we have

$$2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\mu)} \leq \mathcal{G} \cdot \left(\gamma M_{1,\varphi} B_r^{-1} + \gamma M_{1,\varphi} (B_r^{-1}(1-\gamma)^{-1} \gamma M_{1,\varphi} + B_r^{-1}) \right) \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|. \quad (\text{C.33})$$

By simply substituting $\|\cdot\|_{\dot{H}^{-1}(\psi(s;:))}$ for $\|\cdot\|_{\dot{H}^{-1}(\nu)}$ in both (C.32) and (C.33), it also holds for $2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\psi(s;:))}$ that

$$2\|\nu_{\pi_1} - \nu_{\pi_2}\|_{\dot{H}^{-1}(\psi(s;:))} \leq \mathcal{G} \cdot \left(\gamma M_{1,\varphi} B_r^{-1} + \gamma M_{1,\varphi} (B_r^{-1}(1-\gamma)^{-1} \gamma M_{1,\varphi} + B_r^{-1}) \right) \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|. \quad (\text{C.34})$$

Combining (C.30), (C.31), (C.33), and (C.34), we have

$$W_2(\nu_{\pi_1}, \nu_{\pi_2}) \leq \mathcal{G} \cdot \left(\gamma M_{1,\varphi} B_r^{-1} + \gamma M_{1,\varphi} (B_r^{-1}(1-\gamma)^{-1} \gamma M_{1,\varphi} + B_r^{-1}) \right) \cdot \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|. \quad (\text{C.35})$$

Upper Bounding $W_2(\rho_{\pi_1}, \rho_{\pi_2})$. Note that we have $\bar{\rho}_\pi = \nu_\pi \times p_\pi$ in (C.17). By Lemma D.5, we have

$$W_2(\bar{\rho}_{\pi_1}, \bar{\rho}_{\pi_2}) \leq \sqrt{W_2^2(\nu_{\pi_1}, \nu_{\pi_2}) + W_2^2(p_{\pi_1}, p_{\pi_2})} \leq \mathcal{B}' \sup_{s' \in \mathcal{S}} |V^{\pi_1}(s') - V^{\pi_2}(s')|, \quad (\text{C.36})$$

where \mathcal{B}' depends on the discount factor γ and absolute constants $\ell_\beta, B_\beta, B_r, M_{1,\varphi}, \mathcal{G}$ in Assumption 4.2 and 4.3 according to (C.26) and (C.35). By performance difference lemma [33], we have

$$\begin{aligned} |V^{\pi_1}(s) - V^{\pi_2}(s)| &= (1-\gamma)^{-1} \cdot \left| \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\langle A^{\pi_2}(s', \cdot), \pi_1(\cdot | s') - \pi_2(\cdot | s') \rangle_{\mathcal{A}} \right] \right| \\ &= (1-\gamma)^{-1} \cdot \left| \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\langle Q^{\pi_2}(s', \cdot), \pi_1(\cdot | s') - \pi_2(\cdot | s') \rangle_{\mathcal{A}} \right] \right| \\ &\leq (1-\gamma)^{-2} \cdot B_r \cdot \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\|\pi_1(\cdot | s') - \pi_2(\cdot | s')\|_1 \right]. \end{aligned} \quad (\text{C.37})$$

Here the inequality follows from $|Q^\pi(s, a)| \leq (1-\gamma)^{-1} \cdot B_r$. We let $\mathcal{B} = \mathcal{B}' B_r (1-\gamma)^{-2}$. Plugging (C.37) into (C.36), we have

$$W_2(\bar{\rho}_{\pi_1}, \bar{\rho}_{\pi_2}) \leq \mathcal{B} \cdot \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\|\pi_1(\cdot | s') - \pi_2(\cdot | s')\|_1 \right].$$

Recall that we have $\rho_\pi = \bar{\rho}_\pi + (1-\alpha^{-1})(\rho_0 - \bar{\rho}_\pi)$ in (C.18). Then, by Lemma D.6, it holds that

$$W_2(\rho_{\pi_1}, \rho_{\pi_2}) \leq \alpha^{-\frac{1}{2}} W_2(\bar{\rho}_{\pi_1}, \bar{\rho}_{\pi_2}) \leq \alpha^{-\frac{1}{2}} \mathcal{B} \cdot \sup_{s \in \mathcal{S}} \mathbb{E}_{s' \sim \mathcal{E}_s^{\pi_1}} \left[\|\pi_1(\cdot | s') - \pi_2(\cdot | s')\|_1 \right],$$

which completes the proof of Lemma B.1. \square

D Technical Results

In this section, we state and prove some technical results used in the proof of main theorems and lemmas.

Lemma D.1. Under Assumptions 4.3 and 4.2, it holds for any $\rho \in \mathcal{P}_2(\mathbb{R}^D)$ that

$$\sup_{x \in \mathcal{X}} |Q(x; \rho)| \leq \alpha \cdot B_1 \cdot W_2(\rho, \rho_0), \quad (\text{D.1})$$

$$\sup_{\theta \in \mathbb{R}^D} \|\nabla_\theta g(\theta; \rho)\|_{\mathbb{F}} \leq \alpha^{-1} \cdot B_2 \cdot (2\alpha \cdot B_1 \cdot W_2(\rho, \rho_0) + B_r). \quad (\text{D.2})$$

Proof. Following from Assumptions 4.3 and 4.2, we have that $\|\nabla_{\theta}\sigma(x; \theta)\| \leq B_1$ for any $x \in \mathcal{X}$ and $\theta \in \mathbb{R}^D$, which implies that $\text{Lip}(\sigma(x; \cdot)/B_1) \leq 1$ for any $x \in \mathcal{X}$. Note that $Q(x; \rho_0) = 0$ for any $x \in \mathcal{X}$. Thus, by (A.7) and the inequality between W_1 distance and W_2 distance [58], we have for any $\rho \in \mathcal{P}_2(\mathbb{R}^D)$ and $x \in \mathcal{X}$ that

$$|Q(x; \rho)| = \alpha \cdot \left| \int \sigma(x; \theta) \cdot d(\rho - \rho_0)(\theta) \right| \leq \alpha \cdot B_1 \cdot W_1(\rho, \rho_0) \leq \alpha \cdot B_1 \cdot W_2(\rho, \rho_0). \quad (\text{D.3})$$

which completes the proof of (D.1) in Lemma D.1. Following from the definition of g in (3.6), we have for any $x \in \mathcal{X}$ and $\rho \in \mathcal{P}_2(\mathbb{R}^D)$ that

$$\begin{aligned} \|\nabla_{\theta}g(\theta; \rho)\|_{\text{F}} &\leq \alpha^{-1} \cdot \mathbb{E}_{\tilde{\mathcal{D}}} \left[|Q(x; \rho) - r - \gamma \cdot Q(x'; \rho)| \cdot \|\nabla_{\theta\theta}^2\sigma(x; \theta)\|_{\text{F}} \right] \\ &\leq \alpha^{-1} \cdot B_2 \cdot (2\alpha \cdot B_1 \cdot W_2(\rho, \rho_0) + B_r). \end{aligned}$$

Here the last inequality follows from (D.3) and the fact that $\|\nabla_{\theta\theta}^2\sigma(x; \theta)\|_{\text{F}} \leq B_2$ for any $x \in \mathcal{X}$ and $\rho \in \mathcal{P}_2(\mathbb{R}^D)$, which follows from Assumptions 4.3 and 4.2. Thus, we complete the proof of Lemma D.1. \square

Lemma D.2. For $\rho_0, \rho_t, \rho_{\pi_t} \in \mathcal{P}_2(\mathbb{R}^D)$ and the geodesic $\alpha_t^{[0,1]}$ connecting ρ_t and ρ_{π_t} , we have

$$\sup_{s \in [0,1]} W_2(\alpha_t^s, \rho_0) \leq 2 \max\{W_2(\rho_{\pi_t}, \rho_0), W_2(\rho_t, \rho_0)\}. \quad (\text{D.4})$$

Proof. We give a proof by contradiction. Note that α_t^s is the geodesic connecting ρ_{π_t} and ρ_t . Assume there exists t such that

$$\sup_{s \in [0,1]} W_2(\alpha_t^s, \rho_0) > 2 \max\{W_2(\rho_{\pi_t}, \rho_0), W_2(\rho_t, \rho_0)\}.$$

Then, according to the triangle inequality of W_2 metric [59], we have

$$W_2(\rho_t, \alpha_t^s) \geq |W_2(\alpha_t^s, \rho_0) - W_2(\rho_t, \rho_0)| > |2W_2(\rho_t, \rho_0) - W_2(\rho_t, \rho_0)| = W_2(\rho_t, \rho_0),$$

and $W_2(\rho_{\pi_t}, \alpha_t^s) > W_2(\rho_{\pi_t}, \rho_0)$ for the same sake, which conflicts with the definition of geodesic that $W_2(\rho_t, \alpha_t^s) + W_2(\alpha_t^s, \rho_{\pi_t}) = W_2(\rho_t, \rho_{\pi_t}) \leq W_2(\rho_t, \rho_0) + W_2(\rho_{\pi_t}, \rho_0)$. Hence, such t does not exist and (D.4) holds. Thus we complete the proof of Lemma D.2. \square

Lemma D.3. The Chi-squared divergence has the following properties.

- (i) For any probability measure $g \in \mathcal{P}(\Theta)$, function $f : \Theta \rightarrow \mathbb{R}$, and $\alpha \in \mathbb{R}$, we have

$$\chi^2(\alpha f \parallel g) = \alpha^2 \chi^2(f \parallel g) + (1 - \alpha)^2.$$

- (ii) For any probability measures $g_1 \in \mathcal{P}(\Theta_1)$ and $g_2 \in \mathcal{P}(\Theta_2)$, functions $f_1 : \Theta_1 \rightarrow \mathbb{R}$ and $f_2 : \Theta_2 \rightarrow \mathbb{R}$, we have

$$\chi^2(f_1 \times f_2 \parallel g_1 \times g_2) = \chi^2(f_1 \parallel g_1) \cdot \chi^2(f_2 \parallel g_2) + \chi^2(f_1 \parallel g_1) + \chi^2(f_2 \parallel g_2),$$

where $f_1 \times f_2$ is the product of f_1 and f_2 , and $g_1 \times g_2$ is the product measure of g_1 and g_2 , i.e., $(f_1 \times f_2)(\theta_1 \times \theta_2) = f_1(\theta_1)f_2(\theta_2)$ and $(g_1 \times g_2)(\theta_1 \times \theta_2) = g_1(\theta_1)g_2(\theta_2)$, respectively.

- (iii) For any probability measure $g \in \mathcal{P}(\Theta)$, functions $f_1 : \Theta \rightarrow \mathbb{R}$ and $f_2 : \Theta \rightarrow \mathbb{R}$, we have

$$\chi^2(f_1 + f_2 \parallel g) \leq 3(\chi^2(f_1 \parallel g) + \chi^2(f_2 \parallel g) + 1).$$

- (iv) For any probability measure $g \in \mathcal{P}(\Theta)$, function $f : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ and $\alpha : \mathcal{X} \rightarrow \mathbb{R}$, we have

$$\chi^2\left(\int \alpha(x)f(d, \cdot)dx \parallel g\right) \leq \int \alpha(x)^2 dx \cdot \int \chi^2(f(x, \cdot) \parallel g)dx + \left(\int \alpha(x)dx - 1\right)^2.$$

Proof. **Proof of Property (i) of Lemma D.3.** By definition of the Chi-squared divergence, we have

$$\begin{aligned}\chi^2(\alpha f \parallel g) &= \int \left(\frac{\alpha f}{g} - 1 \right)^2 dg \\ &= \int \left(\alpha \left(\frac{f}{g} - 1 \right) + \alpha - 1 \right)^2 dg \\ &= \alpha^2 \chi^2(f \parallel g) + (\alpha - 1)^2.\end{aligned}$$

Thus, we complete the proof of Property (i) of Lemma D.3.

Proof of Property (ii) of Lemma D.3. Let $\tilde{f}_1 = f_1/g_1 - 1$ and $\tilde{f}_2 = f_2/g_2 - 1$. It then holds that

$$\chi^2(f_1 \times f_2 \parallel g_1 \times g_2) = \int \left(\frac{f_1 \times f_2}{g_1 \times g_2} - 1 \right)^2 d(g_1 \times g_2) = \int (\tilde{f}_1 \times \tilde{f}_2 + \tilde{f}_1 + \tilde{f}_2)^2 d(g_1 \times g_2).$$

By further noting that $\int \tilde{f}_1 dg_1 = 0$, $\int \tilde{f}_2 dg_2 = 0$, $\int \tilde{f}_1^2 dg_1 = \chi^2(f_1 \parallel g_1)$, and $\int \tilde{f}_2^2 dg_2 = \chi^2(f_2 \parallel g_2)$, we have

$$\begin{aligned}\chi^2(f_1 \times f_2 \parallel g_1 \times g_2) &= \int (\tilde{f}_1^2 \times \tilde{f}_2^2 + \tilde{f}_1^2 + \tilde{f}_2^2 + 2(\tilde{f}_1)^2 \times \tilde{f}_2 + 2(\tilde{f}_2)^2 \times \tilde{f}_1 + 2\tilde{f}_1 \times \tilde{f}_2) d(g_1 \times g_2) \\ &= \chi^2(f_1 \parallel g_1) \cdot \chi^2(f_2 \parallel g_2) + \chi^2(f_1 \parallel g_1) + \chi^2(f_2 \parallel g_2).\end{aligned}$$

Thus, we complete the proof of Property (ii) of Lemma D.3.

Proof of Property (iii) of Lemma D.3. Let $\tilde{f}_1 = f_1/g_1 - 1$ and $\tilde{f}_2 = f_2/g_2 - 1$. It then holds that

$$\chi^2(f_1 + f_2 \parallel g) = \int \left(\frac{f_1 + f_2}{g} - 1 \right)^2 dg = \int (\tilde{f}_1 + \tilde{f}_2 + 1)^2 dg.$$

By further noting that $\int \tilde{f}_1^2 dg_1 = \chi^2(f_1 \parallel g_1)$ and $\int \tilde{f}_2^2 dg_2 = \chi^2(f_2 \parallel g_2)$, we have

$$\begin{aligned}\chi^2(f_1 + f_2 \parallel g) &\leq 3 \int ((\tilde{f}_1)^2 + (\tilde{f}_2)^2 + 1) dg \\ &= 3(\chi^2(f_1 \parallel g) + \chi^2(f_2 \parallel g) + 1),\end{aligned}$$

where the inequality follows from the Cauchy-Schwarz inequality. Thus, we complete the proof of Property (iii) of Lemma D.3.

Proof of Property (iv) of Lemma D.3. Let $\tilde{f}(x, \cdot) = f(x, \cdot)/g(\cdot) - 1$. It then holds that

$$\begin{aligned}\chi^2\left(\int \alpha(x) f(x, \cdot) dx \parallel g\right) &= \int \left(\int \alpha(x) \tilde{f}(x, \theta) dx \right)^2 g(d\theta) + \left(\int \alpha(x) dx - 1 \right)^2 \\ &\leq \int \left(\int \alpha(x)^2 dx \cdot \int \tilde{f}(x, \theta)^2 dx \right) g(d\theta) + \left(\int \alpha(x) dx - 1 \right)^2 \\ &= \int \alpha(x)^2 dx \cdot \int \chi^2(f(x, \cdot) \parallel g) dx + \left(\int \alpha(x) dx - 1 \right)^2,\end{aligned}$$

where the first equality holds by noting that $\int \tilde{f}(x, \cdot) dg = 0$ and the inequality follows from the Cauchy-Schwarz inequality. Thus, we complete the proof of Property (iv) of Lemma D.3. \square

Lemma D.4. For weighted homogeneous Sobolev norm defined by

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} = \sup \{ |\langle f, \nu_1 - \nu_2 \rangle| \mid \|f\|_{\dot{H}^1(\mu)} \leq 1 \}.$$

we have the following properties.

- (i) For a group of probability measures $\mu_x : \mathcal{X} \rightarrow \mathcal{P}_2(\Theta)$ and $\nu_1, \nu_2 \in \mathcal{P}_2(\Theta)$, if μ is in the convex hull of μ_x , i.e., there exists $\alpha_x \geq 0$ such that both $\int \alpha_x dx = 1$ and $\mu = \int \alpha_x \mu_x dx$ hold, it then holds that

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} \leq \sup_{x \in \mathcal{X}} \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu_x)}.$$

(ii) Assume that we have measures $\mu \in \mathcal{P}_2(\Theta)$ and $\nu_x : \mathcal{X} \rightarrow \mathcal{P}_2(\Theta)$. Let $\beta_1, \beta_2 : \mathcal{X} \rightarrow \mathbb{R}$ be two functions on \mathcal{X} such that $\int \beta_1(x)dx = \int \beta_2(x)dx$. Then, by letting $\nu_1 = \int \nu_x \beta_1(x)dx$ and $\nu_2 = \int \nu_x \beta_2(x)dx$, we have

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} \leq \frac{1}{2} \sup_{(x', x'') \in \mathcal{X} \times \mathcal{X}} \|\nu_{x'} - \nu_{x''}\|_{\dot{H}^{-1}(\mu)} \cdot \int |\beta_1(x) - \beta_2(x)| dx.$$

Proof. Proof of Property (i) of Lemma D.4. By definition of the weighted homogeneous Sobolev norm, we have

$$\begin{aligned} \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} &= \sup_f \left\{ |\langle f, \nu_1 - \nu_2 \rangle| \left| \int |\nabla f|^2 \alpha_x d\mu_x dx \leq 1 \right| \right\} \\ &= \sup_{f, \lambda_x \geq 0} \left\{ |\langle f, \nu_1 - \nu_2 \rangle| \left| \int |\nabla f|^2 d\mu_x \leq \lambda_x, \forall x; \int \alpha_x \lambda_x dx = 1 \right| \right\} \end{aligned} \quad (\text{D.5})$$

$$\leq \sup_{\substack{\int \lambda_x \alpha_x dx = 1, \\ \lambda_x \geq 0}} \inf_x \sup_{f_x} \left\{ |\langle f_x, \nu_1 - \nu_2 \rangle| \left| \int |\nabla f_x|^2 d\mu_x \leq \lambda_x \right| \right\}, \quad (\text{D.6})$$

where the first equality holds by noting that $\mu = \int \alpha_x \mu_x dx$. To illustrate the last equality, we denote by \mathcal{F} and \mathcal{F}_x the allowed function class for f in (D.5) and f_x in (D.6) to choose from, respectively. Note that we have $\mathcal{F} \subseteq \mathcal{F}_x$ for any x , which is because the constraints for each f_x in (D.6) are relaxation of the constraints for f in (D.5). And so, the supremum taken over \mathcal{F} is no larger than the supremum over \mathcal{F}_x for any x . Therefore, the supremum over \mathcal{F} is no larger than the the smallest supremum over \mathcal{F}_x . Thus, (D.6) holds. Furthermore, we have

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} \leq \sup_{\substack{\int \lambda_x \alpha_x dx = 1, \\ \lambda_x \geq 0}} \inf_x \sup_{f_x} \left\{ |\langle f_x, \nu_1 - \nu_2 \rangle| \left| \int |\nabla f_x|^2 d\mu_x \leq \lambda_x \right| \right\} \quad (\text{D.7})$$

$$= \sup_{\substack{\int \lambda_x \alpha_x dx = 1, \\ \lambda_x \geq 0}} \inf_x \left\{ \sqrt{\lambda_x} \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu_x)} \right\} \quad (\text{D.8})$$

$$\leq \sup_{\substack{\int \lambda_x \alpha_x dx = 1, \\ \lambda_x \geq 0}} \inf_x \left\{ \sqrt{\lambda_x} \right\} \cdot \sup_x \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu_x)},$$

where the first equality holds by noting that (D.8) is a rescaling of (D.7) with respect to λ_x in the constraints of (D.7). Here, we let

$$y = \sup_{\substack{\int \lambda_x \alpha_x dx = 1, \\ \lambda_x \geq 0}} \inf_x \left\{ \sqrt{\lambda_x} \right\}.$$

Then, it holds that $y \leq 1$. Otherwise, there must exists λ_x such that $\int \lambda_x \alpha_x dx = 1$ and that $\lambda_x > 1$ holds for any $x \in \mathcal{X}$, which contradicts with our conditions that $\int \alpha_x dx = 1$ and $\alpha_x \geq 0$. Therefore, it further holds that

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} \leq y \cdot \sup_x \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu_x)} \leq \sup_x \|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu_x)}.$$

Thus, we complete the proof of Property (i) of Lemma D.4.

Proof of Property (ii) of Lemma D.4. Let $\alpha = \beta_1 - \beta_2$. Then, we have $\int \alpha(x)dx = 0$. Let $\alpha^+ = \max\{0, \alpha\}$ and $\alpha^- = -\min\{0, \alpha\}$. Then, we have $\alpha^+ - \alpha^- = \alpha = \beta_1 - \beta_2$ and that $\alpha^+ + \alpha^- = |\alpha| = |\beta_1 - \beta_2|$. Since $\int \alpha = 0$, it holds that $\int \alpha^+(x)dx = \int \alpha^-(x)dx = A$, where $A \geq 0$. We further let $\lambda(x, x') = A^{-1} \alpha^+(x) \alpha^-(x')$. Then, it holds that $\alpha^+(x) = \int \lambda(x, x') dx'$ and that $\alpha^-(x') = \int \lambda(x, x') dx$. Therefore, we have

$$\|\nu_1 - \nu_2\|_{\dot{H}^{-1}(\mu)} = \left\| \int_{\mathcal{X}} (\alpha^+ - \alpha^-) \nu_x dx \right\|_{\dot{H}^{-1}(\mu)} = \left\| \int_{\mathcal{X} \times \mathcal{X}} \lambda(x, x') (\nu_x - \nu_{x'}) dx dx' \right\|_{\dot{H}^{-1}(\mu)}.$$

By definition of the weighted homogeneous Sobolev norm, it further holds that

$$\begin{aligned} & \left\| \int_{\mathcal{X} \times \mathcal{X}} \lambda(x, x')(\nu_x - \nu_{x'}) dx dx' \right\|_{\dot{H}^{-1}(\mu)} \\ &= \sup_f \left\{ \left| \left\langle \int_{\mathcal{X} \times \mathcal{X}} \lambda(x, x')(\nu_x - \nu_{x'}) dx dx', f \right\rangle \right| \left| \int |\nabla f|^2 \leq 1 \right\}. \end{aligned} \quad (\text{D.9})$$

We assume the supremum in (D.9) is reached at f^* . Then, we have $\int |\nabla f^*|^2 d\mu \leq 1$ and that

$$\begin{aligned} \left\| \int_{\mathcal{X} \times \mathcal{X}} \lambda(x, x')(\nu_x - \nu_{x'}) dx dx' \right\|_{\dot{H}^{-1}(\mu)} &= \left| \left\langle \int_{\mathcal{X} \times \mathcal{X}} \lambda(x, x')(\nu_x - \nu_{x'}) dx dx', f^* \right\rangle \right| \\ &\leq \sup_{(x, x') \in \mathcal{X} \times \mathcal{X}} |\langle \nu_x - \nu_{x'}, f^* \rangle| \cdot \int \lambda(x, x') dx dx' \\ &\leq \frac{1}{2} \sup_{(x, x') \in \mathcal{X} \times \mathcal{X}} \|\nu_x - \nu_{x'}\|_{\dot{H}^{-1}(\mu)} \cdot \int |\beta_1 - \beta_2| dx, \end{aligned}$$

where the last inequality holds by noting that $\int \lambda(x, x') dx dx' = A = 1/2 \cdot \int |\alpha| dx = 1/2 \cdot \int |\beta_1 - \beta_2| dx$. Thus, we complete the proof of Property (ii) of Lemma D.4. \square

Lemma D.5. For probability measures $\nu_1, \nu_2 \in \mathcal{P}_2(\Theta_1)$ and $p_1, p_2 \in \mathcal{P}_2(\Theta_2)$, it holds that

$$W_2^2(\nu_1 \times p_1, \nu_2 \times p_2) \leq W_2^2(\nu_1, \nu_2) + W_2^2(p_1, p_2).$$

Proof. By the property of optimal transport [6], there exists mapping T_ν and T_p such that

$$\begin{aligned} W_2^2(\nu_1, \nu_2) &= \int \|\theta_1 - T_\nu(\theta_1)\| \nu_1(d\theta_1), \\ W_2^2(p_1, p_2) &= \int \|\theta_1 - T_p(\theta_1)\| p_1(d\theta_1), \end{aligned}$$

and that $(T_\nu)_\# \nu_1 = \nu_2$ and $(T_p)_\# p_1 = p_2$. Note that we have $(T_\nu \times T_p)_\#(\nu_1 \times p_1) = \nu_2 \times p_2$. Thus, by definition of W_2 distance, it holds that

$$\begin{aligned} W_2^2(\nu_1 \times p_1, \nu_2 \times p_2) &\leq \int \left\| (\theta_1, \theta_2) - (T_\nu(\theta_1), T_p(\theta_1)) \right\|^2 \nu_1(d\theta_1) p_1(d\theta_2) \\ &= \int \left(\|\theta_1 - T_\nu(\theta_1)\|^2 + \|\theta_2 - T_p(\theta_2)\|^2 \right) \nu_1(d\theta_1) p_1(d\theta_2) \\ &= W_2^2(\nu_1, \nu_2) + W_2^2(p_1, p_2). \end{aligned}$$

Thus, we complete the proof of Lemma D.5. \square

Lemma D.6. For probability density function $\rho, \rho_1, \rho_2 \in \mathcal{P}_2(\Theta)$, let $\tilde{\rho}_1 = \alpha^{-1} \rho_1 + (1 - \alpha^{-1}) \rho$ and $\tilde{\rho}_2 = \alpha^{-1} \rho_2 + (1 - \alpha^{-1}) \rho$. Then, We have

$$W_2(\tilde{\rho}_1, \tilde{\rho}_2) \leq \alpha^{-1/2} W_2(\rho_1, \rho_2).$$

Proof. Recall the definition of Wasserstain-2 distance that

$$W_2(\rho_1, \rho_2) = \left(\inf_{\gamma \in \Gamma(\rho_1, \rho_2)} \int \|x - y\|^2 d\gamma(x, y) \right)^{1/2},$$

where $\Gamma(\rho_1, \rho_2)$ is the set of all couplings of ρ_1 and ρ_2 . We assume that the infimum is reached by $\gamma^*(x, y) \in \Gamma(\rho_1, \rho_2)$, i.e.,

$$W_2(\rho_1, \rho_2) = \left(\int \|x - y\|^2 d\gamma^*(x, y) \right)^{1/2}.$$

We denote by $\gamma'(x, y)$ the distribution such that

$$\gamma'(x, y) = \alpha^{-1} \gamma^*(x, y) + (1 - \alpha^{-1}) \rho(x) \delta(x - y),$$

where $\delta(x, y)$ is dirac delta function and it holds that $\gamma'(x, y) \in \Gamma(\tilde{\rho}_1, \tilde{\rho}_2)$. Hence, it follows that

$$\begin{aligned} W_2(\tilde{\rho}_1, \tilde{\rho}_2) &\leq \left(\int \|x - y\|^2 d\gamma'(x, y) \right)^{1/2} \\ &= \left(\int \|x - y\|^2 (\alpha^{-1} \gamma^*(x, y) + (1 - \alpha^{-1}) \rho(x) \delta(x - y)) d(x, y) \right)^{1/2} \\ &= \alpha^{-1/2} W_2(\rho_1, \rho_2). \end{aligned}$$

Thus, we complete the proof of Lemma D.6. \square

E Auxiliary Lemmas

We use the definition of absolutely continuous curves in $\mathcal{P}_2(\mathbb{R}^D)$ in [6].

Definition E.1 (Absolutely Continuous Curve). Let $\beta : [a, b] \rightarrow \mathcal{P}_2(\mathbb{R}^D)$ be a curve. Then, we say β is an absolutely continuous curve if there exists a square-integrable function $f : [a, b] \rightarrow \mathbb{R}$ such that

$$W_2(\beta_s, \beta_t) \leq \int_s^t f(\tau) d\tau$$

for any $a \leq s < t \leq b$.

Then, we have the following first variation formula.

Lemma E.2 (First Variation Formula, Theorem 8.4.7 in [6]). Given $\nu \in \mathcal{P}_2(\mathbb{R}^D)$ and an absolutely continuous curve $\mu : [0, T] \rightarrow \mathcal{P}_2(\mathbb{R}^D)$, let $\beta : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^D)$ be the geodesic connecting μ_t and ν . It holds that

$$\frac{d}{dt} \frac{W_2(\mu_t, \nu)^2}{2} = -\langle \dot{\mu}_t, \dot{\beta}_0 \rangle_{\mu_t, W_2},$$

where $\dot{\mu}_t = \partial_t \mu_t$, $\dot{\beta}_0 = \partial_s \beta_s |_{s=0}$, and the inner product is defined in (A.2).

Lemma E.3 (Eulerian Representation of Geodesics, Proposition 5.38 in [58]). Let $\beta : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^D)$ be a geodesic and v be the corresponding vector field such that $\partial_t \beta_t = -\text{div}(\beta_t \cdot v_t)$. It holds that

$$\frac{d}{dt} (\beta_t \cdot v_t) = -\text{div}(\beta_t \cdot v_t \otimes v_t),$$

where \otimes is the outer product between two vectors.

Lemma E.4 (Talagrand's Inequality, Corollary 2.1 in [44]). Let ν be $N(0, \kappa \cdot I_D)$. It holds for any $\mu \in \mathcal{P}_2(\mathbb{R}^D)$ that

$$W_2(\mu, \nu)^2 \leq 2D_{\text{KL}}(\mu \| \nu) / \kappa.$$

Lemma E.5 (Theorem 1 in [46]). Let μ, ν be two probability measures in $\mathcal{P}_2(\theta)$. Then, it holds that

$$W_2(\mu, \nu) \leq 2 \|\mu - \nu\|_{\dot{H}^{-1}(\nu)}.$$

F Conclusions and Limitations

In this work, we study the time evolution of a two-timescale AC represented by a two-layer neural network in the mean-field limit. Specifically, the actor updates its policy via proximal policy optimization, which is closely related to the replicator dynamics, while the critic updates by temporal-difference learning, which is captured by a semigradient flow in the Wasserstein space. By introducing a restarting mechanism, we establish the convergence and optimality of AC with two-layer overparameterized neural network. However, the study has potential limitations. In this work we only study the continuous-time limiting regime, which is an ideal setting with infinitesimal learning rates, and establish finite-time convergence and optimality guarantees. Finite-time results for the more realistic discrete-time setting is left for future research.