
Provably Efficient Neural GTD Algorithm for Off-policy Learning

Hoi-To Wai

The Chinese University of Hong Kong
Shatin, Hong Kong
htwai@se.cuhk.edu.hk

Zhuoran Yang

Princeton University
Princeton, NJ, USA
zy6@princeton.edu

Zhaoran Wang

Northwestern University
Evanston, IL, USA
zhaoranwang@gmail.com

Mingyi Hong

University of Minnesota
Minneapolis, MN, USA
mhong@umn.edu

Abstract

This paper studies a gradient temporal difference (GTD) algorithm using neural network (NN) function approximators to minimize the mean squared Bellman error (MSBE). For off-policy learning, we show that the minimum MSBE problem can be recast into a min-max optimization involving a pair of over-parameterized primal-dual NNs. The resultant formulation can then be tackled using a neural GTD algorithm. We analyze the convergence of the proposed algorithm with a 2-layer ReLU NN architecture using m neurons and prove that it computes an approximate optimal solution to the minimum MSBE problem as $m \rightarrow \infty$.

1 Introduction

Policy evaluation is a key problem in reinforcement learning (RL) whose goal is to estimate the value function of a given policy, i.e., the expected total reward of a discounted Markov decision process (MDP) starting from a given state. Among others, the temporal difference (TD) learning algorithm [Sutton, 1988] has been used to minimize the mean squared (projected) Bellman error (MSBE). For off-policy learning where the behavior policy differs from the target policy, gradient-based TD (GTD) learning algorithms [Sutton et al., 2009a,b] have been proposed with guaranteed convergence. A growing trend is to employ nonlinear approximation such as neural network (NN) functions [e.g., Chung et al., 2019, Haarnoja et al., 2018, Lillicrap et al., 2015, Mnih et al., 2016, Silver, 2012].

The TD/GTD learning algorithms have been analyzed with *linear function approximation* [Bhandari et al., 2018, Dalal et al., 2017]. Meanwhile, nonlinear TD/GTD learning algorithms (as well as other related problems such as Q-learning) are studied in [Bhatnagar et al., 2009, Brandfonbrener and Bruna, 2019, Chung et al., 2019, Dai et al., 2018, Wai et al., 2019], also see [Bertsekas, 2019]. However, these algorithms lack theoretical guarantees related to minimizing the MSBE or MSPBE objective as they may get stuck in a local optimum. Furthermore, these algorithms are designed for arbitrary nonlinear function approximation, whose actual implementations involve computationally intensive steps such as computing the Hessians for the nonlinear functions.

In this paper, we analyze the efficiency of an off-policy GTD learning algorithm with NN function approximation. We consider a simplified setting employing a pair of over-parameterized 2-layer ReLU NNs. A key result proven is that the proposed neural GTD algorithm is guaranteed to *converge globally* to a minimizer of the MSBE problem, despite the corresponding optimization problem is non-convex. Our main contributions are:

- We derive a new formulation for MSBE minimization with a *primal NN* and a *dual NN* as approximators. A neural GTD algorithm, which obeys similar update rules as the classical GTD2 algorithm, is proposed for the resultant min-max optimization problem.
- Under an off-policy learning setting, we analyze the convergence rates of the neural GTD algorithm with various sampling techniques, including population update, stochastic updates with i.i.d. samples and Markov samples.
- We focus on the 2-layer ReLU NN architecture. We show that when the width of the NN employed goes to infinity and the TD error function lies in the NN function class, the proposed neural GTD algorithm is guaranteed to find a *global minimizer* of the MSBE problem. Importantly, the convergence rates measured with the functional distance are independent of NN’s width.

Related Works Recent works have studied the *global convergence to optimal solutions* of different RL algorithms. Examples are [Cai et al., 2019, Wang et al., 2019, Xu and Gu, 2019] which studied neural TD learning, neural policy gradient, and neural Q-learning, respectively. These works rely on that in the limit when the number of neurons approaches infinity, the *nonlinear* NN function is locally approximated by a *linear* function. The present paper follows a similar philosophy, i.e., by reducing the neural GTD algorithm to a primal-dual method for solving a convex-concave problem. We develop new analysis to handle biases in gradients and off-policy learning settings.

Our work is also related to recent works on finite time guarantees for GTD learning algorithms using linear function approximation. For instance, Dalal et al. [2018, 2019], Liu et al. [2015] provide guarantees when the algorithms acquire i.i.d. samples; Du et al. [2017] apply variance reduction technique; Doan [2019], Gupta et al. [2019], Wang et al. [2017], Xu et al. [2019] consider when Markov samples are used. In comparison, we extend these analysis to using NN function approximation and offer finite-time guarantees in the overparameterization limit.

Lastly, the present work is related to recent advances in overparameterized (deep) NNs. Inspired by the classical works [Bartlett, 1998, Rahimi and Recht, 2008] and empirical studies in [Neysshabur and Li, 2019, Zhang et al., 2016], it has been recently shown that overparameterized NNs are efficiently learnable using gradient-type algorithms in [Allen-Zhu et al., 2019a,b, Arora et al., 2019, Jacot et al., 2018, Lee et al., 2019, Mei et al., 2018]. A key insight used is that NN functions exhibit an implicit local linearization, which allows one to ‘convexify’ the corresponding training problem and establish global convergence. This line of analysis is also known as ‘lazy training’ for NNs. Notice that the recent works [Allen-Zhu and Li, 2020, Chizat and Bach, 2018, Daniely, 2017] have suggested stronger theories for the generalization power of deep NNs beyond the ‘lazy training’ characterization. This paper develops the analysis by adapting the above results to neural GTD learning via the implicit local linearization. Unlike the above works, our challenge involves providing dimension-independent bounds in the function space for the off-policy learning algorithm.

2 Markov Decision Process

Consider a discounted markov decision process (MDP) described by (S, A, P^a, R, γ) where S is the possibly infinite state space and A is the action space. Given an action $a \in A$, the operator $P^a : S \times S \rightarrow \mathbb{R}_+$ is a Markov transition kernel such that for any measurable function f on S , we denote for any $s \in S$ that

$$\mathbb{E}[f(s') | s' \sim P^a(s, \cdot)] = P^a f(s) = \int_S f(y) P^a(s, dy), \quad (1)$$

where s' denotes the next state that is transitioned into. The function $R(s, a)$ is the reward received after action a in state s ; lastly, $\gamma \in (0, 1)$ is the discount factor. We assume $\sup_{s,a} R(s, a) \leq \bar{r}$.

The *target policy* $\pi(a|s)$ is the conditional probability of action $a \in A$ given the current state $s \in S$ [Szepesvári, 2010]. This induces a Markov chain with the transition kernel $P^\pi(s, \cdot) := \mathbb{E}_{a \sim \pi(\cdot|s)}[P^a(s, \cdot)]$. The *policy evaluation* problem aims at learning a *value function* $V^\pi : S \rightarrow \mathbb{R}$, defined as the discounted total reward starting from state s :

$$V^\pi(s) := \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \mid s_0 = s, a_t \sim \pi(\cdot|s_t), s_{t+1} \sim P^{a_t}(s_t, \cdot)\right]. \quad (2)$$

Applying the Bellman equation [Van Hasselt, 2012] shows that

$$0 = V^\pi(s) - \mathbb{E}_{a \sim \pi(\cdot|s)}[R(s, a)] - \gamma P^\pi V^\pi(s). \quad (3)$$

The existence of $V^\pi(s)$ follows if the Markov chain driven by P^π is aperiodic and irreducible. From the above, the policy evaluation problem may be solved by finding a value function $V^\pi(s)$ which

satisfies (3). This problem is challenging since (i) the state space \mathcal{S} is large, and (ii) the transition kernel $P^\pi(\cdot, \cdot)$ is unknown. The latter must be learnt from the observed state/action pairs while solving the policy evaluation simultaneously.

2.1 Off-policy Policy Evaluation with Function Approximation

We consider *nonlinear* approximation of the value function using two-layer neural network (NN) of m neurons with rectified linear units (ReLU)s. Each state $s \in \mathcal{S}$ is encoded by a d -dimensional vector, denoted by $\mathbf{x}_s \in \mathbb{R}^d$ with $\|\mathbf{x}_s\|_2 = 1$. The NN is over-parameterized with $m \gg 1$ neurons, each with the weight as $b_r \in \mathbb{R}$, and the parameter described by $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^{md}$ (for simplicity, we may assume $\Theta = \mathbb{R}^{md}$). Our aim is to approximate the value function in (3) as:

$$V^\pi(s) \approx f(s, \boldsymbol{\theta}) := \sum_{r=1}^m \frac{b_r}{\sqrt{m}} \text{ReLU}(\langle \boldsymbol{\theta}^{(r)}, \mathbf{x}_s \rangle) = \sum_{r=1}^m \frac{b_r}{\sqrt{m}} \mathbb{1}\{\langle \boldsymbol{\theta}^{(r)}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}^{(r)}, \mathbf{x}_s \rangle, \quad (4)$$

where $\mathbb{1}\{\cdot\}$ is the 0-1 indicator function with $\mathbb{1}\{\mathcal{E}\} = 1$ if \mathcal{E} is true; otherwise $\mathbb{1}\{\mathcal{E}\} = 0$. Notice that $f(\cdot, \boldsymbol{\theta})$ is differentiable with respect to $\boldsymbol{\theta}$ when $\langle \boldsymbol{\theta}^{(r)}, \mathbf{x}_s \rangle \neq 0$. We assume:

H1. Consider the NN function (4). The weights b_r are generated as $b_r \sim \mathcal{U}(\{-1, 1\})$. The parameters $\boldsymbol{\theta} = (\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(m)})$ are initialized with $\boldsymbol{\theta}_0^{(r)} \sim \mathcal{N}(\mathbf{0}, \frac{1}{d}\mathbf{I}_d)$. All weights/parameters are drawn independently. These initialization parameters are denoted by $\Xi_0 = (\boldsymbol{\theta}_0, b_1, \dots, b_m)$.

Fixing the NN weights at $(b_r)_{r=1}^m$. Let $B > 0$ be a fixed radius, we consider the NN function approximation (4) where $\boldsymbol{\theta}$ is taken from the ball $S_B = \{\boldsymbol{\theta} \in \Theta : \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\| \leq B\}$.

Off-policy Learning. We consider *off-policy learning*, where the observed state-action triplets (s, a, s') are generated from a *behavior policy* $\Pi(a|s) \neq \pi(a|s)$. Similarly, P^Π is the induced Markov kernel for the state sequence (s_0, s_1, \dots) . The Markov chain induced by P^Π is assumed ergodic with the unique stationary distribution μ^Π , and the supports of policies satisfy $\text{supp}\{\pi(\cdot|\cdot)\} \subseteq \text{supp}\{\Pi(\cdot|\cdot)\}$. Let the importance ratio and the temporal difference (TD) error be

$$\rho(a|s) := \frac{\pi(a|s)}{\Pi(a|s)}, \quad \delta(s, a, s'; \boldsymbol{\theta}) := f(s, \boldsymbol{\theta}) - \gamma f(s', \boldsymbol{\theta}) - R(s, a). \quad (5)$$

The following conditional expectation approximates the r.h.s. of (3):

$$\bar{\delta}(s, \boldsymbol{\theta}) := \mathbb{E}_{a \sim \Pi(\cdot|s), s' \sim P^a(s, \cdot)}[\rho(a|s)\delta(s, a, s'; \boldsymbol{\theta})] = \mathbb{E}_{a \sim \pi(\cdot|s), s' \sim P^a(s, \cdot)}[\delta(s, a, s'; \boldsymbol{\theta})]. \quad (6)$$

Our goal is to find the NN parameter $\boldsymbol{\theta} \in S_B$ to minimize the mean squared Bellman error (MSBE). Let $v \geq 0$ be the regularization parameter, we formulate the MSBE minimization problem

$$\min_{\boldsymbol{\theta} \in S_B} J(\boldsymbol{\theta}) := \mathbb{E}_{s \sim \mu^\Pi} \left[\frac{1}{2} \bar{\delta}(s, \boldsymbol{\theta})^2 + \frac{v}{2} f(s, \boldsymbol{\theta})^2 \right], \quad (7)$$

When $v = 0$, one may apply an off-policy modification to the neural TD learning in [Cai et al., 2019], with $\boldsymbol{\theta}' \leftarrow \boldsymbol{\theta} - \beta \rho(a|s) \delta(s, a, s'; \boldsymbol{\theta}) \nabla f(s, \boldsymbol{\theta})$, where $\beta > 0$. However, as observed by [Baird, 1995, Sutton et al., 2016], with off-policy samples the TD algorithm may diverge; also see §3.3. This is due to the non-Hurwitz matrix which forms the mean field of update, since the behavior and target policies are mismatched. Below, we develop an algorithm similar to *gradient-based* TD (GTD) learning which is shown to resolve the non-convergent issue.

Neural GTD Algorithm. Consider rewriting the MSBE function as:

$$\begin{aligned} \mathbb{E}_{s \sim \mu^\Pi} \left[\frac{1}{2} \bar{\delta}(s, \boldsymbol{\theta})^2 \right] &= \int_{\mathcal{S}} \frac{1}{2} \bar{\delta}(s, \boldsymbol{\theta})^2 \mu^\Pi(ds) \stackrel{(a)}{=} \int_{\mathcal{S}} \max_{y(s)} \{y(s) \bar{\delta}(s, \boldsymbol{\theta}) - \frac{1}{2} y(s)^2\} \mu^\Pi(ds) \\ &\stackrel{(b)}{=} \max_{y(\cdot)} \{ \mathbb{E}_{s \sim \mu^\Pi} [y(s) \bar{\delta}(s, \boldsymbol{\theta})] - \frac{1}{2} \mathbb{E}_{s \sim \mu^\Pi} [y(s)^2] \}, \end{aligned} \quad (8)$$

where (a) is due to $\frac{1}{2} \delta^2 = \max_y \{y\delta - \frac{y^2}{2}\}$, and (b) swapped the max and \int operators. Substituting (8) into (7) yields a non-convex min-max problem. Notice that the above reformulation is inspired by those derived in prior works such as [Dai et al., 2017, 2018, Shapiro, 2011].

The optimizer to the maximization (8) is given by $y^*(s, \boldsymbol{\theta}) = \bar{\delta}(s, \boldsymbol{\theta})$, which is the expected TD error. Notice that this leads to another intractable problem since $|\mathcal{S}|$ is large. As such, we consider applying an *additional NN function approximation*. We find an NN parameter $\mathbf{w} \in S_B$ such that

$$\bar{\delta}(s, \boldsymbol{\theta}) \approx f(s, \mathbf{w}), \quad \forall s \in \mathcal{S}. \quad (9)$$

Algorithm 1 Neural GTD algorithms for MSBE

- 1: **Input:** step sizes $(\beta_k)_{k \geq 0}$; maximum number of iterations n .
 - 2: Generate initialization parameters $\theta_0 = \mathbf{w}_0 \in \mathbb{R}^{md}$ with H1, and draw a random integer I_n :

$$\mathbb{P}(I_n = k) = \beta_k / \sum_{\ell=0}^n \beta_\ell, \quad k = 0, \dots, n. \quad (11)$$
 - 3: **for** $k = 0, 1, 2, \dots, I_n$ **do**
 - 4: (*I.i.d. sample*) Draw a state $s \sim \mu^\Pi$ and set $s_k = s$; or (*Markov sample*) set $s_k = s'_{k-1}$.
 - 5: Draw action $a_k \sim \Pi(\cdot | s_k)$, and $s'_k \sim P^{a_k}(s_k, \cdot)$ according to the behavior policy.
 - 6: Compute the gradient at θ_k as $\nabla \tilde{\delta}_k(\theta_k)$, where $\tilde{\delta}_k(\theta_k)$ is the empirical TD error:

$$\tilde{\delta}_k(\theta) := \rho(a_k | s_k) \{f(s_k, \theta) - \gamma f(s'_k, \theta) - R(s_k, a_k)\}. \quad (12)$$
 - 7: Let $\mathcal{P}_{S_B}(\cdot)$ be the Euclidean projection onto S_B , perform the updates as:

$$\begin{aligned} \theta_{k+1} &= \mathcal{P}_{S_B} \{ \theta_k - \beta_k f(s_k, \mathbf{w}_k) \nabla \tilde{\delta}_k(\theta_k) - \beta_k v f(s_k, \theta_k) \nabla f(s_k, \theta_k) \}, \\ \mathbf{w}_{k+1} &= \mathcal{P}_{S_B} \{ \mathbf{w}_k + \beta_k \tilde{\delta}_k(\theta_k) \nabla f(s_k, \mathbf{w}_k) - \beta_k f(s_k, \mathbf{w}_k) \nabla f(s_k, \mathbf{w}_k) \}. \end{aligned} \quad (13)$$
 - 8: **Return:** (approx.) optimal parameters $(\theta_{I_n}, \mathbf{w}_{I_n})$.
-

Using (9), problem (7) is approximated as a *non-convex non-concave* min-max problem:

$$\min_{\theta \in S_B} \max_{\mathbf{w} \in S_B} J_v(\theta, \mathbf{w}) := \mathbb{E}_{s \sim \mu^\Pi} \left[f(s, \mathbf{w}) \tilde{\delta}(s, \theta) - \frac{1}{2} f(s, \mathbf{w})^2 + \frac{v}{2} f(s, \theta)^2 \right]. \quad (10)$$

Observe that problem (10) involves the simultaneous optimization of *two NNs*. In addition to the ‘primal NN’ (θ), we employ a ‘dual NN’ (\mathbf{w}) to approximate the TD error (5). We propose a neural GTD algorithm in Algorithm 1. The algorithm, which shares similar update equations as GTD2 [Sutton et al., 2009a], is essentially a projected primal-dual gradient method on (10). Compared to nonlinear algorithms such as [Bhatnagar et al., 2009], the neural GTD algorithm does not involve computing the Hessians of NN which makes it more practical for the over-parameterized setting.

Remark 1. Dai et al. [2018] derived a similar reformulation of the MSBE function with nonlinear function approximation as (10). However, unlike the neural GTD which is a single loop algorithm, [Dai et al., 2018, Algorithm 1] requires solving an inner maximization for each iteration which possibly requires a simulator. Furthermore, the analysis therein is not suitable for the overparameterized NN setting as it involves bounding the approximation error using an ℓ_∞ norm.

Remark 2. We notice that the GTD2 algorithm [Sutton et al., 2009a] shares a similar update equation as the neural GTD, yet GTD2 was derived through minimizing a projected MSBE objective function. This coincidence can be justified as the approximation in (9) is exact when the expected TD error lies in the space of NN functions. As such, the latter implicitly imposed a function projection. A key difference between our derivation and an algorithm derived from an exact projected MSBE formulation is that the ours avoids a Hessians computation step; e.g., see [Bhatnagar et al., 2009].

3 Finding a Global Minimizer for MSBE

This section summarizes our main results. Our analysis hinges on the following linearized NN function:

$$\hat{f}(s, \theta) := \sum_{r=1}^m \frac{b_r}{\sqrt{m}} \mathbb{1} \{ \langle \theta_0^{(r)}, \mathbf{x}_s \rangle > 0 \} \langle \theta^{(r)}, \mathbf{x}_s \rangle \equiv \langle \theta, \ell(\mathbf{x}_s) \rangle, \quad (14)$$

where $\ell : \mathbb{R}^d \rightarrow \mathbb{R}^{md}$. Compared to (4), the function $\hat{f}(s, \theta)$ is identical to $f(s, \theta)$ except for fixing $\theta = \theta_0$ in the activation function $\mathbb{1}\{\cdot\}$. Moreover, $f(s, \theta)$ is differentiable w.r.t. θ . Note that $\|\ell(\mathbf{x}_s)\|_2^2 \leq 1$ for any $s \in \mathcal{S}$ and initialization Ξ_0 [cf. H1].

Our central idea is to treat $\hat{f}(s, \theta)$ as a surrogate to the *nonlinear* NN function $f(s, \theta)$. In particular, we analyze the convergence of neural GTD for finding a saddle point of the linearized problem:

$$\min_{\theta \in S_B} \max_{\mathbf{w} \in S_B} \hat{J}_v(\theta, \mathbf{w}) := \mathbb{E}_{s \sim \mu^\Pi} \left[\hat{f}(s, \mathbf{w}) \tilde{\delta}(s, \theta) - \frac{1}{2} \hat{f}(s, \mathbf{w})^2 + \frac{v}{2} \hat{f}(s, \theta)^2 \right], \quad (15)$$

where

$$\widehat{\delta}(s, \boldsymbol{\theta}) := \mathbb{E}_{a \sim \Pi(\cdot|s), s' \sim P^a(s, \cdot)} [\rho(a|s) \{ \widehat{f}(s, \boldsymbol{\theta}) - \gamma \widehat{f}(s', \boldsymbol{\theta}) - R(s, a) \}], \quad (16)$$

Problem (15) differs from (10) through linearizing the NN functions. Importantly, problem (15) is a convex-concave optimization with the saddle point denoted as $\widehat{\boldsymbol{z}}(\Xi_0) = (\widehat{\boldsymbol{\theta}}(\Xi_0), \widehat{\boldsymbol{w}}(\Xi_0))$.

Assuming that $m \gg 1$, in Section 3.1, we show that the neural GTD algorithm converges to a saddle point of (15); then in Section 3.2, we show that an optimal solution to (7) can be taken as the primal solution in a saddle point to (15). Finally, Corollary 3.1 shows that the neural GTD algorithm finds a global minimizer of MSBE.

3.1 Convergence to Saddle Point of (15)

We present convergence guarantees for neural GTD in two favors — first we study a population-based method using exact gradients; then we study the stochastic methods using i.i.d. and Markov samples. Before proceeding, we state the following assumption on the stationary distribution:

H2. *There exists a constant c_0 such that for any $\tau > 0$, $\mathbf{y} \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$, it holds almost surely that*

$$\mathbb{E}_{s \sim \mu^\Pi} [\mathbb{1}\{|\langle \mathbf{y}, \mathbf{x}_s \rangle| \leq \tau\} \|\mathbf{y}\|] \leq c_0 \tau / \|\mathbf{y}\|_2. \quad (17)$$

The above condition is a regulatory assumption which requires \mathbf{x}_s to be ‘uniformly’ distributed when s is drawn from μ^Π . To simplify notations, we let $\mathbf{z}_k := (\boldsymbol{\theta}_k^\top \mathbf{w}_k^\top)^\top$. Also, we define the L^2 distance

$$d_{\widehat{f}}(\mathbf{z}_k, \widehat{\boldsymbol{z}}(\Xi_0)) := \mathbb{E}_{s \sim \mu^\Pi} [|\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}}(\Xi_0))|^2 + |\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\boldsymbol{w}}(\Xi_0))|^2], \quad (18)$$

where $\widehat{\boldsymbol{z}}(\Xi_0)$ is a saddle point to (15) given the initial NN parameters Ξ_0 . If $d_{\widehat{f}}(\mathbf{z}_k, \widehat{\boldsymbol{z}}(\Xi_0)) \approx 0$, then \mathbf{z}_k gives the primal-dual NNs that are close in the function space to the NNs parameterized by $\widehat{\boldsymbol{z}}(\Xi_0)$.

Population Neural GTD. We study a population neural GTD algorithm with:

$$\boldsymbol{\theta}_{k+1} = \mathcal{P}_{S_B} \left\{ \boldsymbol{\theta}_k - \beta_k \nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k) \right\}, \quad \mathbf{w}_{k+1} = \mathcal{P}_{S_B} \left\{ \mathbf{w}_k + \beta_k \nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k) \right\}, \quad (19)$$

where we assume that the *exact* gradient of the population objective function $J_v(\boldsymbol{\theta}_k, \mathbf{w}_k)$ is available. This form of update is relevant to a ‘batch’ neural GTD algorithm where samples of state transitions are collected a-priori. We have the following finite-time convergence result for (19):

Theorem 3.1. *Assume that H1, H2 hold, and the step size satisfies $\sup_{k \geq 0} \beta_k \leq \frac{1 \wedge v}{8((1 \vee v)^2 + (1 + \gamma)^2)}$. For the iterates generated by (19) and any $n \geq 1$, it holds:*

$$\min_{k \in \{0, \dots, n\}} \mathbb{E}_{\text{init}} [d_{\widehat{f}}(\mathbf{z}_k, \widehat{\boldsymbol{z}}(\Xi_0))] \leq C_0^p \left(\frac{B^3}{(1 \wedge v) m^{\frac{1}{4}}} \right) + C_1^p \left(\frac{\|\mathbf{z}_0 - \widehat{\boldsymbol{z}}\|_2^2}{(1 \wedge v) \sum_{k=0}^n \beta_k} \right), \quad (20)$$

for some constants C_0^p, C_1^p that are independent of m, B , and $d_{\widehat{f}}(\mathbf{z}_k, \widehat{\boldsymbol{z}})$ is defined in (18). Moreover, the expectation $\mathbb{E}_{\text{init}}[\cdot]$ is taken over the initialization of the NN Ξ_0 .

Stochastic Neural GTD. Next we focus on using stochastic samples in Algorithm 1. First, in the simplest setting, the state pairs (s_k, s'_k) are drawn i.i.d. according to $s_k \sim \mu^\Pi$, $a_k \sim \Pi(\cdot|s_k)$, $s'_k \sim P^{a_k}(s_k, \cdot)$, see line 5 of the pseudo code. We can rewrite the algorithm as

$$\boldsymbol{\theta}_{k+1} = \mathcal{P}_{S_B} \left\{ \boldsymbol{\theta}_k - \beta_k (\nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}_k) + \mathbf{e}_k^{(1)}) \right\}, \quad \mathbf{w}_{k+1} = \mathcal{P}_{S_B} \left\{ \mathbf{w}_k + \beta_k (\nabla_{\mathbf{w}} J_v(\mathbf{z}_k) + \mathbf{e}_k^{(2)}) \right\}, \quad (21)$$

where $\mathbf{e}_k^{(i)}$, $i = 1, 2$ are the noise due to taking i.i.d. samples. Denote $\widetilde{s}_k := (s_k, a_k, s'_k)$, one has

$$\mathbf{e}_k^{(1)} = \nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k; \widetilde{s}_k) - \nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k), \quad \mathbf{e}_k^{(2)} = \nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k; \widetilde{s}_k) - \nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k), \quad (22)$$

where $\nabla J_v(\boldsymbol{\theta}_k, \mathbf{w}_k; \widetilde{s}_k)$ denote the sampled gradients using the state pairs \widetilde{s}_k , see (13). Denote \mathcal{F}_k as the filtration of the random variables $\{\boldsymbol{\theta}_0, \widetilde{s}_0, \widetilde{s}_1, \dots, \widetilde{s}_k\}$. Assume the following holds:

H3. *For any $k \geq 0$, there exists a constant σ such that for any $i = 1, 2$,*

$$\mathbb{E}[\mathbf{e}_k^{(i)} | \mathcal{F}_{k-1}] = 0, \quad \mathbb{E}[\|\mathbf{e}_k^{(i)}\|_2^2 | \mathcal{F}_{k-1}] \leq \sigma^2, \quad (23)$$

In other words, the noise vectors $e_k^{(i)}$, $i = 1, 2$ are martingale differences adapted to the filtration $(\mathcal{F}_k)_{k \geq 0}$. The second condition in (23) can be implied by the boundedness of z_k , in fact, using similar techniques as in Cai et al. [2019], it can be shown that $\sigma^2 = \mathcal{O}(B^2)$.

Theorem 3.2. Assume H1, H2, H3 hold and the step size satisfies $\sup_{k \geq 0} \beta_k \leq \frac{1 \wedge v}{12((1 \vee v)^2 + (1 + \gamma)^2)}$. For the iterates generated by Algorithm 1 and any $n \geq 1$, it holds:

$$\mathbb{E}_{I_n, \text{init}} [d_{\hat{f}}(z_{I_n}, \hat{z}(\Xi_0))] \leq C_0^s \left(\frac{B^3}{(1 \wedge v)m^{\frac{1}{4}}} \right) + C_1^s \left(\frac{\|z_0 - \hat{z}\|_2^2 + \sigma^2 \sum_{k=0}^n \beta_k^2}{(1 \wedge v) \sum_{k=0}^n \beta_k} \right), \quad (24)$$

for some constants C_0^s, C_1^s that are independent of m, B , and the function $d_{\hat{f}}(z_{I_n}, \hat{z}(\Xi_0))$ was defined in (18). The expectation above is taken over the independent r.v. I_n , the initialization to the NN Ξ_0 , and the i.i.d. samples of states drawn from behavior policy during the algorithm.

Discussions of Theorem 3.1 and 3.2. The conclusions (20), (24) imply that neural GTD finds a saddle point to the regularized MSBE problem with linearized NN function (15). For the population neural GTD, if we set β_k to be constant, then the last term in (20) decays to zero at the rate $\mathcal{O}(1/n)$; for the stochastic neural GTD, setting a step size as $\beta_k = \mathcal{O}(1/\sqrt{k})$ and the last term in (24) decays at the rate of $\mathcal{O}(\log n/\sqrt{n})$. These rates are comparable to exact and stochastic primal-dual gradient methods, respectively, see [Chambolle and Pock, 2016, Juditsky et al., 2011]. Meanwhile, the first terms represent the bias controlled by the width of the 2-layer NN, and the bias is in the order of $\mathcal{O}(B^3 m^{-1/4})$. Importantly, we observe that the error bound converges to zero when $n, m \rightarrow \infty$.

For both theorems, the error bound $d_{\hat{f}}(z_k, \hat{z})$ computes the L^2 distance between the linearized NN functions, taken over behavior policy's stationary distribution μ^Π . This optimality measure on the function space is used instead of the Euclidean norm $\|z_k - \hat{z}\|_2^2$ of the parameter space so as to avoid trivial bounds since z_k is $2md$ -dimensional (and we consider $m \rightarrow \infty$).

Markov Samples. An alternative version of Theorem 3.2 is derived when Markov samples are used, i.e., the sampled state-pair $\tilde{s}_k = (s_k, a_k, s_{k+1})$ is drawn from a single sample path of the Markov chain induced by the MDP and behavior policy Π , see line 4 & 7 of Algorithm 1.

Theorem 3.3. Assume H1, the step size satisfies $|\beta_k - \beta_{k+1}| \leq \xi \beta_k^2$ for some constant ξ , $\sup_{k \geq 0} \beta_k \leq \frac{1 \wedge v}{12((1 \vee v)^2 + (1 + \gamma)^2)}$, $\sup_{a,s} \rho(a|s) = \bar{\rho}$. Consider Algorithm 1 with Markov samples and any $n \geq 1$. With probability at least $1 - e^{-\Omega(\log^2 m)}$ over NN initializations, it holds

$$\mathbb{E}_{I_n} [d_{\hat{f}}(z_{I_n}, \hat{z})] = \mathcal{O} \left(\frac{B^{\frac{8}{3}} (\log m)^{\frac{3}{2}} / (1 - \rho)}{(1 \wedge v)m^{\frac{1}{6}}} + \frac{\|z_0 - \hat{z}\|_2^2 + \frac{B^2}{1 - \rho} + (\frac{B^2}{1 - \rho}) \sum_{k=0}^n \beta_k^2}{(1 \wedge v) \sum_{k=0}^n \beta_k} \right), \quad (25)$$

where $d_{\hat{f}}(z_{I_n}, \hat{z}(\Xi_0))$ was defined in (18) and $\rho \in [0, 1)$ is the convergence rate of the Markov chain. The expectation is taken over for r.v. I_n , and the sample path of Markov chain $(\tilde{s}_0, \tilde{s}_1, \dots)$.

Details are in Appendix D where we specify additional conditions on the Markov chain induced by the behavior policy. There are two differences from Theorem 3.2. First, the above holds in high probability w.r.t. the NN initialization. Second, the bias, variance are $\mathcal{O}(B^{\frac{8}{3}} (\log m)^{\frac{3}{2}} (1 - \rho)^{-1} m^{-1/6})$, $\mathcal{O}(B^2 (1 - \rho)^{-1})$, which depends on the mixing time of Markov chain.

3.2 Minimizing the MSBE with Neural GTD Algorithms

Theorems 3.1 & 3.2 show the neural GTD algorithms converge to an optimal solution of (15) when $n, m \rightarrow \infty$. To show that an optimal solution of (15) is also optimal to (7), we consider:

H4. For any $\theta \in S_B$, there exists $w(\theta) \in S_B$ such that $\mathbb{E}_{\text{init}, s \sim \mu^\Pi} [|\bar{\delta}(s, \theta) - f(s, w(\theta))|^2] \leq c_{\text{nn}}$, where $c_{\text{nn}} \geq 0$ and the expectation is taken w.r.t. $s \sim \mu^\Pi$ and the NN initializations Ξ_0 .

The above assumption depends on the reward function $R(s, a)$. Particularly, we have $c_{\text{nn}} = 0$ if the TD error function lies in the function class of 2-layer ReLU NNs. Furthermore, we anticipate that $c_{\text{nn}} \ll 1$ under the overparameterization setting $m \gg 1$. This is due to the representation power of such NNs as demonstrated in the recent works, e.g., [Neyshabur and Li, 2019].

Based on the above results, for any $\theta \in S_B$, we can control the MSBE $J_v(\theta) - J_v(\theta^*)$ as

Theorem 3.4. Assume H1,H2,H4, and the importance ratio is bounded as $\sup_{a,s} \rho(a|s) = \bar{\rho}$. Let $\theta^*(\Xi_0)$ be an optimal solution to (7). For any $\theta \in S_B$,

$$\begin{aligned} \mathbb{E}_{\text{init}} [J_v(\theta) - J_v(\theta^*(\Xi_0))] &\leq C_0^J \mathbb{E}_{\text{init}} [d_{\hat{f}}(z, \hat{z}(\Xi_0))] \\ &+ C_1^J (B + B^{\frac{3}{2}} m^{-\frac{1}{4}}) \sqrt{\mathbb{E}_{\text{init}} [d_{\hat{f}}(z, \hat{z}(\Xi_0))]} + \mathcal{O}(B^3 m^{-1/2} + B^{5/2} m^{-1/4} + c_{\text{nn}}) \end{aligned} \quad (26)$$

for some constants C_0^J, C_1^J that are independent of B, m . In the above, z is defined as the vector $z = (\theta, w)$ for any $w \in S_B$, and $d_{\hat{f}}(z, \hat{z}(\Xi_0))$ was defined in (18).

The difference $J_v(\theta) - J_v(\theta^*(\Xi_0))$ corresponds to the sub-optimality of a given NN parameter θ to the regularized mean squared Bellman error objective function. The above theorem quantifies this sub-optimality in terms of the distance to a saddle point of the problem (15).

Combining Theorem 3.4 with the previous analysis in Theorem 3.1 & 3.2, we obtain the *global convergence* guarantees of MSBE for the neural GTD algorithms as follows:

Corollary 3.1. Assume H1,H2,H4 and the importance ratio is bounded as $\sup_{a,s} \rho(a|s) = \bar{\rho}$. We have the following guarantees for Neural GTD algorithms:

- Consider the population neural GTD algorithm [cf. (19)]. Set $\beta_k = \mathcal{O}(1)$. For any $n \geq 1$,

$$\begin{aligned} \min_{k \in \{0, \dots, n\}} \mathbb{E}_{\text{init}} [J_v(\theta_k) - J_v(\theta^*(\Xi_0))] &= \mathcal{O}(n^{-1} + (B + B^{\frac{3}{2}} m^{-\frac{1}{4}}) n^{-\frac{1}{2}}) \\ &+ \mathcal{O}(B^3 m^{-\frac{1}{4}} + B^{\frac{5}{2}} m^{-\frac{1}{8}} + c_{\text{nn}}), \end{aligned} \quad (27)$$

- Consider the stochastic neural GTD algorithm with i.i.d. samples [cf. Algorithm 1]. Assume in addition H3 and set $\beta_k = \mathcal{O}(1/\sqrt{k})$. For any $n \geq 1$,

$$\begin{aligned} \mathbb{E}_{\text{init}, I_n} [J_v(\theta_{I_n}) - J_v(\theta^*(\Xi_0))] &= \tilde{\mathcal{O}}(\sigma^2 n^{-\frac{1}{2}} + (B + B^{\frac{3}{2}} m^{-\frac{1}{4}}) \sigma n^{-\frac{1}{4}}) \\ &+ \mathcal{O}(B^3 m^{-\frac{1}{4}} + B^{\frac{5}{2}} m^{-\frac{1}{8}} + c_{\text{nn}}), \end{aligned} \quad (28)$$

where the $\tilde{\mathcal{O}}(\cdot)$ notation hides the logarithm terms in the upper bounds.

The expectations above are taken over the NN initialization Ξ_0 and the number of iterations I_n .

When the bias term c_{nn} is small, the above corollary shows that as $n, m \rightarrow \infty$, the neural GTD algorithms find an NN parameter θ_{I_n} which *globally minimizes* the MSBE (7). Moreover, similar conclusions can be drawn for the Markov sample settings.

3.3 Preliminary Numerical Experiment

We perform preliminary experiments to support the above theories on a toy example of off-policy learning. We consider an MDP taken from the Garnet class with $|S| = 500$ states, $|A| = 5$ possible actions per state with uniformly distributed rewards, and the discount factor is $\gamma = 0.9$. We generate two random policies with the same support as the behavior/target policies, respectively. In Fig. 1, we compare the average MSBE against the number of neurons m , using a 2-layer, ReLU NN with random initialization according to H1, after $T = 3 \times 10^5$ iterations of neural GTD and neural TD [Cai et al., 2019] run with Markovian samples [cf. Algorithm 1], from 10 independent runs of state/action.

From the figure, we observe that the average MSBE (solid line) decreases with m stably for neural GTD, as predicted by the above theorems. Meanwhile, the MSBE fluctuates with m with neural TD, indicating that the latter can be unstable in the tested off-policy setting. Note that neural TD algorithm [Cai et al., 2019] has only been analyzed with on-policy data.

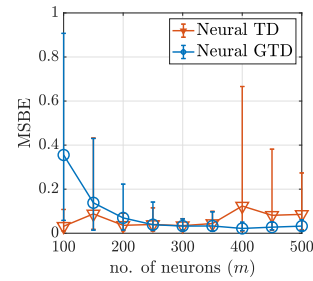


Figure 1: Comparing the averaged MSBE after 10 runs of Neural TD and Neural GTD in off-policy learning.

4 Proof Sketches

This section highlights the major steps involved in showing our main claims from the previous section. We first review on the approximation quality of the NN function, and develop its consequences in

the GTD learning paradigm. This will lead to our main theorems in Theorem 3.1 & 3.2. Then, we perform a perturbation analysis in light of $J_v(\boldsymbol{\theta}; \Xi_0)$ to yield Theorem 3.4.

Approximating an NN function by Linearization. To establish Theorem 3.1 & 3.2, we observe:

Lemma 4.1. [Cai et al., 2019, Lemma 5.1] Under H1, H2, there exists constant c_0 where for any $\boldsymbol{\theta} \in S_B$,

$$\mathbb{E}_{\text{init}, s \sim \mu^\Pi} [|f(s, \boldsymbol{\theta}) - \widehat{f}(s, \boldsymbol{\theta})|^2] \leq c_0 B^3 m^{-1/2}. \quad (29)$$

Essentially, the expected approximation error of the function values $f(\cdot, \boldsymbol{\theta})$ by $\widehat{f}(\cdot, \boldsymbol{\theta})$ decays as $\mathcal{O}(m^{-1/2})$ for any $\boldsymbol{\theta} \in S_B$. As $m \rightarrow \infty$, the NN function behaves like a *linear function*.

Neural GTD as Biased Gradient. Our next step is to analyze the convergence of neural GTD learning [cf. (19) or Algorithm 1]. To this regard, we treat the algorithms as *biased primal-dual gradient* methods for (15), even though the updates have been designed for (10). Concretely, we consider the population neural GTD (19). Observe that

$$\nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k) = \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\boldsymbol{\theta}_k, \mathbf{w}_k) + \widehat{\mathbf{e}}_k^{(1)}, \quad \nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}_k, \mathbf{w}_k) = \nabla_{\mathbf{w}} \widehat{J}_v(\boldsymbol{\theta}_k, \mathbf{w}_k) + \widehat{\mathbf{e}}_k^{(2)}, \quad (30)$$

where $\widehat{\mathbf{e}}_k^{(i)}$, $i = 1, 2$ represent the discrepancies between the gradient of $\widehat{J}_v(\boldsymbol{\theta}, \mathbf{w})$, $J_v(\boldsymbol{\theta}, \mathbf{w})$:

$$\begin{aligned} \widehat{\mathbf{e}}_k^{(1)} &= \mathbb{E} [f(s, \mathbf{w}_k) \nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k) \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k) + v \{ f(s, \boldsymbol{\theta}_k) \nabla f(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \boldsymbol{\theta}_k) \nabla \widehat{f}(s, \boldsymbol{\theta}_k) \}] \\ \widehat{\mathbf{e}}_k^{(2)} &= \mathbb{E} [(\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} \widehat{f}(s, \mathbf{w}_k)], \end{aligned}$$

where the expectations are taken w.r.t. $s \sim \mu^\Pi$. We recall that $\widehat{\delta}(s, \boldsymbol{\theta}_k)$ is the TD error with the linearized NN function parameterized by $\boldsymbol{\theta}_k$. Therefore, each of the above terms represent differences between the NN function and its linear approximation.

As observed in Lemma 4.1, the above discrepancies in the function value diminishes as $m \rightarrow \infty$. This observation also extends to the associated gradients as we prove that

Lemma 4.2. Under H1, H2, it holds for any $k \geq 0$ that

$$\mathbb{E}_{\text{init}} [\|\widehat{\mathbf{e}}_k^{(1)}\|_2^2] \vee \mathbb{E}_{\text{init}} [\|\widehat{\mathbf{e}}_k^{(2)}\|_2^2] \leq C_0 B^3 m^{-1/2}, \quad (31)$$

for some constant C_0 that is independent of B, m , the above expectations are taken with respect to the initialization parameters Ξ_0 of the NN.

From Lemma 4.2, it is clear that if m is sufficiently large, then the population neural GTD algorithm (19) follow closely a primal-dual gradient method for the *convex-concave* problem (15).

If we assume that $\widehat{\mathbf{e}}_k^{(1)} = \widehat{\mathbf{e}}_k^{(2)} = 0$, then the convergence of (19) to a global optimal solution is guaranteed by the classical analysis from, e.g., [Chambolle and Pock, 2016, He and Yuan, 2012]. Fix $n \in \mathbb{N}$, the results from [Chambolle and Pock, 2016] shows that using a slight modification to (19), one can find an $\mathcal{O}(1/n)$ saddle point $(\bar{\boldsymbol{\theta}}_n, \bar{\mathbf{w}}_n)$ in n iterations. However, such result does not immediately relate to a bound on the MSBE $J_v(\boldsymbol{\theta}_n)$ which will be needed later. In addition, the analysis in [Chambolle and Pock, 2016] does not consider bias in the gradient (30). Lastly, although the objective function $\widehat{J}_v(\boldsymbol{\theta}, \mathbf{w})$ admits a quadratic form, we note that the parameters $\boldsymbol{\theta}, \mathbf{w}$ are both md -dimensional vectors. As $m \rightarrow \infty$, the strong convexity/concavity modulus in the Euclidean space associated with the objective function may approach zero. Also see [Lorenz and Pock, 2015] for extensions of the primal-dual algorithm to Hilbert space.

In light of the above challenges, we adopt an error bound metric that adapts to the problem structure at hand. A natural choice is the L^2 distance between $\widehat{f}(\cdot, \boldsymbol{\theta})$, $\widehat{f}(\cdot, \widehat{\boldsymbol{\theta}})$ in the *function space*, as defined in $d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})$ [cf. (18)]. Define the primal-dual gradient operator on $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{w})$ as:

$$\widehat{\Phi}(\mathbf{z}) := (\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z})^\top - \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z})^\top)^\top. \quad (32)$$

Lemma 4.3. Let $\widehat{\mathbf{z}}$ be a saddle point of the problem (15). For any $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{w}) \in S_B \times S_B$. Set $\mu = \min\{1, v\}$ and $L_\Phi = 4((\min\{1, v\})^2 + (1 + \gamma)^2)$, it holds

$$\langle \widehat{\Phi}(\mathbf{z}) - \widehat{\Phi}(\widehat{\mathbf{z}}), \mathbf{z} - \mathbf{z}^* \rangle \geq \mu d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}}), \quad \|\widehat{\Phi}(\mathbf{z}) - \widehat{\Phi}(\widehat{\mathbf{z}})\|_2^2 \leq L_\Phi d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}}). \quad (33)$$

Note that the constants μ, L_Φ are *independent* of the problem dimension m . The condition (33) shows that the primal-dual gradient operator $\widehat{\Phi}(\mathbf{z})$ is a smooth monotone operator.

Lemma 4.2 & 4.3 show that the population neural GTD algorithm is a *biased* primal-dual gradient method on the convex-concave saddle point problem (15). In the appendix, we will show

$$\mathbb{E}[\|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2] \leq \mathbb{E}[\|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\mu\beta_k d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + \beta_k^2 L_\Phi d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}})] + \mathcal{O}(\beta_k B^3/m^{\frac{1}{4}}), \quad (34)$$

where the expectation is taken with respect to the NN initialization. Taking a summation of the above inequalities from $k = 0$ to $k = n$ and canceling terms yield Theorem 3.1.

In the i.i.d. sample case, the stochastic neural GTD algorithm can be analyzed in a similar manner through exploiting the conditional zero mean and bounded variance properties in H3. The latter leads to Theorem 3.2. In the Markov sample case, we utilize the Poisson equation which decomposes the noise terms into a martingale part and a Markov part. We show that the Markov part is small in magnitude. The derivations are similar to [Karimi et al., 2019], yet we have adapted the analysis to the primal-dual gradient method. In addition, inspired by [Gao et al., 2019], we derive similar bounds to Lemma 4.2 which hold in high probability over the initialization. This leads to Theorem 3.3, see Appendix D for the details.

Finding Global Minimizer of MSBE. Our last task is to evaluate the solution quality of the output from neural GTD in terms of deviation from the minimum regularized MSBE (7).

To derive Theorem 3.4, we shall exploit H4 and the Fenchel's conjugation in (8). In particular, it can be shown that

$$J_v(\boldsymbol{\theta}) \stackrel{(a)}{\leq} J_v(\widehat{\boldsymbol{\theta}}) + \mathcal{O}(m^{-1/4} + d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})^{\frac{1}{2}}) \stackrel{(b)}{\leq} J_v(\widehat{\boldsymbol{\theta}}, \mathbf{w}(\widehat{\boldsymbol{\theta}})) + \mathcal{O}(m^{-1/4} + d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})^{\frac{1}{2}} + c_{\text{nn}})$$

where (a) can be derived using Lemma 4.1, and (b) is due to H4 which guarantees the existence of $\mathbf{w}(\widehat{\boldsymbol{\theta}})$. Note that $J_v(\widehat{\boldsymbol{\theta}}, \mathbf{w}(\widehat{\boldsymbol{\theta}}))$ is the primal-dual objective function defined in (10). Subsequently, using Lemma 4.1, we obtain

$$J_v(\widehat{\boldsymbol{\theta}}, \mathbf{w}(\widehat{\boldsymbol{\theta}})) \leq \widehat{J}_v(\widehat{\boldsymbol{\theta}}, \mathbf{w}(\widehat{\boldsymbol{\theta}})) + \mathcal{O}(m^{-1/4}) \leq \widehat{J}_v(\boldsymbol{\theta}^*, \widehat{\mathbf{w}}) + \mathcal{O}(m^{-1/4}) \quad (35)$$

where the last inequality is due to $\widehat{J}_v(\widehat{\boldsymbol{\theta}}, \mathbf{w}(\widehat{\boldsymbol{\theta}})) \leq \widehat{J}_v(\widehat{\boldsymbol{\theta}}, \widehat{\mathbf{w}})$ and the fact $(\widehat{\boldsymbol{\theta}}, \widehat{\mathbf{w}})$ is a saddle point to (15). Applying Lemma 4.1 again yields the inequality

$$\widehat{J}_v(\boldsymbol{\theta}^*, \widehat{\mathbf{w}}) \leq J_v(\boldsymbol{\theta}^*, \widehat{\mathbf{w}}) + \mathcal{O}(m^{-1/4}) \leq J_v(\boldsymbol{\theta}^*) + \mathcal{O}(m^{-1/4}), \quad (36)$$

where the last inequality is due to the optimality of $\bar{\delta}(\cdot; \boldsymbol{\theta}^*)$ for the maximization in (8). We remark that the above inequalities hold in an expectation taken over the initialization Ξ_0 of NN [cf. H1].

Collecting terms in the above leads to Theorem 3.4 which shows

$$J_v(\boldsymbol{\theta}) - J_v(\boldsymbol{\theta}^*) = \mathcal{O}(m^{-1/4} + c_{\text{nn}} + d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})^{\frac{1}{2}}). \quad (37)$$

Combined with Theorem 3.1 & 3.2, we conclude that the neural GTD algorithms find a global minimizer to the regularized MSBE problem (7), i.e., justifying our main claims in Corollary 3.1.

5 Conclusions

We have derived the first neural GTD learning algorithm for off-policy learning and proved its *global convergence* to a minimizer of the regularized MSBE. The main idea is to use a Fenchel conjugate's equivalent formulation to the MSBE objective function and design a novel objective function that involves two NNs. We consider different sample requirements (population, i.i.d. samples and Markov samples), and analyze the convergence rates to a global MSBE minimizer.

Acknowledgement & Funding Disclosure The authors would like to thank Mr. Alan Lun (CUHK) for conducting the preliminary numerical experiments in this paper. H.-T. Wai is supported by the CUHK Direct Grant #4055113. M. Hong is supported in part by NSF under Grant CCF-1651825, CMMI-172775, CIF-1910385 and by AFOSR under grant 19RT0424.

Broader Impact This work does not present any foreseeable societal consequence.

References

- Z. Allen-Zhu and Y. Li. Backward feature correction: How deep learning performs deep learning. *arXiv preprint arXiv:2001.04413*, 2020.
- Z. Allen-Zhu, Y. Li, and Y. Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6155–6166, 2019a.
- Z. Allen-Zhu, Y. Li, and Z. Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019b.
- S. Arora, S. S. Du, W. Hu, Z. Li, and R. Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *ICML*, 2019.
- L. Baird. Residual algorithms: Reinforcement learning with function approximation. In *International Conference on Machine Learning*, pages 30–37, 1995.
- P. Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Transactions on Information Theory*, 44(2):525–536, 1998.
- D. P. Bertsekas. Feature-based aggregation and deep reinforcement learning: A survey and some new implementations. *IEEE/CAA Journal of Automatica Sinica*, 6(1):1–31, 2019.
- J. Bhandari, D. Russo, and R. Singal. A finite time analysis of temporal difference learning with linear function approximation. *arXiv preprint arXiv:1806.02450*, 2018.
- S. Bhatnagar, D. Precup, D. Silver, R. S. Sutton, H. R. Maei, and C. Szepesvári. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Advances in Neural Information Processing Systems*, pages 1204–1212, 2009.
- D. Brandfonbrener and J. Bruna. On the expected dynamics of nonlinear td learning. *arXiv preprint arXiv:1905.12185*, 2019.
- Q. Cai, Z. Yang, J. D. Lee, and Z. Wang. Neural temporal-difference learning converges to global optima. In *Advances in Neural Information Processing Systems*, pages 11312–11322, 2019.
- A. Chambolle and T. Pock. On the ergodic convergence rates of a first-order primal–dual algorithm. *Mathematical Programming*, 159(1-2):253–287, 2016.
- L. Chizat and F. Bach. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 2018.
- W. Chung, S. Nath, A. Joseph, and M. White. Two-timescale networks for nonlinear value function approximation. In *ICLR*, 2019.
- B. Dai, N. He, Y. Pan, B. Boots, and L. Song. Learning from conditional distributions via dual embeddings. In *Artificial Intelligence and Statistics*, pages 1458–1467, 2017.
- B. Dai, A. Shaw, L. Li, L. Xiao, N. He, Z. Liu, J. Chen, and L. Song. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pages 1125–1134, 2018.
- G. Dalal, B. Szorenyi, G. Thoppe, and S. Mannor. Finite sample analysis of two-timescale stochastic approximation with applications to reinforcement learning. *arXiv preprint arXiv:1703.05376*, 2017.
- G. Dalal, G. Thoppe, B. Szörényi, and S. Mannor. Finite sample analysis of two-timescale stochastic approximation with applications to reinforcement learning. In *Conference On Learning Theory*, pages 1199–1233, 2018.
- G. Dalal, B. Szorenyi, and G. Thoppe. A tale of two-timescale reinforcement learning with the tightest finite-time bound. *arXiv preprint arXiv:1911.09157*, 2019.

- A. Daniely. SGD learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*, pages 2422–2430, 2017.
- T. T. Doan. Finite-time analysis and restarting scheme for linear two-time-scale stochastic approximation. *arXiv preprint arXiv:1912.10583*, 2019.
- R. Douc, E. Moulines, P. Priouret, and P. Soulier. *Markov chains*. Springer, 2018.
- S. S. Du, J. Chen, L. Li, L. Xiao, and D. Zhou. Stochastic variance reduction methods for policy evaluation. In *International Conference on Machine Learning*, pages 1049–1058, 2017.
- G. Fort, E. Moulines, P. Priouret, et al. Convergence of adaptive and interacting markov chain monte carlo algorithms. *The Annals of Statistics*, 39(6):3262–3289, 2011.
- R. Gao, T. Cai, H. Li, C.-J. Hsieh, L. Wang, and J. D. Lee. Convergence of adversarial training in overparametrized neural networks. In *Advances in Neural Information Processing Systems*, pages 13009–13020, 2019.
- H. Gupta, R. Srikant, and L. Ying. Finite-time performance bounds and adaptive learning rate selection for two time-scale reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 4706–4715, 2019.
- T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- B. He and X. Yuan. Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective. *SIAM Journal on Imaging Sciences*, 5(1):119–149, 2012.
- A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- B. Karimi, B. Miasojedow, E. Moulines, and H.-T. Wai. Non-asymptotic analysis of biased stochastic approximation scheme. In *Conference on Learning Theory*, pages 1944–1974, 2019.
- J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in neural information processing systems*, pages 8570–8581, 2019.
- T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- B. Liu, J. Liu, M. Ghavamzadeh, S. Mahadevan, and M. Petrik. Finite-sample analysis of proximal gradient td algorithms. In *UAI*, pages 504–513, 2015.
- D. A. Lorenz and T. Pock. An inertial forward-backward algorithm for monotone inclusions. *Journal of Mathematical Imaging and Vision*, 51(2):311–325, 2015.
- S. Mei, A. Montanari, and P.-M. Nguyen. A mean field view of the landscape of two-layer neural networks. *PNAS*, 115(33):E7665–E7671, 2018.
- V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.
- B. Neyshabur and Z. Li. Towards understanding the role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2008.
- A. Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, 2011.
- D. Silver. Gradient temporal difference networks. In *EWRL*, pages 117–130, 2012.
- R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1): 9–44, 1988.
- R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *International Conference on Machine Learning*, pages 993–1000, 2009a.

- R. S. Sutton, H. R. Maei, and C. Szepesvári. A convergent $o(n)$ temporal-difference algorithm for off-policy learning with linear function approximation. In *Advances in Neural Information Processing Systems*, pages 1609–1616, 2009b.
- R. S. Sutton, A. R. Mahmood, and M. White. An emphatic approach to the problem of off-policy temporal-difference learning. *The Journal of Machine Learning Research*, 17(1):2603–2631, 2016.
- C. Szepesvári. Algorithms for reinforcement learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 4(1):1–103, 2010.
- H. Van Hasselt. Reinforcement learning in continuous state and action spaces. In *Reinforcement learning*, pages 207–251. Springer, 2012.
- H.-T. Wai, M. Hong, Z. Yang, Z. Wang, and K. Tang. Variance reduced policy evaluation with smooth function approximation. In *Advances in Neural Information Processing Systems*, pages 5776–5787, 2019.
- L. Wang, Q. Cai, Z. Yang, and Z. Wang. Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*, 2019.
- Y. Wang, W. Chen, Y. Liu, Z.-M. Ma, and T.-Y. Liu. Finite sample analysis of the GTD policy evaluation algorithms in Markov setting. In *Advances in Neural Information Processing Systems*, pages 5504–5513, 2017.
- P. Xu and Q. Gu. A finite-time analysis of q-learning with neural network function approximation. *arXiv preprint arXiv:1912.04511*, 2019.
- T. Xu, S. Zou, and Y. Liang. Two time-scale off-policy td learning: Non-asymptotic analysis over markovian samples. In *Advances in Neural Information Processing Systems*, pages 10633–10643, 2019.
- C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

A Proof of Theorem 3.1

We shall use the notations $(a \vee b) = \max\{a, b\}$ and $(a \wedge b) = \min\{a, b\}$.

Recall that (15) is given by:

$$\min_{\boldsymbol{\theta} \in S_B} \max_{\mathbf{w} \in S_B} \widehat{J}_v(\boldsymbol{\theta}, \mathbf{w}) := \mathbb{E}_{s \sim \mu^\pi} [\widehat{f}(s, \mathbf{w}) \widehat{\delta}(s, \boldsymbol{\theta})] - \frac{1}{2} \mathbb{E}_{s \sim \mu^\pi} [\widehat{f}(s, \mathbf{w})^2] + \frac{v}{2} \mathbb{E}_{s \sim \mu^\pi} [\widehat{f}(s, \boldsymbol{\theta})^2]. \quad (38)$$

To begin our analysis, denote $\mathbf{z}_k := (\boldsymbol{\theta}_k, \mathbf{w}_k)$ and the optimal primal-dual solution to (15) (or (38)) as $\widehat{\mathbf{z}} = (\widehat{\boldsymbol{\theta}}, \widehat{\mathbf{w}})$. In this regard, we can write the population update as:

$$\boldsymbol{\theta}_{k+1} = \mathcal{P}_{S_B} \left\{ \boldsymbol{\theta}_k - \beta_k \left(\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} \right) \right\}, \quad \mathbf{w}_{k+1} = \mathcal{P}_{S_B} \left\{ \mathbf{w}_k + \alpha_k \left(\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} \right) \right\}, \quad (39)$$

with the errors defined as:

$$\begin{aligned} \widehat{\mathbf{e}}_k^{(1)} &= \mathbb{E}_{s \sim \mu^\pi} \left[f(s, \mathbf{w}_k) \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k) + v f(s, \boldsymbol{\theta}_k) \nabla f(s, \boldsymbol{\theta}_k) - (\widehat{f}(s, \mathbf{w}_k) \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k) + v \widehat{f}(s, \boldsymbol{\theta}_k) \nabla \widehat{f}(s, \boldsymbol{\theta}_k)) \right] \\ \widehat{\mathbf{e}}_k^{(2)} &= \mathbb{E}_{s \sim \mu^\pi} \left[(\widehat{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} f(s, \mathbf{w}_k) - ((\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} \widehat{f}(s, \mathbf{w}_k)) \right], \end{aligned} \quad (40)$$

whose magnitude is controlled by the following lemma:

Lemma A.1. *Under H1, H2, it holds that for any $k \geq 0$, we have*

$$\mathbb{E}_{\text{init}} [\|\widehat{\mathbf{e}}_k^{(1)}\|_2^2] \vee \mathbb{E}_{\text{init}} [\|\widehat{\mathbf{e}}_k^{(2)}\|_2^2] \leq C_0 B^3 m^{-1/2}, \quad (41)$$

for some constant C_0 that is independent of B, m, k , the above expectations are taken with respect to the initialization of the NN.

Define the concatenated gradient vectors as $\widehat{\Phi}_k(\mathbf{z}_k) = (\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k)^\top - \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k)^\top)^\top$. Using the fact that $\widehat{\boldsymbol{\theta}} = \mathcal{P}_{S_B} \{\widehat{\boldsymbol{\theta}} - \beta \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}})\}$ for any $\beta \geq 0$ and the non-expansive property of projection, we observe that

$$\begin{aligned} \|\boldsymbol{\theta}_{k+1} - \widehat{\boldsymbol{\theta}}\|_2^2 &\leq \|\boldsymbol{\theta}_k - \beta_k (\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}})) - \widehat{\boldsymbol{\theta}}\|_2^2 \\ &= \|\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}\|_2^2 + \beta_k^2 \|\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 \\ &\quad - 2\beta_k \langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}}) \rangle. \end{aligned} \quad (42)$$

Similarly, using $\widehat{\mathbf{w}} = \mathcal{P}_{S_B} \{\widehat{\mathbf{w}} - \alpha \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}})\}$ for any $\alpha \geq 0$, we get

$$\begin{aligned} \|\mathbf{w}_{k+1} - \widehat{\mathbf{w}}\|_2^2 &\leq \|\mathbf{w}_k + \alpha_k (\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}})) - \widehat{\mathbf{w}}\|_2^2 \\ &= \|\mathbf{w}_k - \widehat{\mathbf{w}}\|_2^2 + \beta_k^2 \|\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 \\ &\quad + 2\beta_k \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}}) \rangle. \end{aligned} \quad (43)$$

Adding up the two inequalities lead to: for any $c_1 > 0$,

$$\begin{aligned} \|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\beta_k \langle \widehat{\Phi}_k(\mathbf{z}_k) - \widehat{\Phi}_k(\widehat{\mathbf{z}}), \mathbf{z}_k - \mathbf{z}^* \rangle + \frac{\beta_k}{c_1} \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 \\ &\quad + 2\beta_k^2 \|\nabla \widehat{J}_v(\mathbf{z}_k) - \nabla \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 + (2\beta_k^2 + c_1 \beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \}, \end{aligned} \quad (44)$$

where we have denoted $\nabla \widehat{J}_v(\mathbf{z}) = (\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z})^\top - \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z})^\top)^\top$. Recall that $\widehat{\mathbf{z}} = (\widehat{\boldsymbol{\theta}}, \widehat{\mathbf{w}})$ is the saddle point solution to (38), we observe the following lemma (which appears in the main paper as Lemma 4.3) to be proven in Appendix A.2.

Lemma A.2. *There exist constants μ, L_Φ such that it holds*

$$\begin{aligned} \frac{1}{\mu} \langle \widehat{\Phi}(\mathbf{z}_k) - \widehat{\Phi}(\widehat{\mathbf{z}}), \mathbf{z}_k - \widehat{\mathbf{z}} \rangle &\geq \mathbb{E}_{s \sim \mu^\pi} [|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})|^2 + |\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}})|^2] =: d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}), \\ \|\nabla \widehat{J}_v(\mathbf{z}_k) - \nabla \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 &\leq L_\Phi d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}), \end{aligned} \quad (45)$$

for any $k \geq 0$, where $\mu = \min\{1, v\} = 1 \wedge v$ and $L_\Phi = 4((1 \vee v)^2 + (1 + \gamma)^2)$,

The above lemma yields:

$$\begin{aligned} \|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\beta_k \mu d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + \frac{\beta_k}{c_1} \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 \\ &\quad + 2\beta_k^2 L_{\Phi} d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + (2\beta_k^2 + c_1 \beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \}. \end{aligned} \quad (46)$$

Notice that $\|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 \leq 4B^2$, we obtain

$$2\beta_k (\mu - \beta_k L_{\Phi}) d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) \leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - \|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 + 4B^2 \frac{\beta_k}{c_1} + (2\beta_k^2 + c_1 \beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \}. \quad (47)$$

Note that we have chosen β_k such that $\mu - \beta_k L_{\Phi} \geq \mu/2$. Setting $c_1 = m^{\frac{1}{4}}$ and using Lemma A.1, we obtain

$$\mu \beta_k d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) \leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - \|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 + 4\beta_k^2 \left(\frac{C_0 B^3}{\sqrt{m}} \right) + 2 \frac{\beta_k}{m^{\frac{1}{4}}} (2B^2 + C_0 B^3). \quad (48)$$

Let $n \geq 1$, summing up both sides of the inequality from $k = 0$ to $k = n$ leads to

$$\min_{k \in \{0, \dots, n\}} d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) \leq \frac{1}{(1 \wedge v) \sum_{k=0}^n \beta_k} \left\{ \|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2 + (4C_0 B^3 / \sqrt{m}) \sum_{k=0}^n \beta_k^2 \right\} + \frac{4B^2 + 2C_0 B^3}{(1 \wedge v) m^{\frac{1}{4}}}. \quad (49)$$

Simplifying the constants conclude the proof.

A.1 Proof of Lemma A.1

In this subsection, we should use the shorthand notation $\mathbb{E}_{\mu}[\cdot]$ to denote that the expectation is taken w.r.t. $s \sim \mu^{\Pi}$. Let us begin by borrowing the following lemmas from [Cai et al., 2019]

Lemma A.3. [Cai et al., 2019, Lemma F.1] Under H1, H2, there exists constant c_1 such that it holds for any $\boldsymbol{\theta} \in S_B$ that

$$\mathbb{E}_{\text{init}, \mu} \left[\frac{1}{m} \sum_{r=1}^m \mathbb{1} \{ |\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2 \} \right] \leq c_1 B m^{-1/2}. \quad (50)$$

The proof of the following lemma can be adapted from [Cai et al., 2019, Lemma F.2], see Appendix A.3 for the proof.

Lemma A.4. Under H1, H2, there exists constant c_2 such that it holds for any $\boldsymbol{\theta} \in S_B$ that

$$\mathbb{E}_{\text{init}, \mu} \left[|\widehat{f}(s, \boldsymbol{\theta}_0)|^2 \frac{1}{m} \sum_{r=1}^m \mathbb{1} \{ |\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2 \} \right] \leq c_2 B m^{-1/2} + m^{-1}. \quad (51)$$

Now, we begin proving the Lemma A.1. First we define a constant \widetilde{C}_0 stated in the lemma as:

$$\widetilde{C}_0 := \max \left\{ \{ 8(1+\gamma^2) + 4v^2 \} \{ c_0 B^3 + 2c_1 B^3 + 2c_2 B + 2 \}, 6c_0 B^3 + 8(c_1 \bar{r}^2 B + 3c_1 B^3 + 3c_2 B + 3) \right\}. \quad (52)$$

Based on the above lemmas, let us observe

$$\begin{aligned} \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 &\leq 2 \mathbb{E}_{\mu} \left[\underbrace{\|f(s, \mathbf{w}_k) \nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k) \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2^2}_{=: E_k^{(1,1)}} \right] \\ &\quad + 2v^2 \mathbb{E}_{\mu} \left[\underbrace{\|f(s, \boldsymbol{\theta}_k) \nabla f(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \boldsymbol{\theta}_k) \nabla \widehat{f}(s, \boldsymbol{\theta}_k)\|_2^2}_{=: E_k^{(1,2)}} \right]. \end{aligned} \quad (53)$$

We have

$$E_k^{(1,1)} \leq 2|f(s, \mathbf{w}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k)\|_2^2 + 2|\widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2^2. \quad (54)$$

Since

$$\nabla \bar{\delta}(s, \boldsymbol{\theta}_k) = \nabla f(s, \boldsymbol{\theta}_k) - \gamma \mathbb{E}_\pi[\nabla f(s', \boldsymbol{\theta}_k) | s], \quad \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k) = \nabla \widehat{f}(s, \boldsymbol{\theta}_k) - \gamma \mathbb{E}_\pi[\nabla \widehat{f}(s', \boldsymbol{\theta}_k) | s], \quad (55)$$

where $\mathbb{E}_\pi[\cdot | s]$ denotes the conditional expectation $\mathbb{E}_{\alpha \sim \pi(\cdot | s), s' \sim \mathcal{P}^\alpha(s, \cdot)}[\cdot]$ using the target policy. With $\|\mathbf{x}_s\|_2 = 1$, it can be shown that $\|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k)\|_2 \vee \|\nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2 \leq (1 + \gamma)$. Moreover,

$$\begin{aligned} \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2^2 &\leq 2\|\nabla f(s, \boldsymbol{\theta}_k) - \nabla \widehat{f}(s, \boldsymbol{\theta}_k)\|_2^2 \\ &\quad + 2\gamma^2 \mathbb{E}_\pi[\|\nabla f(s', \boldsymbol{\theta}_k) - \nabla \widehat{f}(s', \boldsymbol{\theta}_k)\|_2^2 | s]. \end{aligned} \quad (56)$$

For any $s \in \mathcal{S}$, we have

$$\nabla f(s, \boldsymbol{\theta}_k) - \nabla \widehat{f}(s, \boldsymbol{\theta}_k) = \begin{pmatrix} (\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(1)}, \mathbf{x}_s \rangle > 0\} - \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(1)}, \mathbf{x}_s \rangle > 0\}) m^{-1/2} b_1 \mathbf{x}_s \\ \vdots \\ (\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(m)}, \mathbf{x}_s \rangle > 0\} - \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(m)}, \mathbf{x}_s \rangle > 0\}) m^{-1/2} b_m \mathbf{x}_s. \end{pmatrix}. \quad (57)$$

As such,

$$\begin{aligned} \|\nabla f(s, \boldsymbol{\theta}_k) - \nabla \widehat{f}(s, \boldsymbol{\theta}_k)\|_2^2 &= \frac{1}{m} \sum_{r=1}^m |b_r|^2 |\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(r)}, \mathbf{x}_s \rangle > 0\} - \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\}|^2 \|\mathbf{x}_s\|_2^2 \\ &\stackrel{(a)}{\leq} \frac{1}{m} \sum_{r=1}^m |\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(r)}, \mathbf{x}_s \rangle > 0\} - \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\}|^2 \\ &\stackrel{(b)}{\leq} \frac{1}{m} \sum_{r=1}^m \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}_k^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\}, \end{aligned} \quad (58)$$

where (a) uses $|b_r| \leq 1$, $\|\mathbf{x}_s\|_2^2 = 1$, and (b) follows from the fact

$$|\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(r)}, \mathbf{x}_s \rangle > 0\} - \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\}| \leq \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}_k^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\},$$

since $\mathbb{1}\{\langle \boldsymbol{\theta}_k^{(r)}, \mathbf{x}_s \rangle > 0\} \neq \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\}$ implies $|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}_k^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2$. Lastly, we observe

$$|\widehat{f}(s, \mathbf{w}_k)|^2 \leq 2|\widehat{f}(s, \boldsymbol{\theta}_0)|^2 + 2|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \boldsymbol{\theta}_0)|^2 \leq 2|\widehat{f}(s, \boldsymbol{\theta}_0)|^2 + 2B^2. \quad (59)$$

Using Lemma 4.1, A.3, A.4, we obtain

$$\mathbb{E}_{\text{init}, \mu}[E_k^{(1,1)}] \leq 4(1 + \gamma^2) \{c_0 B^3 + 2c_1 B^3 + 2c_2 B\} m^{-1/2} + 8(1 + \gamma^2) m^{-1} \quad (60)$$

Similarly, we observe

$$E_k^{(1,2)} \leq 2|f(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \boldsymbol{\theta}_k)|^2 \|\nabla f(s, \boldsymbol{\theta}_k)\|_2^2 + 2|\widehat{f}(s, \boldsymbol{\theta}_k)|^2 \|\nabla f(s, \boldsymbol{\theta}_k) - \nabla \widehat{f}(s, \boldsymbol{\theta}_k)\|_2^2. \quad (61)$$

Similar to (59), we have $|\widehat{f}(s, \boldsymbol{\theta}_k)|^2 \leq 2|\widehat{f}(s, \boldsymbol{\theta}_0)|^2 + 2B^2$. Applying similar steps as before yields

$$\mathbb{E}_{\text{init}, \mu}[E_k^{(1,2)}] \leq (2c_0 B^3 + 2(2c_1 B^3 + 2c_2 B)) m^{-1/2} + 4m^{-1}. \quad (62)$$

Finally, we get

$$\mathbb{E}_{\text{init}}[\|\widehat{\mathbf{e}}_k^{(1)}\|_2^2] \leq 2 \mathbb{E}_{\text{init}, \mu}[E_k^{(1,1)} + v^2 E_k^{(1,2)}] \leq \widetilde{C}_0 m^{-1/2}. \quad (63)$$

The next step is to bound $\mathbb{E}_{\text{init}}[\|\widehat{\mathbf{e}}_k^{(2)}\|_2^2]$, we observe that

$$\begin{aligned} \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 &\leq \mathbb{E}_\mu \left[\left\| (\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} f(s, \mathbf{w}_k) - ((\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)) \nabla_{\mathbf{w}} \widehat{f}(s, \mathbf{w}_k)) \right\|_2^2 \right] \\ &\leq 2\mathbb{E}_\mu \left[\left\| \bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)) \right\|^2 \|\nabla f(s, \mathbf{w}_k)\|_2^2 \right] \\ &\quad + 2\mathbb{E}_\mu \left[\left\| \widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k) \right\|^2 \|\nabla \widehat{f}(s, \mathbf{w}_k) - \nabla f(s, \mathbf{w}_k)\|_2^2 \right]. \end{aligned} \quad (64)$$

We observe that

$$\begin{aligned} & \mathbb{E}_{\text{init}, \mu} \left[|\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k))|^2 \|\nabla f(s, \mathbf{w}_k)\|_2^2 \right] \\ & \leq \mathbb{E}_{\text{init}, \mu} \left[|\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k))|^2 \right] \leq 3c_0 B^3 m^{-1/2}, \end{aligned} \quad (65)$$

where the last inequality is due to Lemma 4.1 and the fact $\|\nabla f(s, \mathbf{w}_k)\|_2^2 \leq 1$. Moreover, since $|\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \leq 4(\bar{r}^2 + 3B^2 + 3|\widehat{f}(s, \boldsymbol{\theta}_0)|^2)$, we have

$$\begin{aligned} & \mathbb{E}_{\text{init}, \mu} \left[|\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \widehat{f}(s, \mathbf{w}_k) - \nabla f(s, \mathbf{w}_k)\|_2^2 \right] \\ & \stackrel{(a)}{\leq} 4\mathbb{E}_{\text{init}, \mu} \left[(\bar{r}^2 + 3B^2 + 3|\widehat{f}(s, \boldsymbol{\theta}_0)|^2) \frac{1}{m} \sum_{r=1}^m \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\mathbf{w}_k^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\} \right] \\ & \stackrel{(b)}{\leq} (4c_1 \bar{r}^2 B + 12c_1 B^3 + 12c_2 B) m^{-1/2} + 12m^{-1}, \end{aligned} \quad (66)$$

where (a) uses (58) and (b) is due to Lemma A.3. Combining the above inequalities also yields

$$\mathbb{E}_{\text{init}}[\|\widehat{\mathbf{e}}_k^{(2)}\|_2^2] \leq \widetilde{C}_0 m^{-1/2}. \quad (67)$$

The proof is concluded.

A.2 Proof of Lemma A.2

We should use the shorthand notation $\mathbb{E}_\mu[\cdot]$ to denote that the expectation is taken w.r.t. $s \sim \mu^\Pi$, as well as $a \sim \Pi(\cdot|s)$, $s' \sim P^a(s, \cdot)$. Define the following mean field matrix/vectors:

$$\mathbf{A}_0 = \mathbb{E}_\mu[\ell(\mathbf{x}_s)\ell(\mathbf{x}_s)^\top], \quad \mathbf{A}_1 = \mathbb{E}_\mu[\rho(a|s)\ell(\mathbf{x}_s)\ell(\mathbf{x}_{s'})^\top], \quad \mathbf{b} = \mathbb{E}_\mu[\rho(a|s)\mathbf{R}(s, a)\ell(\mathbf{x}_s)],$$

and we further define $\widetilde{\mathbf{A}}_1 := \mathbf{A}_0 - \gamma\mathbf{A}_1$. The objective function can then be written as:

$$\widehat{J}_v(\boldsymbol{\theta}, \mathbf{w}) = \mathbf{w}^\top \widetilde{\mathbf{A}}_1 \boldsymbol{\theta} - \mathbf{w}^\top \mathbf{b} - \frac{1}{2} \mathbf{w}^\top \mathbf{A}_0 \mathbf{w} + \frac{v}{2} \boldsymbol{\theta}^\top \mathbf{A}_0 \boldsymbol{\theta}.$$

A direct computation of the gradient of the above leads to

$$\begin{aligned} \langle \widehat{\Phi}(\mathbf{z}_k) - \widehat{\Phi}(\widehat{\mathbf{z}}), \mathbf{z}_k - \widehat{\mathbf{z}} \rangle &= v(\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})^\top \mathbf{A}_0 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}) + (\mathbf{w}_k - \widehat{\mathbf{w}})^\top \mathbf{A}_0 (\mathbf{w}_k - \widehat{\mathbf{w}}) \\ &= v \mathbb{E}_\mu[(\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})^\top \ell(s)\ell(s)^\top (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})] + \mathbb{E}_\mu[(\mathbf{w}_k - \widehat{\mathbf{w}})^\top \ell(s)\ell(s)^\top (\mathbf{w}_k - \widehat{\mathbf{w}})] \\ &= v \mathbb{E}_\mu[|\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}})|^2] + \mathbb{E}_\mu[|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})|^2] \geq \min\{1, v\} d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}), \end{aligned}$$

where the last equality is due to the definition of $\widehat{f}(\cdot, \cdot)$ in (14). This proves the first inequality of the lemma.

For the second identity, we observe that

$$\begin{pmatrix} \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}}) \\ \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}}) \end{pmatrix} = \begin{pmatrix} \widetilde{\mathbf{A}}_1^\top (\mathbf{w}_k - \widehat{\mathbf{w}}) + v \mathbf{A}_0 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}) \\ \widetilde{\mathbf{A}}_1 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}) - \mathbf{A}_0 (\mathbf{w}_k - \widehat{\mathbf{w}}) \end{pmatrix}.$$

This yields

$$\|\nabla \widehat{J}_v(\mathbf{z}_k) - \nabla \widehat{J}_v(\widehat{\mathbf{z}})\|_2 \leq \|\widetilde{\mathbf{A}}_1^\top (\mathbf{w}_k - \widehat{\mathbf{w}})\|_2 + v \|\mathbf{A}_0 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})\|_2 + \|\widetilde{\mathbf{A}}_1 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})\|_2 + \|\mathbf{A}_0 (\mathbf{w}_k - \widehat{\mathbf{w}})\|_2.$$

We have

$$\begin{aligned} \|\mathbf{A}_0 (\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}})\|_2 &= \|\mathbb{E}_\mu[\ell(\mathbf{x}_s)(\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}}))]\|_2 \leq \mathbb{E}_\mu[\|\ell(\mathbf{x}_s)(\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}}))\|_2] \\ &\leq \mathbb{E}_\mu[|\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}})| \|\ell(\mathbf{x}_s)\|_2] \leq \sqrt{\mathbb{E}_\mu[|\widehat{f}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \widehat{\boldsymbol{\theta}})|^2]} \sqrt{\mathbb{E}_\mu[\|\ell(\mathbf{x}_s)\|_2^2]}, \end{aligned}$$

where we have applied the Cauchy-Schwarz inequality in the last step. Also,

$$\begin{aligned} \|\widetilde{\mathbf{A}}_1^\top (\mathbf{w}_k - \widehat{\mathbf{w}})\|_2 &= \|(\mathbf{A}_0 - \gamma\mathbf{A}_1^\top)(\mathbf{w}_k - \widehat{\mathbf{w}})\|_2 \\ &= \|\mathbb{E}_\mu[(\ell(\mathbf{x}_s) - \gamma\rho(a|s)\ell(\mathbf{x}_{s'}))(\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}}))]\|_2 \\ &\stackrel{(a)}{\leq} \mathbb{E}_{s \sim \mu^\Pi, a \sim \pi(\cdot|s), s' \sim P^a(\cdot|s)}[|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})| \|\ell(\mathbf{x}_s) - \gamma\ell(\mathbf{x}_{s'})\|_2] \\ &\leq \mathbb{E}_\mu[|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})| \|\ell(\mathbf{x}_s)\|_2] + \gamma \mathbb{E}_{s \sim \mu^\Pi, a \sim \pi(\cdot|s), s' \sim P^a(\cdot|s)}[|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})| \|\ell(\mathbf{x}_{s'})\|_2] \\ &\stackrel{(b)}{\leq} (1 + \gamma) \sqrt{\mathbb{E}_\mu[|\widehat{f}(s, \mathbf{w}_k) - \widehat{f}(s, \widehat{\mathbf{w}})|^2]}, \end{aligned}$$

where the expectation in (a) is switched to the target policy through evaluating the expectation with the importance ratio, and (b) uses $\|\ell(\mathbf{x}_s)\|_2^2 \leq 1$.

Similarly, we get

$$\begin{aligned} \|\tilde{\mathbf{A}}_1(\boldsymbol{\theta}_k - \hat{\boldsymbol{\theta}})\|_2 &= \|(\mathbf{A}_0 - \gamma \mathbf{A}_1)(\boldsymbol{\theta}_k - \hat{\boldsymbol{\theta}})\|_2 \leq \|\mathbf{A}_0(\boldsymbol{\theta}_k - \hat{\boldsymbol{\theta}})\|_2 + \gamma \|\mathbf{A}_1(\boldsymbol{\theta}_k - \hat{\boldsymbol{\theta}})\|_2 \\ &= \|\mathbb{E}_\mu[\ell(\mathbf{x}_s)(\hat{f}(s, \boldsymbol{\theta}_k) - \hat{f}(s, \hat{\boldsymbol{\theta}}))]\|_2 + \gamma \|\mathbb{E}_{s \sim \mu^\pi, a \sim \pi(\cdot|s), s' \sim P^a(\cdot|s)}[\ell(\mathbf{x}_s)(\hat{f}(s', \boldsymbol{\theta}_k) - \hat{f}(s', \hat{\boldsymbol{\theta}}))]\|_2 \\ &\leq \mathbb{E}_\mu[|\hat{f}(s, \boldsymbol{\theta}_k) - \hat{f}(s, \hat{\boldsymbol{\theta}})|\|\ell(\mathbf{x}_s)\|_2] + \gamma \mathbb{E}_{s \sim \mu^\pi, a \sim \pi(\cdot|s), s' \sim P^a(\cdot|s)}[|\hat{f}(s', \boldsymbol{\theta}_k) - \hat{f}(s', \hat{\boldsymbol{\theta}})|\|\ell(\mathbf{x}_s)\|_2] \\ &\leq (1 + \gamma) \sqrt{[\hat{f}(s, \boldsymbol{\theta}_k) - \hat{f}(s, \hat{\boldsymbol{\theta}})]^2}. \end{aligned}$$

Therefore,

$$\begin{aligned} \|\nabla \hat{J}_v(\mathbf{z}_k) - \nabla \hat{J}_v(\hat{\mathbf{z}})\|_2^2 &\leq 4 \left((1 + (1 + \gamma)^2) \mathbb{E}_\mu[|\hat{f}(s, \mathbf{w}_k) - \hat{f}(s, \hat{\mathbf{w}})|^2] + (v^2 + (1 + \gamma)^2) \mathbb{E}_\mu[|\hat{f}(s, \boldsymbol{\theta}_k) - \hat{f}(s, \hat{\boldsymbol{\theta}})|^2] \right) \\ &\leq 4 \{ (1 \vee v)^2 + (1 + \gamma)^2 \} d_{\hat{f}}(\mathbf{z}_k, \hat{\mathbf{z}}). \end{aligned}$$

This concludes the proof.

A.3 Proof of Lemma A.4

Observe the following expansion

$$\begin{aligned} |\hat{f}(s, \boldsymbol{\theta}_0)|^2 &= \frac{1}{m} \left(\sum_{r=1}^m b_r \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle \right)^2 \\ &\leq \frac{1}{m} \left(\sum_{r=1}^m \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} |\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle|^2 + \sum_{r=1}^m \sum_{r' \neq r} b_r b_{r'} \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle \right) \\ &\leq \frac{1}{m} \left(\sum_{r=1}^m \|\boldsymbol{\theta}_0^{(r)}\|_2^2 + \sum_{r=1}^m \sum_{r' \neq r} b_r b_{r'} \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle \right). \end{aligned}$$

Together with the above expansion and using the fact that $\mathbb{E}_{\text{init}}[b_r b_{r'}] = 0$ for any $r \neq r'$, the desired expectation with respect to the NN initialization can be computed as

$$\begin{aligned} \mathbb{E}_{\text{init}, \mu} \left[|\hat{f}(s, \boldsymbol{\theta}_0)|^2 \frac{1}{m} \sum_{r=1}^m \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\} \right] \\ \leq \frac{1}{m^2} \mathbb{E}_{\text{init}, \mu} \left[\left(\sum_{r=1}^m \|\boldsymbol{\theta}_0^{(r)}\|_2^2 \right) \left(\sum_{r=1}^m \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\} \right) \right] \\ \leq \frac{1}{m^2} \mathbb{E}_{\text{init}, \mu} \left[\sum_{r=1}^m \|\boldsymbol{\theta}_0^{(r)}\|_2^2 \sum_{r' \neq r} \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r')} - \boldsymbol{\theta}_0^{(r')}\|_2\} + \sum_{r=1}^m \|\boldsymbol{\theta}_0^{(r)}\|_2^2 \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r)} - \boldsymbol{\theta}_0^{(r)}\|_2\} \right]. \end{aligned}$$

Let us study the first term inside the expectation above:

$$\begin{aligned} &\frac{1}{m^2} \mathbb{E}_{\text{init}, \mu} \left[\sum_{r=1}^m \|\boldsymbol{\theta}_0^{(r)}\|_2^2 \sum_{r' \neq r} \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r')} - \boldsymbol{\theta}_0^{(r')}\|_2\} \right] \\ &\stackrel{(a)}{=} \frac{1}{m} \sum_{r=1}^m \mathbb{E}_{\text{init}, \mu} [\|\boldsymbol{\theta}_0^{(r)}\|_2^2] \mathbb{E}_{\text{init}, \mu} \left[\frac{1}{m} \sum_{r' \neq r} \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r')} - \boldsymbol{\theta}_0^{(r')}\|_2\} \right] \\ &\leq \frac{1}{m} \sum_{r=1}^m \mathbb{E}_{\text{init}, \mu} [\|\boldsymbol{\theta}_0^{(r)}\|_2^2] \mathbb{E}_{\text{init}, \mu} \left[\frac{1}{m} \sum_{r'=1}^m \mathbb{1}\{|\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle| \leq \|\boldsymbol{\theta}^{(r')} - \boldsymbol{\theta}_0^{(r')}\|_2\} \right] \stackrel{(b)}{\leq} c_1 B m^{-1/2}, \end{aligned} \tag{68}$$

where (a) is due to the independence between $\theta_0^{(r)}$ and $\theta_0^{(r')}$ with $r' \neq r$ and (b) is due to Lemma A.3. For the second term, we observe

$$\frac{1}{m^2} \mathbb{E}_{\text{init}, \mu} \left[\sum_{r=1}^m \|\theta_0^{(r)}\|_2^2 \mathbb{1}\{\langle \theta_0^{(r)}, \mathbf{x}_s \rangle \leq \|\theta^{(r)} - \theta_0^{(r)}\|_2\} \right] \leq \frac{1}{m^2} \sum_{r=1}^m \mathbb{E}_{\text{init}}[\|\theta_0^{(r)}\|_2^2] = \frac{1}{m}. \quad (69)$$

Combining (68) and (69) leads to the desired result.

B Proof of Theorem 3.2

Note that the stochastic neural GTD algorithm is similar to the population GTD algorithm except for the additional error term $e_k^{(i)}$ in the updates [cf. (22)]. In particular, with respect to (15), the stochastic neural GTD algorithm in Algorithm 1 can be written as:

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \mathcal{P}_{S_B} \left\{ \boldsymbol{\theta}_k - \beta_k \left(\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} + \mathbf{e}_k^{(1)} \right) \right\}, \\ \mathbf{w}_{k+1} &= \mathcal{P}_{S_B} \left\{ \mathbf{w}_k + \beta_k \left(\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} + \mathbf{e}_k^{(2)} \right) \right\}, \end{aligned} \quad (70)$$

We proceed by observing that (42), (43) can be modified as

$$\begin{aligned} \|\boldsymbol{\theta}_{k+1} - \widehat{\boldsymbol{\theta}}\|_2^2 &\leq \|\boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}\|_2^2 + \beta_k^2 \|\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} + \mathbf{e}_k^{(1)} - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 \\ &\quad - 2\beta_k \langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} + \mathbf{e}_k^{(1)} - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\widehat{\mathbf{z}}) \rangle, \\ \|\mathbf{w}_{k+1} - \widehat{\mathbf{w}}\|_2^2 &\leq \|\mathbf{w}_k - \widehat{\mathbf{w}}\|_2^2 + \beta_k^2 \|\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} + \mathbf{e}_k^{(2)} - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 \\ &\quad + 2\beta_k \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} + \mathbf{e}_k^{(2)} - \nabla_{\mathbf{w}} \widehat{J}_v(\widehat{\mathbf{z}}) \rangle. \end{aligned}$$

Note that when conditioned on $\mathcal{F}_{k-1} = \sigma\{\boldsymbol{\theta}_0, \widetilde{s}_0, \dots, \widetilde{s}_{k-1}\}$, the iterate $\mathbf{z}_k = (\boldsymbol{\theta}_k, \mathbf{w}_k)$ is deterministic. As such, using H3 that the noise $e_k^{(i)}$ is zero-mean when conditioned on \mathcal{F}_{k-1} , $\mathbb{E}[\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \mathbf{e}_k^{(1)} \rangle | \mathcal{F}_{k-1}] = \mathbb{E}[\langle \mathbf{w}_k - \widehat{\mathbf{w}}, \mathbf{e}_k^{(2)} \rangle | \mathcal{F}_{k-1}] = 0$.

It follows that for any $c_1 > 0$, we have

$$\begin{aligned} \mathbb{E}[\|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 | \mathcal{F}_{k-1}] &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\beta_k \langle \widehat{\Phi}(\mathbf{z}_k) - \widehat{\Phi}(\widehat{\mathbf{z}}), \mathbf{z}_k - \widehat{\mathbf{z}} \rangle + \frac{\beta_k}{c_1} \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 \\ &\quad + 3\beta_k^2 \|\nabla \widehat{J}_v(\mathbf{z}_k) - \nabla \widehat{J}_v(\widehat{\mathbf{z}})\|_2^2 + (3\beta_k^2 + c_1\beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \} \\ &\quad + 3\beta_k^2 \mathbb{E}[\|\mathbf{e}_k^{(1)}\|_2^2 + \|\mathbf{e}_k^{(2)}\|_2^2 | \mathcal{F}_{k-1}] \\ &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\beta_k \mu d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + \frac{\beta_k}{c_1} 4B^2 + 3\beta_k^2 L_{\Phi} d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) \\ &\quad + (3\beta_k^2 + c_1\beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \} + 6\beta_k^2 \sigma^2, \end{aligned} \quad (71)$$

where the last inequality used $\|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 \leq 4B^2$ and Lemma A.2. Rearranging terms and using our conditions on the step size such that $\mu - \frac{3}{2}\beta_k L_{\Phi} \geq \frac{\mu}{2}$, these lead to

$$\begin{aligned} \beta_k \mu d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - \mathbb{E}[\|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 | \mathcal{F}_{k-1}] + 4B^2 \left(\frac{\beta_k}{c_1} \right) \\ &\quad + (3\beta_k^2 + c_1\beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \} + 6\beta_k^2 \sigma^2. \end{aligned} \quad (72)$$

We take $c_1 = m^{\frac{1}{4}}$ and invoke Lemma A.1, this lead to

$$\beta_k \mu d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}}) \leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - \mathbb{E}[\|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 | \mathcal{F}_{k-1}] + \frac{\beta_k}{m^{\frac{1}{4}}} (2C_0 B^3 + 4B^2) + 2\beta_k^2 \left\{ \frac{3C_0 B^3}{\sqrt{m}} + 3\sigma^2 \right\}. \quad (73)$$

Taking total expectations on both sides and summing up the inequality from $k = 0$ to $k = n$ leads to

$$\mu \sum_{k=0}^n \beta_k \mathbb{E}[d_{\widehat{f}}(\mathbf{z}_k, \widehat{\mathbf{z}})] \leq \mathbb{E}[\|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2] + \frac{2C_0 B^3 + 4B^2}{m^{\frac{1}{4}}} \sum_{k=0}^n \beta_k + 2 \left\{ \frac{3C_0 B^3}{\sqrt{m}} + 3\sigma^2 \right\} \sum_{k=0}^n \beta_k^2. \quad (74)$$

Recall that $I_n \in \{0, \dots, n\}$ is an independent random variable chosen from the distribution $\mathbb{P}(I_n = k) = \beta_k / \sum_{\ell=0}^n \beta_\ell$, we have

$$\mathbb{E}[d_f(\mathbf{z}_{I_n}, \hat{\mathbf{z}})] \leq \frac{2C_0 B^3 + 4B^2}{\mu m^{\frac{1}{4}}} + \frac{\mathbb{E}[\|\mathbf{z}_0 - \hat{\mathbf{z}}\|_2^2] + 2\left\{\frac{3C_0 B^3}{\sqrt{m}} + 3\sigma^2\right\} \sum_{k=0}^n \beta_k^2}{\mu \sum_{k=0}^n \beta_k}. \quad (75)$$

This concludes the proof of our theorem.

C Proof of Theorem 3.4

Note that $\mathbb{E}_\mu[\cdot] = \mathbb{E}_{s \sim \mu^\pi}[\cdot]$. We begin the proof by recalling that, for any $\boldsymbol{\theta} \in S_B$,

$$J_v(\boldsymbol{\theta}) = \frac{1}{2} \mathbb{E}_\mu [\bar{\delta}(s, \boldsymbol{\theta})^2 + v f(s, \boldsymbol{\theta})^2], \quad (76)$$

where we define that (the importance ratio $\rho(a|s)$ has been absorbed in the conditional expectation)

$$\bar{\delta}(s, \boldsymbol{\theta}) = f(s, \boldsymbol{\theta}) - \mathfrak{r}^\pi(s) - \gamma \mathbb{P}^\pi f(s, \boldsymbol{\theta}),$$

and we have defined $\mathfrak{r}^\pi(s) = \mathbb{E}_{a \sim \pi(\cdot|s)}[\mathbb{R}(s, a)]$. Moreover, we define the following using the linearized NN function:

$$\hat{\delta}(s, \boldsymbol{\theta}) = \hat{f}(s, \boldsymbol{\theta}) - \mathfrak{r}^\pi(s) - \gamma \mathbb{P}^\pi \hat{f}(s, \boldsymbol{\theta}).$$

Estimates on the differences between $\bar{\delta}, \hat{\delta}, f, \hat{f}$ We first derive a few upper bounds as follows. Observe from Lemma 4.1 that for any $\boldsymbol{\theta} \in S_B$,

$$\mathbb{E}_{\text{init}, \mu} [|\bar{\delta}(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})|^2] \leq 2(1 + \bar{\rho}^2 \gamma^2) c_0 B^3 m^{-1/2} =: R_\delta^{(0)}$$

as well as $\mathbb{E}_{\text{init}, \mu} [f(s, \boldsymbol{\theta}) - \hat{f}(s, \boldsymbol{\theta})]^2 \leq c_0 B^3 m^{-1/2} =: R_f^{(0)}$. Now, consider that

$$\begin{aligned} |f(s, \boldsymbol{\theta})|^2 &\leq \underbrace{|f(s, \boldsymbol{\theta}) - \hat{f}(s, \boldsymbol{\theta})|^2 + 2|f(s, \boldsymbol{\theta}) - \hat{f}(s, \boldsymbol{\theta})| |\hat{f}(s, \boldsymbol{\theta})| + |\hat{f}(s, \boldsymbol{\theta})|^2}_{=: \bar{R}_f^{(1)}(s, \boldsymbol{\theta})}, \\ |\bar{\delta}(s, \boldsymbol{\theta})|^2 &\leq \underbrace{|\bar{\delta}(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})|^2 + 2|\bar{\delta}(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})| |\hat{\delta}(s, \boldsymbol{\theta})| + |\hat{\delta}(s, \boldsymbol{\theta})|^2}_{=: \bar{R}_\delta^{(1)}(s, \boldsymbol{\theta})}, \\ |f(s, \boldsymbol{\theta})|^2 &\leq \underbrace{|f(s, \boldsymbol{\theta}) - f(s, \hat{\boldsymbol{\theta}})|^2 + 2|f(s, \boldsymbol{\theta}) - f(s, \hat{\boldsymbol{\theta}})| |f(s, \hat{\boldsymbol{\theta}})| + |f(s, \hat{\boldsymbol{\theta}})|^2}_{=: \bar{R}_f^{(2)}(s, \boldsymbol{\theta}, \hat{\boldsymbol{\theta}})}, \\ |\bar{\delta}(s, \boldsymbol{\theta})|^2 &\leq \underbrace{|\bar{\delta}(s, \boldsymbol{\theta}) - \bar{\delta}(s, \hat{\boldsymbol{\theta}})|^2 + 2|\bar{\delta}(s, \boldsymbol{\theta}) - \bar{\delta}(s, \hat{\boldsymbol{\theta}})| |\bar{\delta}(s, \hat{\boldsymbol{\theta}})| + |\bar{\delta}(s, \hat{\boldsymbol{\theta}})|^2}_{=: \bar{R}_\delta^{(2)}(s, \boldsymbol{\theta}, \hat{\boldsymbol{\theta}})}, \end{aligned}$$

Note that $\bar{R}_\delta^{(1)}, \bar{R}_f^{(1)}, \bar{R}_\delta^{(2)}, \bar{R}_f^{(2)}$ are measurable functions on \mathbb{S} and depend on $\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}$. We upper bound these terms as follows.

For $\bar{R}_f^{(1)}$, due to the Cauch-Schwarz inequality, for any $\boldsymbol{\theta} \in S_B$, we have

$$\mathbb{E}_{\text{init}, \mu} [f(s, \boldsymbol{\theta}) - \hat{f}(s, \boldsymbol{\theta}) | \hat{f}(s, \boldsymbol{\theta})] \leq \left(\mathbb{E}_{\text{init}, \mu} [|\hat{f}(s, \boldsymbol{\theta})|^2] \right)^{\frac{1}{2}} \left(\mathbb{E}_{\text{init}, \mu} [f(s, \boldsymbol{\theta}) - \hat{f}(s, \boldsymbol{\theta})]^2 \right)^{\frac{1}{2}}.$$

As we have $|\hat{f}(s, \boldsymbol{\theta})|^2 \leq 2|\hat{f}(s, \boldsymbol{\theta}_0)|^2 + 2B^2$, and

$$\begin{aligned} \mathbb{E}_{\text{init}, \mu} [|\hat{f}(s, \boldsymbol{\theta}_0)|^2] &= \mathbb{E}_{\text{init}, \mu} \left[\frac{1}{m} \sum_{r=1}^m \sum_{r'=1}^m b_r b_{r'} \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle > 0\} \langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle \langle \boldsymbol{\theta}_0^{(r')}, \mathbf{x}_s \rangle \right] \\ &\leq \frac{1}{m} \mathbb{E}_{\text{init}, \mu} \left[\sum_{r=1}^m \mathbb{1}\{\langle \boldsymbol{\theta}_0^{(r)}, \mathbf{x}_s \rangle > 0\} \|\boldsymbol{\theta}_0^{(r)}\|_2^2 \right] \leq 1, \end{aligned}$$

where we have used that b_r is independent from $b_{r'}$ if $r \neq r'$ in the first inequality, also see Appendix A.3. The above implies $\mathbb{E}_{\text{init},\mu}[|\hat{f}(s, \boldsymbol{\theta})|^2] \leq 2(1 + B^2)$, and thus showing that

$$\mathbb{E}_{\text{init},\mu}[\bar{R}_f^{(1)}] \leq c_0 B^3 m^{-1/2} + \sqrt{8c_0(1 + B^2)B^3} m^{-1/4} =: R_f^{(1)}. \quad (77)$$

For $\bar{R}_\delta^{(1)}$, for any $\boldsymbol{\theta} \in S_B$, applying the Cauch-Schwarz inequality yields

$$\mathbb{E}_{\text{init},\mu}[|\delta(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})| |\hat{\delta}(s, \boldsymbol{\theta})|] \leq \left(\mathbb{E}_{\text{init},\mu}[|\hat{\delta}(s, \boldsymbol{\theta})|^2] \right)^{\frac{1}{2}} \left(\mathbb{E}_{\text{init},\mu}[|\delta(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})|^2] \right)^{\frac{1}{2}}. \quad (78)$$

The above implies

$$\begin{aligned} \mathbb{E}_{\text{init},\mu}[|\hat{\delta}(s, \boldsymbol{\theta})|^2] &\leq 3\bar{r}^2 + 3\mathbb{E}_{\text{init},\mu}[|\hat{f}(s, \boldsymbol{\theta})|^2] + 3\gamma^2 \mathbb{E}_{\text{init},\mu^\pi}[|\hat{f}(s', \boldsymbol{\theta})|^2] \\ &\leq 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2)). \end{aligned} \quad (79)$$

Moreover, applying Lemma 4.1 to (78) shows that

$$\mathbb{E}_{\text{init},\mu}[|\delta(s, \boldsymbol{\theta}) - \hat{\delta}(s, \boldsymbol{\theta})| |\hat{\delta}(s, \boldsymbol{\theta})|] \leq \sqrt{3c_0(1 + \bar{\rho}\gamma)(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} B^3 m^{-1/4}.$$

Gathering terms yields

$$\begin{aligned} \mathbb{E}_{\text{init},\mu}[\bar{R}_\delta^{(1)}] &\leq c_0(1 + \bar{\rho}\gamma)B^3 m^{-1/2} + \sqrt{12c_0(1 + \bar{\rho}\gamma)(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} B^3 m^{-1/4} =: R_\delta^{(1)}. \end{aligned}$$

We note that

$$R_f^{(1)} = \mathcal{O}(B^3 m^{-1/2} + B^{3/2} m^{-1/4}), \quad R_\delta^{(1)} = \mathcal{O}(B^3 m^{-1/2} + B^{3/2} m^{-1/4}).$$

To upper bound $\bar{R}_\delta^{(2)}$, we observe that

$$\begin{aligned} \mathbb{E}_\mu[|\delta(s, \boldsymbol{\theta}) - \delta(s, \hat{\boldsymbol{\theta}})|^2] &= \mathbb{E}_\mu \left[|f(s, \boldsymbol{\theta}) - f(s, \hat{\boldsymbol{\theta}}) - \gamma(f(s', \boldsymbol{\theta}) - f(s', \hat{\boldsymbol{\theta}}))|^2 \right] \\ &\leq 3\mathbb{E}_\mu \left[|\hat{f}(s, \boldsymbol{\theta}) - \hat{f}(s, \hat{\boldsymbol{\theta}})|^2 \right] + 3\gamma^2 \mathbb{E}_{\mu^\pi} \left[|\hat{f}(s', \boldsymbol{\theta}) - \hat{f}(s', \hat{\boldsymbol{\theta}})|^2 \right] + 12c_0 B^3 m^{-1/2} \\ &\leq 3(1 + \bar{\rho}^2\gamma^2) d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}}) + 12c_0 B^3 m^{-1/2}, \end{aligned}$$

for $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{w})$ with any $\mathbf{w} \in S_B$, where we have applied Lemma 4.1 in the second inequality.

Moreover, using (79) and the definition of $\bar{R}_\delta^{(1)}$ we have

$$\mathbb{E}_{\text{init},\mu}[|\bar{\delta}(s, \hat{\boldsymbol{\theta}})|^2] \leq R_\delta^{(1)} + 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2)).$$

The above shows that

$$\begin{aligned} \mathbb{E}_{\text{init},\mu} \left[|\delta(s, \boldsymbol{\theta}) - \delta(s, \hat{\boldsymbol{\theta}})| |\delta(s, \hat{\boldsymbol{\theta}})| \right] &\leq \left(\mathbb{E}_{\text{init},\mu}[|\delta(s, \hat{\boldsymbol{\theta}})|^2] \right)^{\frac{1}{2}} \left(\mathbb{E}_{\text{init},\mu}[|\delta(s, \boldsymbol{\theta}) - \delta(s, \hat{\boldsymbol{\theta}})|^2] \right)^{\frac{1}{2}} \\ &\leq \sqrt{3(R_\delta^{(1)} + 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2)))} \sqrt{(1 + \bar{\rho}^2\gamma^2) \mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 4c_0 B^3 m^{-1/2}}. \end{aligned}$$

These yield

$$\begin{aligned} \mathbb{E}_{\text{init},\mu}[\bar{R}_\delta^{(2)}] &\leq 3(1 + \bar{\rho}^2\gamma^2) \mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 12c_0 B^3 m^{-1/2} \\ &\quad + \sqrt{12(R_\delta^{(1)} + 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2)))} \sqrt{(1 + \bar{\rho}^2\gamma^2) \mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 4c_0 B^3 m^{-1/2}} =: R_\delta^{(2)}. \end{aligned}$$

Finally, to upper bound $\bar{R}_f^{(2)}$, applying Lemma 4.1, we observe that

$$\mathbb{E}_\mu[|f(s, \boldsymbol{\theta}) - f(s, \hat{\boldsymbol{\theta}})|^2] \leq 2\mathbb{E}_\mu \left[|\hat{f}(s, \boldsymbol{\theta}) - \hat{f}(s, \hat{\boldsymbol{\theta}})|^2 \right] + 4c_0 B^3 m^{-1/2} \leq 2d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}}) + 4c_0 B^3 m^{-1/2}.$$

Moreover,

$$\mathbb{E}_{\text{init},\mu}[|f(s, \hat{\boldsymbol{\theta}})|^2] \leq R_f^{(1)} + 2(1 + B^2).$$

This shows

$$\mathbb{E}_{\text{init}, \mu} [|f(s, \hat{\boldsymbol{\theta}}) - f(s, \boldsymbol{\theta})|] \leq \sqrt{2R_f^{(1)} + 4(1+B^2)} \sqrt{\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 2c_0 B^3 m^{-1/2}}.$$

As such, we have

$$\begin{aligned} & \mathbb{E}_{\text{init}, \mu} [\bar{R}_f^{(2)}] \\ & \leq 2\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 4c_0 B^3 m^{-1/2} + \sqrt{8R_f^{(1)} + 16(1+B^2)} \sqrt{\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + 2c_0 B^3 m^{-1/2}} =: R_f^{(2)} \end{aligned}$$

We also note that

$$\begin{aligned} R_f^{(2)} &= \mathcal{O} \left(\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + (B + B^{3/2} m^{-1/4}) \sqrt{\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})]} + B^3 m^{-1/2} + B^{5/2} m^{-1/4} \right), \\ R_\delta^{(2)} &= \mathcal{O} \left(\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})] + (B + B^{3/2} m^{-1/4}) \sqrt{\mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})]} + B^3 m^{-1/2} + B^{5/2} m^{-1/4} \right). \end{aligned}$$

Note that $R_f^{(0)}, R_f^{(1)}, R_f^{(2)}, R_\delta^{(0)}, R_\delta^{(1)}, R_\delta^{(2)}$ are deterministic quantities.

Bounding $J_v(\boldsymbol{\theta})$ In the following derivations, the inequalities hold in the expectation taken over the NN initializations. We shall skip $\mathbb{E}_{\text{init}}[\cdot]$ to simplify notations. First observe the following bound using the above definitions

$$J_v(\boldsymbol{\theta}) \leq J_v(\hat{\boldsymbol{\theta}}) + \frac{1}{2} \left(R_\delta^{(2)} + v R_f^{(2)} \right) \quad (80)$$

With a slight abuse in notations, define the following function from the Fenchel's conjugate (8):

$$\bar{J}_v(\boldsymbol{\theta}; y(\cdot)) := \mathbb{E}_\mu [y(s) \bar{\delta}(s, \boldsymbol{\theta}) - \frac{1}{2} y(s)^2 + \frac{v}{2} f(s, \boldsymbol{\theta})^2] \quad (81)$$

From H4, there exists $\mathbf{w}(\hat{\boldsymbol{\theta}}) \in S_B$ such that $\mathbb{E}_\mu [|\bar{\delta}(s, \hat{\boldsymbol{\theta}}) - f(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))|^2] \leq c_{\text{nn}}$. As such, we observe that

$$\begin{aligned} J_v(\hat{\boldsymbol{\theta}}) &= \bar{J}_v(\hat{\boldsymbol{\theta}}; \bar{\delta}(\cdot, \hat{\boldsymbol{\theta}})) \\ &= J_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) + \mathbb{E}_\mu [\bar{\delta}(s, \hat{\boldsymbol{\theta}}) (\bar{\delta}(s, \hat{\boldsymbol{\theta}}) - f(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))) - \frac{1}{2} \{\bar{\delta}(s, \hat{\boldsymbol{\theta}})^2 - f(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))^2\}] \\ &= J_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) + \frac{1}{2} \mathbb{E}_\mu [|\bar{\delta}(s, \hat{\boldsymbol{\theta}}) - f(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))|^2] \leq J_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) + \frac{c_{\text{nn}}}{2}. \end{aligned} \quad (82)$$

Using the bounds developed earlier in this subsection, we have

$$\begin{aligned} J_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) &\leq \hat{J}_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) + \frac{1}{2} (1+v) R_f^{(1)} \\ &\quad + \frac{1}{2} \mathbb{E}_\mu [|\bar{\delta}(s, \hat{\boldsymbol{\theta}}) - \hat{f}(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))| + |\hat{f}(s, \mathbf{w}(\hat{\boldsymbol{\theta}})) - \bar{\delta}(s, \hat{\boldsymbol{\theta}})|]. \end{aligned}$$

Since $(\hat{\boldsymbol{\theta}}, \hat{\mathbf{w}})$ is a saddle point to (15), we have

$$\hat{J}_v(\hat{\boldsymbol{\theta}}, \mathbf{w}(\hat{\boldsymbol{\theta}})) \leq \hat{J}_v(\hat{\boldsymbol{\theta}}, \hat{\mathbf{w}}) \leq \hat{J}_v(\boldsymbol{\theta}^*, \hat{\mathbf{w}}) \quad (83)$$

Notice that

$$\begin{aligned} \hat{J}_v(\boldsymbol{\theta}^*, \hat{\mathbf{w}}) &\leq J_v(\boldsymbol{\theta}^*, \hat{\mathbf{w}}) + \frac{1}{2} (1+v) R_f^{(1)} \\ &\quad + \frac{1}{2} \mathbb{E}_\mu [|\hat{\delta}(s, \boldsymbol{\theta}^*) - f(s, \hat{\mathbf{w}})| + |f(s, \hat{\mathbf{w}}) - \hat{\delta}(s, \boldsymbol{\theta}^*)|]. \end{aligned}$$

Finally, we observe,

$$J_v(\boldsymbol{\theta}^*, \hat{\mathbf{w}}) = \bar{J}_v(\boldsymbol{\theta}^*; f(\cdot, \hat{\mathbf{w}})) \leq \bar{J}_v(\boldsymbol{\theta}^*; \bar{\delta}(\cdot, \boldsymbol{\theta}^*)) = J_v(\boldsymbol{\theta}^*),$$

where the inequality is due to the fact that $\bar{\delta}(\cdot, \boldsymbol{\theta}^*)$ maximizes $\bar{J}_v(\boldsymbol{\theta}^*; y(\cdot))$. Collecting terms, we observe that

$$\begin{aligned}
& J_v(\boldsymbol{\theta}) - J_v(\boldsymbol{\theta}^*) - \frac{1}{2}(c_{\text{nn}} + R_\delta^{(2)} + vR_f^{(2)}) \\
& \leq (1+v)R_f^{(1)} \\
& + \frac{1}{2}\mathbb{E}_\mu \left[|\bar{\delta}(s, \hat{\boldsymbol{\theta}})| |f(s, \mathbf{w}(\hat{\boldsymbol{\theta}})) - \hat{f}(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))| + |\hat{f}(s, \mathbf{w}(\hat{\boldsymbol{\theta}}))| |\bar{\delta}(s, \hat{\boldsymbol{\theta}}) - \hat{\delta}(s, \hat{\boldsymbol{\theta}})| \right] \\
& + \frac{1}{2}\mathbb{E}_\mu \left[|\hat{\delta}(s, \boldsymbol{\theta}^*)| |f(s, \hat{\mathbf{w}}) - \hat{f}(s, \hat{\mathbf{w}})| + |\hat{f}(s, \hat{\mathbf{w}})| |\bar{\delta}(s, \boldsymbol{\theta}^*) - \hat{\delta}(s, \boldsymbol{\theta}^*)| \right] \\
& \leq \frac{1}{2} \left(\sqrt{R_\delta^{(1)} + 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} \sqrt{R_f^{(0)}} + \sqrt{2(1 + B^2)} \sqrt{R_\delta^{(0)}} \right) \\
& + \frac{1}{2} \left(\sqrt{3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} \sqrt{R_f^{(0)}} + \sqrt{R_f^{(1)} + 2(1 + B^2)} \sqrt{R_\delta^{(0)}} \right)
\end{aligned} \tag{84}$$

Finally, we obtain the following bound

$$\begin{aligned}
& J_v(\boldsymbol{\theta}) - J_v(\boldsymbol{\theta}^*) \\
& \leq \frac{1}{2}(c_{\text{nn}} + R_\delta^{(2)} + vR_f^{(2)}) + (1+v)R_f^{(1)} \\
& + \frac{1}{2} \left(\sqrt{R_\delta^{(1)} + 3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} \sqrt{R_f^{(0)}} + \sqrt{2(1 + B^2)} \sqrt{R_\delta^{(0)}} \right) \\
& + \frac{1}{2} \left(\sqrt{3(\bar{r}^2 + 2(1 + \bar{\rho}^2\gamma^2)(1 + B^2))} \sqrt{R_f^{(0)}} + \sqrt{R_f^{(1)} + 2(1 + B^2)} \sqrt{R_\delta^{(0)}} \right)
\end{aligned} \tag{85}$$

The proof is completed by extracting the right orders in terms of $m, B, \mathbb{E}_{\text{init}}[d_{\hat{f}}(\mathbf{z}, \hat{\mathbf{z}})]$ from the previously derived bounds of $R_f^{(0)}, R_f^{(1)}, R_f^{(2)}, R_\delta^{(0)}, R_\delta^{(1)}, R_\delta^{(2)}$.

D Convergence of Neural GTD Algorithm with Markov Samples

We consider a more sophisticated setting when samples used for the neural GTD algorithm are collected along a single sample path of the Markov chain (s_0, s_1, \dots) . At iteration k of the neural GTD algorithm, the state pair is taken as $\tilde{s}_k := (s_k, s_{k+1})$, see Algorithm 1. In this case, the neural GTD update can be written into the same form as (21), (22). However, for finite k , as $(\tilde{s}_k)_{k \geq 0}$ is not drawn from the stationary distribution, denoted as $\tilde{\mu}^\Pi$, H3 does not hold and the analysis leading to Theorem 3.2 cannot be applied. Instead, $(\tilde{s}_k)_{k \geq 0}$ forms a uniformly geometric ergodic Markov chain whose kernel is denoted by \tilde{P}^Π . Let us define $\rho \in [0, 1)$ be its mixing constant. Since the Markov chain is uniformly geometric ergodic, there exists a constant K_P such that

$$\sup_{\tilde{s} \in \mathcal{S} \times \mathcal{S}} \|(\tilde{P}^\Pi)^n(\tilde{s}, \cdot) - \tilde{\mu}^\Pi(\cdot)\| \leq \rho^n K_P, \tag{86}$$

where $\|\cdot\|$ denotes the total variation norm.

We use the standard Bachmann-Landau notations for asymptotic quantities in the following. In particular, consider two non-negative functions $h(x), g(x)$. We say that $h(x) = \mathcal{O}(g(x))$ if there exists $c_0 > 0, x_0 \geq 0$ such that $h(x) \leq c_0 g(x)$ for all $x \geq x_0$; likewise, $h(x) = \Omega(g(x))$ if there exists $c_1 > 0, x_1 \geq 0$ such that $h(x) \geq c_1 g(x)$ for all $x \geq x_0$.

As explained in the main text of the paper, in the Markov sample settings, we prove the convergence of neural GTD with high probability with respect to the random initialization specified in H1. To this end, we need to derive the high probability bounds for certain quantities. Observe the lemmas:

Lemma D.1. *Assume that H1 holds, $m = \Omega(d^{3/2})$, and $B = \mathcal{O}(m^{1/2}(\log m)^{-3})$. For any $(\boldsymbol{\theta}, \mathbf{w}) \in S_B \times S_B$, and $\tilde{s} \in \mathcal{S} \times \mathcal{S}$, with probability at least $1 - e^{-\Omega(\log^2 m)}$ with respect to the random initialization, it holds*

$$\|\nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}, \mathbf{w}; \tilde{s})\|_2 \vee \|\nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}, \mathbf{w}; \tilde{s})\|_2 \leq \mathcal{O}(B \vee \log m) =: \varsigma. \tag{87}$$

Lemma D.2. *Assume that H1 holds, $m = \Omega(d^{3/2})$, and $B = \mathcal{O}(m^{1/2}(\log m)^{-3})$. For any $\mathbf{z}, \mathbf{z}' \in S_B \times S_B$ and $\tilde{s} \in \mathcal{S} \times \mathcal{S}$, there exist constants L_J, C_J such that with probability at least $1 - e^{-\log^2(m)}$ with respect to the random initialization,*

$$\|\nabla J_v(\mathbf{z}; \tilde{s}) - \nabla J_v(\mathbf{z}'; \tilde{s})\|_2 \leq L_J \|\mathbf{z} - \mathbf{z}'\|_2 + C_J B^{4/3} (\log m)^{3/2} m^{-1/6}. \tag{88}$$

Lemma D.3. Assume that *H1* holds, $m = \Omega(d^{3/2})$, and $B = \mathcal{O}(m^{1/2}(\log m)^{-3})$. For any $\boldsymbol{\theta} \in S_B$ and $s \in \mathcal{S}$, with probability at least $1 - e^{-\Omega(B^{2/3}m^{2/3})}$ with respect to the random initialization, it holds

$$|f(s, \boldsymbol{\theta}) - \widehat{f}(s, \boldsymbol{\theta})| = \mathcal{O}(B^{4/3}m^{-1/6}\sqrt{\log m}). \quad (89)$$

We present the main conclusion on the convergence of neural GTD in detail [cf. Theorem 3.3] below:

Theorem D.1. Assume that *H1* hold, $m = \Omega(d^{3/2})$, and $B = \mathcal{O}(m^{1/2}(\log m)^{-3})$. Furthermore, for any $k \geq 0$, there exists a constant ξ such that

$$|\beta_k - \beta_{k+1}| \leq \xi \beta_k^2. \quad (90)$$

For the iterates generated by Algorithm 1 with the Markov chain satisfying (86), and any $n \geq 1$, it holds with probability at least $1 - e^{-\Omega(\log^2 m)}$ over the random initialization that:

$$\mathbb{E}_{I_n} [d_{\widehat{f}}(\mathbf{z}_{I_n}, \widehat{\mathbf{z}})] = \mathcal{O} \left(\frac{B^{\frac{8}{3}} + B^{\frac{7}{3}}(\log m)^{\frac{1}{2}}/(1-\rho)}{(1 \wedge \nu)m^{\frac{1}{6}}/\log m} + \frac{\|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2 + \beta_0 \frac{B^2}{1-\rho} + (B^3 + \frac{B^2}{1-\rho}) \sum_{k=0}^n \beta_k^2}{(1 \wedge \nu) \sum_{k=0}^n \beta_k} \right), \quad (91)$$

where the function $d_{\widehat{f}}(\mathbf{z}_{I_n}, \widehat{\mathbf{z}})$ was defined in (18). Moreover, the expectation above is taken over the independent r.v. I_n , and the i.i.d. samples of states drawn during the algorithm.

Lastly, we remark that the conditions $m = \Omega(d^{3/2})$, and $B = \mathcal{O}(m^{1/2}(\log m)^{-3})$ can be satisfied when m is sufficiently large and we choose $B = \mathcal{O}(1)$.

Proof. Now, let us recall that the neural GTD update can be written as:

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \mathcal{P}_{S_B} \left\{ \boldsymbol{\theta}_k - \beta_k \left(\nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(1)} + \mathbf{e}_k^{(1)} \right) \right\}, \\ \mathbf{w}_{k+1} &= \mathcal{P}_{S_B} \left\{ \mathbf{w}_k + \beta_k \left(\nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k) + \widehat{\mathbf{e}}_k^{(2)} + \mathbf{e}_k^{(2)} \right) \right\}, \end{aligned} \quad (92)$$

and that the errors are expressed as:

$$\begin{aligned} \widehat{\mathbf{e}}_k^{(1)} &= \nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}_k) - \nabla_{\boldsymbol{\theta}} \widehat{J}_v(\mathbf{z}_k), & \mathbf{e}_k^{(1)} &= \nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}_k; \widetilde{\mathbf{s}}_k) - \mathbb{E}_{\widetilde{\mu}}[\nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}_k; \widetilde{\mathbf{s}})], \\ \widehat{\mathbf{e}}_k^{(2)} &= \nabla_{\mathbf{w}} J_v(\mathbf{z}_k) - \nabla_{\mathbf{w}} \widehat{J}_v(\mathbf{z}_k), & \mathbf{e}_k^{(2)} &= \nabla_{\mathbf{w}} J_v(\mathbf{z}_k; \widetilde{\mathbf{s}}_k) - \mathbb{E}_{\widetilde{\mu}}[\nabla_{\mathbf{w}} J_v(\mathbf{z}_k; \widetilde{\mathbf{s}})], \end{aligned} \quad (93)$$

where the expectation $\mathbb{E}_{\widetilde{\mu}}[\cdot]$ is evaluated with respect to the stationary distribution $\widetilde{\mu}^{\Pi}$ of the Markov chain $(\widetilde{\mathbf{s}}_k)_{k \geq 0}$. We observe the following lemma which is analogous to Lemma A.1:

Lemma D.4. Assume that *H1*, *H2* hold. For any $\mathbf{z} \in S_B \times S_B$, with probability at least $1 - e^{-\Omega(\log^2 m)}$ with respect to the random initialization, it holds

$$\|\mathbf{e}_k^{(1)}\|_2^2 \vee \|\mathbf{e}_k^{(2)}\|_2^2 \leq C^e B^{8/3} m^{-1/3} (\log m)^2. \quad (94)$$

Using [Douc et al., 2018, Proposition 21.2.3], together with Lemma D.1 and some regulatory conditions, there exists measurable functions $\widehat{\mathbf{d}}J_v^{(i)} : S_B \times \mathcal{S} \rightarrow \mathbb{R}^{md}$, $i = 1, 2$, satisfying

$$\mathbf{e}_k^{(i)} = \widehat{\mathbf{d}}J_v^{(i)}(\mathbf{z}_k; \widetilde{\mathbf{s}}_k) - \widetilde{\mathbf{P}}^{\Pi} \widehat{\mathbf{d}}J_v^{(i)}(\mathbf{z}_k; \widetilde{\mathbf{s}}_k), \quad i = 1, 2. \quad (95)$$

This is also known as the Poisson equation. Again with Lemma D.1, it can be shown that (e.g., using [Fort et al., 2011])

$$\sup_{\mathbf{z} \in S_B \times S_B, \widetilde{\mathbf{s}} \in \mathcal{S} \times \mathcal{S}} \|\widehat{\mathbf{d}}J_v^{(i)}(\mathbf{z}; \widetilde{\mathbf{s}})\|_2 \leq \varsigma K_P / (1 - \rho) =: \widehat{\varsigma},$$

Furthermore, for any $\widetilde{\mathbf{s}} \in \mathcal{S} \times \mathcal{S}$ and $\mathbf{z}, \mathbf{z}' \in S_B \times S_B$, using Lemma D.2,

$$\|\widehat{\mathbf{d}}J_v^{(i)}(\mathbf{z}; \widetilde{\mathbf{s}}) - \widehat{\mathbf{d}}J_v^{(i)}(\mathbf{z}'; \widetilde{\mathbf{s}})\|_2 \leq L_J K_P / (1 - \rho) \|\mathbf{z} - \mathbf{z}'\|_2 + (C_J K_P / (1 - \rho)) B^{4/3} (\log m)^{3/2} m^{-1/6}, \quad (96)$$

for brevity, we denote $\widehat{L}_J := L_J K_P / (1 - \rho)$, $\widehat{C}_J := C_J K_P / (1 - \rho)$.

We proceed to proving the convergence of neural GTD by following the analysis done in previous sections. Since the neural GTD update follows the same form as in (21), (22), we observe that every steps in the proof of Theorem 3.2 follows. Also, the second inequality in H3 holds with σ^2 replaced by $4\zeta^2$. In particular, applying Lemma A.2 and we observe that the following (which is analogous to (71)) holds for any $k \geq 0$, $c_1 > 0$

$$\begin{aligned} \mathbb{E}[\|\mathbf{z}_{k+1} - \widehat{\mathbf{z}}\|_2^2 | \mathcal{F}_{k-1}] &\leq \|\mathbf{z}_k - \widehat{\mathbf{z}}\|_2^2 - 2\beta_k \mu d_{\widehat{\mathcal{F}}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + 3\beta_k^2 L_{\Phi} d_{\widehat{\mathcal{F}}}(\mathbf{z}_k, \widehat{\mathbf{z}}) + \frac{\beta_k}{c_1} 4B^2 \\ &\quad + (3\beta_k^2 + c_1\beta_k) \{ \|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 + \|\widehat{\mathbf{e}}_k^{(2)}\|_2^2 \} + 24\beta_k^2 \zeta^2 \\ &\quad - 2\beta_k \mathbb{E}[\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \mathbf{e}_k^{(1)} \rangle - \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \mathbf{e}_k^{(2)} \rangle | \mathcal{F}_{k-1}], \end{aligned}$$

where the last term is new. Using the step size condition, Lemma D.4 and setting $c_1 = m^{1/6}/\log m$, we can derive the following inequality:

$$\begin{aligned} \mathbb{E}[d_{\widehat{\mathcal{F}}}(\mathbf{z}_{I_n}, \widehat{\mathbf{z}})] &\leq \frac{2C^e B^{\frac{8}{3}} + 4B^2}{\mu m^{\frac{1}{6}} / \log m} + \frac{\mathbb{E}[\|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2] + 2\{3C^e B^{\frac{8}{3}} m^{-1/3} (\log m)^2 + 12\zeta^2\} \sum_{k=0}^n \beta_k^2}{\mu \sum_{k=0}^n \beta_k} \\ &\quad + \frac{2|\sum_{k=0}^n \beta_k \mathbb{E}[\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \mathbf{e}_k^{(1)} \rangle - \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \mathbf{e}_k^{(2)} \rangle]}{\mu \sum_{k=0}^n \beta_k}. \end{aligned} \tag{97}$$

Our remaining task is to bound the weighted sum of inner products in the last term. Observe

$$\beta_k \langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \mathbf{e}_k^{(1)} \rangle = \beta_k \left\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k) - \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) + \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) - \widetilde{\mathbf{P}}^{\Pi} \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k) \right\rangle.$$

By (95), we observe that $\widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) - \widetilde{\mathbf{P}}^{\Pi} \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k)$ is a martingale difference such that $\mathbb{E}[\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) - \widetilde{\mathbf{P}}^{\Pi} \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k) \rangle | \mathcal{F}_{k-1}] = 0$, where we recall $\mathcal{F}_k = \sigma\{\boldsymbol{\theta}_0, \widetilde{s}_0, \widetilde{s}_1, \dots, \widetilde{s}_k\}$. Furthermore, the remaining inner product can be decomposed into four terms:

$$\begin{aligned} &\beta_k \left\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k) - \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) \right\rangle \\ &= \underbrace{\beta_k \left\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_k) \right\rangle}_{=A_k^1} - \underbrace{\beta_{k+1} \left\langle \boldsymbol{\theta}_{k+1} - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_{k+1}; \widetilde{s}_{k+1}) \right\rangle}_{=A_k^2} \\ &\quad + \underbrace{(\beta_{k+1} - \beta_k) \left\langle \boldsymbol{\theta}_{k+1} - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_{k+1}; \widetilde{s}_{k+1}) \right\rangle}_{=A_k^3} + \underbrace{\beta_k \left\langle \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_{k+1}; \widetilde{s}_{k+1}) \right\rangle}_{=A_k^4} \\ &\quad + \underbrace{\beta_k \left\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_{k+1}; \widetilde{s}_{k+1}) - \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_k; \widetilde{s}_{k+1}) \right\rangle}_{=A_k^4}. \end{aligned}$$

In particular, we have

$$\left| \sum_{k=0}^n A_k^1 \right| = \beta_0 \left\langle \boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_0; \widetilde{s}_0) \right\rangle - \beta_{n+1} \left\langle \boldsymbol{\theta}_{n+1} - \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{d}}_{\mathbf{J}_v}^{(1)}(\mathbf{z}_{n+1}; \widetilde{s}_{n+1}) \right\rangle \leq (\beta_0 + \beta_{n+1}) B \zeta,$$

$$\begin{aligned} \left| \sum_{k=0}^n A_k^2 \right| &\stackrel{(a)}{\leq} \xi B \zeta \sum_{k=0}^n \beta_k^2, \quad \left| \sum_{k=0}^n A_k^3 \right| \stackrel{(b)}{\leq} \widehat{\zeta} \sum_{k=0}^n \beta_k^2, \\ \left| \sum_{k=0}^n A_k^4 \right| &\stackrel{(c)}{\leq} 2\widehat{L}_J B \zeta \sum_{k=0}^n \beta_k^2 + \widehat{C}_J \frac{B^{7/3} (\log m)^{3/2}}{m^{1/6}} \sum_{k=0}^n \beta_k, \end{aligned}$$

where (a) is due to the additional condition (90) on the step size β_k ; (b) is due to the uniform boundedness on sampled gradient, i.e., Lemma D.1, as well as the non-expansive property of

projection; (c) is due to the Lipschitz gradient condition in Lemma D.2 and (96). Notice that the same analysis is applicable to the sum of inner product $\beta_k \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \mathbf{e}_k^{(1)} \rangle$. The above analysis shows that

$$2 \left| \sum_{k=0}^n \beta_k \mathbb{E}[\langle \boldsymbol{\theta}_k - \widehat{\boldsymbol{\theta}}, \mathbf{e}_k^{(1)} \rangle - \langle \mathbf{w}_k - \widehat{\mathbf{w}}, \mathbf{e}_k^{(2)} \rangle] \right| \leq 4 \left\{ (\beta_0 + \beta_{n+1}) B \widehat{\varsigma} + (\xi B \widehat{\varsigma} + \widehat{\varsigma} \varsigma + 2 \widehat{L}_J B \varsigma) \sum_{k=0}^n \beta_k^2 + \widehat{C}_J \frac{B^{7/3} (\log m)^{3/2}}{m^{1/6}} \sum_{k=0}^n \beta_k \right\}. \quad (98)$$

Plugging the above into (97) gives the convergence rate as

$$\begin{aligned} \mathbb{E}[d_f(\mathbf{z}_{I_n}, \mathbf{z}^*)] &\leq \frac{2C^e B^{\frac{8}{3}} + 4B^2}{\mu m^{\frac{1}{6}} / \log m} + \frac{\mathbb{E}[\|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2] + 2\{3C^e B^{\frac{8}{3}} m^{-1/3} (\log m)^2 + 12\varsigma^2\} \sum_{k=0}^n \beta_k^2}{\mu \sum_{k=0}^n \beta_k} \\ &\quad + \frac{2(\beta_0 + \beta_{n+1}) B \widehat{\varsigma} + 2(\xi B \widehat{\varsigma} + \widehat{\varsigma} \varsigma + 2 \widehat{L}_J B \varsigma) \sum_{k=0}^n \beta_k^2}{\mu \sum_{k=0}^n \beta_k} + \frac{4 \widehat{C}_J B^{7/3} (\log m)^{3/2}}{m^{1/6}} \\ &= \mathcal{O} \left(\frac{B^{\frac{8}{3}} + B^{\frac{7}{3}} (\log m)^{\frac{1}{2}} / (1 - \rho)}{(1 \wedge \nu) m^{\frac{1}{6}} / \log m} + \frac{\|\mathbf{z}_0 - \widehat{\mathbf{z}}\|_2^2 + \beta_0 \frac{B^2}{1 - \rho} + \frac{B^2}{1 - \rho} \sum_{k=0}^n \beta_k^2}{(1 \wedge \nu) \sum_{k=0}^n \beta_k} \right), \end{aligned} \quad (99)$$

where we have recalled that $\mu = 1 \wedge \nu$ and assumed $B = \Omega(\log m)$ in estimating the upper bound for $\varsigma, \widehat{\varsigma}$.

We thus conclude that the neural GTD algorithm converges when using Markov samples, yet the rate holds with high probability with respect to the random initialization. Moreover, the asymptotic bias of $\mathcal{O}(B^{\frac{8}{3}} (\log m)^{\frac{3}{2}} / ((1 - \rho)(1 \wedge \nu) m^{\frac{1}{6}}))$ is higher than that of the i.i.d. case. \square

Lastly, we present the following result that is akin to Theorem 3.4, but holds with high probability with respect to the random initialization of NN.

Corollary D.1. *Assume H1, H4. For any $\boldsymbol{\theta} \in S_B$, with probability at least $1 - e^{-\Omega(\log^2 m)}$, it holds*

$$\begin{aligned} J_v(\boldsymbol{\theta}) - J_v(\boldsymbol{\theta}^*) &\leq \mathcal{O} \left(c_{\text{nn}} + d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}}) + B \sqrt{d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})} + B^{8/3} m^{-1/3} \log m + B^{7/3} m^{-1/6} \sqrt{\log m} \right), \end{aligned}$$

where \mathbf{z} is defined as the vector $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{w})$ for any $\mathbf{w} \in S_B$, and $d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}}(\Xi_0))$ was defined in (18).

Proof. The proof follows from repeating the steps from the previous section using the high probability bounds in this appendix. Following the previous notations, it can be shown that it holds with probability at least $1 - e^{-\Omega(\log^2 m)}$ for

$$\text{r.h.s. of (84)} = \mathcal{O} \left(B^{8/3} m^{-1/3} \log m + B^{7/3} m^{-1/6} \sqrt{\log m} \right),$$

where we have repeatedly applied Lemma D.3 and [Allen-Zhu et al., 2019b, Theorem 1], see (101).

Moreover, for $R_\delta^{(2)}, R_f^{(2)}$, it can be shown that

$$R_\delta^{(2)} \vee R_f^{(2)} = \mathcal{O} \left(d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}}) + B \sqrt{d_{\widehat{f}}(\mathbf{z}, \widehat{\mathbf{z}})} \right).$$

Combining the terms yields the desired corollary. \square

D.1 Proof of Lemma D.1

First we observe that the gradient is given by:

$$\nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}, \mathbf{w}; \widetilde{\mathbf{s}}) = f(s, \mathbf{w}) \nabla \delta(s, a, s'; \boldsymbol{\theta}) + v f(s, \boldsymbol{\theta}) \nabla f(s, \boldsymbol{\theta}), \quad (100)$$

where we recall the definition of $\delta(\cdot)$ from (5). Using [Gao et al., 2019, Lemma A.5], it can be shown that with probability at least $1 - e^{-\Omega(B^{2/3} m^{2/3})}$, we have for any $s, a, s' \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$,

$$\|\nabla f(s, \boldsymbol{\theta})\|_2 = \mathcal{O}(1), \quad \|\nabla \delta(s, a, s'; \boldsymbol{\theta})\|_2 = \mathcal{O}(1),$$

where we have also used $\sup_{a,s} \rho(a|s) = \mathcal{O}(1)$. Furthermore, from [Allen-Zhu et al., 2019b, Theorem 1], with probability at least $1 - e^{-\Omega(\log^2 m)}$, we have

$$|f(s, \boldsymbol{\theta})| \leq |f(s, \boldsymbol{\theta}_0)| + |f(s, \boldsymbol{\theta}) - f(s, \boldsymbol{\theta}_0)| = \mathcal{O}(B \vee \log m), \quad (101)$$

where we have used the Lipschitz property of the NN function where $|f(s, \boldsymbol{\theta}) - f(s, \boldsymbol{\theta}_0)| = \mathcal{O}(B)$. Combining the above observations gives $\|\nabla_{\boldsymbol{\theta}} J_v(\boldsymbol{\theta}, \mathbf{w}; \tilde{s})\|_2 = \mathcal{O}(B \vee \log m)$ with probability at least $1 - e^{-\Omega(\log^2 m)}$. On the other hand, we observe

$$\nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}, \mathbf{w}; \tilde{s}) = (\delta(s, a, s'; \boldsymbol{\theta}) - f(s, \mathbf{w})) \nabla_{\mathbf{w}} f(s, \mathbf{w}). \quad (102)$$

Following the same arguments as before yields $\|\nabla_{\mathbf{w}} J_v(\boldsymbol{\theta}, \mathbf{w}; \tilde{s})\|_2 = \mathcal{O}(B \vee \log m)$ with high probability. The proof is concluded.

D.2 Proof of Lemma D.2

Observe that

$$\begin{aligned} \|\nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}; \tilde{s}) - \nabla_{\boldsymbol{\theta}} J_v(\mathbf{z}'; \tilde{s})\| &\leq \|f(s, \mathbf{w}) \nabla \delta(s, a, s'; \boldsymbol{\theta}) - f(s, \mathbf{w}') \nabla \delta(s, a, s'; \boldsymbol{\theta}')\|_2 \\ &\quad + v \|f(s, \boldsymbol{\theta}) \nabla f(s, \boldsymbol{\theta}) - f(s, \boldsymbol{\theta}') \nabla f(s, \boldsymbol{\theta}')\|_2. \end{aligned}$$

As we have

$$\begin{aligned} &\|f(s, \mathbf{w}) \nabla \delta(s, a, s'; \boldsymbol{\theta}) - f(s, \mathbf{w}') \nabla \delta(s, a, s'; \boldsymbol{\theta}')\|_2 \\ &\leq |f(s, \mathbf{w}) - f(s, \mathbf{w}')| \|\nabla \delta(s, a, s'; \boldsymbol{\theta})\|_2 + |f(s, \mathbf{w}')| \|\nabla \delta(s, a, s'; \boldsymbol{\theta}) - \nabla \delta(s, a, s'; \boldsymbol{\theta}')\|_2 \\ &\stackrel{(a)}{\leq} \mathcal{O}(1) \|\mathbf{w} - \mathbf{w}'\|_2 + \mathcal{O}(B^{4/3} m^{-1/6} (\log m)^{3/2}), \end{aligned}$$

where (a) holds with probability at least $1 - e^{-\Omega(\log^2 m)}$ by using [Allen-Zhu et al., 2019b, Theorem 1] on the first term and [Gao et al., 2019, Lemma A.5] on the second term. Similarly, we have

$$\|f(s, \boldsymbol{\theta}) \nabla f(s, \boldsymbol{\theta}) - f(s, \boldsymbol{\theta}') \nabla f(s, \boldsymbol{\theta}')\|_2 \leq \mathcal{O}(1) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2 + \mathcal{O}(B^{4/3} m^{-1/6} (\log m)^{3/2}).$$

On the other hand, we have

$$\begin{aligned} &\|\nabla_{\mathbf{w}} J_v(\mathbf{z}; \tilde{s}) - \nabla_{\mathbf{w}} J_v(\mathbf{z}'; \tilde{s})\|_2 \\ &= \|(\delta(s, a, s'; \boldsymbol{\theta}) - f(s, \mathbf{w})) \nabla f(s, \mathbf{w}) - (\delta(s, a, s'; \boldsymbol{\theta}') - f(s, \mathbf{w}')) \nabla f(s, \mathbf{w}')\|_2 \\ &\leq |\delta(s, a, s'; \boldsymbol{\theta}) - \delta(s, a, s'; \boldsymbol{\theta}') - f(s, \mathbf{w}) + f(s, \mathbf{w}')| \|\nabla f(s, \mathbf{w})\|_2 \\ &\quad + |\delta(s, a, s'; \boldsymbol{\theta}') - f(s, \mathbf{w}')| \|\nabla f(s, \mathbf{w}') - \nabla f(s, \mathbf{w}')\|_2. \end{aligned} \quad (103)$$

Using the same arguments as before, we can show

$$\|\nabla_{\mathbf{w}} J_v(\mathbf{z}; \tilde{s}) - \nabla_{\mathbf{w}} J_v(\mathbf{z}'; \tilde{s})\|_2 \leq \mathcal{O}(1) \|\mathbf{z} - \mathbf{z}'\|_2 + \mathcal{O}(B^{4/3} m^{-1/6} (\log m)^{3/2}).$$

This concludes our proof.

D.3 Proof of Lemma D.3

Observe that

$$\begin{aligned} f(s, \boldsymbol{\theta}) - \widehat{f}(s, \boldsymbol{\theta}) &= \int_0^1 \langle \nabla f(s, (1-t)\boldsymbol{\theta}_0 + t\boldsymbol{\theta}) - \nabla \widehat{f}(s, \boldsymbol{\theta}_0), \boldsymbol{\theta} - \boldsymbol{\theta}_0 \rangle dt \\ &\leq B \int_0^1 \|\nabla f(s, (1-t)\boldsymbol{\theta}_0 + t\boldsymbol{\theta}) - \nabla f(s, \boldsymbol{\theta}_0)\|_2 dt = \mathcal{O}(B^{4/3} m^{-1/6} \sqrt{\log m}), \end{aligned}$$

where the last equality is due to [Gao et al., 2019, Lemma A.5] and it holds with probability at least $1 - e^{-\Omega(B^{2/3} m^{2/3})}$. This concludes the proof of the lemma.

D.4 Proof of Lemma D.4

Following the proof from our previous Lemma A.1, we observe that

$$\|\widehat{\mathbf{e}}_k^{(1)}\|_2^2 \leq 2\mathbb{E}_\mu[E_k^{(1,1)} + v^2 E_k^{(1,2)}],$$

with

$$E_k^{(1,1)} \leq 2|f(s, \mathbf{w}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k)\|_2^2 + 2|\widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2^2,$$

$$E_k^{(1,2)} \leq 2|f(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \boldsymbol{\theta}_k)|^2 \|\nabla f(s, \boldsymbol{\theta}_k)\|_2^2 + 2|\widehat{f}(s, \boldsymbol{\theta}_k)|^2 \|\nabla f(s, \boldsymbol{\theta}_k) - \nabla \widehat{f}(s, \boldsymbol{\theta}_k)\|_2^2.$$

Using Lemma D.3 and [Gao et al., 2019, Lemma A.5], with probability at least $1 - e^{-\Omega(B^{2/3}m^{2/3})}$, we have

$$|f(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \boldsymbol{\theta}_k)|^2 \|\nabla f(s, \boldsymbol{\theta}_k)\|_2^2 = \mathcal{O}(B^{8/3}m^{-1/3} \log m).$$

Moreover, as $\nabla \widehat{f}(s, \boldsymbol{\theta}_k) = \nabla \widehat{f}(s, \boldsymbol{\theta}_0)$ and $\nabla \widehat{\delta}(s, \boldsymbol{\theta}_k) = \nabla \widehat{\delta}(s, \boldsymbol{\theta}_0)$, using (101), we have

$$|\widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \bar{\delta}(s, \boldsymbol{\theta}_k) - \nabla \widehat{\delta}(s, \boldsymbol{\theta}_k)\|_2^2 = \mathcal{O}((B \vee \log m)B^{2/3}m^{-1/3} \log m),$$

where the equality holds with probability at least $1 - e^{-\Omega(\log^2 m)}$ due to [Allen-Zhu et al., 2019b, Theorem 1] and [Gao et al., 2019, Lemma A.5]. Combining the above yields

$$\|e_k^{(1)}\|_2^2 = \mathcal{O}(B^{8/3}m^{-1/3}(\log m)^2). \quad (104)$$

On the other hand,

$$\begin{aligned} \|e_k^{(2)}\|_2^2 &\leq 2\mathbb{E}_\mu \left[|\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k))|^2 \|\nabla f(s, \mathbf{w}_k)\|_2^2 \right] \\ &\quad + 2\mathbb{E}_\mu \left[|\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \widehat{f}(s, \mathbf{w}_k) - \nabla f(s, \mathbf{w}_k)\|_2^2 \right]. \end{aligned}$$

Likewise, we can bound

$$|\bar{\delta}(s, \boldsymbol{\theta}_k) - f(s, \mathbf{w}_k) - (\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k))|^2 \|\nabla f(s, \mathbf{w}_k)\|_2^2 = \mathcal{O}(B^{8/3}m^{-1/3} \log m)$$

$$|\widehat{\delta}(s, \boldsymbol{\theta}_k) - \widehat{f}(s, \mathbf{w}_k)|^2 \|\nabla \widehat{f}(s, \mathbf{w}_k) - \nabla f(s, \mathbf{w}_k)\|_2^2 = \mathcal{O}((B \vee \log m)B^{2/3}m^{-1/3} \log m),$$

with probability at least $1 - e^{-\Omega(\log^2 m)}$. Thus, we also get

$$\|e_k^{(1)}\|_2^2 = \mathcal{O}(B^{8/3}m^{-1/3}(\log m)^2). \quad (105)$$

This concludes our proof.