
Improved Sample Complexity Analysis of Natural Policy Gradient Algorithm with General Parameterization for Infinite Horizon Discounted Reward Markov Decision Processes

Washim Uddin Mondal
Purdue University, USA

Vaneet Aggarwal
Purdue University, USA

Abstract

We consider the problem of designing sample efficient learning algorithms for infinite horizon discounted reward Markov Decision Process. Specifically, we propose the Accelerated Natural Policy Gradient (ANPG) algorithm that utilizes an accelerated stochastic gradient descent process to obtain the natural policy gradient. ANPG achieves $\mathcal{O}(\epsilon^{-2})$ sample complexity and $\mathcal{O}(\epsilon^{-1})$ iteration complexity with general parameterization where ϵ defines the optimality error. This improves the state-of-the-art sample complexity by a $\log(\frac{1}{\epsilon})$ factor. ANPG is a first-order algorithm and unlike some existing literature, does not require the unverifiable assumption that the variance of importance sampling (IS) weights is upper bounded. In the class of Hessian-free and IS-free algorithms, ANPG beats the best-known sample complexity by a factor of $\mathcal{O}(\epsilon^{-\frac{1}{2}})$ and simultaneously matches their state-of-the-art iteration complexity.

1 INTRODUCTION

Reinforcement Learning (RL) is a sequential decision-making framework that finds its applications in a wide range of areas, ranging from epidemic control to transportation to online marketing (Ling et al., 2023; Singh et al., 2023; Al-Abbasi et al., 2019). The objective of an RL agent is to obtain a policy that maximizes the discounted sum of expected rewards that are generated due to its interaction with the environment. One way to solve this problem is called the policy gradient (PG) approach which performs the optimization directly in

the policy space. Value function-based approaches, on the contrary, optimize the Q -functions which are then used to filter out the policies.

Paired with the general function approximation, PG-based methods turn out to be an effective tool in handling large state spaces. Specifically, deep neural network (DNN)-based PG has shown tremendous success empirically (Schulman et al., 2015, 2017). The idea of function approximation is to parameterize the policies by a d -dimensional parameter, θ so that the optimization can be performed over \mathbb{R}^d . In this case, however, the objective function, $J(\theta)$, appears to be non-convex. As a result, many articles primarily analyze the first-order stationary (FOS) convergence properties of $J(\theta)$. For example, (Yuan et al., 2022) derives $\tilde{\mathcal{O}}(\epsilon^{-4})$ sample complexity for Vanilla-PG algorithm to achieve an ϵ -FOS error.

In this paper, however, we are interested in the global convergence property of $J(\theta)$. It can be shown that under the assumption of Fisher Non-Degeneracy (FND), $J(\theta)$ satisfies the gradient domination property which implies that there does not exist any FOS sub-optimal points of $J(\theta)$. As a result, one can achieve any ϵ global optimum (GO) gap, despite the objective function being non-convex. For example, (Liu et al., 2020) derives $\mathcal{O}(\epsilon^{-4})$ and $\mathcal{O}(\epsilon^{-3})$ sample complexities of the PG and the Natural PG (NPG) algorithms respectively corresponding to an ϵ GO gap. In recent years, importance sampling (IS) has garnered popularity as an important sample complexity-reducing method. (Liu et al., 2020) utilizes IS to propose variance-reduction (VR) versions of PG and NPG algorithms, each achieving a sample complexity of $\tilde{\mathcal{O}}(\epsilon^{-3})$. The caveat of IS-based methods is that one needs to impose a bound on the variance of the IS weights. This is a strong assumption and might not be verifiable for most practical cases. Alternatives to the IS-based method include the momentum-based N-PG-IGT algorithm of (Fatkhullin et al., 2023) and the Hessian-based SCRN algorithm of (Masiha et al., 2022), both of which achieve $\tilde{\mathcal{O}}(\epsilon^{-2.5})$ sample complexity. Hessian-based (second-order) approaches typically

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

Algorithm	Sample Complexity	Iteration Complexity	Hessian-free	IS-free
Vanilla-PG (Yuan et al., 2022)	$\tilde{\mathcal{O}}(\epsilon^{-3})$	$\mathcal{O}(\epsilon^{-3})$	Yes	Yes
STORM-PG-F (Ding et al., 2022)	$\tilde{\mathcal{O}}((1+W)^{\frac{3}{2}}\epsilon^{-3})$	$\mathcal{O}(\epsilon^{-3})$	Yes	No
SCRN (Masiha et al., 2022)	$\tilde{\mathcal{O}}(\epsilon^{-2.5})$	$\mathcal{O}(\epsilon^{-0.5})$	No	Yes
VR-SCRN (Masiha et al., 2022)	$\mathcal{O}(\epsilon^{-2} \log(\frac{1}{\epsilon}))$	$\mathcal{O}(\epsilon^{-0.5})$	No	No
NPG (Liu et al., 2020)	$\mathcal{O}(\epsilon^{-3})$	$\mathcal{O}(\epsilon^{-1})$	Yes	Yes
SRVR-NPG (Liu et al., 2020)	$\mathcal{O}(W\epsilon^{-2.5} + \epsilon^{-3})$	$\mathcal{O}(\epsilon^{-1})$	Yes	No
SRVR-PG (Liu et al., 2020)	$\mathcal{O}(W\epsilon^{-3})$	$\mathcal{O}(\epsilon^{-2})$	Yes	No
N-PG-IGT (Fatkhullin et al., 2023)	$\tilde{\mathcal{O}}(\epsilon^{-2.5})$	$\mathcal{O}(\epsilon^{-2.5})$	Yes	Yes
HARPG (Fatkhullin et al., 2023)	$\mathcal{O}(\epsilon^{-2} \log(\frac{1}{\epsilon}))$	$\mathcal{O}(\epsilon^{-2})$	No	Yes
ANPG (This work)	$\mathcal{O}(\epsilon^{-2})$	$\mathcal{O}(\epsilon^{-1})$	Yes	Yes
Lower Bound (Azar et al., 2017)	$\mathcal{O}(\epsilon^{-2})$	–	–	–

Table 1: Summary of recent sample complexity results for global convergence in discounted reward MDPs with general parameterized policies. The term, ‘‘Hessian-free’’ states that the underlying algorithm only utilizes first-order information. The ‘‘IS-free’’ column states whether a bound (W) on the variance of the importance sample (IS) weights is assumed. It is worthwhile to mention here that although the HARPG algorithm is Hessian-aided, its per-iteration computational and memory requirements are similar to that of Hessian-free methods (Fatkhullin et al., 2023). We have explicitly written the logarithmic terms in the sample complexities of HARPG and VR-SCRN to highlight the superiority of our result.

have more memory and computational requirements in comparison to the first-order algorithms and thus, are less preferred in practice. The state-of-the-art (SOTA) $\tilde{\mathcal{O}}(\epsilon^{-2})$ sample complexity is achieved by two recently proposed algorithms: the VR-SCRN by (Masiha et al., 2022) and the HARPG proposed by (Fatkhullin et al., 2023). The first one uses a combination of Hessian and IS-based methods whereas the other one solely applies a Hessian-aided technique. This leads to the following question.

Does there exist an IS-free and Hessian-free algorithm that either achieves or improves the SOTA $\tilde{\mathcal{O}}(\epsilon^{-2})$ sample complexity?

1.1 Our Contributions and Challenges

In this paper, we provide an affirmative answer to the above question. In particular,

- We propose an acceleration-based NPG (ANPG) algorithm that uses accelerated stochastic gradient descent (ASGD) to determine an estimate of the natural policy gradient.
- We prove that ANPG achieves a sample complexity of $\mathcal{O}(\epsilon^{-2})$ and an iteration complexity of $\mathcal{O}(\epsilon^{-1})$.
- ANPG is IS-free and does not need any second-order (Hessian-related) computation.

Note that the best known sample complexity is $\tilde{\mathcal{O}}(\epsilon^{-2}) = \mathcal{O}(\epsilon^{-2} \log(\frac{1}{\epsilon}))$. Clearly, our proposed ANPG algo-

rithm improves the SOTA sample complexity by a factor of $\log(\frac{1}{\epsilon})$. Additionally, in the class of Hessian-free and IS-free algorithms, it improves the SOTA sample complexity by a factor of $\epsilon^{-0.5}$ (see Table 1). Note that our sample complexity matches the lower bound and hence, cannot be improved further. Moreover, the iteration complexity of our algorithm matches the SOTA iteration complexity of Hessian-free and IS-free algorithms and beats that of the HARPG algorithm by a factor of $\mathcal{O}(\epsilon^{-1})$.

The NPG algorithm updates the policy parameter, θ for K iterations and at each iteration, $k \in \{0, \dots, K-1\}$, at estimate, ω_k , of the true natural gradient, ω_k^* is calculated via an H -step SGD method. The iteration complexity of the NPG, as shown in (Liu et al., 2020) is $\mathcal{O}(K) = \mathcal{O}(\epsilon^{-1})$. This is difficult to improve since even with exact natural gradients and softmax parameterization, we obtain the same result (Agarwal et al., 2021, Theorem 16). On the other hand, the estimation error bound $\mathbb{E}\|\omega_k - \omega_k^*\| = \mathcal{O}(1/\sqrt{H})$ is standard and hard to change. This leads to a choice of $H = \mathcal{O}(\epsilon^{-2})$, resulting in a sample complexity of $\mathcal{O}(\epsilon^{-3})$. Therefore, the sample complexity of the NPG, in its current form, seems to be hard to break. We introduce the following modifications to surpass this apparent barrier. Firstly, the intermediate direction finding SGD routine is substituted by the ASGD procedure. Next, by improving the global convergence analysis, we show that the first-order estimation error can be more accurately written as $\mathbf{E}\|(\mathbf{E}[\omega_k|\theta_k] - \omega_k^*)\|$ where θ_k is the policy parameter at the k th iteration. Finally, we prove that the stated

term can be interpreted as the approximation error of a noiseless (deterministic) accelerated gradient descent process and hence can be upper bounded as $\exp(-\kappa H)$ for some H -independent parameter κ . This allows us to take $H = \mathcal{O}(\epsilon^{-1})$, and resulting in $\mathcal{O}(\epsilon^{-2})$ sample complexity.

1.2 Related Works

• **Convergence with Exact Gradients:** A series of papers in the literature explore the global convergence rate (or equivalently, iteration complexity) of PG-type algorithms assuming access to the exact gradients. For example, (Agarwal et al., 2021) shows $\mathcal{O}(\epsilon^{-2})$ iteration complexity of the exact PG with softmax parameterization and log barrier regularization. (Mei et al., 2020) obtains an improved $\mathcal{O}(\epsilon^{-1})$ result without regularization and $\mathcal{O}(\log(\frac{1}{\epsilon}))$ result with regularization. (Shani et al., 2020) proves $\mathcal{O}(\epsilon^{-1})$ iteration complexity for the exact regularized NPG. (Agarwal et al., 2021) proves an identical result for the unregularized NPG. Finally, $\mathcal{O}(\log(\frac{1}{\epsilon}))$ iteration complexity is recently shown both for regularized (Cen et al., 2022; Lan, 2023; Zhan et al., 2023) and unregularized (Bhandari and Russo, 2021; Xiao, 2022) NPG.

• **FOS Sample Complexity:** Many variance reduction methods exist in the literature that reaches the ϵ -FOS point faster than the Vanilla-PG algorithm (Yuan et al., 2022). The majority of these works incorporate some form of the importance sampling (IS) technique. Works that fall into this category include (Papini et al., 2018; Xu et al., 2019, 2020; Gargiani et al., 2022; Pham et al., 2020; Yuan et al., 2020; Huang et al., 2020), etc. As an alternative to the IS-based method, some papers (Salehkaleybar et al., 2022; Shen et al., 2019) resort to second-order (Hessian-related) information. The techniques stated above are sometimes used in conjunction with momentum-based updates (Salehkaleybar et al., 2022; Huang et al., 2020).

• **GO Sample Complexity:** Most of the works that target global optimal (GO) convergence using stochastic gradients have been discussed earlier. Alongside, a plethora of works has emerged that investigate the GO convergence via actor-critic methods with linear function approximation (Wu et al., 2020; Chen et al., 2022; Chen and Maguluri, 2022; Khodadadian et al., 2022; Qiu et al., 2021). These works assume the underlying MDP to be ergodic. Fortunately, such an assumption is not a requirement for our analysis. As mentioned before, we assume Fisher Non-Degenerate (FND) policies and the compatible function approximation framework (defined in section 4) to arrive at our result. A similar analysis of the PG algorithm with general parameterization has been recently explored for average reward MDPs (Bai et al., 2023b).

2 PROBLEM FORMULATION

We consider a Markov Decision Process (MDP) characterized by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, P, \gamma, \rho)$ where \mathcal{S}, \mathcal{A} are the state space and the action space respectively, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{|\mathcal{S}|}$ is the state transition kernel (where $\Delta^{|\mathcal{S}|}$ indicates a probability simplex of size $|\mathcal{S}|$), $\gamma \in (0, 1)$ denotes the reward discount factor, and $\rho \in \Delta^{|\mathcal{S}|}$ is the initial state distribution. A policy $\pi : \mathcal{S} \rightarrow \Delta^{|\mathcal{A}|}$ characterizes the distribution of action to be executed given the current state. For a state-action pair (s, a) , the Q -function of a policy π is defined as follows.

$$Q^\pi(s, a) \triangleq \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \middle| s_0 = s, a_0 = a \right]$$

where the expectation is computed over all π -induced trajectories $\{(s_t, a_t)\}_{t=0}^{\infty}$ where $s_t \sim P(s_{t-1}, a_{t-1})$, and $a_t \sim \pi(s_t)$, $\forall t \in \{1, 2, \dots\}$. Similarly, the V -function of the policy π is defined as,

$$\begin{aligned} V^\pi(s) &\triangleq \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \middle| s_0 = s \right] \\ &= \sum_{a \in \mathcal{A}} \pi(a|s) Q^\pi(s, a) \end{aligned}$$

Finally, the advantage function is defined as,

$$A^\pi(s, a) \triangleq Q^\pi(s, a) - V^\pi(s), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}$$

Our objective is to maximize the function defined below over the class of all policies.

$$J_\rho^\pi \triangleq \mathbf{E}_{s \sim \rho} [V^\pi(s)] = \frac{1}{1-\gamma} \sum_{s,a} d_\rho^\pi(s, a) \pi(a|s) r(s, a)$$

where the state occupancy $d_\rho^\pi \in \Delta^{|\mathcal{S}|}$ is defined as,

$$d_\rho^\pi(s) \triangleq (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s | s_0 \sim \rho, \pi), \quad \forall s \in \mathcal{S}$$

We define the state-action occupancy $\nu_\rho^\pi \in \Delta^{|\mathcal{S} \times \mathcal{A}|}$ as $\nu_\rho^\pi(s, a) \triangleq d_\rho^\pi(s) \pi(a|s)$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$.

In many application scenarios, the state space is quite large which calls for the parameterization of the policy functions by a deep neural network with d -dimensional parameters. Let, π_θ indicate the policy corresponding to the parameter $\theta \in \mathbb{R}^d$. In this case, our maximization problem can be restated as shown below.

$$\max_{\theta \in \mathbb{R}^d} J_\rho^{\pi_\theta} \quad (1)$$

For notational convenience, we denote $J_\rho^{\pi_\theta}$ as $J_\rho(\theta)$ in the rest of the paper.

3 ALGORITHM

One way to solve the maximization (1) is via updating the policy parameters by applying the gradient ascent: $\theta_{k+1} = \theta_k + \eta \nabla_{\theta} J_{\rho}(\theta_k)$, $k \in \{0, 1, \dots\}$ starting with an initial parameter, θ_0 . Here, $\eta > 0$ denotes the learning rate, and the policy gradients (PG) are given as follows (Sutton et al., 1999).

$$\nabla_{\theta} J_{\rho}(\theta) = \frac{1}{1-\gamma} H_{\rho}(\theta), \text{ where} \quad (2)$$

$$H_{\rho}(\theta) \triangleq \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi_{\theta}}} [A^{\pi_{\theta}}(s, a) \nabla_{\theta} \log \pi_{\theta}(a|s)] \quad (3)$$

In contrast, this paper uses the natural policy gradient (NPG) to update the policy parameters. In particular, $\forall k \in \{1, 2, \dots\}$, we have,

$$\theta_{k+1} = \theta_k + \eta F_{\rho}(\theta_k)^{\dagger} \nabla_{\theta} J_{\rho}(\theta_k) \quad (4)$$

where \dagger is the Moore-Penrose pseudoinverse operator, and F_{ρ} is called the Fisher information function which is defined as follows $\forall \theta \in \mathbb{R}^d$.

$$F_{\rho}(\theta) \triangleq \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi_{\theta}}} [\nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s)] \quad (5)$$

where \otimes denotes the outer product. In particular, for any $\mathbf{a} \in \mathbb{R}^d$, $\mathbf{a} \otimes \mathbf{a} = \mathbf{a}\mathbf{a}^T$.

Note that, in comparison to a PG-based iteration, the advantage of an NPG-based iteration is that the latter modulates the learning parameter, η , according to the rate of change of the policy function itself. This results in a lower sample complexity as shown in the previous study (Liu et al., 2020). As mentioned before, our goal in this paper is to show that the sample complexity of the NPG-based update can be reduced further.

Let $\omega_{\theta}^* \triangleq F_{\rho}(\theta)^{\dagger} \nabla_{\theta} J_{\rho}(\theta)$. Notice that we have removed the dependence of ω_{θ}^* on ρ for notational convenience. Invoking (2), the term, ω_{θ}^* can be written as a solution of a quadratic optimization (Peters and Schaal, 2008). Specifically, $\omega_{\theta}^* = \arg \min_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$ where $L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$ is the compatible function approximation error and it is mathematically defined as,

$$L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) \triangleq \frac{1}{2} \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi_{\theta}}} \left[\frac{1}{1-\gamma} A^{\pi_{\theta}}(s, a) - \omega^T \nabla_{\theta} \log \pi_{\theta}(a|s) \right]^2 \quad (6)$$

Using the above notations, NPG updates can be written as $\theta_{k+1} = \theta_k + \eta \omega_k^*$, $k \in \{1, 2, \dots\}$ where $\omega_k^* = \omega_{\theta_k}^*$. However, in most practical scenarios, the learner does not have knowledge of the state transition probabilities which makes it difficult to directly calculate ω_k^* . Below

we clarify how one can get its sample-based estimates. Note that the gradient of $L_{\nu_{\rho}^{\pi_{\theta}}}(\cdot, \theta)$ can be obtained as shown below for any arbitrary θ .

$$\nabla_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) = F_{\rho}(\theta) \omega - \frac{1}{1-\gamma} H_{\rho}(\theta) \quad (7)$$

Although, in general, it is difficult to obtain $F_{\rho}(\theta)$ and $H_{\rho}(\theta)$ exactly, their (unbiased) estimates can be easily obtained via Algorithm 1. This is similar to Algorithm 3 stated in (Agarwal et al., 2021). The Algorithm first runs the MDP for T time instances following the policy π_{θ} starting from $s_0 \sim \rho$ where T is a geometric random variable with success probability $1-\gamma$. The tuple (s_T, a_T) is taken as the state-action pair (\hat{s}, \hat{a}) sampled from the distribution $\nu_{\rho}^{\pi_{\theta}}$. With probability $\frac{1}{2}$, the algorithm then generates another π_{θ} -induced trajectory of length $T \sim \text{Geo}(1-\gamma)$ starting from the pair (\hat{s}, \hat{a}) , and assigns twice of the total reward of this trajectory as the estimate $\hat{Q}^{\pi_{\theta}}(\hat{s}, \hat{a})$. In its complementary event, the trajectory initiates with \hat{s} , and the twice of its total reward is assigned as $\hat{V}^{\pi_{\theta}}(\hat{s})$. Its associated advantage estimate is obtained as $\hat{A}^{\pi_{\theta}}(\hat{s}, \hat{a}) = \hat{Q}^{\pi_{\theta}}(\hat{s}, \hat{a}) - \hat{V}^{\pi_{\theta}}(\hat{s})$. Finally, the estimate of the gradient is computed as,

$$\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) = \hat{F}_{\rho}(\theta) \omega - \frac{1}{1-\gamma} \hat{H}_{\rho}(\theta) \quad (8)$$

where $\hat{F}_{\rho}(\theta)$ and $\hat{H}_{\rho}(\theta)$ are given as follows.

$$\begin{aligned} \hat{F}_{\rho}(\theta) &= \nabla_{\theta} \log \pi_{\theta}(\hat{a}|\hat{s}) \otimes \nabla_{\theta} \log \pi_{\theta}(\hat{a}|\hat{s}) \\ \hat{H}_{\rho}(\theta) &= \hat{A}^{\pi_{\theta}}(\hat{s}, \hat{a}) \nabla_{\theta} \log \pi_{\theta}(\hat{a}|\hat{s}) \end{aligned} \quad (9)$$

The following lemma ensures that the estimate of the gradient generated by Algorithm 1 is indeed unbiased.

Lemma 1. *If $\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$ denotes the gradient estimate yielded by Algorithm 1, then the following holds.*

$$\mathbf{E} \left[\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) | \omega, \theta \right] = \nabla_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) \quad (10)$$

Thanks to Algorithm 1, we can utilize a gradient-based iterative process to minimize $L_{\nu_{\rho}^{\pi_{\theta}}}(\cdot, \theta_k)$ to generate an estimate of ω_k^* . Here we use the accelerated stochastic gradient descent (ASGD) process (Jain et al., 2018b), to achieve this goal. This is in contrast with the SGD-based algorithms typically used in the NPG literature (Liu et al., 2020). The entire procedure is described in Algorithm 2.

Algorithm 2 can be broadly segregated into two parts. In its *outer loop*, the policy parameters are updated K number of times following (16) where $\{\omega_k\}_{k=0}^{K-1}$ denote the estimates of $\{\omega_k^*\}_{k=0}^{K-1}$. These estimates are calculated by iterating the ASGD algorithm H number of times in the *inner loop*. Each ASGD iteration consists of four steps described by (11) – (14) where $(\alpha, \beta, \xi, \delta)$

Algorithm 1 Unbiased Sampling

1: **Input:** Parameters (θ, ω) , Initial Distribution ρ
 2: $T \sim \text{Geo}(1 - \gamma)$
 3: Sample $s_0 \sim \rho, a_0 \sim \pi_\theta(s_0)$
 4: **for** $j \in \{0, \dots, T - 1\}$ **do**
 5: Execute a_j and observe $s_{j+1} \sim P(s_j, a_j)$
 6: Sample $a_{j+1} \sim \pi_\theta(s_{j+1})$
 7: $(\hat{s}, \hat{a}) \leftarrow (s_T, a_T)$
 8: $(s_0, a_0) \leftarrow (s_T, a_T)$
 9: $T \sim \text{Geo}(1 - \gamma), X \sim \text{Bernoulli}(\frac{1}{2})$
 10: **if** $X = 1$ **then**
 11: Sample $a_0 \sim \pi_\theta(s_0)$
 12: **for** $j \in \{0, \dots, T - 1\}$ **do**
 13: Execute a_j and observe $s_{j+1} \sim P(s_j, a_j)$
 14: Sample $a_{j+1} \sim \pi_\theta(s_{j+1})$
 15: $\hat{Q}^{\pi_\theta}(\hat{s}, \hat{a}) \leftarrow 2(1 - X) \sum_{j=0}^T r(s_j, a_j)$
 16: $\hat{V}^{\pi_\theta}(\hat{s}) \leftarrow 2X \sum_{j=0}^T r(s_j, a_j)$
 17: $\hat{A}^{\pi_\theta}(\hat{s}, \hat{a}) \leftarrow \hat{Q}^{\pi_\theta}(\hat{s}, \hat{a}) - \hat{V}^{\pi_\theta}(\hat{s})$
 18: $\hat{F}_\rho(\theta) \leftarrow \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s}) \otimes \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s})$
 19: $\hat{H}_\rho(\theta) \leftarrow \hat{A}^{\pi_\theta}(\hat{s}, \hat{a}) \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s})$
 20: **Output:** $\hat{F}_\rho(\theta)\omega - \frac{1}{1 - \gamma} \hat{H}_\rho(\theta)$

are the associated learning rates. Note that in an SGD setup, the function argument is updated based only on the gradient observed at its current value. In contrast, the convoluted ASGD steps (11) – (14) ensure that the gradients obtained in the past iterations are also taken into account while updating the argument. The idea is essentially similar to Nesterov’s acceleration algorithm for deterministic gradients (Nesterov, 2012). The estimate, ω_k , is finally given by (15) where the outcomes of the last $H/2$ iterations are averaged.

4 ANALYSIS

In this section, we shall discuss the convergence properties of Algorithm 2. Before delving into the mathematical details, we would like to state the assumptions that are needed for the analysis.

Assumption 1. *The score function is G -Lipschitz and B -smooth. Mathematically, the following relations hold $\forall \theta, \theta_1, \theta_2 \in \mathbb{R}^d, \forall (s, a) \in \mathcal{S} \times \mathcal{A}$.*

- (a) $\|\nabla_\theta \log \pi_\theta(a|s)\| \leq G$
 (b) $\|\nabla_\theta \log \pi_{\theta_1}(a|s) - \nabla_\theta \log \pi_{\theta_2}(a|s)\| \leq B\|\theta_1 - \theta_2\|$ (17)

where B and G are some positive reals.

The Lipschitz continuity and smoothness of the score function are common assumptions in the literature and

Algorithm 2 Accelerated Natural Policy Gradient

1: **Input:** Policy Parameter θ_0 , State distribution ρ , Run-time Parameters K, H , Learning Parameters $(\eta, \alpha, \beta, \xi, \delta)$
 2: **for** $k \in \{0, \dots, K - 1\}$ **do** ▷ Outer Loop
 3: $\mathbf{x}_0, \mathbf{v}_0 \leftarrow \mathbf{0}$
 4: **for** $h \in \{0, \dots, H - 1\}$ **do** ▷ Inner Loop
 5: ▷ Accelerated Stochastic Gradient Descent

$$\mathbf{y}_h \leftarrow \alpha \mathbf{x}_h + (1 - \alpha) \mathbf{v}_h$$
 (11)
 6: Get $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta_k)|_{\omega=\mathbf{y}_h}$ via Algorithm 1

$$\mathbf{x}_{h+1} \leftarrow \mathbf{y}_h - \delta \hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta_k)|_{\omega=\mathbf{y}_h}$$
 (12)

$$\mathbf{z}_h \leftarrow \beta \mathbf{y}_h + (1 - \beta) \mathbf{v}_h$$
 (13)

$$\mathbf{v}_{h+1} \leftarrow \mathbf{z}_h - \xi \hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta_k)|_{\omega=\mathbf{y}_h}$$
 (14)
 7: Tail Averaging:

$$\omega_k \leftarrow \frac{2}{H} \sum_{\frac{H}{2} < h \leq H} \mathbf{x}_h$$
 (15)
 8: Policy Parameter Update:

$$\theta_{k+1} \leftarrow \theta_k + \eta \omega_k$$
 (16)

9: **Output:** $\{\theta_k\}_{k=0}^{K-1}$

these can be verified for simple parameterized policies such as Gaussian policies (Liu et al., 2020; Zhang et al., 2021; Aggarwal et al., 2020).

Note that Assumption 1 implies the following result.

Lemma 2. *If Assumption 1 holds, then $J_\rho(\cdot)$ defined in (1) satisfies the following properties $\forall \theta \in \mathbb{R}^d$.*

- (a) $\|\nabla_\theta J_\rho(\theta)\| \leq \frac{G}{(1 - \gamma)^2}$
 (b) $J_\rho(\cdot)$ is L -smooth, $L \triangleq \frac{B}{(1 - \gamma)^2} + \frac{2G^2}{(1 - \gamma)^3}$

Proof. Statement (a) can be proven by combining Assumption 1(a) with (2) whereas statement (b) follows directly from Proposition 4.2 of (Xu et al., 2019). \square

Lemma 2 establishes that the function $J_\rho(\cdot)$ is Lipschitz continuous and smooth with appropriate parameters. These two properties will be useful in our further analysis.

Assumption 2. *The compatible function approximation error defined in (6) satisfies the following $\forall \theta \in \mathbb{R}^d$.*

$$L_{\nu_\rho^{\pi^*}}(\omega_\theta^*, \theta) \leq \epsilon_{\text{bias}} \quad (18)$$

where π^* is the optimal policy i.e., $\pi^* = \arg \max_{\pi} J_{\rho}^{\pi}$ and ω_{θ}^* is defined as follows.

$$\omega_{\theta}^* \triangleq \arg \max_{\omega \in \mathbb{R}^d} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta) \quad (19)$$

Assumption 2 dictates that the parameterized policy class is rich enough such that for each policy parameter, θ , the transferred compatible function approximation error, symbolized by the LHS of (18), is bounded by a certain positive number, ϵ_{bias} . One can show that $\epsilon_{\text{bias}} = 0$ for softmax parameterization (Agarwal et al., 2021). Furthermore, ϵ_{bias} turns out to be small for rich neural network-based parameterization (Wang et al., 2019).

Assumption 3. *There exists a constant $\mu_F > 0$ such that $F_{\rho}(\theta) - \mu_F I_d$ is positive semidefinite where I_d denotes an identity matrix of dimension d and $F_{\rho}(\theta)$ is defined in (5). Equivalently, $F_{\rho}(\theta) \succcurlyeq \mu_F I_d$.*

Assumption 3 states that the Fisher information function cannot be too small. This is also commonly used in the policy gradient-based literature (Liu et al., 2020; Bai et al., 2023a; Zhang et al., 2020). This is obeyed by Gaussian policies with linearly parameterized means. The set of policies that follows Assumption 3 is called a Fisher Non-degeneracy (FND) set. Note that, $F_{\rho}(\theta) = \nabla_{\omega}^2 L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$. Hence, Assumption 3 essentially states that $L_{\nu_{\rho}^{\pi_{\theta}}}(\cdot, \theta)$ is strongly convex with parameter, μ_F . (Mondal et al., 2023) lays out a concrete example of a class of policies that satisfies assumptions 1 – 3.

4.1 Outer Loop Analysis

Recall that in the outer loop, the policy parameter θ_k 's are updated via (16) where ω_k 's are estimates of ω_k^* 's, the natural policy gradients. The following lemma establishes how θ_k 's convergence is intimately connected with the convergence of ω_k 's.

Lemma 3. *Let, the parameters $\{\theta_k\}_{k=0}^{K-1}$ be updated via (16), π^* be the optimal policy and J_{ρ}^* be the optimal value of the function $J_{\rho}(\cdot)$. If assumption 1 and 2 hold, then the following inequality is satisfied.*

$$\begin{aligned} J_{\rho}^* - \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}[J_{\rho}(\theta_k)] &\leq \sqrt{\epsilon_{\text{bias}}} \\ &+ \frac{G}{K} \sum_{k=0}^{K-1} \mathbf{E}[\|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|] + \frac{B\eta}{2K} \sum_{k=0}^{K-1} \mathbf{E}\|\omega_k\|^2 \\ &+ \frac{1}{\eta K} \mathbf{E}_{s \sim d^{\pi^*}} [KL(\pi^*(\cdot|s) \| \pi_{\theta_0}(\cdot|s))] \end{aligned} \quad (20)$$

where $KL(\cdot \| \cdot)$ is the Kullback-Leibler divergence.

Lemma 3 is almost identical to Proposition 4.5 of (Liu et al., 2020). However, there is an important difference

between these results. The first order term in the cited paper is $\mathbf{E}\|\omega_k - \omega_k^*\|$ whereas in (20), we have improved it to $\mathbf{E}\|(\mathbf{E}[\omega_k | \theta_k] - \omega_k^*)\|$. Such a seemingly innocuous change turns out to be crucial to improving the sample complexity.

Proof. (Outline) To get an idea of how the first order term is improved, note that we can write the following inequality using the smoothness property of the score function (Assumption 1) and the update rule (16).

$$\begin{aligned} \mathbf{E}_{s \sim d^{\pi^*}} [KL(\pi^*(\cdot|s) \| \pi_{\theta_k}(\cdot|s)) - KL(\pi^*(\cdot|s) \| \pi_{\theta_{k+1}}(\cdot|s))] \\ \leq \eta \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi^*}} [\nabla_{\theta} \log \pi_{\theta_k}(a|s) \cdot \omega_k^*] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ + \eta \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi^*}} [\nabla_{\theta} \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] \end{aligned} \quad (21)$$

The first term on the RHS of (21) can be connected to ϵ_{bias} via Assumption 2 and the performance difference lemma (Agarwal et al., 2021) stated as,

$$J_{\rho}^* - J_{\rho}(\theta_k) = \frac{1}{1-\gamma} \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi^*}} [A^{\pi_{\theta_k}}(s, a)]$$

Note that the second term on the RHS of (21) is connected to the second-order error in (20). Before bounding the third term, we take expectations on both sides of (21). Then a bound is obtained as follows.

$$\begin{aligned} &\left| \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi^*}} \mathbf{E} [\nabla_{\theta} \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] \right| \\ &= \left| \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi^*}} \mathbf{E} [\nabla_{\theta} \log \pi_{\theta_k}(a|s) \cdot (\mathbf{E}[\omega_k | \theta_k] - \omega_k^*)] \right| \\ &\stackrel{(a)}{\leq} G \times \mathbf{E}[\|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|] \end{aligned}$$

where (a) follows from Assumption 1. Rearranging the terms and taking a sum over $k \in \{0, 1, \dots, K-1\}$, we finally prove Lemma 3. \square

Observe that, if the estimate ω_k produced by the inner loop is a good approximation of ω_k^* , then the first order term in (20) will be small. To bound the second-order term, note the following inequality.

$$\begin{aligned} &\frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}\|\omega_k\|^2 \\ &\leq \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}\|\omega_k - \omega_k^*\|^2 + \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}\|\omega_k^*\|^2 \\ &\stackrel{(a)}{\leq} \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}\|\omega_k - \omega_k^*\|^2 + \frac{\mu_F^{-2}}{K} \sum_{k=0}^{K-1} \mathbf{E}\|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 \end{aligned} \quad (22)$$

where (a) follows from Assumption 3 and the definition that $\omega_k^* = F_{\rho}(\theta_k)^{\dagger} \nabla_{\theta} J_{\rho}(\theta_k)$. Notice that the first term in (22) will be small if ω_k well approximates ω_k^* , $\forall k$. To bound the second term, we use the following lemma.

Lemma 4. Let $\{\theta_k\}_{k=0}^{K-1}$ be recursively defined by (16). If Assumption 1 and 3 hold, then the following inequality is satisfied with learning rate $\eta = \frac{\mu_F^2}{4G^2L}$.

$$\begin{aligned} \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 &\leq \left(\frac{8G^4L}{\mu_F^2(1-\gamma)} \right) \frac{1}{K} \\ &+ (2G^4 + \mu_F^2) \left(\frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E} \|\omega_k - \omega_k^*\|^2 \right) \end{aligned} \quad (23)$$

Lemma 4 reveals that the second term in (22) can be bounded above by the approximation error of the inner loop. Combining Lemma 3, (22), and Lemma 4, we get the following result.

Corollary 1. Consider the setup described in Lemma 3. If assumptions 1-3 hold, then the following inequality is satisfied for $\eta = \frac{\mu_F^2}{4G^2L}$.

$$\begin{aligned} J_{\rho}^* - \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}[J_{\rho}(\theta_k)] \\ \leq \sqrt{\epsilon_{\text{bias}}} + \frac{G}{K} \sum_{k=0}^{K-1} \mathbf{E} \|(\mathbf{E}[\omega_k|\theta_k] - \omega_k^*)\| \\ + \frac{B}{4L} \left(\frac{\mu_F^2}{G^2} + G^2 \right) \left(\frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E} \|\omega_k - \omega_k^*\|^2 \right) \\ + \frac{G^2}{\mu_F^2 K} \left(\frac{B}{1-\gamma} + 4L \mathbf{E}_{s \sim d_{\pi^*}} [KL(\pi^*(\cdot|s) \|\pi_{\theta_0}(\cdot|s))] \right) \end{aligned} \quad (24)$$

Observe that both the first-order approximation error, $\mathbf{E} \|(\mathbf{E}[\omega_k|\theta] - \omega_k^*)\|$ and the second-order error, $\mathbf{E} \|\omega_k - \omega_k^*\|^2$ must be functions of H , the number of iterations of the inner loop. In section 4.2, we point out the exact nature of this dependence. Furthermore, the last term in (24) is a function of K , the length of the outer loop. Section 4.3 discusses how H and K can be judiciously chosen to achieve the desired sample complexity. Note that due to the presence of ϵ_{bias} in (24), the optimality gap of $J_{\rho}(\cdot)$ cannot be made arbitrarily small. This is an unavoidable error for any parameterized policy class that is incomplete i.e., does not contain all stochastic policies. However, for rich classes of policies, this error can be assumed to be insignificant.

4.2 Inner Loop Analysis

Recall that the inner loop produces estimates of $\omega_k^* = \arg \min_{\omega \in \mathbb{R}^d} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)|_{\theta=\theta_k}$ using stochastic gradients obtained from Algorithm 1. Lemma 1 showed that the gradient estimates are unbiased. The following lemma characterizes their variance.

Lemma 5. Let $\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$ denote the gradient estimate generated by Algorithm 1. If assumptions 1 and

3 hold, the following semidefinite inequality is satisfied for any $\theta \in \mathbb{R}^d$.

$$\mathbf{E} \left[\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega_{\theta}^*, \theta) \otimes \hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega_{\theta}^*, \theta) \right] \preceq \sigma^2 F_{\rho}(\theta) \quad (25)$$

where $\omega_{\theta}^* = \arg \min_{\omega \in \mathbb{R}^d} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega, \theta)$, $F_{\rho}(\theta)$ is given by (5), and σ^2 is defined by the following expression.

$$\sigma^2 \triangleq \frac{2G^4}{\mu_F^2(1-\gamma)^4} + \frac{32}{(1-\gamma)^4} \quad (26)$$

Note that the term σ^2 can be interpreted as the scaled variance of the gradient estimates. In a noiseless scenario (deterministic estimates), the LHS of (25) turns out to be zero which allows us to choose $\sigma^2 = 0$. This information will be useful in proving one of our subsequent results. Merging Lemma 5 with the convergence result of ASGD as stated in (Jain et al., 2018b, Corollary 2), we arrive at the following lemma.

Lemma 6. Let ω_k be an estimate of ω_k^* generated by Algorithm 2 at the k th iteration of the outer loop, $k \in \{0, 1, \dots, K-1\}$. If assumptions 1 and 3 hold, then the following inequality is satisfied with learning rates $\alpha = \frac{3\sqrt{5}G^2}{\mu_F + 3\sqrt{5}G^2}$, $\beta = \frac{\mu_F}{9G^2}$, $\xi = \frac{1}{3\sqrt{5}G^2}$, and $\delta = \frac{1}{5G^2}$ provided that the inner loop length obeys $H > \bar{C} \frac{G^2}{\mu_F} \log \left(\sqrt{d} \frac{G^2}{\mu_F} \right)$ for some universal constant, \bar{C} .

$$\begin{aligned} \mathbf{E} \|\omega_k - \omega_k^*\|^2 &\leq 22 \frac{\sigma^2 d}{\mu_F H} \\ &+ C \exp \left(-\frac{\mu_F}{20G^2} H \right) \left[\frac{1}{\mu_F(1-\gamma)^4} \right] \end{aligned} \quad (27)$$

where C denotes a universal constant and σ^2 is defined in (26).

Lemma 6 implies that the second-order inner loop approximation error can be bounded above as $\mathcal{O}(1/H)$. Interestingly, however, the leading factor $1/H$ appears due to the presence of noise variance, σ^2 . Hence, in a noiseless scenario ($\sigma^2 = 0$), exponentially fast convergence should be viable. This idea has been exploited in the following result.

Lemma 7. Consider the same setup and the choice of parameters as dictated in Lemma 6. If assumptions 1 and 3 hold, then we have the following inequality.

$$\begin{aligned} \mathbf{E} \|(\mathbf{E}[\omega_k|\theta_k] - \omega_k^*)\| \\ \leq \sqrt{C} \exp \left(-\frac{\mu_F}{40G^2} H \right) \left(\frac{1}{\sqrt{\mu_F(1-\gamma)^2}} \right) \end{aligned} \quad (28)$$

Lemma 7 dictates that the first-order inner loop approximation error exponentially decreases with H . To intuitively understand the proof of this result, consider the ASGD update rules (11)–(14) in Algorithm 2 that

are used to determine ω_k via (15), $\forall k \in \{0, \dots, K-1\}$. Note the following property of the gradient estimator produced by Algorithm 1.

$$\begin{aligned} & \mathbf{E} \left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta) \mid \theta \right] \\ &= \mathbf{E} \left[\mathbf{E} \left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta) \mid \omega, \theta \right] \mid \theta \right] \quad (29) \\ &\stackrel{(a)}{=} \mathbf{E} \left[\nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta) \mid \theta \right] \stackrel{(b)}{=} \nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\mathbf{E}[\omega \mid \theta], \theta) \end{aligned}$$

where (a) uses the unbiasedness (Lemma 1) of the estimate and (b) follows from the linearity of the gradient. Taking expectation conditioned on θ_k on both sides of (15), we obtain the following relation.

$$\mathbf{E}[\omega_k \mid \theta_k] = \frac{2}{H} \sum_{\frac{H}{2} < h \leq H} \mathbf{E}[\mathbf{x}_h \mid \theta_k] \quad (30)$$

where the conditional expectation of $\{\mathbf{x}_h\}_{h=1}^H$ are recursively computed as follows starting with the initial condition, $\bar{\mathbf{x}}_0 \triangleq \mathbf{E}[\mathbf{x}_0 \mid \theta_k] = \mathbf{0}$ and $\bar{\mathbf{v}}_0 \triangleq \mathbf{E}[\mathbf{v}_0 \mid \theta_k] = \mathbf{0}$.

$$\bar{\mathbf{y}}_h = \alpha \bar{\mathbf{x}}_h + (1 - \alpha) \bar{\mathbf{v}}_h \quad (31)$$

$$\bar{\mathbf{x}}_{h+1} = \bar{\mathbf{y}}_h - \delta \nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta_k) \Big|_{\omega=\bar{\mathbf{y}}_h} \quad (32)$$

$$\bar{\mathbf{z}}_h = \beta \bar{\mathbf{y}}_h + (1 - \beta) \bar{\mathbf{v}}_h \quad (33)$$

$$\bar{\mathbf{v}}_{h+1} = \bar{\mathbf{z}}_h - \xi \nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta_k) \Big|_{\omega=\bar{\mathbf{y}}_h} \quad (34)$$

where $\bar{\mathbf{x}}_h \triangleq \mathbf{E}[\mathbf{x}_h \mid \theta_k]$, $\bar{\mathbf{y}}_h \triangleq \mathbf{E}[\mathbf{y}_h \mid \theta_k]$, $\bar{\mathbf{v}}_h \triangleq \mathbf{E}[\mathbf{v}_h \mid \theta_k]$, $\bar{\mathbf{z}}_h \triangleq \mathbf{E}[\mathbf{z}_h \mid \theta_k]$ and $h \in \{0, 1, \dots, H\}$. The above equations can be derived by taking conditional expectation on both sides of (11)-(14) and invoking (29). Note that (30) together with (31)–(34) can be interpreted as the update rules of a noiseless (deterministic) accelerated gradient descent procedure. This allows us to derive a bound similar to (27) but with $\sigma^2 = 0$.

4.3 Final Result

Lemma 6 and 7 together dictate that the second and third terms in (24) can be bounded above as $\mathcal{O}(1/H)$. Additionally, the last term in (24) is $\mathcal{O}(1/K)$. Thus by taking $H = \mathcal{O}(1/\epsilon)$ and $K = \mathcal{O}(1/\epsilon)$, we guarantee the optimality gap to be at most $\sqrt{\epsilon_{\text{bias}}} + \epsilon$. This results in a sample complexity of $\mathcal{O}(HK) = \mathcal{O}(1/\epsilon^2)$ and an iteration complexity of $\mathcal{O}(K) = \mathcal{O}(1/\epsilon)$. The result is formalized in the following theorem.

Theorem 1. *Let $\{\theta_k\}_{k=0}^{K-1}$ be the policy parameters generated by Algorithm 2, π^* be the optimal policy and J_ρ^* denote the optimal value of $J_\rho(\cdot)$ corresponding to an initial distribution ρ . Assume that assumptions 1–3 hold and the learning parameters are chosen as stated in Corollary 1 and Lemma 6. For all sufficiently small*

ϵ , and $H = \mathcal{O}(1/\epsilon)$, $K = \mathcal{O}(1/\epsilon)^1$, the following holds.

$$J_\rho^* - \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}[J_\rho(\theta_k)] \leq \sqrt{\epsilon_{\text{bias}}} + \epsilon \quad (35)$$

This results in $\mathcal{O}((1-\gamma)^{-6}\epsilon^{-2})$ sample complexity and $\mathcal{O}((1-\gamma)^{-3}\epsilon^{-1})$ iteration complexity.

A few remarks are in order. Note the importance of the first-order approximation error, $\mathbf{E}\|(\mathbf{E}[\omega_k \mid \theta_k] - \omega_k^*)\|$, in establishing the sample complexity result. If we follow existing analysis such as that given in (Liu et al., 2020), then this term would turn out to be $\mathbf{E}\|\omega_k - \omega_k^*\|$ which is bounded above as $\mathcal{O}(H^{-\frac{1}{2}})$ following Lemma 6. This requires $H = \mathcal{O}(\epsilon^{-2})$ and $K = \mathcal{O}(\epsilon^{-1})$ to guarantee an optimality gap of $\sqrt{\epsilon_{\text{bias}}} + \epsilon$. Unfortunately, the sample complexity, in this case, worsens to $\mathcal{O}(\epsilon^{-3})$.

Secondly, note the importance of ASGD in our analysis. If we use SGD (instead of ASGD) in the inner loop and utilize its convergence result such as that provided in (Bach and Moulines, 2013), the first-order error still ends up being $\mathcal{O}(H^{-\frac{1}{2}})$. With an improved analysis of SGD (Jain et al., 2018a), one can make the first-order term an exponentially decreasing function of H . However, it requires the learner to have an exact knowledge of the occupancy measure, $\nu_\rho^{\pi_{\theta_k}}$ in order to choose an appropriate learning rate. This directly contrasts with the basic premise of reinforcement learning.

5 CONCLUSION

This paper proposes the Accelerated Natural Policy Gradient (ANPG) algorithm that uses an acceleration-based stochastic gradient descent procedure for finding natural gradients. By introducing an improved global convergence analysis, and utilizing a novel observation that a first-order term can be interpreted as the error generated by a noiseless accelerated gradient descent procedure, we establish a sample complexity of $\mathcal{O}(\epsilon^{-2})$ and an iteration complexity of $\mathcal{O}(\epsilon^{-1})$ for the ANPG. This improves the SOTA sample complexity by a factor of $\log\left(\frac{1}{\epsilon}\right)$. Within the class of all Hessian-free and importance sample-free algorithms, it beats the SOTA sample complexity by a $\mathcal{O}(\epsilon^{-\frac{1}{2}})$ factor and matches the best known iteration complexity.

This paper unleashes several promising directions for future research. For example, our analysis can be used to improve the sample complexities in constrained RL, concave utility RL, etc. Another important task would be improving the iteration complexity while maintaining the sample complexity reported in our paper.

¹The exact values of these parameters are stated in the proof of the theorem in the appendix.

References

- A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. Optimality and approximation with policy gradient methods in markov decision processes. In J. Abernethy and S. Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 64–66. PMLR, 09–12 Jul 2020.
- A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *The Journal of Machine Learning Research*, 22(1):4431–4506, 2021.
- A. O. Al-Abbasi, A. Ghosh, and V. Aggarwal. Deep-pool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20(12):4714–4727, 2019.
- M. G. Azar, I. Osband, and R. Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- F. Bach and E. Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate $o(1/n)$. *Advances in neural information processing systems*, 26, 2013.
- Q. Bai, A. S. Bedi, and V. Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via conservative natural policy gradient primal-dual algorithm. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6737–6744, 2023a.
- Q. Bai, W. U. Mondal, and V. Aggarwal. Regret analysis of policy gradient algorithm for infinite horizon average reward markov decision processes. *arXiv preprint arXiv:2309.01922*, 2023b.
- J. Bhandari and D. Russo. On the linear convergence of policy gradient methods for finite mdps. In *International Conference on Artificial Intelligence and Statistics*, pages 2386–2394. PMLR, 2021.
- S. Cen, C. Cheng, Y. Chen, Y. Wei, and Y. Chi. Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 70(4):2563–2578, 2022.
- Z. Chen and S. T. Maguluri. Sample complexity of policy-based methods under off-policy sampling and linear function approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 11195–11214. PMLR, 2022.
- Z. Chen, S. Khodadadian, and S. T. Maguluri. Finite-sample analysis of off-policy natural actor-critic with linear function approximation. *IEEE Control Systems Letters*, 6:2611–2616, 2022.
- Y. Ding, J. Zhang, and J. Lavaei. On the global optimum convergence of momentum-based policy gradient. In *International Conference on Artificial Intelligence and Statistics*, pages 1910–1934. PMLR, 2022.
- I. Fatkhullin, A. Barakat, A. Kireeva, and N. He. Stochastic policy gradient methods: Improved sample complexity for fisher-non-degenerate policies. *arXiv preprint arXiv:2302.01734*, 2023.
- M. Gargiani, A. Zanelli, A. Martinelli, T. Summers, and J. Lygeros. Page-pg: A simple and loopless variance-reduced policy gradient method with probabilistic gradient estimation. In *International Conference on Machine Learning*, pages 7223–7240. PMLR, 2022.
- F. Huang, S. Gao, J. Pei, and H. Huang. Momentum-based policy gradient methods. In *International conference on machine learning*, pages 4422–4433. PMLR, 2020.
- P. Jain, S. M. Kakade, R. Kidambi, P. Netrapalli, and A. Sidford. Parallelizing stochastic gradient descent for least squares regression: Mini-batching, averaging, and model misspecification. *Journal of Machine Learning Research*, 18(223):1–42, 2018a.
- P. Jain, S. M. Kakade, R. Kidambi, P. Netrapalli, and A. Sidford. Accelerating stochastic gradient descent for least squares regression. In *Conference On Learning Theory*, pages 545–604. PMLR, 2018b.
- S. Khodadadian, T. T. Doan, J. Romberg, and S. T. Maguluri. Finite sample analysis of two-time-scale natural actor-critic algorithm. *IEEE Transactions on Automatic Control*, 2022.
- G. Lan. Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *Mathematical programming*, 198(1):1059–1106, 2023.
- L. Ling, W. U. Mondal, and S. V. Ukkusuri. Cooperating graph neural networks with deep reinforcement learning for vaccine prioritization. *arXiv preprint arXiv:2305.05163*, 2023.
- Y. Liu, K. Zhang, T. Basar, and W. Yin. An improved analysis of (variance-reduced) policy gradient and natural policy gradient methods. *Advances in Neural Information Processing Systems*, 33:7624–7636, 2020.
- S. Masiha, S. Salehkaleybar, N. He, N. Kiyavash, and P. Thiran. Stochastic second-order methods improve best-known sample complexity of SGD for gradient-dominated functions. *Advances in Neural Information Processing Systems*, 35:10862–10875, 2022.

- J. Mei, C. Xiao, C. Szepesvari, and D. Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR, 2020.
- W. U. Mondal, V. Aggarwal, and S. Ukkusuri. Mean-field control based approximation of multi-agent reinforcement learning in presence of a non-decomposable shared global state. *Transactions on Machine Learning Research*, 2023.
- Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- M. Papini, D. Binaghi, G. Canonaco, M. Pirodda, and M. Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pages 4026–4035. PMLR, 2018.
- J. Peters and S. Schaal. Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190, 2008.
- N. Pham, L. Nguyen, D. Phan, P. H. Nguyen, M. Dijk, and Q. Tran-Dinh. A hybrid stochastic policy gradient algorithm for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 374–385. PMLR, 2020.
- S. Qiu, Z. Yang, J. Ye, and Z. Wang. On finite-time convergence of actor-critic algorithm. *IEEE Journal on Selected Areas in Information Theory*, 2(2):652–664, 2021.
- S. Salehkaleybar, S. Khorasani, N. Kiyavash, N. He, and P. Thiran. Adaptive momentum-based policy gradient with second-order information. *arXiv preprint arXiv:2205.08253*, 2022.
- J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- L. Shani, Y. Efroni, and S. Mannor. Adaptive trust region policy optimization: Global convergence and faster rates for regularized mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5668–5675, 2020.
- Z. Shen, A. Ribeiro, H. Hassani, H. Qian, and C. Mi. Hessian aided policy gradient. In *International conference on machine learning*, pages 5729–5738. PMLR, 2019.
- V. Singh, B. Nanavati, A. K. Kar, and A. Gupta. How to maximize clicks for display advertisement in digital marketing? a reinforcement learning approach. *Information Systems Frontiers*, 25(4):1621–1638, 2023.
- R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.
- L. Wang, Q. Cai, Z. Yang, and Z. Wang. Neural policy gradient methods: Global optimality and rates of convergence. In *International Conference on Learning Representations*, 2019.
- Y. F. Wu, W. Zhang, P. Xu, and Q. Gu. A finite-time analysis of two time-scale actor-critic methods. *Advances in Neural Information Processing Systems*, 33:17617–17628, 2020.
- L. Xiao. On the convergence rates of policy gradient methods. *The Journal of Machine Learning Research*, 23(1):12887–12922, 2022.
- P. Xu, F. Gao, and Q. Gu. Sample efficient policy gradient methods with recursive variance reduction. In *International Conference on Learning Representations*, 2019.
- P. Xu, F. Gao, and Q. Gu. An improved convergence analysis of stochastic variance-reduced policy gradient. In *Uncertainty in Artificial Intelligence*, pages 541–551. PMLR, 2020.
- H. Yuan, X. Lian, J. Liu, and Y. Zhou. Stochastic recursive momentum for policy gradient methods. *arXiv preprint arXiv:2003.04302*, 2020.
- R. Yuan, R. M. Gower, and A. Lazaric. A general sample complexity analysis of vanilla policy gradient. In *International Conference on Artificial Intelligence and Statistics*, pages 3332–3380. PMLR, 2022.
- W. Zhan, S. Cen, B. Huang, Y. Chen, J. D. Lee, and Y. Chi. Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *SIAM Journal on Optimization*, 33(2):1061–1091, 2023.
- J. Zhang, C. Ni, Z. Yu, C. Szepesvari, and M. Wang. On the convergence and sample efficiency of variance-reduced policy gradient method, 2021.
- K. Zhang, A. Koppel, H. Zhu, and T. Basar. Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612, 2020.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]

- (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
- (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A PROOF OF LEMMA 1

Note that the following relation holds $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$.

$$\Pr(\hat{s} = s, \hat{a} = a | \rho, \pi_\theta) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s, a_t = a | s_0 \sim \rho, \pi_\theta) = \nu_\rho^{\pi_\theta}(s, a) \quad (36)$$

Moreover, for a given pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, the following equality holds.

$$\begin{aligned} \mathbf{E} \left[\hat{Q}^{\pi_\theta}(s, a) \right] &= \frac{1}{2} \times 2\mathbf{E} \left[(1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \sum_{j=0}^t r(s_j, a_j) \middle| s_0 = s, a_0 = a, \pi_\theta \right] \\ &\stackrel{(a)}{=} \mathbf{E} \left[(1 - \gamma) \sum_{j=0}^{\infty} r(s_j, a_j) \sum_{t=j}^{\infty} \gamma^t \middle| s_0 = s, a_0 = a, \pi_\theta \right] \\ &= \mathbf{E} \left[\sum_{j=0}^{\infty} \gamma^j r(s_j, a_j) \middle| s_0 = s, a_0 = a, \pi_\theta \right] = Q^{\pi_\theta}(s, a) \end{aligned} \quad (37)$$

where (a) is obtained by exchanging the order of the sum. Similarly, one can exhibit that $\mathbf{E} \left[\hat{V}^{\pi_\theta}(s) \right] = V^{\pi_\theta}(s)$, and therefore $\mathbf{E} \left[\hat{A}^{\pi_\theta}(s, a) \right] = A^{\pi_\theta}(s, a)$, $\forall s \in \mathcal{S}$. Combining this result with (36), we establish the lemma.

B PROOF OF LEMMA 3

Using the definition of KL divergence, we obtain the following.

$$\begin{aligned} &\mathbf{E}_{s \sim d_\rho^{\pi^*}} [KL(\pi^*(\cdot|s) \| \pi_{\theta_k}(\cdot|s)) - KL(\pi^*(\cdot|s) \| \pi_{\theta_{k+1}}(\cdot|s))] \\ &= \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} \left[\log \frac{\pi_{\theta_{k+1}}(a|s)}{\pi_{\theta_k}(a|s)} \right] \\ &\stackrel{(a)}{\geq} \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\theta_{k+1} - \theta_k)] - \frac{B}{2} \|\theta_{k+1} - \theta_k\|^2 \\ &= \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot \omega_k] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ &= \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot \omega_k^*] + \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ &= \eta [J_\rho^* - J_\rho(\theta_k)] + \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot \omega_k^*] - \eta [J_\rho^* - J_\rho(\theta_k)] \\ &+ \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ &\stackrel{(b)}{=} \eta [J_\rho^* - J_\rho(\theta_k)] + \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} \left[\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot \omega_k^* - \frac{1}{1-\gamma} A^{\pi_{\theta_k}}(s, a) \right] \\ &+ \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ &\stackrel{(c)}{\geq} \eta [J_\rho^* - J_\rho(\theta_k)] - \eta \sqrt{\mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} \left[\left(\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot \omega_k^* - \frac{1}{1-\gamma} A^{\pi_{\theta_k}}(s, a) \right)^2 \right]} \\ &+ \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] - \frac{B\eta^2}{2} \|\omega_k\|^2 \\ &\stackrel{(d)}{\geq} \eta [J_\rho^* - J_\rho(\theta_k)] - \eta \sqrt{\epsilon_{\text{bias}}} + \eta \mathbf{E}_{s \sim d_\rho^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_\theta \log \pi_{\theta_k}(a|s) \cdot (\omega_k - \omega_k^*)] - \frac{B\eta^2}{2} \|\omega_k\|^2 \end{aligned} \quad (38)$$

where (a) and (b) are implied by Assumption 1(b) and Lemma 8 respectively. Inequality (c) follows from the convexity of the function $f(x) = x^2$. Finally, step (d) is a consequence of Assumption 2. Taking expectation on

both sides, we obtain,

$$\begin{aligned}
 & \mathbf{E} \left[\mathbf{E}_{s \sim d_{\rho}^*} [KL(\pi^*(\cdot|s) \|\pi_{\theta_k}(\cdot|s)) - KL(\pi^*(\cdot|s) \|\pi_{\theta_{k+1}}(\cdot|s))] \right] \\
 & \geq \eta [J_{\rho}^* - \mathbf{E} [J_{\rho}(\theta_k)]] - \eta \sqrt{\epsilon_{\text{bias}}} \\
 & \quad + \eta \mathbf{E} \left[\mathbf{E}_{s \sim d_{\rho}^*} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\nabla_{\theta} \log \pi_{\theta_k}(a|s) \cdot (\mathbf{E}[\omega_k | \theta_k] - \omega_k^*)] \right] - \frac{B\eta^2}{2} \mathbf{E} [\|\omega_k\|^2] \\
 & \geq \eta [J_{\rho}^* - \mathbf{E} [J_{\rho}(\theta_k)]] - \eta \sqrt{\epsilon_{\text{bias}}} \\
 & \quad - \eta \mathbf{E} \left[\mathbf{E}_{s \sim d_{\rho}^*} \mathbf{E}_{a \sim \pi^*(\cdot|s)} [\|\nabla_{\theta} \log \pi_{\theta_k}(a|s)\| \|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|] \right] - \frac{B\eta^2}{2} \mathbf{E} [\|\omega_k\|^2] \\
 & \stackrel{(a)}{\geq} \eta [J_{\rho}^* - \mathbf{E} [J_{\rho}(\theta_k)]] - \eta \sqrt{\epsilon_{\text{bias}}} - \eta G \mathbf{E} [\|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|] - \frac{B\eta^2}{2} \mathbf{E} [\|\omega_k\|^2]
 \end{aligned} \tag{39}$$

where (a) follows from Assumption 1(a). Rearranging the terms, we get,

$$\begin{aligned}
 J_{\rho}^* - \mathbf{E} [J_{\rho}(\theta_k)] & \leq \sqrt{\epsilon_{\text{bias}}} + G \mathbf{E} [\|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|] + \frac{B\eta}{2} \mathbf{E} [\|\omega_k\|^2] \\
 & \quad + \frac{1}{\eta} \mathbf{E} \left[\mathbf{E}_{s \sim d_{\rho}^*} [KL(\pi^*(\cdot|s) \|\pi_{\theta_k}(\cdot|s)) - KL(\pi^*(\cdot|s) \|\pi_{\theta_{k+1}}(\cdot|s))] \right]
 \end{aligned} \tag{40}$$

Summing from $k = 0$ to $K - 1$, using the non-negativity of KL divergence and dividing the resulting expression by K , we get the desired result.

C PROOF OF LEMMA 4

Using the L -smoothness property of the function, $J_{\rho}(\cdot)$ (Lemma 2), we obtain the following.

$$\begin{aligned}
 & J_{\rho}(\theta_{k+1}) \\
 & \geq J_{\rho}(\theta_k) + \langle \nabla_{\theta} J_{\rho}(\theta_k), \theta_{k+1} - \theta_k \rangle - \frac{L}{2} \|\theta_{k+1} - \theta_k\|^2 \\
 & = J_{\rho}(\theta_k) + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k \rangle - \frac{\eta^2 L}{2} \|\omega_k\|^2 \\
 & = J_{\rho}(\theta_k) + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k^* \rangle + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k - \omega_k^* \rangle - \frac{\eta^2 L}{2} \|\omega_k - \omega_k^* + \omega_k^*\|^2 \\
 & \stackrel{(a)}{\geq} J_{\rho}(\theta_k) + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), F_{\rho}(\theta_k)^{\dagger} \nabla_{\theta} J_{\rho}(\theta_k) \rangle + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k - \omega_k^* \rangle - \eta^2 L \|\omega_k - \omega_k^*\|^2 - \eta^2 L \|\omega_k^*\|^2 \\
 & \stackrel{(b)}{\geq} J_{\rho}(\theta_k) + \frac{\eta}{G^2} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 + \eta \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k - \omega_k^* \rangle - \eta^2 L \|\omega_k - \omega_k^*\|^2 - \eta^2 L \|\omega_k^*\|^2 \\
 & = J_{\rho}(\theta_k) + \frac{\eta}{2G^2} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 + \frac{\eta}{2G^2} [\|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 + 2G^2 \langle \nabla_{\theta} J_{\rho}(\theta_k), \omega_k - \omega_k^* \rangle + G^4 \|\omega_k - \omega_k^*\|^2] \\
 & \quad - \left(\frac{\eta G^2}{2} + \eta^2 L \right) \|\omega_k - \omega_k^*\|^2 - \eta^2 L \|\omega_k^*\|^2 \\
 & = J_{\rho}(\theta_k) + \frac{\eta}{2G^2} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 + \frac{\eta}{2G^2} \|\nabla_{\theta} J_{\rho}(\theta_k) + G^2(\omega_k - \omega_k^*)\|^2 - \left(\frac{\eta G^2}{2} + \eta^2 L \right) \|\omega_k - \omega_k^*\|^2 - \eta^2 L \|\omega_k^*\|^2 \\
 & \geq J_{\rho}(\theta_k) + \frac{\eta}{2G^2} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 - \left(\frac{\eta G^2}{2} + \eta^2 L \right) \|\omega_k - \omega_k^*\|^2 - \eta^2 L \|F_{\rho}(\theta_k)^{\dagger} \nabla_{\theta} J_{\rho}(\theta_k)\|^2 \\
 & \stackrel{(c)}{\geq} J_{\rho}(\theta_k) + \left(\frac{\eta}{2G^2} - \frac{\eta^2 L}{\mu_F^2} \right) \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 - \left(\frac{\eta G^2}{2} + \eta^2 L \right) \|\omega_k - \omega_k^*\|^2
 \end{aligned} \tag{41}$$

where (a) utilizes the Cauchy-Schwarz inequality and the definition that $\omega_k^* = F_{\rho}(\theta_k)^{\dagger} \nabla_{\theta} J_{\rho}(\theta_k)$. Inequalities (b), and (c) follow from Assumption 1(a) and 3 respectively. Choosing $\eta = \frac{\mu_F^2}{4G^2L}$, taking a sum over $k = 0, \dots, K - 1$,

and rearranging the resulting terms, we get the following.

$$\begin{aligned} \frac{\mu_F^2}{16G^4L} \left(\frac{1}{K} \sum_{k=0}^{K-1} \|\nabla_{\theta} J_{\rho}(\theta_k)\|^2 \right) &\leq \frac{J_{\rho}(\theta_K) - J_{\rho}(\theta_0)}{K} + \left(\frac{\mu_F^2}{8L} + \frac{\mu_F^4}{16G^4L} \right) \left(\frac{1}{K} \sum_{k=0}^{K-1} \|\omega_k - \omega_k^*\|^2 \right) \\ &\stackrel{(a)}{\leq} \frac{2}{(1-\gamma)K} + \left(\frac{\mu_F^2}{8L} + \frac{\mu_F^4}{16G^4L} \right) \left(\frac{1}{K} \sum_{k=0}^{K-1} \|\omega_k - \omega_k^*\|^2 \right) \end{aligned} \quad (42)$$

where (a) uses the fact that $J_{\rho}(\cdot)$ is absolutely bounded above by $1/(1-\gamma)$. This proves the desired result.

D PROOF OF LEMMA 5

Note that,

$$\begin{aligned} &\mathbf{E} \left[\hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega_{\theta}^*, \theta) \otimes \hat{\nabla}_{\omega} L_{\nu_{\rho}^{\pi_{\theta}}}(\omega_{\theta}^*, \theta) \right] \\ &= \mathbf{E}_{s \sim d_{\rho}^{\pi^*}} \mathbf{E}_{a \sim \pi^*(\cdot|s)} \left[\mathbf{E} \left[\underbrace{\nabla_{\theta} \log \pi_{\theta}(a|s) \cdot \omega_{\theta}^* - \frac{1}{1-\gamma} \hat{A}^{\pi_{\theta}}(s, a)}_{\triangleq \zeta_{\theta}(s, a)} \right]^2 \nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s) \right] \end{aligned} \quad (43)$$

Therefore, it is sufficient to show that $\mathbf{E}[\zeta_{\theta}(s, a)] \leq \sigma^2, \forall (s, a) \in \mathcal{S} \times \mathcal{A}$. Observe the following chain of inequalities,

$$\begin{aligned} &\mathbf{E} \left[\nabla_{\theta} \log \pi_{\theta}(a|s) \cdot \omega_{\theta}^* - \frac{1}{1-\gamma} \hat{A}^{\pi_{\theta}}(s, a) \right]^2 \\ &\leq 2 \left[\nabla_{\theta} \log \pi_{\theta}(a|s) \cdot \omega_{\theta}^* \right]^2 + \frac{2}{(1-\gamma)^2} \mathbf{E} \left[\hat{A}^{\pi_{\theta}}(s, a) \right]^2 \\ &\stackrel{(a)}{\leq} 2 \|\nabla_{\theta} \log \pi_{\theta}(a|s)\|^2 \|\omega_{\theta}^*\|^2 + \frac{4}{(1-\gamma)^2} \mathbf{E} \left[\hat{Q}^{\pi_{\theta}}(s, a) \right]^2 + \frac{4}{(1-\gamma)^2} \mathbf{E} \left[\hat{V}^{\pi_{\theta}}(s) \right]^2 \\ &\stackrel{(b)}{\leq} 2G^2 \|F_{\rho}(\theta)^{\dagger} \nabla_{\theta} J_{\rho}(\theta)\|^2 + \frac{32}{(1-\gamma)^4} \\ &\stackrel{(c)}{\leq} \frac{2G^4}{\mu_F^2(1-\gamma)^4} + \frac{32}{(1-\gamma)^4} \end{aligned} \quad (44)$$

Relation (a) is a consequence of Cauchy-Schwarz inequality and the fact that $(a-b)^2 \leq 2(a^2+b^2)$ for any two reals a, b . Inequality (b) uses Assumption 1, the definition that $\omega_{\theta}^* = F_{\rho}(\theta)^{\dagger} \nabla_{\theta} J_{\rho}(\theta)$, and the following two trivial bounds whereas (c) follows from Assumption 3 and Lemma 2.

$$\mathbf{E} \left[\hat{Q}^{\pi_{\theta}}(s, a) \right]^2 \leq \frac{4}{(1-\gamma)^2}, \text{ and } \mathbf{E} \left[\hat{V}^{\pi_{\theta}}(s) \right]^2 \leq \frac{4}{(1-\gamma)^2}, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (45)$$

To establish the first bound, note that $\hat{Q}^{\pi_{\theta}}(s, a)$ is assigned a zero value with probability $1/2$ and a positive value of at most $2(j+1)$ with probability $\frac{1}{2}(1-\gamma)\gamma^j$. Therefore,

$$\mathbf{E} \left[\hat{Q}^{\pi_{\theta}}(s, a) \right]^2 \leq \frac{1}{2} \sum_{j=0}^{\infty} 4(1-\gamma)(j+1)^2 \gamma^j = \frac{2(1+\gamma)}{(1-\gamma)^2} < \frac{4}{(1-\gamma)^2} \quad (46)$$

The second bound in (45) can be proven similarly. This concludes the lemma.

E PROOFS OF LEMMA 6 AND 7

We shall prove Lemma 6 and 7 using Corollary 2 of (Jain et al., 2018b). Note the following statements.

S1 : The following quantities exist and are finite $\forall \theta \in \mathbb{R}^d$.

$$F_{\rho}(\theta) \triangleq \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi_{\theta}}} \left[\nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s) \right], \quad (47)$$

$$G_{\rho}(\theta) \triangleq \mathbf{E}_{(s,a) \sim \nu_{\rho}^{\pi_{\theta}}} \left[\nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s) \otimes \nabla_{\theta} \log \pi_{\theta}(a|s) \right] \quad (48)$$

S2 : There exists $\sigma^2 > 0$ such that the following is satisfied $\forall \theta \in \mathbb{R}^d$ where ω_θ^* is a maximizer of $L_{\nu_\rho^{\pi_\theta}}(\cdot, \theta)$.

$$\mathbf{E} \left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega_\theta^*, \theta) \otimes \hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega_\theta^*, \theta) \right] \preceq \sigma^2 F_\rho(\theta) \quad (49)$$

S3 : There exists $\mu_F, G > 0$ such that the following statements hold $\forall \theta \in \mathbb{R}^d$.

$$(a) F_\rho(\theta) \succeq \mu_F I_d, \quad (50)$$

$$(b) \mathbf{E}_{(s,a) \sim \nu_\rho^{\pi_\theta}} \left[\|\nabla_\theta \log \pi_\theta(a|s)\|^2 \nabla_\theta \log \pi_\theta(a|s) \otimes \nabla_\theta \log \pi_\theta(a|s) \right] \preceq G^2 F_\rho(\theta), \quad (51)$$

$$(c) \mathbf{E}_{(s,a) \sim \nu_\rho^{\pi_\theta}} \left[\|\nabla_\theta \log \pi_\theta(a|s)\|_{F_\rho(\theta)^\dagger}^2 \nabla_\theta \log \pi_\theta(a|s) \otimes \nabla_\theta \log \pi_\theta(a|s) \right] \preceq \frac{G^2}{\mu_F} F_\rho(\theta) \quad (52)$$

Statement **S1** follows from Assumption 1(a) whereas **S2** is established in Lemma 5. Statement **S3(a)** is identical to Assumption 3, **S3(b)** follows from Assumption 1(a), and **S3(c)** is a result of Assumption 1(a) and 3. Therefore, we can apply Corollary 2 of (Jain et al., 2018b) with $\kappa = \tilde{\kappa} = G^2/\mu_F$ and obtain the convergence result stated below with $H > \bar{C}\sqrt{\kappa\tilde{\kappa}} \log(\sqrt{d}\sqrt{\kappa\tilde{\kappa}})$ iterations of the inner loop, and the following learning rates $\alpha = \frac{3\sqrt{5}\sqrt{\kappa\tilde{\kappa}}}{1+3\sqrt{5}\sqrt{\kappa\tilde{\kappa}}}$, $\beta = \frac{1}{9\sqrt{\kappa\tilde{\kappa}}}$, $\xi = \frac{1}{3\sqrt{5}\mu_F\sqrt{\kappa\tilde{\kappa}}}$, and $\delta = \frac{1}{5G^2}$.

$$\mathbf{E} [l_k(\omega_k)] - l_k(\omega_k^*) \leq \frac{C}{2} \exp\left(-\frac{H}{20\sqrt{\kappa\tilde{\kappa}}}\right) [l_k(\mathbf{0}) - l_k(\omega_k^*)] + 11 \frac{\sigma^2 d}{H}, \quad (53)$$

$$\text{where } l_k(\omega) \triangleq L_{\nu_\rho^{\pi_\theta}}(\omega, \theta)|_{\theta=\theta_k}, \quad \forall \omega \in \mathbb{R}^d$$

The term, C is a universal constant. Note that $l_k(\omega_k^*) \geq 0$ and $l_k(\mathbf{0})$ is bounded above as follows.

$$l_k(\mathbf{0}) = \frac{1}{2} \mathbf{E}_{(s,a) \sim \nu_\rho^{\pi_{\theta_k}}} \left[\frac{1}{1-\gamma} A^{\pi_{\theta_k}}(s, a) \right]^2 \stackrel{(a)}{\leq} \frac{1}{2(1-\gamma)^4} \quad (54)$$

where (a) follows from the fact that $|A^{\pi_\theta}(s, a)| \leq 1/(1-\gamma)$, $\forall (s, a)$, $\forall \theta$. Combining (53), (54), and the fact that $l_k(\cdot)$ is μ_F -strongly convex, we establish,

$$\mathbf{E} \|\omega_k - \omega_k^*\|^2 \leq \frac{2}{\mu_F} [\mathbf{E} [l_k(\omega_k)] - l_k(\omega_k^*)] \leq 22 \frac{\sigma^2 d}{\mu_F H} + C \exp\left(-\frac{\mu_F}{20G^2} H\right) \left[\frac{1}{\mu_F(1-\gamma)^4} \right] \quad (55)$$

This concludes Lemma 6. For noiseless ($\sigma^2 = 0$) gradient updates, we get the following.

$$\mathbf{E} \|\mathbf{E}[\omega_k | \theta_k] - \omega_k^*\|^2 \leq C \exp\left(-\frac{\mu_F}{20G^2} H\right) \left[\frac{1}{\mu_F(1-\gamma)^4} \right] \quad (56)$$

Lemma 7, therefore, can be established from (56) by applying Jensen's inequality on the function $f(x) = x^2$.

F PROOF OF THEOREM 1

Combining Lemma 6 and 7 with Corollary 1, we get the following bound for appropriate choices of the learning rates as mentioned in the theorem.

$$\begin{aligned} J_\rho^* - \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}[J_\rho(\theta_k)] &\leq \sqrt{\epsilon_{\text{bias}}} + G \left\{ \sqrt{C} \exp\left(-\frac{\mu_F}{40G^2} H\right) \left[\frac{1}{\sqrt{\mu_F(1-\gamma)^2}} \right] \right\} \\ &\quad + \frac{B}{4L} \left(\frac{\mu_F^2}{G^2} + G^2 \right) \left\{ C \exp\left(-\frac{\mu_F}{20G^2} H\right) \left[\frac{1}{\mu_F(1-\gamma)^4} \right] + 22 \frac{\sigma^2 d}{\mu_F H} \right\} \\ &\quad + \frac{G^2}{\mu_F^2 K} \left(\frac{B}{1-\gamma} + 4L \mathbf{E}_{s \sim d_\rho^*} [KL(\pi^*(\cdot|s) \| \pi_{\theta_0}(\cdot|s))] \right) \end{aligned} \quad (57)$$

Using the fact that $\exp(-x) < \frac{1}{x}$, and $\exp(-x) < 1, \forall x > 0$, we get the following result.

$$\begin{aligned}
 J_\rho^* - \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{E}[J_\rho(\theta_k)] &\leq \sqrt{\epsilon_{\text{bias}}} + \underbrace{\left[\frac{40G^3\sqrt{C}}{\mu_F^{3/2}(1-\gamma)^2} + \frac{B}{4L} \left(\frac{\mu_F^2}{G^2} + G^2 \right) \left\{ \frac{20G^2C}{\mu_F^2(1-\gamma)^4} + 22\frac{\sigma^2 d}{\mu_F} \right\} \right]}_{\triangleq P} \frac{1}{H} \\
 &\quad + \underbrace{\left\{ \frac{G^2}{\mu_F^2} \left(\frac{B}{1-\gamma} + 4L\mathbf{E}_{s \sim d_p^*} [KL(\pi^*(\cdot|s) \parallel \pi_{\theta_0}(\cdot|s))] \right) \right\}}_{\triangleq Q} \frac{1}{K}
 \end{aligned} \tag{58}$$

To get an optimality gap of $\sqrt{\epsilon_{\text{bias}}}$, we must choose $H = \frac{2P}{\epsilon}$ and $K = \frac{2Q}{\epsilon}$. Note that, $H > \bar{C} \frac{G^2}{\mu_F} \log \left(\sqrt{d} \frac{G^2}{\mu_F} \right) \triangleq H_0$ can be achieved if $\epsilon < \frac{2P}{H_0}$. In other words, ϵ must be sufficiently small for the theorem to be true. To determine the sample and iteration complexities as functions of γ , note the following observations.

- The terms B, G, C, μ_F, d , and $\mathbf{E}_{s \sim d_p^*} KL(\pi^*(\cdot|s) \parallel \pi_{\theta_0}(\cdot|s))$ are constants i.e., independent of γ .
- Lemma 5 demonstrates that $\sigma^2 = \mathcal{O}((1-\gamma)^{-4})$.
- According to Lemma 2, $L = \mathcal{O}((1-\gamma)^{-3})$ and $1/L = \mathcal{O}((1-\gamma)^3)$.

Thus, $H = \mathcal{O}((1-\gamma)^{-2}\epsilon^{-1})$ and $K = \mathcal{O}((1-\gamma)^{-3}\epsilon^{-1})$. Note that the expected number of steps executed by Algorithm 1 is $\mathcal{O}((1-\gamma)^{-1})$. Therefore, the sample complexity is $\mathcal{O}((1-\gamma)^{-1}HK) = \mathcal{O}((1-\gamma)^{-6}\epsilon^{-2})$. Moreover, the iteration complexity is $\mathcal{O}(K) = \mathcal{O}((1-\gamma)^{-3}\epsilon^{-1})$.

G AUXILIARY LEMMA

Lemma 8. (Agarwal et al., 2021, Lemma 2) *The following relation holds for any two policies π_1, π_2 .*

$$J_\rho^{\pi_1} - J_\rho^{\pi_2} = \frac{1}{1-\gamma} \mathbf{E}_{(s,a) \sim \nu_\rho^{\pi_1}} [A^{\pi_2}(s,a)] \tag{59}$$