arXiv:2405.10624v3 [cs.LG] 31 Oct 2024

Sample-Efficient Constrained Reinforcement Learning with General Parameterization

Washim Uddin Mondal
Department of Electrical Engineering
Indian Institute of Technology Kanpur
Kanpur, UP, India 208016
wmondal@iitk.ac.in

Vaneet Aggarwal
School of IE and ECE
Purdue University
West Lafayette, IN, USA 47906
vaneet@purdue.edu

Abstract

We consider a constrained Markov Decision Problem (CMDP) where the goal of an agent is to maximize the expected discounted sum of rewards over an infinite horizon while ensuring that the expected discounted sum of costs exceeds a certain threshold. Building on the idea of momentum-based acceleration, we develop the Primal-Dual Accelerated Natural Policy Gradient (PD-ANPG) algorithm that ensures an ϵ global optimality gap and ϵ constraint violation with $\tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$ sample complexity for general parameterized policies where γ denotes the discount factor. This improves the state-of-the-art sample complexity in general parameterized CMDPs by a factor of $\mathcal{O}((1-\gamma)^{-1}\epsilon^{-2})$ and achieves the theoretical lower bound in ϵ^{-1}.

1 Introduction

Reinforcement learning (RL) is a framework where an agent repeatedly interacts with an unknown Markovian environment to find a policy that maximizes the expected discounted sum of its observed rewards. Such problems, often modeled via Markov Decision Processes (MDPs), find applications in many areas, including transportation [1], communication networks [2], robotics [3], etc. In many applications, however, the agents must also obey certain constraints. For example, in a food delivery network, the orders must be delivered within a stipulated time window; the marketing decisions of a firm must satisfy its budget constraints, etc. Such constraints are incorporated into RL by introducing a cost function. In constrained MDPs (CMDPs), the agents not only maximize the expected sum of discounted rewards but also ensure that the expected sum of discounted costs does not cross a predefined boundary.

Finding an optimal policy for a CMDP is a challenging problem, especially when the environment, i.e., the state transition function, is unknown. The efficiency of a solution to a CMDP is measured by its sample complexity, which essentially states how many state-transition samples it takes to yield a policy that is ϵ close to the optimal one while also ensuring that the expected sum of discounted costs does not violate the imposed boundary by more than ϵ amount. Many articles in the literature solve the CMDP with an unknown environment. Most of these works, however, focus on the tabular case where the number of states is finite. These solutions cannot be applied to many real-life scenarios where the state space is either large or infinite. To tackle this issue, the concept of policy parameterization must be invoked. Unfortunately, as exhibited in Table 1, only a few works are available on CMDPs with parameterized policies. While, for softmax parameterization, the state-of-the-art (SOTA) sample complexity is $\mathcal{O}(\epsilon^{-2})$, the same for the general parameterization is $\tilde{\mathcal{O}}(\epsilon^{-4})$ which is far from the lower bound $\Omega(\epsilon^{-2})$. It should be noted that the number of parameters needed in softmax parameterization is $\mathcal{O}(SA)$ where S, A are the sizes of the state and action spaces of the underlying CMDP. On the other hand, general parameterization uses $\mathrm{d} \ll SA$ number of parame-

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

Algorithm	Sample Complexity	Parameterization
PMD-PD [9]	$\mathcal{O}(\epsilon^{-3})$	Softmax
PD-NAC [10]	$\mathcal{O}(\epsilon^{-6})$	Softmax
NPG-PD [5]	$\mathcal{O}((1-\gamma)^{-5}\epsilon^{-2})$	Softmax
CRPO [6]	$\mathcal{O}((1-\gamma)^{-7}\epsilon^{-4})$	Softmax
NPG-PD [5]	$\mathcal{O}((1-\gamma)^{-8}\epsilon^{-6})$	General
CRPO [6]	$\mathcal{O}((1-\gamma)^{-13}\epsilon^{-6})$	General
C-NPG-PDA [4]	$\tilde{\mathcal{O}}((1-\gamma)^{-8}\epsilon^{-4})$[1]	General
PD-ANPG (This Work)	$\tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$	General
Lower Bound [11]	$\Omega((1-\gamma)^{-5}\epsilon^{-2})$	–

Table 1: Summary of sample complexity results on CMDP with parameterized policies. The parameter γ indicates the discount factor. The dependence of the sample complexities of PMD-PD and PD-NAC on γ is not depicted in [9, 10].

ters, which makes it appropriate for dealing with large or infinite states. Given the importance of general parameterization for large state space CMDPs, the following question naturally arises: *"Is it possible to solve CMDPs with general parameterization and achieve a sample complexity better than the SOTA $\tilde{\mathcal{O}}(\epsilon^{-4})$ bound?"*

In this article, we provide an affirmative answer to the above question. We propose a Primal-Dual-based Accelerated Natural Policy Gradient (PD-ANPG) algorithm to solve γ-discounted CMDPs with general parameterization. We theoretically prove that PD-ANPG achieves ϵ optimality gap and ϵ constraint violation with $\tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$ sample complexity (Theorem 1) that improves the SOTA $\tilde{\mathcal{O}}((1-\gamma)^{-8}\epsilon^{-4})$ sample complexity result of [4]. It closes the gap between the theoretical upper and lower bounds of sample complexity in general parameterized CMDPs (in terms of ϵ^{-1}), which was an open problem for quite some time (see the results of [5], [6], [4] in Table 1).

1.1 Challenges and Key Insights

Our algorithm builds upon the idea of primal dual-based NPG [7, 4]. However, unlike the previous works, we use accelerated stochastic gradient descent (ASGD) in the inner loop to compute the estimate of the NPG. The improvement in the sample complexity results from two key observations. Firstly, we establish a global-to-local convergence lemma (Lemma 3), which dictates how the global convergence of the Lagrange function is related to the first and second-order estimation error of the NPG. Here, via careful analysis, we show that the first-order term can be written as the expected bias of the NPG estimator (i.e., the difference between the true NPG and the expectation of its estimate). Secondly, we show (Lemma 5 and its subsequent discussion) that the bias of the NPG estimate can be interpreted as the convergence error of an ASGD program with non-stochastic (i.e., deterministic) gradients. These, combined with the ASGD convergence result provided by [8], lead to a convergence result of the Lagrange function (Corollary 1).

Finally, Theorem 1 segregates the objective and constraint violation rates from Lagrange convergence. Corollary 1 shows that Lagrange convergence error is bounded by an independent function of ζ (the dual learning rate). One might, hence, be tempted to make ζ arbitrarily small. However, our analysis shows that although small ζ leads to better objective convergence, it worsens the constraint violation rate. We demonstrate how to optimally choose ζ to reach the middle ground which eventually leads us to $\tilde{\mathcal{O}}(\epsilon^{-2})$ sample complexity.

1.2 Related Works

Unconstrained RL: Many algorithms solve MDPs with exact gradients [12, 13, 14, 15, 16]. Moreover, many works use generative models to show either first-order [17, 18, 19, 20, 21] or global convergence [22, 23, 24, 25, 7, 26, 27].

[1] We would like to point out that the sample complexity of C-NPG-PDA reported in [4] is $\tilde{\mathcal{O}}((1-\gamma)^{-6}\epsilon^{-4})$. However, the result is erroneous. The authors have subsequently corrected their result, and the sample complexity has been modified to $\tilde{\mathcal{O}}((1-\gamma)^{-8}\epsilon^{-4})$ in the arXiv version (updated May 2024).

Constrained RL: The tabular setting is well investigated, and many model-based [28, 29, 30, 31] and model-free [6, 30, 32, 33] algorithms are available in the literature. In comparison, there are relatively fewer works on parameterized policies. Policy mirror descent-primal dual (PMD-PD) algorithm was proposed by [9] that achieves $\mathcal{O}(\epsilon^{-3})$ sample complexity for softmax policies. For the same parameterization, [10] achieved $\mathcal{O}(\epsilon^{-6})$ sample complexity via their proposed Online Primal-Dual Natural Actor-Critic Algorithm. The primal-dual Natural Policy Gradient algorithm suggested by [5] yields $\mathcal{O}(\epsilon^{-2})$ and $\mathcal{O}(\epsilon^{-6})$ sample complexities for softmax and general parametrization respectively. [6] also proposed a primal policy-based algorithm that works for both the softmax and general function approximation cases. The state of the art for the general parameterization is given by [4], where the $\tilde{\mathcal{O}}(\epsilon^{-4})$ sample complexity is obtained. This work improves upon this direction to obtain $\tilde{\mathcal{O}}(\epsilon^{-2})$ sample complexity. The comparisons are summarized in Table 1.

2 Formulation

Let us consider a constrained Markov Decision Process (CMDP) characterized by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, c, P, \gamma, \rho)$ where $\mathcal{S}, \mathcal{A}$ denote the state space and the action space respectively, $r : \mathcal{S} \times \mathcal{A} \to [0, 1]$ defines the reward function, $c : \mathcal{S} \times \mathcal{A} \to [-1, 1]$ is the cost function, $P : \mathcal{S} \times \mathcal{A} \to \Delta(\mathcal{S})$ indicates the state transition kernel (where $\Delta(\mathcal{S})$ is the collection of all probability distributions over $\mathcal{S}$), γ is the discount factor, and $\rho \in \Delta(\mathcal{S})$ is the initial state distribution. Note that the state space $\mathcal{S}$, in our setting, can potentially be a compact set of infinite size. However, for simplicity, we assume it to be countable. The action space, $\mathcal{A}$, is assumed to be of finite size. The range of the cost function is chosen to be $[-1, 1]$, rather than $[0, 1]$, to ensure that the constraint in our central optimization problem (defined later in (2)) is non-trivial. A policy, $\pi : \mathcal{S} \to \Delta(\mathcal{A})$ is defined as a distribution over the action space for a given state of the environment. For a given policy, π, and a state-action pair (s, a), we define the Q-value associated with $g \in \{r, c\}$ as follows.

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

where $\mathbf{E}_\pi$ is the expectation computed over all π-induced trajectories $\{(s_t, a_t)\}_{t=0}^{\infty}$ where $s_{t+1} \sim P(s_t, a_t)$ and $a_t \sim \pi(s_t), \forall t \in \{0, 1, \cdots\}$. Similarly, the V-value associated with policy π, state s, and $g \in \{r, c\}$ is defined below.

$$V_g^\pi(s) \triangleq \mathbf{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t g(s_t, a_t) \middle| s_0 = s \right] = \sum_a \pi(a|s) Q_g^\pi(s, a)$$

Below we define the advantage value for a policy π, a state-action pair (s, a), and $g \in \{r, c\}$.

$$A_g^\pi(s, a) \triangleq Q_g^\pi(s, a) - V_g^\pi(s)$$

Define a function $J_{g,\rho}^\pi, \forall g \in \{r, c\}$ as follows.

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

where $d_\rho^\pi \in \Delta(\mathcal{S})$ is the state occupancy measure given by,

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

Similarly, the state-action occupancy measure is defined as,

$$\nu_\rho^\pi(s, a) = d_\rho^\pi(s) \pi(a|s), \ \forall (s, a) \in \mathcal{S} \times \mathcal{A} \tag{1}$$

Our goal is to maximize the function $J_{r,\rho}^\pi$ over all policies π while ensuring that $J_{c,\rho}^\pi$ does not lie below a predefined threshold. Without loss of generality, we can formally express this problem as,

$$\max_\pi J_{r,\rho}^\pi \quad \text{subject to: } J_{c,\rho}^\pi \geq 0 \tag{2}$$

If the state space, $\mathcal{S}$, is large or infinite (which is the case in many application scenarios), the policies can no longer be represented in the tabular format; rather, they are indexed by a parameter, $\theta \in \Theta$. In

this paper, we assume $\Theta = \mathbb{R}^{\mathrm{d}}$. Such indexing can be done via, for example, neural networks (NNs). Let $J_{g,\rho}(\theta) \triangleq J_{g,\rho}^{\pi_\theta}$. This allows us to redefine the constrained optimization problem as follows.

$$\max_{\theta \in \Theta} J_{r,\rho}(\theta) \quad \text{subject to: } J_{c,\rho}(\theta) \geq 0 \tag{3}$$

We assume the existence of at least one interior point solution of the above optimization. This is also known as Slater condition which can be formally expressed as follows.

Assumption 1. *There exists $\bar{\theta}$ such that $J_{c,\rho}(\bar{\theta}) \geq c_{\text{slater}}$ for some $c_{\text{slater}} \in (0, \frac{1}{1-\gamma}]$.*

3 Algorithm

The dual problem associated with the constraint optimization (3) can be written as follows.

$$\min_{\lambda \geq 0} \max_{\theta \in \Theta} J_{\mathrm{L},\rho}(\theta, \lambda) \;\; \text{where } J_{\mathrm{L},\rho}(\theta, \lambda) \triangleq J_{r,\rho}(\theta) + \lambda J_{c,\rho}(\theta) \tag{4}$$

The function, $J_{\mathrm{L},\rho}(\cdot,\cdot)$ is called the Lagrange function while λ is said to be the Lagrange multiplier. The above problem can be solved by iteratively applying the following update rule $\forall k \in \{0, \cdots, K-1\}$, starting with (θ_0, λ_0) where $\theta_0 \in \Theta$ is arbitrary and $\lambda_0 = 0$.

$$\theta_{k+1} = \theta_k + \eta F_\rho(\theta_k)^\dagger \nabla_\theta J_{\mathrm{L},\rho}(\theta_k, \lambda_k) \tag{5}$$

$$\lambda_{k+1} = \mathcal{P}_\Lambda \left[\lambda_k - \zeta J_{c,\rho}(\theta_k)\right] \tag{6}$$

where η, ζ are learning rates, $\mathcal{P}_\Lambda$ denotes the projection function onto the set, $\Lambda \triangleq [0, \lambda_{\max}]$, and $\dagger$ is the Moore-Penrose pseudoinverse operator. The choice of $\lambda_{\max}$ will be specified later. Note that the update rule of θ is similar to that of the standard policy gradient method except here the learning rate, η is modulated by the inverse of the Fisher matrix, $F_\rho(\theta)$ which is defined below.

$$F_\rho(\theta) \triangleq \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}} \left[\nabla_\theta \log \pi_\theta(a|s) \otimes \nabla_\theta \log \pi_\theta(a|s)\right] \tag{7}$$

where $\otimes$ indicates the outer product. Using a variation of the classical policy gradient theorem [34], one can obtain the gradient of the Lagrange function as follows.

$$\begin{aligned} \nabla_\theta J_{\mathrm{L},\rho}(\theta, \lambda) = \frac{1}{1-\gamma} H_\rho(\theta, \lambda), \text{ where } H_\rho(\theta, \lambda) \triangleq \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}} \left[A_{\mathrm{L},\lambda}^{\pi_\theta}(s,a) \nabla_\theta \log \pi_\theta(a|s)\right] \\ \text{and } A_{\mathrm{L},\lambda}^{\pi_\theta}(s,a) \triangleq A_r^{\pi_\theta}(s,a) + \lambda A_c^{\pi_\theta}(s,a) \end{aligned} \tag{8}$$

In most application scenarios, the learner is unaware of the state transition function, P, and thereby, of the advantage function, $A_{\mathrm{L},\lambda}^{\pi_\theta}$ and the occupancy measure, $\nu_\rho^{\pi_\theta}$. This makes the exact computation of $F_\rho(\theta)$ and $H_\rho(\theta, \lambda)$ an impossible task. Fortunately, there is a way to obtain an approximate value of the *natural policy gradient* $\omega_{\theta,\lambda}^* \triangleq F_\rho(\theta)^\dagger \nabla_\theta J_{\mathrm{L},\rho}(\theta, \lambda)$ that does not require the knowledge of P. Invoking (8), one can prove that $\omega_{\theta,\lambda}^*$ is a solution of a quadratic optimization. Formally, we have,

$$\begin{aligned} \omega_{\theta,\lambda}^* \in \arg\min_{\omega \in \mathbb{R}^{\mathrm{d}}} L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda), \\ \text{where } L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda) \triangleq \frac{1}{2} \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}} \left[\left(\frac{1}{1-\gamma} A_{\mathrm{L},\lambda}^{\pi_\theta}(s,a) - \omega^{\mathrm{T}} \nabla_\theta \log \pi_\theta(a|s)\right)^2\right] \end{aligned} \tag{9}$$

The above reformulation opens up the possibility to compute $\omega_{\theta,\lambda}^*$ via a gradient descent-type iterative procedure. Observe that the gradient of $L_{\nu_\rho^{\pi_\theta}}(\cdot, \theta, \lambda)$ can be calculated as follows.

$$\nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda) = F_\rho(\theta)\omega - \frac{1}{1-\gamma} H_\rho(\theta, \lambda) \tag{10}$$

Algorithm 1 describes a procedure to obtain unbiased estimates of this gradient. This is inspired by Algorithm 3 of [12]. Additionally, observe from (6) that the update of the Lagrange variable, λ requires the computation of $J_{c,\rho}(\theta)$ which is also difficult to accomplish without having an explicit knowledge about P. Algorithm 1 also provides an unbiased estimation of the above quantity.

Algorithm 1 first samples a horizon length, T, from the geometric distribution with success probability $(1-\gamma)$ and executes the CMDP for T instances following the policy, π_θ, starting from a state

Algorithm 1 Unbiased Sampling

Input: Parameters $(\theta, \omega, \lambda, \gamma)$, Initial Distribution ρ

$T \sim \mathrm{Geo}(1-\gamma)$, $s_0 \sim \rho$, $a_0 \sim \pi_\theta(s_0)$
for $j \in \{0, \cdots, T-1\}$ **do**
    Execute a_j, observe $s_{j+1} \sim P(s_j, a_j)$ and sample $a_{j+1} \sim \pi_\theta(s_{j+1})$
$\hat{J}_{c,\rho}(\theta) \leftarrow \sum_{j=0}^{T} c(s_j, a_j)$, $(\hat{s}, \hat{a}) \leftarrow (s_T, a_T)$

$T \sim \mathrm{Geo}(1-\gamma)$, $(s_0, a_0) \leftarrow (\hat{s}, \hat{a})$ ▷ Q-function Estimation
for $j \in \{0, \cdots, T-1\}$ **do**
    Execute a_j, observe $s_{j+1} \sim P(s_j, a_j)$, and sample $a_{j+1} \sim \pi_\theta(s_{j+1})$
for $g \in \{r, c\}$ **do**
    $\hat{Q}_g^{\pi_\theta}(\hat{s}, \hat{a}) \leftarrow \sum_{j=0}^{T} g(s_j, a_j)$

$T \sim \mathrm{Geo}(1-\gamma)$, $s_0 \leftarrow \hat{s}$, $a_0 \sim \pi_\theta(s_0)$ ▷ V-function Estimation
for $j \in \{0, \cdots, T-1\}$ **do**
    Execute a_j, observe $s_{j+1} \sim P(s_j, a_j)$, and sample $a_{j+1} \sim \pi_\theta(s_{j+1})$
for $g \in \{r, c\}$ **do**
    $\hat{V}_g^{\pi_\theta}(\hat{s}) \leftarrow \sum_{j=0}^{T} g(s_j, a_j)$
    $\hat{A}_g^{\pi_\theta}(\hat{s}, \hat{a}) \leftarrow \hat{Q}_g^{\pi_\theta}(\hat{s}, \hat{a}) - \hat{V}_g^{\pi_\theta}(\hat{s})$

▷ Estimation of Relevant functions

$$\hat{A}_{\mathrm{L},\lambda}^{\pi_\theta}(\hat{s}, \hat{a}) \leftarrow \hat{A}_r^{\pi_\theta}(\hat{s}, \hat{a}) + \lambda \hat{A}_c^{\pi_\theta}(\hat{s}, \hat{a}) \tag{11}$$

$$\hat{F}_\rho(\theta) \leftarrow \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s}) \otimes \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s}) \tag{12}$$

$$\hat{H}_\rho(\theta, \lambda) \leftarrow \hat{A}_{\mathrm{L},\lambda}^{\pi_\theta}(\hat{s}, \hat{a}) \nabla_\theta \log \pi_\theta(\hat{a}|\hat{s}) \tag{13}$$

▷ Gradient Estimate

$$\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda) \leftarrow \hat{F}_\rho(\theta)\omega - \frac{1}{1-\gamma}\hat{H}_\rho(\theta, \lambda) \tag{14}$$

Output: $\hat{J}_{c,\rho}(\theta)$, $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda)$

$s_0 \sim \rho$. The total cost observed in the resulting trajectory is assigned as the estimate $\hat{J}_{c,\rho}(\theta)$. The state-action pair (s_T, a_T) can be assumed to be an arbitrary sample $(\hat{s}, \hat{a})$ chosen from the occupancy measure $\nu_\rho^{\pi_\theta}$. The algorithm then generates a π_θ-induced trajectory of length $T \sim \mathrm{Geo}(1-\gamma)$, taking $(\hat{s}, \hat{a})$ as the starting point. The total reward and cost observed in this trajectory are assigned as $\hat{Q}_r^{\pi_\theta}(\hat{s}, \hat{a})$ and $\hat{Q}_c^{\pi_\theta}(\hat{s}, \hat{a})$ respectively. Next, another π_θ-induced trajectory of length $T \sim \mathrm{Geo}(1-\gamma)$ is generated assuming the state, $\hat{s}$ as the initiation point. The total reward and cost of this trajectory are assigned as $\hat{V}_r^{\pi_\theta}(\hat{s})$ and $\hat{V}_c^{\pi_\theta}(\hat{s})$ respectively. For $g \in \{r, c\}$, an estimate of the advantage value is computed as $\hat{A}_g^{\pi_\theta}(\hat{s}, \hat{a}) = \hat{Q}_g^{\pi_\theta}(\hat{s}, \hat{a}) - \hat{V}_g^{\pi_\theta}(\hat{s})$. Finally, the estimates of $F_\rho(\theta)$ and $H_\rho(\theta, \lambda)$ are obtained via (12) and (13) respectively which produces an estimation of the desired gradient in (14). The following Lemma demonstrates that the estimates produced by Algorithm 1 are unbiased.

Lemma 1. *Let $\hat{J}_{c,\rho}(\theta)$, $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda)$ be the estimates produced by Algorithm 1 for a predefined set of parameters $(\omega, \theta, \lambda)$. The following equations hold.*

$$\mathbf{E}\left[\hat{J}_{c,\rho}(\theta)\middle|\theta\right] = J_{c,\rho}(\theta) \;\; \textit{and} \;\; \mathbf{E}\left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda)\middle|\omega, \theta, \lambda\right] = \nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda)$$

In the absence of knowledge about the transition model, P, one can utilize the estimates generated by Algorithm 1 as good proxies for their true values. In particular, one can obtain an approximate value of the natural policy gradient $\omega^*_{\theta,\lambda}$ by iteratively minimizing the function $L_{\nu_\rho^{\pi_\theta}}(\cdot, \theta, \lambda)$ using the gradient estimate $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega, \theta, \lambda)$. On the other hand, using the estimate $\hat{J}_{c,\rho}(\theta)$, an approximate update equation of the Lagrange parameter can be formed. Algorithm 2 uses these two ideas to obtain a policy that is close to the optimal one.

Algorithm 2 Primal-Dual Accelerated Natural Policy Gradient (PD-ANPG)

Input: Parameters (θ_0, λ_0), Distribution ρ, Run-time Parameters K, H, Learning Parameters $\eta, \zeta, \alpha, \beta, \xi, \delta$

for $k \in \{0, \cdots, K-1\}$ **do** ▷ Outer Loop

  $\mathbf{x}_0, \mathbf{v}_0 \leftarrow \mathbf{0}$

  for $h \in \{0, \cdots, H-1\}$ **do** ▷ Inner Loop

    ▷ Accelerated Stochastic Gradient Descent

$$\mathbf{y}_h \leftarrow \alpha \mathbf{x}_h + (1-\alpha)\mathbf{v}_h \tag{15}$$

$$\hat{G} \leftarrow \hat{\nabla}_\omega L_{\nu_\rho^{\pi_{\theta_k}}}(\omega, \theta_k, \lambda_k)\big|_{\omega=\mathbf{y}_h} \text{ (Algorithm 1)}$$

$$\mathbf{x}_{h+1} \leftarrow \mathbf{y}_h - \delta \hat{G} \tag{16}$$

$$\mathbf{z}_h \leftarrow \beta \mathbf{y}_h + (1-\beta)\mathbf{v}_h \tag{17}$$

$$\mathbf{v}_{h+1} \leftarrow \mathbf{z}_h - \xi \hat{G} \tag{18}$$

  Tail Averaging:

$$\omega_k \leftarrow \frac{2}{H} \sum_{\frac{H}{2} < h \leq H} \mathbf{x}_h \tag{19}$$

  Obtain $\hat{J}_{c,\rho}(\theta_k)$ via Algorithm 1.

  Parameter Updates:

$$\theta_{k+1} \leftarrow \theta_k + \eta \omega_k \tag{20}$$

$$\lambda_{k+1} \leftarrow \mathcal{P}_\Lambda[\lambda_k - \zeta \hat{J}_{c,\rho}(\theta_k)] \tag{21}$$

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

Algorithm 2 has a nested loop structure. The *outer loop* runs K number of times. At a given instance, k, of the outer loop, the policy parameter θ_k and the Lagrange parameter, λ_k are updated via (20) and (21). The estimate, $\hat{J}_{c,\rho}(\theta_k)$ is computed via Algorithm 1. On the other hand, ω_k, the approximate value of the natural policy gradient $\omega^*_{\theta_k,\lambda_k}$ is obtained by iteratively minimizing $L_{\nu_\rho^{\pi_{\theta_k}}}(\cdot, \theta_k, \lambda_k)$ in H number of *inner loop* steps via the Accelerated Stochastic Gradient Descent (ASGD) procedure as stated in [8]. ASGD comprises the iterative updates (15)−(18) with tunable learning parameters $(\alpha, \beta, \xi, \delta)$ followed by a tail-averaging step (19). The gradient estimate utilized in (16) and (18) is obtained via Algorithm 1. It is worth mentioning that existing NPG algorithms such as that given in [7] typically apply the SGD, rather than the ASGD procedure, to obtain ω_k. The difference between these subroutines is that while SGD uses only the current gradient estimate to update ω_k, ASGD considers the contribution of all previous gradient estimates (momentum) using its convoluted iteration and tail-averaging steps.

4 Analysis

Our goal in this section is to characterize the rate of convergence of the objective function and the constraint violation if policy parameters are generated via Algorithm 2. We start by stating a few assumptions needed for the analysis.

Assumption 2. *The log-likelihood function is G-Lipschitz and B-smooth where $B, G > 0$. Formally, the following relations hold $\forall \theta, \theta_1, \theta_2 \in \Theta$, and $\forall (s,a) \in \mathcal{S} \times \mathcal{A}$.*

$$\|\nabla_\theta \log \pi_\theta(a|s)\| \leq G, \text{ and } \|\nabla_\theta \log \pi_{\theta_1}(a|s) - \nabla_\theta \log \pi_{\theta_2}(a|s)\| \leq B\|\theta_1 - \theta_2\|$$

Remark 1. *Assumption 2 is commonly applied in proving convergence guarantees of policy gradient-type algorithms [35, 12, 7]. This assumption is obeyed by many widely used policy classes such as the class of neural networks with bounded weights.*

Assumption 2 implies the boundedness of the gradient of the Lagrange function. This can be formally expressed as follows.

Lemma 2. *If Assumption 2 holds, then the following inequality is true* $\forall\theta \in \Theta$ *and* $\forall\lambda \in \Lambda$.

$$\|\nabla_\theta J_{\mathrm{L},\rho}(\theta,\lambda)\| \leq \frac{G(1+\lambda_{\max})}{(1-\gamma)^2}$$

Proof. Statement (a) can be proven using (8) along with Assumption 2 and the facts that $|A^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)|$ is bounded by $(1+\lambda)/(1-\gamma)$ and $\lambda \leq \lambda_{\max}$, $\forall\lambda \in \Lambda$. □

The result established by Lemma 2 will be pivotal in our further analysis.

Assumption 3. *The compatible function approximation error defined in* (9) *satisfies the inequality* $L_{\nu_\rho^{\pi^*}}(\omega^*_{\theta,\lambda},\theta,\lambda) \leq \epsilon_{\mathrm{bias}}/2$, $\forall\theta \in \Theta$ *and* $\forall\lambda \in \Lambda$ *where* π^* *is a solution to the original constrained optimization problem* (2) *and* $\omega^*_{\theta,\lambda}$ *is defined in* (9)*. The term* ϵ_{bias} *is a non-negative constant. The factor* 2 *is used for notational convenience.*

Remark 2. *The term* ϵ_{bias} *quantifies the expressivity of the parameterized policy class. For example, if the parameterization is complete i.e., includes all possible policies (such as in direct or softmax parameterization), then* $\epsilon_{\mathrm{bias}} = 0$ *[12]. A similar result can be proven for linear MDPs [36]. For incomplete policy classes, we have* $\epsilon_{\mathrm{bias}} > 0$*. However, if the class is sufficiently rich (such as neural networks with a large number of parameters),* ϵ_{bias} *can be assumed to be negligibly small [37].*

Assumption 4. *There exists a positive constant* μ_F *such that* $F_\rho(\theta) - \mu_F I_{\mathrm{d}}$ *is positive semidefinite i.e.,* $F_\rho(\theta) \succeq \mu_F I_{\mathrm{d}}$, $\forall\theta \in \Theta$ *where* I_d *is a* $\mathrm{d} \times \mathrm{d}$ *identity matrix and* $F_\rho(\cdot)$ *is defined in* (7)*.*

The property of the policy classes laid out in Assumption 4 is called Fisher Non-Degeneracy (FND) which essentially ensures that $\forall\theta \in \Theta$, the Fisher matrix $F_\rho(\theta)$ is away from the zero matrix by a certain amount. Observe that the Hessian of the function, $l_{\theta,\lambda}(\cdot) \triangleq L_{\nu_\rho^{\pi_\theta}}(\cdot,\theta,\lambda)$ is $F_\rho(\theta)$. Therefore, Assumption 4 also indicates that $l_{\theta,\lambda}$ is μ_F-strongly convex. This assumption is commonly applied in analyzing policy-gradient algorithms [4, 38, 7]. Assumption 4 also ensures that the matrix $F_\rho(\theta)$ is invertible, which, in turn, implies the uniqueness of the maximizer $\omega^*_{\theta,\lambda} = \arg\min_{\omega\in\mathbb{R}^{\mathrm{d}}} l_{\theta,\lambda}(\omega)$. [39] describes a concrete set of policies that obeys Assumption $2-4$.

4.1 Local-to-Global Convergence Lemma

Recall that our goal is to establish the global convergence rates. Lemma 3 (stated below) is the first step in that direction. Specifically, it demonstrates how the average optimality gap of the Lagrange function can be bounded by the first and second-order error of the gradient estimates.

Lemma 3. *If the parameters* $\{\theta_k,\lambda_k\}_{k=0}^{K-1}$ *are updated via* (20) *and* (21) *and assumptions* $2-4$ *hold, then the following inequality holds for any* K.

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big(J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)\Big) \leq \sqrt{\epsilon_{\mathrm{bias}}} + \frac{G}{K}\sum_{k=0}^{K-1}\mathbf{E}\|(\mathbf{E}\left[\omega_k\big|\theta_k,\lambda_k\right] - \omega_k^*)\| \\ + \frac{B\eta}{2K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k\|^2 + \frac{1}{\eta K}\mathbf{E}_{s\sim d_\rho^{\pi^*}}[KL(\pi^*(\cdot|s)\|\pi_{\theta_0}(\cdot|s))]\end{aligned} \tag{22}$$

where $\omega_k^* \triangleq \omega^*_{\theta_k,\lambda_k}$, $\omega^*_{\theta_k,\lambda_k}$ *is the natural policy gradient defined in* (9)*, and* π^* *is the solution to the constrained optimization* (2)*. Finally,* ω_k *is the approximation of* ω_k^* *given by* (19) *and* $KL(\cdot\|\cdot)$ *is the KL-divergence.*

Note the presence of the term, ϵ_{bias} in (22). It shows that due to the incompleteness of the parameterized policy class, the average optimality error cannot be made arbitrarily small. It is worth mentioning that many existing CMDP analyses (such as [4]) follow a path similar to that of Lemma 3. However, while the first order term in those works turns out to be $\mathbf{E}\|\omega_k - \omega_k^*\|$, we improved it to $\mathbf{E}\|\mathbf{E}[\omega_k|\theta_k,\lambda_k] - \omega_k^*\|$. Such seemingly insignificant improvement has important ramifications for

our analysis as explained later in the paper. The second order term in (22) can be expanded as,

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k\|^2 &\leq \frac{2}{K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k-\omega_k^*\|^2+\frac{2}{K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k^*\|^2\\ &\overset{(a)}{\leq} \frac{2}{K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k-\omega_k^*\|^2+\frac{2}{\mu_F^2 K}\sum_{k=0}^{K-1}\mathbf{E}\left\|\nabla_\theta J_{\mathrm{L},\rho}(\theta_k,\lambda_k)\right\|^2\\ &\overset{(b)}{\leq} \frac{2}{K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k-\omega_k^*\|^2+\frac{2G^2(1+\lambda_{\max})^2}{\mu_F^2(1-\gamma)^4}\end{aligned} \tag{23}$$

where (a) utilises $\omega_k^* = F_\rho(\theta_k)^\dagger \nabla_\theta J_{\mathrm{L},\rho}(\theta_k,\lambda_k)$ and Assumption 4. The second inequality applies Lemma 2 together with the fact that $\lambda_k \in \Lambda$. Our next subsection provides a bound on $\mathbf{E}\|\omega_k-\omega_k^*\|^2$ and the first order term $\mathbf{E}\|\mathbf{E}[\omega_k|\theta_k,\lambda_k]-\omega_k^*\|$.

4.2 Local Convergence of the Natural Policy Gradient

To deliver the promised bounds, we first provide some characterization of the gradient estimate.

Lemma 4. *Let $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega,\theta,\lambda)$ be the estimate produced by Algorithm 1. Under assumptions 2 and 4, the following semidefinite inequality holds for any $\theta\in\Theta$ and $\lambda\in\Lambda$.*

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

where $\omega_{\theta,\lambda}^*$, $F_\rho(\theta)$ are given by (9) *and* (7) *respectively and σ^2 is defined below.*

$$\sigma^2 \triangleq \frac{1}{(1-\gamma)^4}\left[\frac{2G^4}{\mu_F^2}+32\right](1+\lambda_{\max})^2 \tag{24}$$

The term σ^2 defined in Lemma 4 can be described as the scaled variance of the gradient estimate, $\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega,\theta,\lambda)$. Note that if the estimates were non-stochastic (i.e., deterministic), we would have $\sigma^2=0$ since $\nabla_\omega L_{\nu_\rho^{\pi_\theta}}(\omega_{\theta,\lambda}^*,\theta,\lambda)=0$. The last equation can be proved using the definition of $\omega_{\theta,\lambda}^*$ given in (9), the gradient expression provided in (10), and observing that the Fisher matrix, $F_\rho(\theta)$ is invertible due to Assumption 4. The above information is crucial in bounding the first-order error, as stated in the following lemma.

Lemma 5. *If assumptions 2 and 4 hold, then the following relations are satisfied $\forall k\in\{0,\cdots,K-1\}$ 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 of Algorithm 2 obeys $H>\bar{C}\frac{G^2}{\mu_F}\log\left(\sqrt{\mathrm{d}}\frac{G^2}{\mu_F}\right)$ for some universal constant, $\bar{C}$.*

$$\mathbf{E}\|\omega_k-\omega_k^*\|^2\leq 22\frac{\sigma^2\mathrm{d}}{\mu_F H}+C\exp\left(-\frac{\mu_F}{20G^2}H\right)\left[\frac{(1+\lambda_{\max})^2}{\mu_F(1-\gamma)^4}\right], \tag{25}$$

$$\mathbf{E}\|\mathbf{E}\left[\omega_k\left|\theta_k,\lambda_k\right.\right]-\omega_k^*\|\leq\sqrt{C}\exp\left(-\frac{\mu_F}{40G^2}H\right)\left[\frac{1+\lambda_{\max}}{\sqrt{\mu_F}(1-\gamma)^2}\right] \tag{26}$$

where C denotes a universal constant, ω_k is given by (19)*, and σ^2 is defined in* (24)*.*

The first bound (25) is a consequence of Lemma 4 and the ASGD convergence result provided in [8] (Corollary 2). To gain intuition about the second result (26), note that by taking the conditional expectation $\mathbf{E}[\cdot|\theta_k,\lambda_k]$ on both sides of the ASGD iterations (15)−(18), and applying the unbiasedness of the gradient estimate (Lemma 1) we obtain the following $\forall h\in\{0,\cdots,H-1\}$.

$$\begin{aligned}\bar{\mathbf{y}}_h &= \alpha\bar{\mathbf{x}}_h+(1-\alpha)\bar{\mathbf{v}}_h,\\ \bar{\mathbf{x}}_{h+1} &= \bar{\mathbf{y}}_h-\delta\nabla_\omega L_{\nu_\rho^{\pi_{\theta_k}}}(\omega,\theta_k,\lambda_k)\big|_{\omega=\bar{\mathbf{y}}_h}\\ \bar{\mathbf{z}}_h &= \beta\bar{\mathbf{y}}_h+(1-\beta)\bar{\mathbf{v}}_h,\\ \bar{\mathbf{v}}_{h+1} &= \bar{\mathbf{z}}_h-\xi\nabla_\omega L_{\nu_\rho^{\pi_{\theta_k}}}(\omega,\theta_k,\lambda_k)\big|_{\omega=\bar{\mathbf{y}}_h}\end{aligned} \tag{27}$$

where $\bar{l}_h = \mathbf{E}[l_h|\theta_k, \lambda_k]$, $l \in \{\mathbf{v}, \mathbf{x}, \mathbf{y}, \mathbf{z}\}$. Moreover, taking conditional expectation on both sides of the tail averaging process (19), we arrive at the following.

$$\bar{\omega}_k \triangleq \mathbf{E}\left[\omega_k \middle| \theta_k, \lambda_k\right] = \frac{2}{H} \sum_{\frac{H}{2} < h \leq H} \bar{\mathbf{x}}_h \tag{28}$$

Note that the steps (27)−(28) resemble the iterative updates of a deterministic ASGD. This allows us to obtain $\mathbf{E}\|\bar{\omega}_k - \omega_k^*\|$ by substituting $\sigma^2 = 0$ in (25) and applying the Cauchy-Schwarz inequality.

4.3 Global Convergence of the Lagrange

Combining Lemma 3, (23) and using the expected gradient errors provided by Lemma 5, we bound the average Lagrange optimality gap as a function of tunable parameters H and K as stated in the following corollary.

Corollary 1. *Consider the same setup and the choice of parameters described in Lemma 3− 5. The following inequality holds if assumptions 2−4 are met.*

$$\begin{aligned}
&\frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\Big(J_{\mathrm{L},\rho}(\pi^*, \lambda_k) - J_{\mathrm{L},\rho}(\theta_k, \lambda_k)\Big) \leq \sqrt{\epsilon_{\text{bias}}} + \left[\frac{G\sqrt{C}(1+\lambda_{\max})}{\sqrt{\mu_F}(1-\gamma)^2}\right] \exp\left(-\frac{\mu_F}{40G^2}H\right) \\
&+ B\eta\left[\frac{22\sigma^2 \mathrm{d}}{\mu_F H} + \exp\left(-\frac{\mu_F}{20G^2}H\right)\left[\frac{C(1+\lambda_{\max})^2}{\mu_F(1-\gamma)^4}\right] + \frac{G^2(1+\lambda_{\max})^2}{\mu_F^2(1-\gamma)^4}\right] \\
&+ \frac{1}{\eta K}\mathbf{E}_{s\sim d_\rho^{\pi^*}}\left[KL(\pi^*(\cdot|s)\|\pi_{\theta_0}(\cdot|s))\right]
\end{aligned} \tag{29}$$

Corollary 1 bounds the optimality error of the Lagrange function as $\mathcal{O}(\sqrt{\epsilon_{\text{bias}}} + \exp(-C_0 H) + \eta + \frac{1}{\eta K})$ where C_0 is some problem specific constant. Interestingly, the dual learning parameter, ζ does not appear in this bound. However, the next section shows that ζ plays a pivotal role in deciding the objective and constraint violation rates.

4.4 Decoupling the Objective and the Constraint Violation Rates

The goal of the following theorem is to choose optimal values of the tunable parameters and decouple the objective and constraint violation rates from the Lagrange convergence result given in (29).

Theorem 1. *Consider the same setup and the choice of parameters given in Lemma 3− 5. Assume* $\eta = (1-\gamma)^2(1+\lambda_{\max})^{-1}/\sqrt{K}$, $\zeta = \lambda_{\max}(1-\gamma)/\sqrt{K}$, *and* $\lambda_{\max} = 2/[(1-\gamma)c_{\text{slater}}]$. *For sufficiently small* $\epsilon > 0$, *the following inequalities hold*

$$\begin{aligned}
&\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\left[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\right] \leq \sqrt{\epsilon_{\text{bias}}} + \epsilon, \\
&\mathbf{E}\left[\frac{1}{K}\sum_{k=0}^{K-1} -J_{c,\rho}(\theta_k)\right] \leq (1-\gamma)c_{\text{slater}}\sqrt{\epsilon_{\text{bias}}} + \epsilon
\end{aligned} \tag{30}$$

whenever assumptions 1−4 are met, $H = \mathcal{O}(\log(\epsilon^{-1}))$ *and* $K = \mathcal{O}((1-\gamma)^{-6}\epsilon^{-2})$. *Therefore, the sample complexity to ensure* (30) *is* $\mathcal{O}((1-\gamma)^{-1}HK) = \tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$. *It is to be clarified that the* $(1-\gamma)^{-1}$ *factor in the sample complexity calculation appears due to the fact that it requires* $\mathcal{O}((1-\gamma)^{-1})$ *samples on an average to obtain a gradient estimate via Algorithm 1.*

Theorem 1 dictates that, with appropriate choice of the parameters, the rate of convergence of the objective and that of constraint violation can be bounded as $\mathcal{O}(\sqrt{\epsilon_{\text{bias}}} + \epsilon)$ with $\tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$ samples. This beats the SOTA $\tilde{\mathcal{O}}(\epsilon^{-4})$ sample complexity of [4] and achieves the theoretical lower bound. Our derived sample complexity is also dependent on c_{slater}. However, this is not explicitly mentioned in Theorem 1. Interested readers can find such details in the appendix.

Remark 3. *Note the importance of the nested expectation in the first-order term* $\mathbf{E}\|\mathbf{E}[\omega_k|\theta_k, \lambda_k] - \omega_k^*\|$ *in Lemma 3. Lemma 5 bounds this term as* $\mathcal{O}(\exp(-C_0 H))$ *where* C_0 *is a problem dependent constant. Other terms in the Lagrange optimality bound (Lemma 3) can be expressed as* $\mathcal{O}(\eta + \frac{1}{\eta K})$*. Moreover, following the proof of Theorem 1, one sees that decoupling the objective optimality error and constraint violation bounds incurs additional* $\mathcal{O}(\zeta + \frac{1}{\zeta K})$ *terms. Choosing* η, ζ *as prescribed in Theorem 1, makes both the objective and constraint violation errors as* $\mathcal{O}(\sqrt{\epsilon_{\text{bias}}} + \exp(-C_0 H) + K^{-0.5})$*. This allows us to take* $H = \tilde{\mathcal{O}}(1)$ *and* $K = \mathcal{O}(\epsilon^{-2})$*, leading to* $\tilde{\mathcal{O}}(\epsilon^{-2})$ *sample complexity. Had the first-order term been* $\mathbf{E}\|\omega_k - \omega_k^*\|$*, it would have resulted in* $\mathcal{O}(\sqrt{\epsilon_{\text{bias}}} + H^{-0.5} + K^{-0.5})$ *objective and constraint violation errors which would have lead to* $\mathcal{O}(\epsilon^{-4})$ *sample complexity.*

5 Conclusions and Limitations

This paper considers the problem of learning a CMDP where the goal is to maximize the objective value function while guaranteeing that the cost value exceeds a predefined threshold. We propose an acceleration-based primal-dual natural policy gradient algorithm that ensures ϵ optimality gap and ϵ constraint violation with $\tilde{\mathcal{O}}(\epsilon^{-2})$ sample complexity. This improves upon the previous state-of-the-art sample complexity of $\mathcal{O}(\epsilon^{-4})$ and achieves the theoretical lower bound. Future works include applying the idea of acceleration-based NPG to improve sample complexities in other related domains of constrained reinforcement learning, e.g., non-linear CMDP, average reward CMDP, etc.

References

[1] Al-Abbasi, A. O., A. Ghosh, V. Aggarwal. Deeppool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20(12):4714–4727, 2019.

[2] Geng, N., T. Lan, et al. A multi-agent reinforcement learning perspective on distributed traffic engineering. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–11. IEEE, 2020.

[3] Gonzalez, G., M. Balakuntala, M. Agarwal, et al. Asap: A semi-autonomous precise system for telesurgery during communication delays. *IEEE Transactions on Medical Robotics and Bionics*, 5(1):66–78, 2023.

[4] Bai, Q., A. S. Bedi, 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*, vol. 37, pages 6737–6744. 2023.

[5] Ding, D., K. Zhang, T. Basar, M. Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33:8378–8390, 2020.

[6] Xu, T., Y. Liang, G. Lan. Crpo: A new approach for safe reinforcement learning with convergence guarantee. In *International Conference on Machine Learning*, pages 11480–11491. 2021.

[7] Liu, Y., K. Zhang, T. Basar, 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.

[8] Jain, P., S. M. Kakade, R. Kidambi, P. Netrapalli, A. Sidford. Accelerating stochastic gradient descent for least squares regression. In *Conference On Learning Theory*, pages 545–604. 2018.

[9] Liu, T., R. Zhou, D. Kalathil, P. Kumar, C. Tian. Policy optimization for constrained mdps with provable fast global convergence. *arXiv preprint arXiv:2111.00552*, 2021.

[10] Zeng, S., T. T. Doan, J. Romberg. Finite-time complexity of online primal-dual natural actor-critic algorithm for constrained markov decision processes. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4028–4033. IEEE, 2022.

[11] Vaswani, S., L. Yang, C. Szepesvári. Near-optimal sample complexity bounds for constrained mdps. *Advances in Neural Information Processing Systems*, 35:3110–3122, 2022.

[12] Agarwal, A., S. M. Kakade, J. D. Lee, 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.

[13] Bhandari, J., D. Russo. On the linear convergence of policy gradient methods for finite mdps. In *International Conference on Artificial Intelligence and Statistics*, pages 2386–2394. 2021.

[14] Cen, S., C. Cheng, Y. Chen, Y. Wei, Y. Chi. Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 70(4):2563–2578, 2022.

[15] Lan, G. Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *Mathematical programming*, 198(1):1059–1106, 2023.

[16] Zhan, W., S. Cen, B. Huang, Y. Chen, J. D. Lee, Y. Chi. Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *SIAM Journal on Optimization*, 33(2):1061–1091, 2023.

[17] Xu, P., F. Gao, Q. Gu. An improved convergence analysis of stochastic variance-reduced policy gradient. In *Uncertainty in Artificial Intelligence*, pages 541–551. 2020.

[18] Gargiani, M., A. Zanelli, A. Martinelli, T. Summers, 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. 2022.

[19] Huang, F., S. Gao, J. Pei, H. Huang. Momentum-based policy gradient methods. In *International conference on machine learning*, pages 4422–4433. 2020.

[20] Salehkaleybar, S., M. Khorasani, N. Kiyavash, N. He, P. Thiran. Momentum-based policy gradient with second-order information. *Transactions on Machine Learning Research*, 2024.

[21] Shen, Z., A. Ribeiro, H. Hassani, H. Qian, C. Mi. Hessian aided policy gradient. In *International conference on machine learning*, pages 5729–5738. 2019.

[22] Chen, Z., S. Khodadadian, S. T. Maguluri. Finite-sample analysis of off-policy natural actor–critic with linear function approximation. *IEEE Control Systems Letters*, 6:2611–2616, 2022.

[23] Chen, Z., 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. 2022.

[24] Khodadadian, S., T. T. Doan, J. Romberg, S. T. Maguluri. Finite sample analysis of two-time-scale natural actor-critic algorithm. *IEEE Transactions on Automatic Control*, 2022.

[25] Fatkhullin, I., A. Barakat, A. Kireeva, N. He. Stochastic policy gradient methods: Improved sample complexity for fisher-non-degenerate policies. In *International Conference on Machine Learning*, pages 9827–9869. 2023.

[26] Masiha, S., S. Salehkaleybar, N. He, N. Kiyavash, 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.

[27] Mondal, W. U., V. Aggarwal. Improved sample complexity analysis of natural policy gradient algorithm with general parameterization for infinite horizon discounted reward Markov decision processes. In *International Conference on Artificial Intelligence and Statistics*, pages 3097–3105. 2024.

[28] Efroni, Y., S. Mannor, M. Pirotta. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.

[29] Liu, T., R. Zhou, D. Kalathil, P. Kumar, C. Tian. Learning policies with zero or bounded constraint violation for constrained mdps. *Advances in Neural Information Processing Systems*, 34:17183–17193, 2021.

[30] Ding, D., X. Wei, Z. Yang, Z. Wang, M. Jovanovic. Provably efficient safe exploration via primal-dual policy optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3304–3312. 2021.

[31] He, J., D. Zhou, Q. Gu. Nearly minimax optimal reinforcement learning for discounted mdps. *Advances in Neural Information Processing Systems*, 34:22288–22300, 2021.

[32] Wei, H., X. Liu, L. Ying. Triple-q: A model-free algorithm for constrained reinforcement learning with sublinear regret and zero constraint violation. In *International Conference on Artificial Intelligence and Statistics*, pages 3274–3307. 2022.

[33] Bai, Q., A. S. Bedi, M. Agarwal, A. Koppel, V. Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via primal-dual approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pages 3682–3689. 2022.

[34] Sutton, R. S., D. McAllester, S. Singh, Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

[35] Zhang, J., C. Ni, C. Szepesvari, M. Wang. On the convergence and sample efficiency of variance-reduced policy gradient method. *Advances in Neural Information Processing Systems*, 34:2228–2240, 2021.

[36] Jin, C., Z. Yang, Z. Wang, M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In J. Abernethy, S. Agarwal, eds., *Proceedings of Thirty Third Conference on Learning Theory*, vol. 125 of *Proceedings of Machine Learning Research*, pages 2137–2143. PMLR, 2020.

[37] Wang, L., Q. Cai, Z. Yang, Z. Wang. Neural policy gradient methods: Global optimality and rates of convergence. In *International Conference on Learning Representations*. 2019.

[38] Zhang, K., A. Koppel, H. Zhu, T. Basar. Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612, 2020.

[39] Mondal, W. U., V. Aggarwal, 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.

[40] Ding, D., K. Zhang, J. Duan, T. Başar, M. R. Jovanović. Convergence and sample complexity of natural policy gradient primal-dual methods for constrained mdps. *arXiv preprint arXiv:2206.02346*, 2022.

A Helper Lemma

Lemma 6. *With slight abuse of notation, define $J_{\mathrm{L},\rho}(\pi,\lambda) \triangleq J^{\pi}_{r,\rho} + \lambda J^{\pi}_{c,\rho}$. The following relation holds for any two policies π_1, π_2 and $\lambda \in \Lambda$.*

$$J_{\mathrm{L},\rho}(\pi_1,\lambda) - J_{\mathrm{L},\rho}(\pi_2,\lambda) = \frac{1}{1-\gamma}\mathbf{E}_{(s,a)\sim\nu^{\pi_1}_{\rho}}\left[A^{\pi_2}_{\mathrm{L},\lambda}(s,a)\right] \tag{31}$$

Proof. This can be proved using Lemma 2 of [12] and the definition of the Lagrange function. □

B Proof of Lemma 1

Proof. Fix arbitrary θ and λ. Note that the following equation holds due to the definition of $\hat{J}_{c,\rho}(\theta)$.

$$\begin{aligned}
\mathbf{E}\left[\hat{J}_{c,\rho}(\theta)\middle|\theta\right] &= (1-\gamma)\mathbf{E}\left[\sum_{t=0}^{\infty}\gamma^t\sum_{j=0}^{t}c(s_j,a_j)\middle|s_0\sim\rho,\pi_\theta\right] \\
&= (1-\gamma)\mathbf{E}\left[\sum_{j=0}^{\infty}c(s_j,a_j)\sum_{t=j}^{\infty}\gamma^t\middle|s_0\sim\rho,\pi_\theta\right] \\
&= \mathbf{E}\left[\sum_{j=0}^{\infty}\gamma^j c(s_j,a_j)\middle|s_0\sim\rho,\pi_\theta\right] = J_{c,\rho}(\theta)
\end{aligned} \tag{32}$$

To prove the unbiasedness of the gradient, we first prove that the distribution of the sample pair $(\hat{s},\hat{a})$ produced by Algorithm 1 is indeed $\nu^{\pi_\theta}_{\rho}$. Observe the following.

$$\Pr(\hat{s}=s,\hat{a}=a|\rho,\pi_\theta) = (1-\gamma)\sum_{t=0}^{\infty}\gamma^t\Pr(s_t=s,a_t=a|s_0\sim\rho,\pi_\theta) = \nu^{\pi_\theta}_{\rho}(s,a) \tag{33}$$

Next, we show that for a given pair (s,a), the estimate $\hat{Q}^{\pi_\theta}_{g}(s,a)$, $g\in\{r,c\}$ is unbiased.

$$\begin{aligned}
\mathbf{E}\left[\hat{Q}^{\pi_\theta}_{g}(s,a)\middle|\theta,s,a\right] &= (1-\gamma)\mathbf{E}\left[\sum_{t=0}^{\infty}\gamma^t\sum_{j=0}^{t}g(s_i,a_i)\middle|s_0=s,a_0=a,\pi_\theta\right] \\
&= (1-\gamma)\mathbf{E}\left[\sum_{j=0}^{\infty}g(s_i,a_i)\sum_{t=j}^{t}\gamma^t\middle|s_0=s,a_0=a,\pi_\theta\right] \\
&= \mathbf{E}\left[\sum_{j=0}^{\infty}\gamma^i g(s_i,a_i)\middle|s_0=s,a_0=a,\pi_\theta\right] = Q^{\pi_\theta}_{g}(s,a)
\end{aligned} \tag{34}$$

In a similar fashion, one can establish that $\mathbf{E}[\hat{V}^{\pi_\theta}_{g}(s)\big|\theta,s] = V^{\pi_\theta}_{g}(s)$, $\forall g\in\{r,c\}$. Combining this with (34) leads to: $\mathbf{E}[\hat{A}^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)|\theta,\lambda,s,a] = A^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)$. We arrive at,

$$\mathbf{E}\left[\hat{F}_\rho(\theta)\middle|\theta\right] = \mathbf{E}_{(\hat{s},\hat{a})\sim\nu^{\pi_\theta}_{\rho}}\left[\nabla_\theta\log\pi_\theta(\hat{a}|\hat{s})\otimes\nabla_\theta\log\pi_\theta(\hat{a}|\hat{s})\middle|\theta\right] = F_\rho(\theta), \tag{35}$$

$$\begin{aligned}
\mathbf{E}\left[\hat{H}_\rho(\theta,\lambda)\middle|\theta,\lambda\right] &= \mathbf{E}_{(\hat{s},\hat{a})\sim\nu^{\pi_\theta}_{\rho}}\left[\mathbf{E}\left[\hat{A}^{\pi_\theta}_{\mathrm{L},\lambda}(\hat{s},\hat{a})\middle|\theta,\lambda,\hat{s},\hat{a}\right]\nabla_\theta\log\pi_\theta(\hat{a}|\hat{s})\middle|\theta,\lambda\right] \\
&= \mathbf{E}_{(\hat{s},\hat{a})\sim\nu^{\pi_\theta}_{\rho}}\left[A^{\pi_\theta}_{\mathrm{L},\lambda}(\hat{s},\hat{a})\nabla_\theta\log\pi_\theta(\hat{a}|\hat{s})\middle|\theta,\lambda\right] = H_\rho(\theta,\lambda)
\end{aligned} \tag{36}$$

Finally, we arrive at the following by utilizing the definitions (10) and (14).

$$\begin{aligned}
\mathbf{E}\left[\hat{\nabla}_\omega L_{\nu^{\pi_\theta}_{\rho}}(\omega,\theta,\lambda)\middle|\omega,\theta,\lambda\right] &= \mathbf{E}\left[\hat{F}_\rho(\theta)\middle|\theta\right]\omega - \frac{1}{1-\gamma}\mathbf{E}\left[\hat{H}_\rho(\theta,\lambda)\middle|\theta,\lambda\right] \\
&= F_\rho(\theta)\omega - \frac{1}{1-\gamma}H_\rho(\theta,\lambda) = \nabla_\omega L_{\nu^{\pi_\theta}_{\rho}}(\omega,\theta,\lambda)
\end{aligned} \tag{37}$$

This concludes the proof. □

C Proof of Lemma 3

Proof. Invoking the definition of KL divergence, we arrive at the following series of inequalities.

$$
\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,a)\sim\nu_\rho^{\pi^*}}\left[\log\frac{\pi_{\theta_{k+1}}(a|s)}{\pi_{\theta_k}(a|s)}\right] \\
&\overset{(a)}{\geq} \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\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,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^*] + \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot(\omega_k-\omega_k^*)] - \frac{B\eta^2}{2}\|\omega_k\|^2 \\
&= \eta[J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)] + \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot\omega_k^*] \\
&- \eta[J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)] + \eta\mathbf{E}_{(s,a)\sim\nu^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot(\omega_k-\omega_k^*)] - \frac{B\eta^2}{2}\|\omega_k\|^2 \\
&\overset{(b)}{=} \eta[J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)] + \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}\left[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot\omega_k^* - \frac{1}{1-\gamma}A_{\mathrm{L},\lambda_k}^{\pi_{\theta_k}}(s,a)\right] \\
&+ \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot(\omega_k-\omega_k^*)] - \frac{B\eta^2}{2}\|\omega_k\|^2 \\
&\overset{(c)}{\geq} \eta[J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)] - \eta\sqrt{\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}\left[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot\omega_k^* - \frac{1}{1-\gamma}A_{\mathrm{L},\lambda_k}^{\pi_{\theta_k}}(s,a)\right]^2} \\
&+ \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot(\omega_k-\omega_k^*)] - \frac{B\eta^2}{2}\|\omega_k\|^2 \\
&\overset{(d)}{\geq} \eta[J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)] - \eta\sqrt{\epsilon_{\mathrm{bias}}} \\
&\qquad + \eta\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi^*}}[\nabla_\theta \log\pi_{\theta_k}(a|s)\cdot(\omega_k-\omega_k^*)] - \frac{B\eta^2}{2}\|\omega_k\|^2
\end{aligned}
\tag{38}
$$

where the step (a) holds by Assumption 2 and step (b) holds by Lemma 6. Step (c) uses the convexity of the function $f(x) = x^2$. Finally, step (d) comes from the Assumption 3. Rearranging the terms and taking expectations on both sides, we have,

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

where (a) follows from Assumption 2. Summing from $k = 0$ to $K-1$, using the non-negativity of KL divergence and dividing the resulting expression by K, we obtain,

$$
\begin{aligned}
\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big(J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)\Big) &\leq \sqrt{\epsilon_{\mathrm{bias}}} + \frac{G}{K}\sum_{k=0}^{K-1}\mathbf{E}\|(\mathbf{E}\left[\omega_k\middle|\theta_k,\lambda_k\right] - \omega_k^*)\| \\
&+ \frac{B\eta}{2K}\sum_{k=0}^{K-1}\mathbf{E}\|\omega_k\|^2 + \frac{1}{\eta K}\mathbf{E}_{s\sim d_\rho^{\pi^*}}[KL(\pi^*(\cdot|s)\|\pi_{\theta_0}(\cdot|s))]
\end{aligned}
\tag{40}
$$

This concludes the proof. □

D Proof of Lemma 4

Proof. Fix a $\theta \in \Theta$ and a $\lambda \in \Lambda$. Observe the following equation.

$$\mathbf{E}\left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega^*_{\theta,\lambda},\theta,\lambda)\otimes\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega^*_{\theta,\lambda},\theta,\lambda)\right]$$
$$=\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}}\left[\mathbf{E}\Bigg[\underbrace{\nabla_\theta\log\pi_\theta(a|s)\cdot\omega^*_{\theta,\lambda}-\frac{1}{1-\gamma}\hat{A}^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)}_{\triangleq\zeta_{\theta,\lambda}(s,a)}\Bigg]^2\nabla_\theta\log\pi_\theta(a|s)\otimes\nabla_\theta\log\pi_\theta(a|s)\right]$$

To prove the lemma, it is sufficient to demonstrate that $\mathbf{E}[\zeta_{\theta,\lambda}(s,a)]\leq\sigma^2,\forall(s,a)$. Notice the chain of inequalities stated below.

$$\begin{aligned}
&\mathbf{E}\left[\nabla_\theta\log\pi_\theta(a|s)\cdot\omega^*_{\theta,\lambda}-\frac{1}{1-\gamma}\hat{A}^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)\right]^2\\
&\leq 2\left[\nabla_\theta\log\pi_\theta(a|s)\cdot\omega^*_{\theta,\lambda}\right]^2+\frac{2}{(1-\gamma)^2}\mathbf{E}\left[\hat{A}^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)\right]^2\\
&\overset{(a)}{\leq} 2\|\nabla_\theta\log\pi_\theta(a|s)\|^2\|\omega^*_{\theta,\lambda}\|^2+\frac{4(1+\lambda)^2}{(1-\gamma)^2}\max_{g\in\{r,c\}}\left\{\mathbf{E}\left[\hat{A}^{\pi_\theta}_g(s,a)\right]^2\right\}\\
&\overset{(b)}{\leq} 2G^2\|F_\rho(\theta)^\dagger\nabla_\theta J_{\mathrm{L},\rho}(\theta,\lambda)\|^2+\frac{8(1+\lambda)^2}{(1-\gamma)^2}\max_{g\in\{r,c\}}\left\{\mathbf{E}\left[\hat{Q}^{\pi_\theta}_g(s,a)\right]^2+\mathbf{E}\left[\hat{V}^{\pi_\theta}_g(s,a)\right]^2\right\}\\
&\overset{(c)}{\leq}\frac{2G^4(1+\lambda)^2}{\mu_F^2(1-\gamma)^4}+\frac{32(1+\lambda)^2}{(1-\gamma)^4}\leq\frac{1}{(1-\gamma)^4}\left[\frac{2G^4}{\mu_F^2}+32\right](1+\lambda_{\max})^2
\end{aligned}\tag{41}$$

Inequality (a) follows from the Cauchy-Schwarz inequality and the fact that $(a+b)^2\leq 2(a^2+b^2)$ for any two numbers a,b. The same argument is also applied in inequality (b). Additionally, it uses Assumption 2 and the definition of $\omega^*_{\theta,\lambda}$. Finally, (c) is a consequence of Assumption 4, Lemma 2, and the following two bounds.

$$\mathbf{E}\left[\hat{Q}^{\pi_\theta}_g(s,a)\right]^2\leq\frac{2}{(1-\gamma)^2},\ \text{and}\ \mathbf{E}\left[\hat{V}^{\pi_\theta}_g(s)\right]^2\leq\frac{2}{(1-\gamma)^2},\ \ \forall(s,a)\in\mathcal{S}\times\mathcal{A},\forall g\in\{r,c\}\tag{42}$$

To establish the first bound, note that $|\hat{Q}^{\pi_\theta}_g(s,a)|$ is assigned a value of at most $(j+1)$ with probability $(1-\gamma)\gamma^j,\forall g\in\{r,c\}$. Therefore,

$$\mathbf{E}\left[\hat{Q}^{\pi_\theta}_g(s,a)\right]^2\leq\sum_{j=0}^{\infty}(1-\gamma)(j+1)^2\gamma^j=\frac{(1+\gamma)}{(1-\gamma)^2}<\frac{2}{(1-\gamma)^2}\tag{43}$$

The second bound in (42) can be proven similarly. This concludes the lemma. □

E Proof of Lemma 5

We establish Lemma 5 applying Corollary 2 of [8]. Note the following statements.

S1 : The following quantities exist and are finite $\forall\theta\in\Theta$.

$$F_\rho(\theta)\triangleq\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}}\left[\nabla_\theta\log\pi_\theta(a|s)\otimes\nabla_\theta\log\pi_\theta(a|s)\right],\tag{44}$$

$$G_\rho(\theta)\triangleq\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}}\left[\nabla_\theta\log\pi_\theta(a|s)\otimes\nabla_\theta\log\pi_\theta(a|s)\otimes\nabla_\theta\log\pi_\theta(a|s)\otimes\nabla_\theta\log\pi_\theta(a|s)\right]\tag{45}$$

S2 : There exists σ^2 such that the following is obeyed $\forall\theta\in\Theta$ where $\omega^*_{\theta,\lambda}$ minimizes $L_{\nu_\rho^{\pi_\theta}}(\cdot,\theta,\lambda)$.

$$\mathbf{E}\left[\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega^*_{\theta,\lambda},\theta,\lambda)\otimes\hat{\nabla}_\omega L_{\nu_\rho^{\pi_\theta}}(\omega^*_{\theta,\lambda},\theta,\lambda)\right]\preccurlyeq\sigma^2F_\rho(\theta)\tag{46}$$

S3 : There exists $\mu_F, G > 0$ such that the following statements hold $\forall \theta \in \Theta$.

$$(a)\ F_\rho(\theta) \succcurlyeq \mu_F I_{\mathrm{d}}, \tag{47}$$

$$(b)\ \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}}\left[\|\nabla_\theta \log \pi_\theta(a|s)\|^2 \nabla_\theta \log \pi_\theta(a|s) \otimes \nabla_\theta \log \pi_\theta(a|s)\right] \preccurlyeq G^2 F_\rho(\theta), \tag{48}$$

$$(c)\ \mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_\theta}}\left[\|\nabla_\theta \log \pi_\theta(a|s)\|^2_{F_\rho(\theta)^\dagger} \nabla_\theta \log \pi_\theta(a|s) \otimes \nabla_\theta \log \pi_\theta(a|s)\right] \preccurlyeq \frac{G^2}{\mu_F} F_\rho(\theta) \tag{49}$$

Statement **S1** follows from Assumption 2 whereas **S2** is a consequence of Lemma 4. Statement **S3**(a) is identical to Assumption 4, **S3**(b) results from Assumption 2, and finally, **S3**(c) follows from Assumption 2 and 4. We can, therefore, apply Corollary 2 of [8] with $\kappa = \tilde{\kappa} = G^2/\mu_F$ and deduce the following convergence result whenever $H > \bar{C}\sqrt{\kappa\tilde{\kappa}}\log(\sqrt{\mathrm{d}}\sqrt{\kappa\tilde{\kappa}})$ and the learning rates are set as $\alpha = \frac{3\sqrt{5}\sqrt{\kappa\tilde{\kappa}}}{1+3\sqrt{5\kappa\tilde{\kappa}}}$, $\beta = \frac{1}{9\sqrt{\kappa\tilde{\kappa}}}$, $\xi = \frac{1}{3\sqrt{5}\mu_F\sqrt{\kappa\tilde{\kappa}}}$, and $\delta = \frac{1}{5G^2}$.

$$\begin{aligned}&\mathbf{E}\left[l_k(\omega_k)\right] - l_k(\omega_k^*) \leq \frac{C}{2}\exp\left(-\frac{H}{20\sqrt{\kappa\tilde{\kappa}}}\right)\left[l_k(\mathbf{0}) - l_k(\omega_k^*)\right] + 11\frac{\sigma^2\mathrm{d}}{H},\\ &\text{where } l_k(\omega) \triangleq L_{\nu_\rho^{\pi_{\theta_k}}}(\omega, \theta_k, \lambda_k),\ \forall \omega \in \mathbb{R}^{\mathrm{d}}\end{aligned} \tag{50}$$

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.

$$l_k(\mathbf{0}) = \frac{1}{2}\mathbf{E}_{(s,a)\sim\nu_\rho^{\pi_{\theta_k}}}\left[\frac{1}{1-\gamma}A^{\pi_{\theta_k}}_{\mathrm{L},\lambda_k}(s,a)\right]^2 \overset{(a)}{\leq} \frac{(1+\lambda_{\max})^2}{2(1-\gamma)^4} \tag{51}$$

where (a) is a result of the fact that $|A^{\pi_\theta}_{\mathrm{L},\lambda}(s,a)| \leq (1+\lambda_{\max})/(1-\gamma)$, $\forall (s,a) \in \mathcal{S}\times\mathcal{A}$, $\forall \theta \in \Theta$, and $\forall \lambda \in \Lambda$. Combining (50), (51), and the fact that $l_k(\cdot)$ is μ_F-strongly convex, we establish,

$$\mathbf{E}\|\omega_k - \omega_k^*\|^2 \leq \frac{2}{\mu_F}\left[\mathbf{E}\left[l_k(\omega_k)\right] - l_k(\omega_k^*)\right] \leq 22\frac{\sigma^2\mathrm{d}}{\mu_F H} + C\exp\left(-\frac{\mu_F}{20G^2}H\right)\left[\frac{(1+\lambda_{\max})^2}{\mu_F(1-\gamma)^4}\right] \tag{52}$$

This proves the first statement. We get the following for noiseless ($\sigma^2 = 0$) gradient updates.

$$\mathbf{E}\|(\mathbf{E}[\omega_k|\theta_k] - \omega_k^*)\|^2 \leq C\exp\left(-\frac{\mu_F}{20G^2}H\right)\left[\frac{(1+\lambda_{\max})^2}{\mu_F(1-\gamma)^4}\right] \tag{53}$$

The second statement can be established from (53) by applying Jensen's inequality on the function $f(x) = x^2$.

F Proof of Theorem 1

Applying the inequalities $H \geq 1$, $\exp(-(\mu_F/20G^2)H) \leq 1$ and substituting the values of σ^2 and η (as stated in Lemma 4 and Theorem 1 respectively), we can rewrite (29) as follows.

$$\begin{aligned}&\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\left(J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k)\right)\\ &\qquad\leq \sqrt{\epsilon_{\mathrm{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\left(-\frac{\mu_F}{40G^2}H\right) + f_1\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}}\end{aligned} \tag{54}$$

The terms f_0 and f_1 are defined below.

$$\begin{aligned}f_0 &\triangleq \frac{G\sqrt{C}}{\sqrt{\mu_F}},\\ f_1 &\triangleq B\left[\frac{44\mathrm{d}}{\mu_F}\left(\frac{G^4}{\mu_F^2} + 16\right) + \frac{C}{\mu_F} + \frac{G^2}{\mu_F^2}\right] + \mathbf{E}_{s\sim d_\rho^{\pi^*}}[KL(\pi^*(\cdot|s)\|\pi_{\theta_0}(\cdot|s))]\end{aligned} \tag{55}$$

Using the definition of $J_{\mathrm{L},\rho}(\cdot,\cdot)$, one can write,

$$\begin{aligned}J_{\mathrm{L},\rho}(\pi^*,\lambda_k) - J_{\mathrm{L},\rho}(\theta_k,\lambda_k) &= (J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)) + \lambda_k(J_{c,\rho}^{\pi^*} - J_{c,\rho}(\theta_k))\\ &\overset{(a)}{\geq} (J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)) + \lambda_k(-J_{c,\rho}(\theta_k))\end{aligned} \tag{56}$$

where (a) follows from the fact that $\lambda_k \geq 0$ and $J_{c,\rho}^{\pi^*} \geq 0$ due to feasibility. Combining (54) and (56), we obtain the following.

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\Big] + \frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[-\lambda_k J_{c,\rho}(\theta_k)\Big] \\ \leq \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) + f_1\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}}\end{aligned} \tag{57}$$

F.1 Convergence Rate of the Objective Function

Note that (57) can be alternatively written as,

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\Big] \leq \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) \\ + f_1\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}} + \frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[\lambda_k J_{c,\rho}(\theta_k)\Big]\end{aligned} \tag{58}$$

To obtain the convergence rate of the objective function, we need to bound the last term in the above expression. Observe the following chain of inequalities.

$$\begin{aligned}0 \leq &(\lambda_K)^2 \overset{(a)}{=} \sum_{k=0}^{K-1}\Big((\lambda_{k+1})^2 - (\lambda_k)^2\Big) \\ &\overset{(b)}{\leq} \sum_{k=0}^{K-1}\Big(\big[\lambda_k - \zeta\hat{J}_{c,\rho}(\theta_k)\big]^2 - (\lambda_k)^2\Big) = -2\zeta\sum_{k=0}^{K-1}\lambda_k\hat{J}_{c,\rho}(\theta_k) + \zeta^2\sum_{k=0}^{K-1}\hat{J}_{c,\rho}^2(\theta_k)\end{aligned} \tag{59}$$

where (a) uses $\lambda_0 = 0$ and (b) follows from the contraction property of the projection operator, $\mathcal{P}_\Lambda$. Using the above inequality, one can write,

$$\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[\lambda_k\hat{J}_{c,\rho}(\theta_k)\Big] \leq \frac{\zeta}{2K}\sum_{k=0}^{K-1}\mathbf{E}\Big[\hat{J}_{c,\rho}^2(\theta_k)\Big] \tag{60}$$

Using the unbiasedness of $\hat{J}_{c,\rho}(\theta_k)$ (Lemma 1), we deduce the following.

$$\mathbf{E}\Big[\lambda_k\hat{J}_{c,\rho}(\theta_k)\Big] \overset{(a)}{=} \mathbf{E}\Big[\lambda_k\mathbf{E}\Big[\hat{J}_{c,\rho}(\theta_k)\big|\theta_k\Big]\Big] = \mathbf{E}\Big[\lambda_k J_{c,\rho}(\theta_k)\Big] \tag{61}$$

where (a) is a consequence of the fact that $\hat{J}_{c,\rho}(\theta_k)$ and λ_k are conditionally independent given θ_k. Note that, $\hat{J}_{c,\rho}^2(\theta_k)$ is assigned a value of at most $(j+1)^2$ with probability $(1-\gamma)\gamma^j$, $j \in \{0, 1, \cdots\}$. Therefore, the RHS of (60) can be bounded as follows.

$$\mathbf{E}\Big[\hat{J}_{c,\rho}^2(\theta_k)\Big] \leq \sum_{j=0}^{\infty}(1-\gamma)(j+1)^2\gamma^j = \frac{1+\gamma}{(1-\gamma)^2} < \frac{2}{(1-\gamma)^2} \tag{62}$$

Combining (58), (60), (61), and (62), we finally obtain,

$$\begin{aligned}&\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\Big[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\Big] \\ &\leq \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) + f_1\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}} + \frac{\zeta}{(1-\gamma)^2} \\ &\overset{(a)}{=} \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) + f_1\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}} + \frac{\lambda_{\max}}{(1-\gamma)}\frac{1}{\sqrt{K}} \\ &\leq \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) + (f_1+1)\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}}\end{aligned} \tag{63}$$

where (a) uses the substitution $\zeta = \lambda_{\max}(1-\gamma)/\sqrt{K}$.

F.2 Rate of Constraint Violation

The following inequality is satisfied for any $k \in \{0, 1, \cdots, K-1\}$.

$$\begin{aligned} &|\lambda_{k+1} - \lambda_{\max}|^2 = |\mathcal{P}_\Lambda(\lambda_k - \zeta \hat{J}_{c,\rho}(\theta_k)) - \lambda_{\max}|^2 \\ &\overset{(a)}{\leq} \left|\lambda_k - \zeta \hat{J}_{c,\rho}(\theta_k) - \lambda_{\max}\right|^2 = \left|\lambda_k - \lambda_{\max}\right|^2 - 2\zeta \hat{J}_{c,\rho}(\theta_k)\big(\lambda_k - \lambda_{\max}\big) + \zeta^2 \hat{J}^2_{c,\rho}(\theta_k) \end{aligned} \tag{64}$$

where (a) is due to the contractive property of $\mathcal{P}_\Lambda$. Performing an average over $k \in \{0, \cdots, K-1\}$, and applying expectations on both sides, we get,

$$\begin{aligned} \frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\left[(\lambda_k - \lambda_{\max})\hat{J}_{c,\rho}(\theta_k)\right] &\overset{(a)}{\leq} \frac{|\lambda_{\max}|^2 - |\lambda_K - \lambda_{\max}|^2}{2\zeta K} + \frac{\zeta}{2K}\sum_{k=0}^{K-1} \mathbf{E}\left[\hat{J}^2_{c,\rho}(\theta_k)\right] \\ &\overset{(b)}{\leq} \frac{\lambda^2_{\max}}{2\zeta K} + \frac{\zeta}{(1-\gamma)^2} \overset{(c)}{=} \frac{3\lambda_{\max}}{2(1-\gamma)}\frac{1}{\sqrt{K}} \end{aligned} \tag{65}$$

where (a) utilises $\lambda_0 = 0$, (b) applies (62), and (c) is derived using $\zeta = \lambda_{\max}(1-\gamma)/\sqrt{K}$. Note that one can write the following using Lemma 1 and the observation that $\hat{J}_{c,\rho}(\theta_k)$ and λ_k are conditionally independent given θ_k.

$$\mathbf{E}\left[(\lambda_k - \lambda_{\max})\hat{J}_{c,\rho}(\theta_k)\right] = \mathbf{E}\left[(\lambda_k - \lambda_{\max})\mathbf{E}\left[\hat{J}_{c,\rho}(\theta_k)\big|\theta_k\right]\right] = \mathbf{E}\Big[(\lambda_k - \lambda_{\max})J_{c,\rho}(\theta_k)\Big] \tag{66}$$

Combining (65) and (66), we get,

$$\frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\Big[(\lambda_k - \lambda_{\max})J_{c,\rho}(\theta_k)\Big] \leq \frac{3\lambda_{\max}}{2(1-\gamma)}\frac{1}{\sqrt{K}} \leq \frac{2(1+\lambda_{\max})}{(1-\gamma)^2\sqrt{K}} \tag{67}$$

Finally, combining (58) and (67), we arrive at,

$$\begin{aligned} &\frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\Big[J^{\pi^*}_{r,\rho} - J_{r,\rho}(\theta_k)\Big] + \lambda_{\max}\frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\Big[-J_{c,\rho}(\theta_k)\Big] \\ &\qquad \leq \sqrt{\epsilon_{\text{bias}}} + f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) + (f_1+2)\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}} \end{aligned} \tag{68}$$

Since the functions $\{J_{g,\rho}(\theta_k)\}$, $g \in \{r, c\}$, $k \in \{0, \cdots, K-1\}$ are linear in occupancy measure, there exists a policy $\bar{\pi}$ such that the following holds $\forall g \in \{r, c\}$.

$$\frac{1}{K}\sum_{k=0}^{K-1} J_{g,\rho}(\theta_k) = J^{\bar{\pi}}_{g,\rho}$$

This allows us to rewrite (68) as,

$$\begin{aligned} J^{\pi^*}_{r,\rho} - \mathbf{E}\big[J^{\bar{\pi}}_{r,\rho}\big] + \lambda_{\max}\mathbf{E}\big[-J^{\bar{\pi}}_{c,\rho}\big] \leq \sqrt{\epsilon_{\text{bias}}} &+ f_0\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) \\ &+ (f_1+2)\frac{(1+\lambda_{\max})}{(1-\gamma)^2}\frac{1}{\sqrt{K}} \end{aligned} \tag{69}$$

Applying Lemma 9 and verifying (via Lemma 8) that $\lambda_{\max} \geq 2\lambda^*$ where λ^* is the non-negative minimizer of the dual function corresponding to the unparameterized constrained optimization (2), we can write the constraint violation rate as follows.

$$\begin{aligned} \frac{1}{K}\sum_{k=0}^{K-1} \mathbf{E}\left[-J_{c,\rho}(\theta_k)\right] \leq \frac{2\sqrt{\epsilon_{\text{bias}}}}{\lambda_{\max}} &+ 2f_0\frac{(1+\lambda_{\max})}{\lambda_{\max}(1-\gamma)^2}\exp\Big(-\frac{\mu_F}{40G^2}H\Big) \\ &+ 2(f_1+2)\frac{(1+\lambda_{\max})}{\lambda_{\max}(1-\gamma)^2}\frac{1}{\sqrt{K}} \end{aligned} \tag{70}$$

F.3 Final Result

Substituting $\lambda_{\max} = 2/[(1-\gamma)c_{\text{slater}}] \geq 2$ in (63), we get the rate of convergence of the objective as follows.

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\left[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\right] \\ \leq \sqrt{\epsilon_{\text{bias}}} + \frac{3f_0 c_{\text{slater}}^{-1}}{(1-\gamma)^3}\exp\left(-\frac{\mu_F}{40G^2}H\right) + \frac{3(f_1+1)c_{\text{slater}}^{-1}}{(1-\gamma)^3}\frac{1}{\sqrt{K}}\end{aligned} \tag{71}$$

Similarly, we obtain the constraint violation rate as,

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\left[-J_{c,\rho}(\theta_k)\right] \\ \leq (1-\gamma)c_{\text{slater}}\sqrt{\epsilon_{\text{bias}}} + \frac{3f_0}{(1-\gamma)^2}\exp\left(-\frac{\mu_F}{40G^2}H\right) + \frac{3(f_1+2)}{(1-\gamma)^2}\frac{1}{\sqrt{K}}\end{aligned} \tag{72}$$

Let, H and K are chosen as follows for an arbitrary $\epsilon > 0$.

$$\begin{aligned}H &= \frac{40G^2}{\mu_F}\log\left(\max\left\{\frac{6f_0 c_{\text{slater}}^{-1}}{(1-\gamma)^3}\epsilon^{-1}, \frac{6f_0}{(1-\gamma)^2}\epsilon^{-1}\right\}\right) = \mathcal{O}(\log(\epsilon^{-1})), \\ K &= \max\left\{\frac{36(f_1+1)^2 c_{\text{slater}}^{-2}}{(1-\gamma)^6}\epsilon^{-2}, \frac{36(f_1+2)^2}{(1-\gamma)^4}\epsilon^{-2}\right\} = \mathcal{O}((1-\gamma)^{-6}\epsilon^{-2})\end{aligned} \tag{73}$$

For the above choice of H and K, we have,

$$\begin{aligned}\frac{1}{K}\sum_{k=0}^{K-1}\mathbf{E}\left[J_{r,\rho}^{\pi^*} - J_{r,\rho}(\theta_k)\right] &\leq \sqrt{\epsilon_{\text{bias}}} + \epsilon, \\ \mathbf{E}\left[\frac{1}{K}\sum_{k=0}^{K-1} -J_{c,\rho}(\theta_k)\right] &\leq (1-\gamma)c_{\text{slater}}\sqrt{\epsilon_{\text{bias}}} + \epsilon\end{aligned} \tag{74}$$

Note that the expected number of steps required in Algorithm 1 is $\mathcal{O}((1-\gamma)^{-1})$. Therefore, the sample complexity required to ensure (74) is $\mathcal{O}((1-\gamma)^{-1}HK) = \tilde{\mathcal{O}}((1-\gamma)^{-7}\epsilon^{-2})$. Finally, note that Lemma 5 requires $H > \bar{C}\frac{G^2}{\mu_F}\log\left(\sqrt{\mathrm{d}}\frac{G^2}{\mu_F}\right)$. This can be ensured if ϵ is sufficiently small.

G Strong Duality and Related Lemmas

Define the dual function associated with the unparameterized constrained optimization (2) as follows.

$$J_{\mathrm{D},\rho}^{\lambda} = \max_{\pi}\left\{J_{r,\rho}^{\pi} + \lambda J_{c,\rho}^{\pi}\right\} \tag{75}$$

The following lemma formally describes the strong duality result.

Lemma 7. *[40] [Lemma 3] If $\lambda^* \triangleq \arg\min_{\lambda\geq 0} J_{\mathrm{D},\rho}^{\lambda}$ and π^* is a solution of* (2)*, then the following holds whenever Assumption 1 is true.*

$$J_{r,\rho}^{\pi^*} = J_{\mathrm{D},\rho}^{\lambda^*} \tag{76}$$

It is to be mentioned that strong duality, in general, does not hold for the parameterized optimization (3). The following lemma established a bound on λ^* which becomes the foundation for choosing the value of $\lambda_{\max}$ in Algorithm 2.

Lemma 8. *[40][Lemma 3] Let λ^* be the optimal dual variable as defined in Lemma 7. The following inequalities hold where c_{slater} is defined in Assumption 1.*

$$0 \leq \lambda^* \leq \frac{1}{(1-\gamma)c_{\text{slater}}}$$

The following lemma is the main tool in decoupling the objective and constraint violation rates.

Lemma 9. *Let Slater's condition (Assumption 1) hold. If $C \geq 2\lambda^*$ where λ^* is defined in Lemma 7 and $\bar{\pi}$ is a policy such that $J_{r,\rho}^{\pi^*} - J_{r,\rho}^{\bar{\pi}} + C[-J_{c,\rho}^{\bar{\pi}}] \leq \zeta$ for some $\zeta > 0$ then*

$$-J_{c,\rho}^{\bar{\pi}} \leq \frac{2\zeta}{C} \tag{77}$$

Proof. Define the function $v(\cdot)$ as follows.

$$v(\tau) = \max_{\pi} \left\{ J_{r,\rho}^{\pi} \middle| J_{c,\rho}^{\pi} \geq \tau \right\}, \ \tau \in \mathbb{R}^{\mathrm{d}} \tag{78}$$

Let $\tau = J_{c,\rho}^{\bar{\pi}}$. Therefore, one can write the following chain of inequalities.

$$\begin{aligned} J_{r,\rho}^{\bar{\pi}} \overset{(a)}{\leq} v(\tau) = \max_{\pi} \left\{ J_{r,\rho}^{\pi} \middle| J_{c,\rho}^{\pi} \geq \tau \right\} &\overset{(b)}{\leq} \max_{\pi} \left\{ J_{r,\rho}^{\pi} + \lambda^*(J_{c,\rho}^{\pi} - \tau) \middle| J_{c,\rho}^{\pi} \geq \tau \right\} \\ &= \max_{\pi} \left\{ J_{r,\rho}^{\pi} + \lambda^* J_{c,\rho}^{\pi} \middle| J_{c,\rho}^{\pi} \geq \tau \right\} - \tau\lambda^* \\ &\overset{(c)}{\leq} J_{\mathrm{D},\rho}^{\lambda^*} - \tau\lambda^* \overset{(d)}{=} J_{r,\rho}^{\pi^*} - \tau\lambda^* \end{aligned} \tag{79}$$

where (a) follows from the definition of $v(\tau)$, (b) uses $\lambda^* \geq 0$, (c) follows from the definition of the dual function, and (d) is a result of the strong duality (Lemma 7). Utilizing (79), we get,

$$(C - \lambda^*)(-\tau) \leq J_{r,\rho}^{\pi^*} - J_{r,\rho}^{\bar{\pi}} + C(-\tau) \leq \zeta \tag{80}$$

This leads to $-\tau \leq \zeta/(C - \lambda^*) \leq 2\zeta/C$. □