

Variance-Reduced Conservative Policy Iteration

Naman Agarwal

Google AI Princeton

NAMANAGARWAL@GOOGLE.COM

Brian Bullins

Department of Computer Science, Purdue University

BBULLINS@PURDUE.EDU

Karan Singh

Tepper School of Business, Carnegie Mellon University

KARANSINGH@CMU.EDU

Editors: Shipra Agrawal and Francesco Orabona

Abstract

We study the sample complexity of reducing reinforcement learning to a sequence of empirical risk minimization problems over the policy space. Such reductions-based algorithms exhibit local convergence in the function space, as opposed to the parameter space for policy gradient algorithms, and thus are unaffected by the possibly non-linear or discontinuous parameterization of the policy class. We propose a variance-reduced variant of Conservative Policy Iteration that improves the sample complexity of producing a ε -functional local optimum from $O(\varepsilon^{-4})$ to $O(\varepsilon^{-3})$. Under state-coverage and policy-completeness assumptions, the algorithm enjoys ε -global optimality after sampling $O(\varepsilon^{-2})$ times, improving upon the previously established $O(\varepsilon^{-3})$ sample requirement.

Keywords: variance reduction, reinforcement learning, non-convex optimization

1. Introduction

Reinforcement learning agents interact with the environment by adaptively executing actions with the goal of maximizing a cumulative long-term reward. A persistent challenge for such agents is operating in situations that involve large or continuous state spaces. Such large-scale Markov Decision Processes (MDPs) are accompanied by both the statistical challenge of generalization across states and the computational challenge of working with a large decision set, since often the policy class used in conjunction is too large to enumerate efficiently.

One approach to deal with these issues is to reduce reinforcement learning (RL) to a sequence of better-understood and easier learning problems like supervised learning (SL) or empirical risk minimization (ERM). Indeed, such an approach is well-studied (e.g. [Kakade and Langford \(2002\)](#); [Kakade \(2003\)](#); [Langford and Zadrozny \(2003\)](#); see [Agarwal et al. \(2021\)](#) for a modern treatment), and has since inspired widely-used practical variants ([Schulman et al., 2015, 2017](#); [Vieillard et al., 2020](#)). A typical algorithm of this kind incrementally updates the policy in small steps, using the solutions obtained from a SL/ERM oracle, to arrive at a policy that is a local minimum in the space of policies as opposed to a local minimum in the space of policy parameters. This *functional* local optimality is a stronger notion of local optimality which holds independently of how the policy class itself may be parameterized, and is hence distinct from local convergence in parameter space. See below for a more detailed comparison of these notions. Existing structural results (e.g. Theorem 14.4 in [Agarwal et al. \(2021\)](#)) then translate such functional local optimality to global optimality guarantees under various appropriate state-coverage assumptions.

In this paper we study and improve the sample complexity of such reductions-based algorithms. Specifically, consider the Conservative Policy Iteration (CPI) ([Kakade and Langford, 2002](#)) algorithm

which reduces RL to weighted multi-class classification problems over the policy space. For any policy class Π , the CPI algorithm guarantees upon sampling $O\left(\frac{\log|\Pi|}{\varepsilon^4}\right)$ ¹ transitions outputs a policy $\bar{\pi}$ which satisfies ε -local optimality in the function space. The notion of functional local optimality over a policy space Π can be defined by considering the gradient of the value function ∇V^π with respect to the policy π ². A policy $\bar{\pi}$ is an ε -functional local optimum if $\max_{\pi' \in \Pi} \langle \nabla V^{\bar{\pi}}, \pi' - \bar{\pi} \rangle \leq \varepsilon$.

Our main result (formalized as Theorem 2) is an algorithm that retains the above-stated functional local optimality guarantee while drawing $\tilde{O}\left(\frac{\log|\Pi|}{\varepsilon^3}\right)$ transitions. This reduction in sample complexity happens due to careful reuse of samples across successive calls to an ERM oracle, much like Experience Replay (Schaul et al., 2015). A crucial feature here is that this reuse of samples does not employ trajectory-wise (or marginal) importance weights whose size may be uncontrollable, but instead uses a momentum-like weighting scheme inspired by variance-reduction methods (Cutkosky and Orabona, 2019) for non-convex optimization.

Contrasting local optimality in parameter and function spaces. Typically, policy gradient methods guarantee convergence to local optima in the parameter space. We contrast these notions below.

1. Local optimality in function space is a stronger notion. Consider the policy class $\Pi = \{\pi(\theta) : \theta \in \Theta\}$ where $\pi(\theta) : \Theta \times \mathbb{R}^{S \times A}$ is a differentiable function and a possibly non-convex loss function $V : \mathcal{F} \rightarrow \mathbb{R}$. Via the application of a chain rule we see that $\nabla_\theta V(\pi(\theta)) = \nabla V^\pi \frac{\partial \pi(\theta)}{\partial \theta}$. It is now easy to verify that for any θ

$$\max_{\pi' \in \Pi} \langle \nabla V^{\pi(\theta)}, \pi' - \pi \rangle \leq \varepsilon \implies \|\nabla_\theta V(\pi(\theta))\| \leq O(\varepsilon)$$

for smoothly parameterized policy classes. Yet as the following example demonstrates the reverse implications does not hold true.

Example 1 Consider a singleton dataset $(x, y) = (1, 1)$, loss function $l(y, \hat{y}) = -(\hat{y} - y)^2/2$, and function class $\mathcal{F} = \{f_\theta(x) = \theta^2 x : \theta \in [-10, 10]\}$. $\theta = 0$ is a local optima in the parameter space, i.e. $\nabla_\theta l(y, f_\theta(x)) = -2(\theta^2 x - 1)\theta x = 0$, yet $f_\theta = 0$ is not a local optima in the function space. In particular, $\nabla_f l(y, f_0(x)) = 1$ and $\nabla l(y, f_0(x)) \times (f_1(x) - f_0(x)) = 1 \neq 0$.

2. For any composite optimization problem $\min_\theta V(\pi(\theta))$, local optimality in function space decouples the non-convexity induced by the policy parameterization $\pi(\theta)$ from the intrinsic non-convexity of the evaluation function V . Consequently, function-space algorithms (Mason et al., 1999) are unaffected when the policy class is non-linearly (like neural nets) or discontinuously (like decision trees) parameterized. Meanwhile, parameter-space algorithms like SGD have to contend with the local optima introduced by non-linear parameterizations of the function class even for convex loss functions.
3. In the context of RL, unlike parameter-space local convergence, functional local optimality guarantees readily translate to global guarantees under appropriate state-coverage assumptions. See below.

1. We handle infinite policy classes via the notion of covering numbers defined further in the paper.

2. As defined later in the paper we consider policies to be in the space $\mathbb{R}^{S \times A}$ and thus the gradient lives in that space too

	Without variance reduction	With variance reduction	Remarks
Local optimality in parameter space (Policy Gradient)	$\frac{1}{\varepsilon^4}$ (Ghadimi and Lan, 2013)	$\frac{1}{\varepsilon^3}$ (Xu et al., 2019; Shen et al., 2019)	Assuming a smoothly parameterized policy class and possibly a variance bound on trajectory-wise importance weights.
Local optimality in function space	$\frac{\log \Pi }{\varepsilon^4}$ (CPI; Kakade and Langford (2002))	$\frac{\log \Pi }{\varepsilon^3}$ (Theorem 2)	Holds regardless of the parameterization of the policy class Π ; without any state-coverage completeness assumption.
Global optimality given a restart distribution ν	$\frac{\log \Pi }{\varepsilon^4}$ (CPI; Kakade and Langford (2002))	$\frac{\log \Pi }{\varepsilon^3}$ (Corollary 5)	Assuming $D_\infty = \left\ \frac{d^{\pi^*}}{\nu} \right\ _\infty < \infty$; upto a policy-completeness residual term (Definition 3).
Global optimality given an exploratory policy class	$\frac{\log \Pi }{\varepsilon^3}$ (CPI; Kakade and Langford (2002), Brukhim et al. (2022))	$\frac{\log \Pi }{\varepsilon^2}$ (Theorem 8)	$C_\infty = \max_{\text{CH}(\pi) \in \Pi} \left\ \frac{d^{\pi^*}}{d^\pi} \right\ _\infty < \infty$; upto a policy-completeness residual term (Definition 6).

Table 1: The sample complexity of various optimality objectives, suppressing $\text{poly}(|A|, 1/(1-\gamma))$. Furthermore in the paper we don't assume that the policy class is finite but rather prove our results with $|\Pi_\varepsilon|$ which is the ε -covering number of Π . For exploratory policy classes, the stated sample complexity results from an improved iteration complexity of the CPI algorithm (Kakade and Langford, 2002) as observed in Brukhim et al. (2022), by directly appealing to the gradient domination characteristics of the value function.

Functional Local Optimality to Global Optimality Using standard local-to-global results in RL literature (e.g. Lemma 4 & Lemma 7 in this work, or Theorem 14.4 in Agarwal et al. (2021)), we translate our functional local optimality result to a $\tilde{O}\left(\frac{\log|\Pi|}{\varepsilon^3}\right)$ sample complexity for global optimality (upto a policy-completeness term) when the learner has access to an exploratory distribution μ with sufficient state coverage. This is formalized in Theorem 5 and the coverage assumptions are formalized in Definition 3.

Furthermore, when the policy class under consideration has sufficient state coverage by itself, we show that it is possible to skip this local-to-global approach, and directly guarantee an improved sample complexity of $\tilde{O}\left(\frac{\log|\Pi|}{\varepsilon^2}\right)$ for global convergence, which is tight even for the subcase of binary classification. This is formalized as Theorem 8 and the policy class coverage is formalized as Definition 6.

2. Related work

Reduction from reinforcement learning to supervised learning has been studied extensively, beginning with Kakade and Langford (2002); Kakade (2003); Langford and Zadrozny (2003); Bagnell et al. (2004). The local functional descent view of Conservative Policy Iteration (CPI) (Kakade and Langford, 2002) is presented in detail in Kakade (2003); Scherrer and Geist (2014). In particular, CPI runs for $1/\varepsilon^2$ rounds and requires an ε -accurate supervised learning oracle in each round, thereby needing $O(1/\varepsilon^4)$ samples in total.

In recent years, the technique of variance reduction (Schmidt et al., 2017; Defazio et al., 2014; Johnson and Zhang, 2013) has led to faster rates for convex optimization in the finite-sum setting. The idea behind these methods is to use a (mini-)batch of gradients at a point to help reduce the variance of the stochastic estimator at subsequent points. These techniques have since been combined with ideas such as acceleration (Shalev-Shwartz and Zhang, 2014; Lin et al., 2015; Allen-Zhu, 2017), and they have even led to improvements for non-convex problems (Allen-Zhu and Hazan, 2016; Reddi et al., 2016a) in terms of reaching ε -approximate stationary points, i.e., points with gradient norm bounded by ε . For the stochastic non-convex setting, variance reduced approaches have, in a series of works (Xu et al., 2018; Allen-Zhu and Li, 2018; Nguyen et al., 2017; Fang et al., 2018; Cutkosky and Orabona, 2019), led to improvements over the classical rate of $O(1/\varepsilon^4)$ (Ghadimi and Lan, 2013), ultimately reaching a rate of $O(1/\varepsilon^3)$, which is tight under a mean-squared smoothness property of the noisy gradient (Arjevani et al., 2019). Furthermore, such applications have proven useful in the context of constrained non-convex optimization, whereby similar improvements have been shown in (Reddi et al., 2016b; Zhang et al., 2020).

Following the analysis of SGD for nonconvex objectives (Ghadimi and Lan, 2013), it follows that REINFORCE (Sutton et al., 1999) converges to a local minimum in parameter space after $\mathcal{O}(1/\varepsilon^4)$ samples. A recent thread (Shen et al., 2019; Xu et al., 2019; Zhang et al., 2021; Papini et al., 2018) of research applies variance-reduction techniques to improve upon this to $O(1/\varepsilon^3)$.

3. Problem setting

We setup a few notations of common use, and then proceed to delineate the interaction model.

3.1. Basic definitions

A differentiable function $f : \mathcal{K} \rightarrow \mathbb{R}$ is defined to be L -smooth, over a domain \mathcal{K} with respect to a norm $\|\cdot\|$ and an inner product $\langle \cdot, \cdot \rangle$, if for every $x, y \in \mathcal{K}$ the inequality stated below holds. We note that the norm $\|\cdot\|$ need not necessarily be the canonical norm associated with the inner product $\langle \cdot, \cdot \rangle$.

$$|f(y) - f(x) - \langle \nabla f(x), y - x \rangle| \leq \frac{L}{2} \|x - y\|^2$$

For any $N \in \mathbb{N}$, let Δ_N be the unit simplex in N dimensions, i.e. $\Delta_N = \{x \in \mathbb{R}_+^N : \|x\|_1 = 1\}$. For matrices in $\mathbb{R}^{M \times N}$, define the dual-norm³ pair $\|\cdot\|_{1,\infty}$ and $\|\cdot\|_{\infty,1}$ as follows. Note that the latter norm $\|\cdot\|_{1,\infty}$ is not a vector-induced matrix norm.

$$\|A\|_{\infty,1} = \max_{1 \leq i \leq M} \sum_{j=1}^N |A_{i,j}|, \quad \|A\|_{1,\infty} = \sum_{i=1}^M \max_{1 \leq j \leq N} |A_{i,j}|.$$

Let $\mathcal{B}_{\infty,1}, \mathcal{B}_{1,\infty} \subset \mathbb{R}^{S \times A}$ (S, A correspond to the size of the state space and action space defined in the next section) denote the unit norm balls for the two norms over $S \times A$ matrices. For any element $\pi \in \mathcal{B}_{\infty,1}$, and any $s \in S$, we will use the notation $\pi(s) \in \mathbb{R}^{|A|}$ to denote the row corresponding to the state s . In subsequent sections we consider the decision set to be $\mathcal{K} = \Delta_A^S \subset \mathcal{B}_{\infty,1}$, which is the cartesian product of S unit simplices Δ_A , to represent the space of stochastic policies over state space \mathcal{S} and action set \mathcal{A} . Note that for any $\pi \in \Delta_A^S$, $\|\pi\|_{\infty,1} = \max_{s \in S} \sum_{a \in \mathcal{A}} \pi(a|s) = 1$.

For any $\varepsilon \geq 0$ and policy class Π , the ε -covering number $|\Pi_\varepsilon|$ is the minimum size of a set Π_ε such that for every $\pi \in \Pi$, there exists a $\pi' \in \Pi_\varepsilon$ such that $\|\pi - \pi'\|_{\infty,1} \leq \varepsilon$. Such covering numbers are typically scale logarithmically as function of ε and henceforth we assume the same. We define $\text{CH}(\Pi)$ to be the convex hull of Π .

3.2. Interaction model

A Markov Decision Process (MDP) is a decision making framework specified by state space \mathcal{S} with $|\mathcal{S}| \triangleq S$, action space \mathcal{A} with $|\mathcal{A}| \leq A$, a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, a transition kernel $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_A$, discount factor $\gamma \in [0, 1)$, and an initial state distribution $\rho \in \Delta_S$. Define an *effective* horizon of $H_\gamma = \frac{1}{1-\gamma}$. A stochastic policy $\pi \in \Delta_A^S$ prescribes a choice of actions at each state s as $a \sim \pi(\cdot|s)$. The execution of such a policy on the MDP induces a distribution over the space of trajectories, where each trajectory $\tau = (s_0, a_0, \dots)$ is a random variable sampled as

$$s_0 \sim \rho, a_i \sim \pi(\cdot|s_i), s_{i+1} = P(\cdot|s_i, a_i)$$

Thus, averaging with respect to this distribution, it is possible to ascribe a value to every state-action pair as

$$Q^\pi(s, a) \triangleq \mathbb{E}_\tau \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \middle| \pi, s_0 = s, a_0 = a \right] \text{ and } V^\pi(s) \triangleq \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(s, a) | \pi, s].$$

Further we will also consider a setting wherein the start state of the MDP, s_0 could be sampled from an arbitrary distribution μ . Under this setting we define the following quantities,

$$V_\mu^\pi \triangleq \mathbb{E}_{s \sim \mu} [V^\pi(s) | \pi] \text{ and } d_\mu^\pi(s, a) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s_0 \sim \mu} \mathbb{E} [\mathbf{1}_{s_t = s \wedge a_t = a} | \pi, s_0].$$

3. See Lemma 20 for a proof of the duality

V_μ^π and d_μ^π above captures the expected return and the (discounted) state distribution associated with a policy π when the initial state is sampled from a distribution μ . To ease the presentation, through the paper for quantities that depend on start state distributions, if the start state distribution is not explicitly specified, it is assumed to be the canonical start state distribution ρ . Thus, $V^\pi = V_\rho^\pi$ and $d^\pi = d_\rho^\pi$. Given a single argument s , we shall use the unary function $d_\mu^\pi(s)$ to denote the state-space marginal of $d_\mu^\pi(s, a)$, i.e. $d_\mu^\pi(s) = \sum_{a \in \mathcal{A}} d_\mu^\pi(s, a)$. Given a single state $s \in \mathcal{S}$, d_s^π represents the (discounted) state distribution achieved when starting from a fixed state s .

In subsequent sections, we consider two distinct ways in which a learner might interact with the MDP. In the μ -**reset model**, the learner has access to an exploratory restart distribution $\mu \sim \Delta_S$, and can draw a trajectory of finite length from the MDP starting with an initial state $s_0 \sim \mu$ sampled from μ . Nevertheless, the objective for the learner still remains to maximize $V^\pi = V_\rho^\pi$.

The other alternative considered is the **episodic model** where the learner can draw trajectories from the MDP with respect to the canonical start state distribution ρ .

3.3. Computational model

Since this work considers a reduction-based approach to RL, we assume that the learner is aided by the following computational oracles. Note that the tolerance parameter in the following definition scales naturally with the maximum possible loss.

Definition 1 (ERM Oracle) *Let $\mathcal{L} = \{x \rightarrow l^\top x : l \in \mathbb{R}^A\}$ be the class of linear loss function over the decision set Δ_A . Given a dataset $D = \{(s_i, l_i)\}_{i=1}^N$ where each example is in $\mathcal{S} \times \mathcal{L}$, and a tolerance $\varepsilon_{\text{ERM}} > 0$, the Empirical Risk Minimization (ERM) oracle ERM outputs a policy $\pi \in \Pi$ such that*

$$\sum_{i=1}^N l_i^\top \pi(s_i) \geq \max_{\pi^* \in \Pi} \sum_{i=1}^N l_i^\top \pi^*(s_i) - \varepsilon_{\text{ERM}} \cdot \sum_{i=1}^N \|l_i\|_\infty.$$

Previous approaches based on reducing RL to better understood subroutines sometimes also make use of a supervised learning oracle where in the dataset is replaced by a distribution over examples. Note that it is always possible to construct an ERM oracle (for any tolerance ε_{ERM}) using a supervised learning oracle, by resampling with replacement.

Through the main paper we use $\tilde{O}()$ to contain polynomial factors in problem constants in particular including $(1 - \gamma)^{-1}$, A . In the appendix we restate our main results including all such dependency.

4. Main results

Our main result is a variance-reduced algorithm (Algorithm 1) which we formally describe and explain in the next section. In this section we present the guarantees we prove. We begin by stating our main result in the μ -reset model for any start state distribution μ . Note that the episodic model is a natural sub-case whence the start state distribution is set to be ρ , the canonical start state distribution.

Theorem 2 (Local optimality in function space) *There exists an algorithm such that given any start state distribution μ input to the algorithm and any given ε, δ , the algorithm produces a policy $\bar{\pi}$ which satisfies the following with probability $1 - \delta$,*

$$\max_{\pi' \in \Pi} \langle \nabla V_\mu^{\bar{\pi}}, \pi' - \bar{\pi} \rangle \leq \varepsilon$$

Furthermore the algorithm samples at most $\tilde{O}(\log(|\Pi_\epsilon|)/\epsilon^3)$ episodes of expected length $O\left(\frac{1}{1-\gamma}\right)$ from the MDP starting from the start state distribution μ .

The above theorem improves the sample complexity of such functional local convergence from the best known $\tilde{O}(\epsilon^{-4})$ for the CPI algorithm to $\tilde{O}(\epsilon^{-3})$. As highlighted in Table 1, such improvements only exist for local optimality in parameter space which can be a significantly weaker guarantee.

Global Optimality Results Next, we demonstrate how such functional local optimality may be translated to a statement on global convergence with certain state-coverage and policy-completeness assumptions which we define next. These local-to-global translation lemmas (Lemma 4; proven in the appendix for completeness) were first noted in Kakade and Langford (2002), along with Scherrer and Geist (2014) who formally introduced the policy completeness notion. Our contribution (Corollary 5) here is an improvement in the dependence of the sample complexity on ϵ as a consequence of an improved local functional optimality result (Theorem 2).

Definition 3 Let π^* be an optimal policy for the MDP in consideration. Given μ , a start-state distribution the learner can draw from, define the distribution mismatch coefficient D_∞^μ as stated below. Further, given a policy class Π , define $\epsilon_{\Pi,\mu}$ as a quantitative measure of policy completeness.

$$D_\infty^\mu \triangleq \left\| \frac{d^{\pi^*}}{\mu} \right\|_\infty, \quad \epsilon_{\Pi,\mu} = \max_{\pi \in \text{CH}(\Pi)} \min_{\pi^* \in \Pi} \mathbb{E}_{s \sim d_\mu^\pi} \left[\max_{a \in A} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi^*(s) \right].$$

The distribution mismatch coefficient measures how exploratory the restart distribution μ is, and the associated policy completeness notion is a policy analogue of inherent bellman error (Munos and Szepesvári, 2008). The latter measures the degree to which a policy in Π can best approximate the bellman optimality operator in an average sense with respect to the state distribution. Under the above definitions, we show the following lemma, which relates the global optimality gap of a policy to the local optimality measure of the policy.

Lemma 4 For any state distribution μ and any policy $\pi \in \text{CH}(\Pi)$, the following holds

$$V^* - V^\pi \leq \frac{D_\infty^\mu}{1-\gamma} \left(\max_{\pi' \in \Pi} \langle \nabla V_\mu^\pi, \pi' - \pi \rangle + \frac{1}{1-\gamma} \epsilon_{\Pi,\mu} \right).$$

In particular the above lemma shows that if one has access to a start state distribution with state coverage (i.e. D_∞^μ is small) local optimality implies global optimality upto the policy completeness measure. We now provide sample complexity bounds for achieving global optimality via the following corollary of Theorem 2 which translates the functional local optimality guarantee stated in the latter to the associated global optimality guarantee under appropriate policy coverage assumptions when the learner has access to an exploratory restart distribution μ , with the aid of Lemma 4.

Corollary 5 (Global optimality) There exists an algorithm such that given any start state distribution μ input to the algorithm and any given ϵ, δ , the algorithm produces a policy $\bar{\pi}$ which satisfies the following with probability $1 - \delta$,

$$V^* - V^{\bar{\pi}} \leq \epsilon + \frac{D_\infty^\mu \epsilon_{\Pi,\mu}}{(1-\gamma)^2}.$$

Furthermore the algorithm samples at most $\tilde{O}\left(\frac{(D_\infty^\mu)^3 \log |\Pi_\epsilon|}{\epsilon^3}\right)$ episodes of expected length $O\left(\frac{1}{1-\gamma}\right)$ from the MDP starting from the start state distribution μ .

Improved rates with an exploratory policy class Next, we provide improved rates of convergence to global optimality under the assumption that the underlying policy class by itself has sufficient overlap with the state distribution of an optimal policy. *Under this assumption we no longer require access to a state-distribution with sufficient coverage, so we state the results in the more restricted episodic setting, where every sample episode begins at a state sampled from the MDP’s canonical start state distribution ρ .* For brevity we define $\epsilon_\Pi = \epsilon_{\Pi, \rho}$ to be the policy completeness with respect to the canonical start state distribution ρ . Formally, the policy class’ overlap condition is captured by the following definition.

Definition 6 Consider a policy class Π . Let π^* be an optimal policy for the MDP in consideration. Define the distribution mismatch coefficient as $C_\infty = \max_{\pi \in \text{CH}(\Pi)} \left\| \frac{d\pi^*}{d\pi} \right\|_\infty$.

In such cases, the value function is approximately gradient dominated as the following lemma shows and it is possible to forego the procedure of arriving at a local optimality guarantee before making claims on the global suboptimality.

Lemma 7 For any $\pi \in \text{CH}(\Pi)$, the following holds

$$V^* - V^\pi \leq C_\infty \left(\max_{\pi' \in \Pi} \langle \nabla V^\pi, \pi' - \pi \rangle + \frac{1}{1-\gamma} \epsilon_\Pi \right).$$

A more direct analysis yields the following result which improves the sample complexity to be scaling as $\tilde{O}(\varepsilon^{-2})$.

Theorem 8 (Faster global optimality) There exists an algorithm such that given any ε, δ , the algorithm produces a policy $\bar{\pi}$ which satisfies the following with probability $1 - \delta$,

$$V^* - V^{\bar{\pi}} \leq \varepsilon + \frac{C_\infty \epsilon_\Pi}{1-\gamma}.$$

Furthermore the algorithm samples at most $\tilde{O}(C_\infty^2 \log(|\Pi_\varepsilon|)/\varepsilon^2)$ episodes of expected length $O\left(\frac{1}{1-\gamma}\right)$ sampled from the MDP starting at the canonical start state distribution ρ .

5. Algorithms

5.1. Variance-reduced Conservative Policy Iteration

In this section we present our main algorithm **Variance-reduced Conservative Policy Iteration**. The algorithm is formally described as Algorithm 1. The algorithm takes as input a start state distribution μ for the MDP and assumes access to an ERM oracle (Definition 1) over the policy class Π . The algorithm is parameterized by parameters $\eta, \lambda, \gamma, T, \varepsilon_{\text{ERM}}$, wherein notably ε_{ERM} is the accuracy target for the ERM oracle calls and the algorithm samples a total of $3T$ episodes from the MDP starting from the start state distribution μ . The algorithm makes use of two important subroutines the Q-sampler (Algorithm 2) and H-sampler (Algorithm 3) to compute unbiased estimates of linear forms in the functional policy gradient and quadratic forms in the functional policy Hessian respectively. These sub-routines are formally stated in the appendix (Algorithms 2 and 3) and we state their properties in this section. We now proceed to explain the construction of Algorithm 1.

Conservative Policy Iteration The core structure of our proposed algorithm follows the idea proposed by the Conservative Policy Iteration(CPI) algorithm (Kakade and Langford, 2002) which maintains a policy π_t . At every step a new candidate policy π'_t is obtained via the following

$$\pi'_t = \operatorname{argmin} \langle \nabla V^{\pi_t}, \pi \rangle. \quad (5.1)$$

Note that ∇V^{π_t} is the functional policy gradient. The following lemma from Sutton et al. (1999) provides a concrete estimator for the functional policy gradient.

Lemma 9 [Sutton et al. (1999)] For any policy pair π, π' and start-state distribution μ , we have

$$\langle \nabla V_{\mu}^{\pi}, \pi' \rangle = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\mu}^{\pi}} \left[Q^{\pi}(s, \cdot)^{\top} \pi'(s) \right].$$

Using the above, it can be observed that the minimization problem (5.1) can be cast as a supervised learning problem over the policy space. CPI obtains the next policy via a convex combination $\pi_{t+1} = \eta\pi_t + (1-\eta)\pi'_{t+1}$. This step can be interpreted as a step of the Frank-Wolfe algorithm (Frank and Wolfe, 1956) in the policy space. Via standard convergence analyses (see eg. Hazan et al. (2016)) one can expect to require a number of iterations scaling with ε^{-2} . To solve the per-step optimization problem, the sample complexity scales with ε^{-2} , leading to a total sample complexity scaling with ε^{-4} . In order to improve sample complexity, we instead employ a variance-reduced estimator of the gradient ∇V^{π_t} .

Variance Reduction The variance-reduced estimator of the functional policy gradient we use is an adaptation of the momentum based low-variance estimator STORM proposed by Cutkosky and Orabona (2019). For a stochastic function $F(x) = \mathbb{E}[f(x, \xi)|x]$, Cutkosky and Orabona (2019) propose the following estimator

$$v_t \triangleq (1-\lambda)v_{t-1} + \lambda \nabla f(x_t, \xi_t) + (1-\lambda)(\nabla f(x_t, \xi_t) - \nabla f(x_{t-1}, \xi_t)) \quad (5.2)$$

The proposed estimator runs a running average of the previous stochastic gradient estimates (akin to momentum) with a key addition of a term accounting for gradient differences viz. $\nabla f(x_t, \xi_t) - \nabla f(x_{t-1}, \xi_t)$. It is critical that the random variable ξ_t is the same on both gradient estimates in the above expression. It is therefore tricky to estimate the gradient difference term for RL applications as the distribution of the stochastic gradient depends on the current policy (trajectories are sampled from the current policy), an issue highlighted in the previous works (Xu et al., 2019; Shen et al., 2019). We avoid the issue by noting that this term can also be estimated via the following expression involving the Hessian of f ,

$$\nabla f(x_t, \xi_t) - \nabla f(x_{t-1}, \xi_t) \sim \nabla^2 f((1-b)x_t + bx_{t-1}, \xi_t)(x_t - x_{t-1}), \quad (5.3)$$

where b is uniformly sampled from $[0, 1]$. We employ this Hessian-based approach for our problem. We note that this *correction* term (either computed via gradient difference or a Hessian-vector product) is necessary in a sense for reducing the complexity of stochastic optimization in general (Arjevani et al., 2019, 2020). The Hessian based estimator requires building estimates for the functional Hessian-vector products (in the policy space) for which we provide a sub-routine H-Sampler (Algorithm 3). We now proceed to describe the construction of our algorithm.

Input: Initial state distribution μ , ERM oracle ERM , parameters $\eta, \lambda, \gamma, T, \varepsilon_{\text{ERM}}$.

Choose an initial policy $\pi_0 = \pi_1 \in \Pi$ arbitrarily.

Create an empty dataset D_0 supported over state-linear-loss pairs $\{(s_i, \widehat{l}_i) : s_i \in S, \widehat{l}_i \in \mathbb{R}^A\}$.

for $t = 1$ **to** T **do**

$(s_t, \widehat{Q}_t) \in S \times \mathbb{R}^A \leftarrow$ Q-sampler (Algorithm 2) with inputs $\pi_t, \mu, \varepsilon_Q$.

Sample $b \sim \text{Unif}([0, 1])$, and define a policy $\bar{\pi}_t = (1 - b)\pi_t + b\pi_{t-1}$.

$(s'_t, s''_t, \widehat{H}_t) \in S \times S \times \mathbb{R}^{A \times A} \leftarrow$ H-sampler (Algorithm 3) with inputs $\bar{\pi}_t, \mu, \varepsilon_H$.

Create a new dataset D_t by first multiplying each existing loss vector in D_{t-1} by $(1 - \lambda)$:

$$D_t = \{(s_i, (1 - \lambda)l_i) : (s_i, l_i) \in D_{t-1}\},$$

and then appending the following three tuples to the said new dataset D_t .

$$\left(s_t, \widehat{l}_{t,1} = \frac{\lambda}{1 - \gamma} \widehat{Q}_t \right), \quad \left(s'_t, \widehat{l}_{t,2} = \frac{\gamma(1 - \lambda)}{(1 - \gamma)^2} \widehat{H}_t (\pi_t(s''_t) - \pi_{t-1}(s''_t)) \right), \\ \left(s''_t, \widehat{l}_{t,3} = \frac{\gamma(1 - \lambda)}{(1 - \gamma)^2} \widehat{H}_t (\pi_t(s'_t) - \pi_{t-1}(s'_t)) \right).$$

Call the ERM oracle with tolerance ε_{ERM} on the dataset D_t to obtain $\pi'_t = \text{ERM}(D_t)$.

Update $\pi_{t+1} = (1 - \eta)\pi_t + \eta\pi'_t$.

Empirically compute local improvement of π'_t in comparison to π_t with respect to D_t as

$$\widehat{A}_t = \sum_{(s, \widehat{l}) \in D_t} \widehat{l}^\top (\pi'_t(s) - \pi_t(s)).$$

end

Option 1: return $\bar{\pi} = \pi_{t'}$ where $t' = \text{argmin}_{T/2 \leq t \leq T} \widehat{A}_t$.

Option 2: return $\bar{\pi} = \pi_T$.

Algorithm 1: Variance-reduced Conservative Policy Iteration

Construction of the algorithm Our overall algorithm is based on building an estimator of ∇V^{π_t} akin to the v_t estimator from 5.2. With such an estimator we wish to solve the minimization problem

$$\pi'_t = \text{argmin}_{\pi \in \Pi} \langle v^t, \pi \rangle. \quad (5.4)$$

We approach the above problem via reduction to ERM problem over the policy class and solve it via mapping the estimators to datasets of (s, l) pairs over which we solve the ERM problem. Concretely at every step of Algorithm 1 we call the Q-sampler which we show to have the following guarantee

Lemma 10 *Q-sampler (Algorithm 2) when run with a policy π and start-state distribution μ produces a random tuple $(s, \widehat{Q}) \in S \times \mathbb{R}^A$ with the distribution $s \sim d_\mu^\pi$, and $\mathbb{E} [\widehat{Q} | s, \pi] = Q^\pi(s, \cdot)$. This in particular implies that for any π' ,*

$$\langle \nabla V_\mu^\pi, \pi' \rangle = \frac{1}{1 - \gamma} \mathbb{E}_{(s, \widehat{Q}) \sim \mathcal{A}} [\widehat{Q}^\top \pi'(s)].$$

Furthermore, $\|\widehat{Q}\|_1 \leq \frac{|A|}{1 - \gamma}$, and the expected length of sampled episode is at most $\frac{3}{1 - \gamma}$.

Input: Policy π , and start-state distribution μ , tolerance ε_Q .

Sample the initial state $s_0 \sim \mu$.

for $t = 0, \dots, \infty$ **do**

 With probability $1 - \gamma$, record the current state as s_t and exit the loop.

 Else, draw an action $a_t \sim \pi(s_t)$ and transition to the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$.

end

Draw the t^{th} action $a_t \sim \text{Unif}(A)$ uniformly, and observe the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$.

for $t' = t + 1, \dots, \infty$ **do**

 With probability $1 - \gamma$, record $R = \frac{r(s_{t'}, a_{t'})}{1 - \gamma}$, and construct the vector in \mathbb{R}^A as

$$\widehat{Q}(a) = \mathbb{I}_{a=a_t} AR$$

 and return the tuple (s_t, \widehat{Q}) ending the algorithm's execution.

 Else, draw an action $a_{t'} \sim \pi(s_{t'})$ to transition to the next state $s_{t'+1} \sim P(\cdot | s_{t'}, a_{t'})$.

end

Algorithm 2: Q-sampler

Therefore we include s_t and a scaled version of \widehat{Q}_t in our dataset. To see why the scaling is $\frac{\lambda}{(1-\gamma)}$, note that the $\frac{1}{(1-\gamma)}$ from the scaling in the gradient expression above and the scaling of λ comes from the usage in the estimator (5.2). Similarly to estimate the Hessian term akin to (5.3) in the estimator, we mix π_{t-1}, π_t to obtain $\bar{\pi}_t$ and we call the H-Sampler subroutine. For the H-Sampler subroutine we show the following guarantee

Lemma 11 *H-sampler (Algorithm 3) when run with a policy π and start-state distribution μ produces a random tuple $(s, s', \widehat{H}) \in \mathcal{S} \times \mathcal{S} \times \mathbb{R}^{A \times A}$ such that for any policy pair π', π'' ,*

$$\langle \pi'', \nabla V_{\mu}^{\pi} \pi' \rangle = \frac{\gamma}{(1-\gamma)^2} \mathbb{E}_{(s, s', \widehat{H}) \sim \mathcal{A}} \left[\pi''(s)^\top \widehat{H} \pi'(s') + \pi'(s)^\top \widehat{H} \pi''(s') \right].$$

Furthermore, $\sum_{a, a' \in \mathcal{A} \times \mathcal{A}} |\widehat{H}(a, a')| \leq \frac{A^2}{1-\gamma}$ and the expected episode length is at most $\frac{5}{1-\gamma}$.

Further we account for the scalings arising both from the Hessian guarantee and the usage in the estimator when including the \widehat{H}_t into our dataset. As a result we see that at all times t our dataset construction satisfies the following lemma

Lemma 12 *For all $t \geq 0$, let $v_t \in \mathbb{R}^{S \times A}$ be a vector defined recursively such that for any $\pi \in \mathcal{B}_{\infty, 1}$, such that $\langle v_0, \pi \rangle = 0$, and*

$$\begin{aligned} \langle v_t, \pi \rangle &\triangleq (1 - a) \langle v_{t-1}, \pi \rangle + \frac{\lambda}{1 - \gamma} \widehat{Q}_t^\top \pi(s_t) \\ &+ \frac{\gamma(1 - \lambda)}{(1 - \gamma)^2} \left((\widehat{H}_t(\pi_t(s'_t)) - \pi_{t-1}(s'_t))^\top \pi(s'_t) + (\widehat{H}_t(\pi_t(s'_t)) - \pi_{t-1}(s'_t))^\top \pi(s''_t) \right). \end{aligned}$$

We have that for all times $t \leq T$, and for any policy π , the ERM loss for π on D_t can be expressed as

$$\langle v_t, \pi \rangle = \sum_{(s, l) \in D_t} \widehat{l}^\top \pi(s).$$

The above lemma ensures that all times t , π'_t is the solution of the minimization problem described in (5.4) for an estimator v_t which by using Lemma 10 and Lemma 11 can readily be seen to be the same as the STORM estimator 5.2 but in the functional space.

5.2. H-Sampler

Input: Policy π , and start-state distribution μ , tolerance ε_H .

Sample the initial state $s_0 \sim \mu$.

for $t = 0, \dots, \infty$ **do**

 With probability $1 - \gamma$, record the current state as s_t and exit the loop.

 Else, draw an action $a_t \sim \pi(s_t)$ and transition to the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$.

end

Draw the t^{th} action $a_t \sim \text{Unif}(A)$ uniformly, and observe the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$.

for $t' = t + 1, \dots, \infty$ **do**

 With probability $1 - \gamma$, record the current state as $s_{t'}$ and exit the loop.

 Else, draw an action $a_{t'} \sim \pi(s_{t'})$ and transition to the next state $s_{t'+1} \sim P(\cdot | s_{t'}, a_{t'})$.

end

Draw the t'^{th} action $a_{t'} \sim \text{Unif}(A)$ uniformly, and observe the next state $s_{t'+1} \sim P(\cdot | s_{t'}, a_{t'})$.

for $t'' = t' + 1, \dots, \infty$ **do**

 With probability $1 - \gamma$, record $R = \frac{r(s_{t''), a_{t''}}}{1 - \gamma}$ construct the matrix in $\mathbb{R}^{|A| \times |A|}$ as

$$\widehat{H}(a_1, a_2) = \mathbb{I}_{a_1 = a_t \wedge a_2 = a_{t'}} |A|^2 R$$

 and return the tuple $(s_t, s_{t'}, \widehat{H})$ ending the algorithm's execution.

 Else, draw an action $a_{t''} \sim \pi(s_{t''})$ to transition to the next state $s_{t''+1} \sim P(\cdot | s_{t''}, a_{t''})$.

end

Algorithm 3: H-sampler

The Q-Sampler and H-Sampler, which are importance sampling based estimators, to are stated as algorithms 2 and 3 respectively. We highlight the salient aspects of the proposed H-Sampler which we believe to be of independent interest. To obtain the H-Sampler, similar to the case of gradient in Lemma 9, we provide an explicit characterization of the functional policy Hessian. To define the Hessian, we make use of the notion of a future advantage $F^\pi(s, a | \pi')$ of a policy π' with respect to a baseline policy π , when starting from some state s and action a . Intuitively, it represents the value of playing one step of a candidate policy π' at a random step (geometrically distributed) in the future when starting from a state-action pair (s, a) , all the while executing a baseline policy π .

Definition 13 For any policy pair π, π' , define the future advantage $F^\pi(\cdot | \pi') : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ of a policy π' with respect to a baseline policy π as

$$F^\pi(s, a | \pi') = \mathbb{E}_{s' \sim P(\cdot | s, a)} \mathbb{E}_{s'' \sim d_{s'}^\pi} \left[Q^\pi(s'', \cdot)^\top \pi'(s'') \right].$$

Note that the future advantage $F^\pi(\cdot | \pi')$ is linear in π' . Lemma 14 provides a characterization of the functional Hessian of the value function as a bi-linear operator over π', π'' . The interchangeability of the roles of π' and π'' ensures the symmetry of the bi-linear form.

Lemma 14 For any policy triplet π, π', π'' and start-state distribution μ , we have

$$\langle \pi'', \nabla^2 V_\mu^\pi \pi' \rangle = \frac{\gamma}{(1 - \gamma)^2} \left(\mathbb{E}_{s \sim d_\mu^\pi} \left[F^\pi(s, \cdot | \pi')^\top \pi''(s) + F^\pi(s, \cdot | \pi'')^\top \pi'(s) \right] \right).$$

The following lemma shows the main guarantee for our H-Sampler, which immediately implies Lemma 11 using Lemma 14.

Lemma 15 *H-sampler (Algorithm 3) when run with a policy π and start-state distribution μ produces a random tuple $(s, s', \widehat{H}) \in \mathcal{S} \times \mathcal{S} \times \mathbb{R}^{A \times A}$ such that $s \sim d_\mu^\pi$, and for any π'*

$$\mathbb{E} \left[\widehat{H}^\top \pi'(s') \middle| s, \pi \right] = F^\pi(s, \cdot | \pi').$$

6. Overview of Analysis

Due to space constraints, we defer the analysis and proofs of the theorem entirely to the appendix, where the theorems are restated with the correct parameter instantiations. We provide a high level summary of the analysis approach here. Overall, the core of our algorithmic approach and analysis resembles the one-sample stochastic Frank-Wolfe algorithm proposed by Zhang et al. (2020) which also employs the STORM estimator for variance reduction in stochastic optimization. However the RL setting and especially performing the variance reduction in functional space brings some unique challenges which we tackle in our analysis. In particular, the functional(policy) space is bounded in ∞ -norm with gradients bounded in 1-norm. Thereby, the variance reduction properties of the STORM algorithm which are naturally stated in ℓ_2 norms need to be extended to $\infty, 1$ norms. To this end we provide an alternative analysis of STORM which bounds the deviation between the estimator and the true gradient with *high probability*(as opposed to smaller variance) over a covering set of the policy space. We believe this alternative analysis extending STORM to $\infty, 1$ norms and establishing high probability guarantees can be of independent interest. Furthermore to construct the estimator in the functional settings for RL, we devise novel functional Hessian-vector product oracle, which requires developing a sampling based expression for the Hessian-vector product (summarized in Lemma 14).

7. Conclusion

We revisit the problem of reducing reinforcement learning to a sequence of ERM problems. Using ideas from variance reduction in stochastic optimization, we improve the sample complexity of achieving a functional local optimum in policy space from $O(\varepsilon^{-4})$ to $O(\varepsilon^{-3})$. As we discuss, functional local optimum guarantees can be significantly stronger than parameter space local optimum guarantees, which we demonstrate by translating our improved sample complexity results for functional local optimum to improved bounds for global optimality under state coverage assumptions.

References

- Alekh Agarwal, Nan Jiang, Sham M. Kakade, and Wen Sun. *Reinforcement Learning: Theory and Algorithms*. 2021.
- Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. *The Journal of Machine Learning Research*, 18(1):8194–8244, 2017.
- Zeyuan Allen-Zhu and Elad Hazan. Variance reduction for faster non-convex optimization. In *International conference on machine learning*, pages 699–707. PMLR, 2016.

- Zeyuan Allen-Zhu and Yuanzhi Li. Neon2: Finding local minima via first-order oracles. *Advances in Neural Information Processing Systems*, 31, 2018.
- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365*, 2019.
- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Ayush Sekhari, and Karthik Sridharan. Second-order information in non-convex stochastic optimization: Power and limitations. In *Conference on Learning Theory*, pages 242–299. PMLR, 2020.
- J. Bagnell, Sham M Kakade, Jeff Schneider, and Andrew Ng. Policy search by dynamic programming. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2004. URL <https://proceedings.neurips.cc/paper/2003/file/3837a451cd0abc5ce4069304c5442c87-Paper.pdf>.
- Nataly Brukhim, Elad Hazan, and Karan Singh. A boosting approach to reinforcement learning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=GNHyNOR8Sn>.
- Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex sgd. *Advances in neural information processing systems*, 32, 2019.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in Neural Information Processing Systems*, 27, 2014.
- Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. *Advances in Neural Information Processing Systems*, 31, 2018.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems*, 26, 2013.
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer, 2002.
- Sham Machandranath Kakade. *On the sample complexity of reinforcement learning*. University of London, University College London (United Kingdom), 2003.
- John Langford and Bianca Zadrozny. Reducing t-step reinforcement learning to classification. 2003.

- Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. *Advances in Neural Information Processing Systems*, 28, 2015.
- Llew Mason, Jonathan Baxter, Peter Bartlett, and Marcus Frean. Boosting algorithms as gradient descent. *Advances in neural information processing systems*, 12, 1999.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(5), 2008.
- Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pages 2613–2621. PMLR, 2017.
- Matteo Papini, Damiano Binaghi, Giuseppe Canonaco, Matteo Pirodda, and Marcello Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pages 4026–4035. PMLR, 2018.
- Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pages 314–323. PMLR, 2016a.
- Sashank J Reddi, Suvrit Sra, Barnabás Póczos, and Alex Smola. Stochastic frank-wolfe methods for nonconvex optimization. In *2016 54th annual Allerton conference on communication, control, and computing (Allerton)*, pages 1244–1251. IEEE, 2016b.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- Bruno Scherrer and Matthieu Geist. Local policy search in a convex space and conservative policy iteration as boosted policy search. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer, 2014.
- Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1):83–112, 2017.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *International Conference on Machine Learning*, pages 64–72. PMLR, 2014.
- Zebang Shen, Alejandro Ribeiro, Hamed Hassani, Hui Qian, and Chao Mi. Hessian aided policy gradient. In *International conference on machine learning*, pages 5729–5738. PMLR, 2019.

- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.
- Nino Vieillard, Olivier Pietquin, and Matthieu Geist. Deep conservative policy iteration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6070–6077, 2020.
- Pan Xu, Felicia Gao, and Quanquan Gu. Sample efficient policy gradient methods with recursive variance reduction. *arXiv preprint arXiv:1909.08610*, 2019.
- Yi Xu, Rong Jin, and Tianbao Yang. First-order stochastic algorithms for escaping from saddle points in almost linear time. *Advances in neural information processing systems*, 31, 2018.
- Junyu Zhang, Chengzhuo Ni, Zheng Yu, Csaba Szepesvari, and Mengdi Wang. On the convergence and sample efficiency of variance-reduced policy gradient method. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=Re_VXF0yyO.
- Mingrui Zhang, Zebang Shen, Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. One sample stochastic frank-wolfe. In *International Conference on Artificial Intelligence and Statistics*, pages 4012–4023. PMLR, 2020.

Appendix A. Detailed statements of theorems

A.1. Theorem 2

Theorem 16 (Theorem 2 detailed) For a given ε, δ define $\varepsilon_{\text{cover}} \triangleq \frac{\varepsilon(1-\gamma)^2}{80A}$ and define the function

$$C(T) \triangleq \frac{8A^{3/2} \sqrt{\gamma \log(2T|\Pi_{\varepsilon_{\text{cover}}}|/\delta)}}{(1-\gamma)^{5/2}}.$$

Then Algorithm 1 when run with any parameters η, T satisfying the following equations,

$$\eta \leq \frac{\varepsilon(1-\gamma)^3}{40\gamma} \quad \eta T \geq \frac{(1-\gamma)}{2\gamma A} \log\left(\frac{1}{20\varepsilon(1-\gamma)^2}\right) \quad \frac{2}{(1-\gamma)\eta T} + 5C(T) \cdot \sqrt{\eta} \leq 3\varepsilon/10 \quad (\text{A.1})$$

and $\varepsilon_{\text{ERM}} = \frac{\varepsilon(1-\gamma)^2}{60A}$, then given any start state distribution μ input to the algorithm, the algorithm produces a policy $\bar{\pi}$ which satisfies the following with probability $1 - \delta$,

$$\max_{\pi' \in \Pi} \langle \nabla V_{\mu}^{\bar{\pi}}, \pi' - \bar{\pi} \rangle \leq \varepsilon.$$

Further there exists a setting of $T = \tilde{O}(\log(|\Pi_{\varepsilon}|)A^3(1-\gamma)^{-6}\varepsilon^{-3})$ such that the conditions (A.1) can be satisfied and therefore the algorithm samples at most $\tilde{O}(\log(|\Pi_{\varepsilon}|)A^3(1-\gamma)^{-6}\varepsilon^{-3})$ episodes of expected length $O\left(\frac{1}{1-\gamma}\right)$ from the MDP starting from the start state distribution μ .

In the above theorem, \tilde{O} hides polylogarithmic factors in the relevant parameters.

A.2. Theorem 8

Theorem 17 (Theorem 8 detailed) For a given ε, δ define $\varepsilon_{\text{cover}} \triangleq \frac{\varepsilon(1-\gamma)^2}{80A}$ and define the function

$$C(T) \triangleq \frac{8A^{3/2} \sqrt{\gamma \log(2T|\Pi_{\varepsilon_{\text{cover}}}|/\delta)}}{(1-\gamma)^{5/2}}.$$

Then Algorithm 1 when run with any parameters η, T satisfying the following equations,

$$\eta T \geq 2C_{\infty} \log\left(\frac{10}{\varepsilon(1-\gamma)}\right) \quad \eta \log\left(\frac{2T\Pi_{\varepsilon_{\text{cover}}}}{\delta}\right) \leq \frac{\varepsilon^2(1-\gamma)^5}{6400A^3} \quad \eta \leq \frac{\varepsilon(1-\gamma)^3}{40\gamma C_{\infty}} \quad (\text{A.2})$$

and $\varepsilon_{\text{ERM}} = \frac{\varepsilon(1-\gamma)^2}{20AC_{\infty}}$, then starting from the canonical start state distribution ρ , the algorithm produces a policy $\bar{\pi}$ which satisfies the following with probability $1 - \delta$,

$$V^* - V^{\bar{\pi}} \leq \varepsilon + \frac{C_{\infty}\varepsilon\Pi}{1-\gamma}.$$

Further there exists a setting of $T = \tilde{O}(C_{\infty}^2 \log(|\Pi_{\varepsilon}|)A^3(1-\gamma)^{-5}\varepsilon^{-2})$ such that the conditions (A.1) can be satisfied and therefore the algorithm samples at most $\tilde{O}(C_{\infty}^2 \log(|\Pi_{\varepsilon}|)A^3(1-\gamma)^{-5}\varepsilon^{-2})$ episodes of expected length $O\left(\frac{1}{1-\gamma}\right)$ from the MDP starting from the start state distribution ρ .

In the above theorem, \tilde{O} hides polylogarithmic factors in the relevant parameters.

Appendix B. Proofs of the Main Results

B.1. Proof of Theorem 16

Proof [Proof of Theorem 16] Let us first observe that value function V_μ^π is a smooth over the space of policies. The following statement holds independently of how the policy class itself may be parameterized.

Lemma 18 *For any start-state distribution μ , V_μ^π is $\frac{\gamma}{(1-\gamma)^3}$ -smooth in the $\|\cdot\|_{\infty,1}$ norm, i.e. for any two policies π, π' , we have that*

$$|V_\mu^{\pi'} - V_\mu^\pi - \langle \nabla V_\mu^\pi, \pi' - \pi \rangle| \leq \frac{\gamma}{(1-\gamma)^3} \|\pi' - \pi\|_{\infty,1}^2$$

Further for any two policies π, π' and any starting distribution μ we have that

$$\langle \nabla V_\mu^{\pi'} - \nabla V_\mu^{\pi''}, \pi \rangle \leq \frac{2\gamma}{(1-\gamma)^3} \|\pi' - \pi''\|_{\infty,1}$$

We invoke smoothness of V_μ^π , as Lemma 18 certifies, to observe that since successive iterates are close in the $\|\cdot\|_{\infty,1}$ norm, we have

$$\begin{aligned} V_\mu^{\pi_{t+1}} &= V_\mu^{\pi_t + \eta(\pi'_t - \pi_t)} \\ &\geq V_\mu^{\pi_t} + \eta \langle \nabla V_\mu^{\pi_t}, \pi'_t - \pi_t \rangle - \frac{\gamma\eta^2}{(1-\gamma)^3} \|\pi'_t - \pi_t\|_{\infty,1}^2. \end{aligned} \quad (\text{B.1})$$

Next, we wish to use the fact that π'_t was chosen by a ERM oracle, and therefore approximately maximizes the inner product with the gradient of the value function. To do this, we first relate the ERM objective (as in Algorithm 1) to the said gradient. This result supplants Lemma 12.

Theorem 19 *For all $t \geq 0$, define a sequence of vectors $v_t \in \mathbb{R}^{S \times A}$ recursively as follows. Let v_0 be any vector such that for all $\pi \in \mathcal{B}_{\infty,1}$, we have that $\langle v_0, \pi \rangle = 0$. Further for any $\pi \in \mathcal{B}_{\infty,1}$, $t > 0$, let v_t be an vector satisfying,*

$$\begin{aligned} \langle v_t, \pi \rangle &\triangleq (1-\lambda) \langle v_{t-1}, \pi \rangle + \frac{\lambda}{1-\gamma} \widehat{Q}_t^\top \pi(s_t) \\ &\quad + \frac{\gamma(1-\lambda)}{(1-\gamma)^2} \left((\widehat{H}_t(\pi_t(s'_t)) - \pi_{t-1}(s'_t))^\top \pi(s'_t) + (\widehat{H}_t(\pi_t(s'_t)) - \pi_{t-1}(s'_t))^\top \pi(s''_t) \right). \end{aligned}$$

Here s_t, s'_t, s''_t are the sequence of states produced by the algorithm. We have that for all $t \geq 0$ the dataset D_t maintained by Algorithm 1 satisfies the property that for any policy π , the ERM loss for π on D_t can be expressed as

$$\sum_{(s,l) \in D_t} l^\top \pi(s) = \langle v_t, \pi \rangle.$$

Further for any $\eta < \frac{(1-\gamma)}{4A\gamma}$ setting $\lambda = \frac{4\eta A\gamma}{(1-\gamma)}$, we have that for any ε, δ , with probability at least $1 - \delta$, for all policies $\pi \in \text{CH}(\Pi)$ and time $t \leq T$ the following holds

$$|\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \leq \frac{1}{(1-\gamma)^2} (1-\lambda)^t + \frac{8A^{3/2} \sqrt{\gamma \log(2T|\Pi_\varepsilon|/\delta)}}{(1-\gamma)^{5/2}} \cdot \sqrt{\eta} + \frac{4A\varepsilon}{(1-\gamma)^2}.$$

Furthermore it holds with probability 1, that for all t , $\|v_t\|_{1,\infty} \leq \frac{2A}{(1-\gamma)^2}$.

For the rest of the proof define $\varepsilon_{\text{cover}} \triangleq \frac{\varepsilon(1-\gamma)^2}{80A}$ and define the function $C(T) \triangleq \frac{8A^{3/2}\sqrt{\gamma \log(2T|\Pi_{\varepsilon_{\text{cover}}}|/\delta)}}{(1-\gamma)^{5/2}}$. Invoking Theorem 19 using $\varepsilon_{\text{cover}}$ and using the conditions on T in the statement of Theorem 16, it can now be checked that for any δ and for any $t \in [T/2, T]$ we have that

$$|\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \leq \frac{\varepsilon}{10} + C(T) \cdot \sqrt{\eta}. \quad (\text{B.2})$$

Now the ε_{ERM} -tolerant Empirical Risk Minimization Oracle, by its definition, guarantees for any t

$$\begin{aligned} \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi'_t \rangle &\leq \max_{\pi^* \in \Pi} \langle v_t, \pi^* - \pi'_t \rangle + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \\ &= \max_{\pi^* \in \Pi} \sum_{i=0}^{3t} \widehat{l}_i^\top (\pi^*(s_i) - \pi'_t(s_i)) + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \\ &\leq \|v_t\|_{1,\infty} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \end{aligned}$$

Using (B.2), with probability $1 - \delta$, for all $t \in [T/2, T]$, the inequality concerning successive iterates may thus be written as

$$\begin{aligned} \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi_t \rangle &= \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi'_t \rangle + \langle \nabla V_\mu^{\pi_t}, \pi'_t - \pi_t \rangle \\ &\leq \|v_t\|_{1,\infty} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| + \langle \nabla V_\mu^{\pi_t}, \pi'_t - \pi_t \rangle \\ &\leq \|v_t\|_{1,\infty} \varepsilon_{\text{ERM}} + \frac{\varepsilon}{10} + C(T) \cdot \sqrt{\eta} + \frac{V_\mu^{\pi_{t+1}} - V_\mu^{\pi_t}}{\eta} + \frac{4\gamma\eta}{(1-\gamma)^3} \\ &\leq \|v_t\|_{1,\infty} \varepsilon_{\text{ERM}} + \frac{\varepsilon}{5} + C(T) \cdot \sqrt{\eta} + \frac{V_\mu^{\pi_{t+1}} - V_\mu^{\pi_t}}{\eta}. \end{aligned}$$

where the second last inequality uses (B.1) and . that for any policy π , $\|\pi\|_{\infty,1} = 1$, and therefore, $\|\pi'_t - \pi_t\|_{\infty,1} \leq 2$. The last inequality follows from the condition on η in the theorem. Further since value functions are always bounded by $\frac{1}{1-\gamma}$, we average the inequality over iterations via telescoping to observe that with probability $1 - \delta$,

$$\frac{2}{T} \sum_{t=T/2}^T \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi_t \rangle \leq \frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + \frac{\varepsilon}{5} + C(T) \cdot \sqrt{\eta} + \frac{2}{(1-\gamma)\eta T}. \quad (\text{B.3})$$

Now, finally, we move from an average to a guarantee on a specific iterate. From, we have that for any $t \in [T/2, T]$,

$$\begin{aligned} &\left| \widehat{A}_t - \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi_t \rangle \right| \\ &= \left| \langle v_t, \pi'_t - \pi_t \rangle - \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi_t \rangle \right| \\ &= \left| \max_{\pi^* \in \Pi} \langle v_t, \pi^* - \pi_t \rangle - \max_{\pi^* \in \Pi} \langle \nabla V_\mu^{\pi_t}, \pi^* - \pi_t \rangle \right| + \varepsilon_{\text{ERM}} \\ &\leq 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| + \|v_t\|_{1,\infty} \varepsilon_{\text{ERM}} \end{aligned}$$

Let $\bar{t} = \operatorname{argmin}_{t \in [T/2, T]} \widehat{A}_t$ and $t' = \operatorname{argmin}_{t \in [T/2, T]} \max_{\pi^*} \langle \nabla V_{\mu}^{\pi_t}, \pi^* - \pi_t \rangle$. Then it follows by definition of t' and \bar{t} and the above inequality that

$$\begin{aligned}
 & \max_{\pi^* \in \Pi} \langle \nabla V_{\mu}^{\pi_{\bar{t}}}, \pi^* - \pi_{\bar{t}} \rangle \\
 & \leq \widehat{A}_{\bar{t}} + \frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V_{\mu}^{\pi_t}, \pi \rangle| \\
 & \leq \widehat{A}_{t'} + \frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V_{\mu}^{\pi_t}, \pi \rangle| \\
 & \leq \min_{t \in [T/2, T]} \max_{\pi^* \in \Pi} \langle \nabla V_{\mu}^{\pi_t}, \pi^* - \pi_t \rangle + \frac{4A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 4 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V_{\mu}^{\pi_t}, \pi \rangle| \\
 & \leq \frac{2}{T} \sum_{t=T/2}^T \max_{\pi^* \in \Pi} \langle \nabla V_{\mu}^{\pi_t}, \pi^* - \pi_t \rangle + \frac{4A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 4 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V_{\mu}^{\pi_t}, \pi \rangle|
 \end{aligned}$$

where the second last inequality follows from the average iterate guarantee in Equation B.3. Now combining the above, (B.3) and (B.2) we get that for any δ the following holds with probability $1 - \delta$,

$$\begin{aligned}
 \max_{\pi^* \in \Pi} \langle \nabla V_{\mu}^{\pi_{\bar{t}}}, \pi^* - \pi_{\bar{t}} \rangle & \leq \frac{2}{(1-\gamma)\eta T} + \frac{6A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + \frac{6\varepsilon}{10} + 5C(T) \cdot \sqrt{\eta} \\
 & \leq \varepsilon
 \end{aligned}$$

■

Appendix C. Faster global convergence - Proof of Theorem 17

Proof [Proof of Theorem 8] We invoke smoothness of V^{π} , as Lemma 18 certifies, to observe that since successive iterates are close in the $\|\cdot\|_{\infty,1}$ norm, we have

$$\begin{aligned}
 V^{\pi_{t+1}} & = V^{\pi_t + \eta(\pi'_t - \pi_t)} \\
 & \geq V^{\pi_t} + \eta \langle \nabla V^{\pi_t}, \pi'_t - \pi_t \rangle - \frac{\gamma \eta^2}{(1-\gamma)^3} \|\pi'_t - \pi_t\|_{\infty,1}^2.
 \end{aligned}$$

Using Theorem 19 and the definition of ε_{ERM} -tolerant Empirical Risk Minimization Oracle, we have that for any t ,

$$\begin{aligned}
 & \max_{\pi^* \in \Pi} \langle \nabla V^{\pi_t}, \pi^* - \pi_t \rangle \\
 & = \max_{\pi^* \in \Pi} \langle \nabla V^{\pi_t}, \pi^* - \pi'_t \rangle + \langle \nabla V^{\pi_t}, \pi'_t - \pi_t \rangle \\
 & \leq \frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| + \langle \nabla V^{\pi_t}, \pi'_t - \pi_t \rangle
 \end{aligned}$$

Continuing on, using the above inequality and Lemma 7, we have

$$\begin{aligned}
 V^* - V^{\pi_{t+1}} &\leq V^* - V^{\pi_t} - \eta \langle \nabla V^{\pi_t}, \pi'_t - \pi_t \rangle + \frac{\gamma \eta^2}{(1-\gamma)^3} \|\pi'_t - \pi_t\|_{\infty,1}^2 \\
 &\leq V^* - V^{\pi_t} - \eta \max_{\pi^* \in \Pi} \langle \nabla V^{\pi_t}, \pi^* - \pi_t \rangle + \eta \left(\frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| \right) + \frac{4\gamma \eta^2}{(1-\gamma)^3} \\
 &\leq V^* - V^{\pi_t} - \eta \frac{V^* - V^{\pi_t}}{C_\infty} + \eta \left(\frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| + \frac{\varepsilon_\Pi}{1-\gamma} \right) + \frac{4\gamma \eta^2}{(1-\gamma)^3} \\
 &= \left(1 - \frac{\eta}{C_\infty}\right) V^* - V^{\pi_t} + \eta \left(\frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| + \frac{\varepsilon_\Pi}{1-\gamma} \right) + \frac{4\gamma \eta^2}{(1-\gamma)^3}.
 \end{aligned}$$

Unrolling the above inequality from $t = T/2$ to $t = T$ and noting that for any policy π , $V^* - V^\pi \leq \frac{1}{1-\gamma}$ we get that

$$\begin{aligned}
 V^* - V^{\pi_T} &\leq \left(1 - \frac{\eta}{C_\infty}\right)^{T/2} \frac{1}{1-\gamma} + C_\infty \left(\frac{2A}{(1-\gamma)^2} \varepsilon_{\text{ERM}} + 2 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| + \frac{\varepsilon_\Pi}{1-\gamma} \right) + \frac{4\gamma \eta C_\infty}{(1-\gamma)^3} \\
 &\leq \frac{3\varepsilon}{10} + C_\infty \cdot 2 \max_{t \in [T/2, T]} \max_{\pi \in \Pi} |\langle v_t - \nabla V^{\pi_t}, \pi \rangle| + \frac{C_\infty \varepsilon_\Pi}{(1-\gamma)} \\
 &\leq \frac{3\varepsilon}{10} + 2C_\infty C(T) \sqrt{\eta} + \frac{C_\infty \varepsilon_\Pi}{(1-\gamma)}
 \end{aligned}$$

where the second last inequality follows from the conditions of the theorem and the last inequality follows from (B.2) which holds with probability $1 - \delta$ for any δ . Now using the conditions in the theorem give the requisite statement for any δ with probability at least $1 - \delta$. \blacksquare

Appendix D. High Probability Bound for Gradient Estimator - Proof of Theorem 19

Proof The first part of the theorem follows immediately via the definition of v_t and the definition of the datasets D_t in Algorithm 1. We now proceed with the bound on the deviation. Remember that v_t is defined in a recursive fashion by satisfying the following for any $\pi \in \mathcal{B}_{\infty,1}$,

$$\begin{aligned}
 \langle v_t, \pi \rangle &\triangleq (1-\lambda) \langle v_{t-1}, \pi \rangle + \lambda \cdot \frac{\widehat{Q}_t^\top \pi(s_t)}{1-\gamma} \\
 &\quad + (1-\lambda) \left(\frac{\gamma}{(1-\gamma)^2} \left((\widehat{H}_t(\pi_t(s_t'') - \pi_{t-1}(s_t''))^\top \pi(s_t') + (\widehat{H}_t(\pi_t(s_t') - \pi_{t-1}(s_t'))^\top \pi(s_t'')) \right) \right).
 \end{aligned}$$

For brevity in the proof we define the following random functions defined over all $\pi \in \mathcal{B}_{\infty,1}$ and for all $t \geq 0$

$$\begin{aligned}
 \psi_t(\pi) &\triangleq \frac{\widehat{Q}_t^\top \pi(s_t)}{1-\gamma} \\
 \zeta_t(\pi) &\triangleq \frac{\gamma}{(1-\gamma)^2} \left((\widehat{H}_t(\pi_t(s_t'') - \pi_{t-1}(s_t''))^\top \pi(s_t') + (\widehat{H}_t(\pi_t(s_t') - \pi_{t-1}(s_t'))^\top \pi(s_t'')) \right)
 \end{aligned}$$

Therefore by definition we have that for all $\pi \in \mathcal{B}_{\infty,1}$ and for all t ,

$$\langle v_t, \pi \rangle = (1-\lambda) \langle v_{t-1}, \pi \rangle + \lambda \psi_t(\pi) + (1-\lambda) \zeta_t(\pi).$$

Before moving onto the proof we will provide some simple upper bounds on the random variables $\psi_t(\pi), \zeta_t(\pi)$ for any $\pi \in \mathcal{B}_{\infty,1}$. Using Claim 10 we get that for any $t \leq T, \pi \in \mathcal{B}_{\infty,1}$,

$$|\psi_t(\pi)| \leq \frac{\|\hat{Q}_t\|_1 \|\pi(s_t)\|_\infty}{1-\gamma} \leq \frac{A}{(1-\gamma)^2}. \quad (\text{D.1})$$

Further using Claim 11 we get that for any $t \leq T, \pi \in \mathcal{B}_{\infty,1}$,

$$\begin{aligned} |\zeta_t(\pi)| &\leq \frac{\gamma}{(1-\gamma)^2} \left(\sum_{a,a' \in \mathcal{A} \times \mathcal{A}} |H(a, a')| \right) \|\pi(s'_t)\|_\infty (\|\pi_t(s'_t) - \pi_{t-1}(s'_t)\|_\infty + \|\pi_t(s''_t) - \pi_{t-1}(s''_t)\|_\infty) \\ &\leq \frac{2A^2\gamma}{(1-\gamma)^3} \|\pi_t - \pi_{t-1}\|_\infty \leq \frac{2A^2\gamma}{(1-\gamma)^3} \|\pi_t - \pi_{t-1}\|_{\infty,1}. \end{aligned} \quad (\text{D.2})$$

We now move on to the main proof. Let \mathbb{E}_t represent expectation fixing all the randomness upto and including time t . Then we have using Lemma 10 that $\pi \in \mathcal{B}_{\infty,1}$

$$\mathbb{E}_{t-1}[\psi_t(\pi)] = \langle \nabla V_\mu^{\pi_t}, \pi \rangle.$$

Further using Lemma 11 and the inputs to the H-sampler from Algorithm 1, it follows that

$$\mathbb{E}_{t-1}[\zeta_t(\pi)] = \mathbb{E}_{b_t} [\langle \pi_t - \pi_{t-1}, \nabla^2 V_\mu^{\bar{\pi}_t} \pi \rangle] = \mathbb{E}_{b_t} [\langle \nabla^2 V_\mu^{\bar{\pi}_t} (\pi_t - \pi_{t-1}), \pi \rangle] = \langle \nabla V_\mu^{\pi_t} - \nabla V_\mu^{\pi_{t-1}}, \pi \rangle.$$

Next consider the definitions of the following sequences for every $t \geq 0$ and $\pi \in \mathcal{B}_{\infty,1}$,

$$\epsilon_t \triangleq v_t - \nabla V_\mu^{\pi_t} \quad y_t(\pi) \triangleq \frac{\langle \epsilon_t, \pi \rangle}{(1-\lambda)^t}.$$

We next show that for any $\pi \in \mathcal{B}_{\infty,1}$, the sequence $\{y_t(\pi)\}$ is a martingale sequence over time t . To see this consider the following derivation,

$$\begin{aligned} \mathbb{E}_{t-1}[y_t(\pi)] &= \frac{\mathbb{E}_{t-1}[\langle \epsilon_t, \pi \rangle]}{(1-\lambda)^t} = \frac{\mathbb{E}_{t-1}[\langle (v_t - \nabla V_\mu^{\pi_t}), \pi \rangle]}{(1-\lambda)^t} \\ &= \frac{1}{(1-\lambda)^t} ((1-\lambda)\mathbb{E}_{t-1}[\langle (v_{t-1} - \nabla V_\mu^{\pi_{t-1}}), \pi \rangle] + \lambda (\mathbb{E}_{t-1}[\psi_t(\pi)] - \langle \nabla V_\mu^{\pi_t}, \pi \rangle) \\ &\quad + (1-\lambda) (\mathbb{E}_{t-1}[\zeta_t(\pi)] - \langle (\nabla V_\mu^{\pi_t} - \nabla V_\mu^{\pi_{t-1}}), \pi \rangle)) \\ &= \frac{(1-\lambda)\langle \epsilon_{t-1}, \pi \rangle}{(1-\lambda)^t} = y_{t-1}(\pi). \end{aligned}$$

We now wish to use Azuma's inequality to show concentration for the martingale sequence. Note that for any policies π, π' Lemma 9 implies that $\langle \nabla V_\mu^\pi, \pi' \rangle \leq \frac{1}{(1-\gamma)^2}$. Using the above derivations and Lemma 18, we can now bound the differences of the martingale sequences as follows which

holds for any $\pi \in \mathcal{B}_{\infty,1}$ and all t ,

$$\begin{aligned}
 x_t(\pi) &\triangleq |y_t(\pi) - y_{t-1}(\pi)| = \frac{|\langle \epsilon_t, \pi \rangle - (1-\lambda)\langle \epsilon_{t-1}, \pi \rangle|}{(1-\lambda)^t} \\
 &= \frac{|\lambda(\psi_t - \langle \nabla V_\mu^{\pi_t}, \pi \rangle) + (1-\lambda)(\zeta_t - \langle (\nabla V_\mu^{\pi_t} - \nabla V_\mu^{\pi_{t-1}}), \pi \rangle)|}{(1-\lambda)^t} \\
 &\leq \frac{\lambda \cdot \frac{2A}{(1-\gamma)^2} + (1-\lambda) \cdot \frac{2A^2\gamma}{(1-\gamma)^3} \|\pi_t - \pi_{t-1}\|_{\infty,1}}{(1-\lambda)^t} \\
 &\leq \frac{2A}{(1-\gamma)^2} \cdot \frac{\lambda + \eta \cdot \frac{4A\gamma}{(1-\gamma)}}{(1-\lambda)^t}.
 \end{aligned}$$

To use Azuma's inequality we need to control the sum of the worst-case differences. To this end consider the following which holds with probability 1 for any $\pi \in \mathcal{B}_{\infty,1}$,

$$\begin{aligned}
 \sum_{\tau=1}^t x_\tau^2(\pi) &\leq \sum_{\tau=1}^t \frac{4A^2}{(1-\gamma)^4} \cdot \frac{\left(\lambda + \eta \cdot \frac{4A\gamma}{(1-\gamma)}\right)^2}{(1-\lambda)^{2\tau}} \\
 &\leq \sum_{\tau=1}^t \frac{4A^2}{(1-\gamma)^4} \cdot \frac{\left(\lambda + \eta \cdot \frac{4A\gamma}{(1-\gamma)}\right)^2}{(1-\lambda)^{2\tau}} \\
 &\leq \frac{4A^2}{(1-\gamma)^4} \cdot \frac{\left(\lambda + \eta \cdot \frac{4A\gamma}{(1-\gamma)}\right)^2}{\lambda(1-\lambda)^{2t}}.
 \end{aligned}$$

A direct application of Azuma's inequality implies that for any $\pi \in \mathcal{B}_{\infty,1}$ and any t, δ with probability at least $1 - \delta$ the following holds,

$$|y_t(\pi)| \leq |y_0(\pi)| + \frac{4A}{(1-\gamma)^2} \cdot \frac{\sqrt{(\sqrt{\lambda} + \frac{\eta}{\sqrt{\lambda}} \cdot \frac{4A\gamma}{(1-\gamma)})^2 \log(2/\delta)}}{(1-\lambda)^t}$$

Setting $\lambda = \frac{4\eta A\gamma}{(1-\gamma)}$, we get that for any $\pi \in \mathcal{B}_{\infty,1}, t, \delta$, with probability at least $1 - \delta$ the following holds

$$|y_t(\pi)| \leq |y_0(\pi)| + \frac{8A^{3/2} \sqrt{\gamma \log(2/\delta)}}{(1-\gamma)^{5/2}} \cdot \frac{\sqrt{\eta}}{(1-\lambda)^t}$$

Replacing the definition of $y_t(\pi)$ we get that for any $\pi \in \mathcal{B}_{\infty,1}, \delta, t$, with probability least $1 - \delta$, the following holds

$$\begin{aligned}
 |\langle \epsilon_t, \pi \rangle| &\leq |\langle \epsilon_0, \pi \rangle| (1-\lambda)^t + \frac{8A^{3/2} \sqrt{\gamma \log(2/\delta)}}{(1-\gamma)^{5/2}} \cdot \sqrt{\eta} \\
 &\leq \frac{1}{(1-\gamma)^2} (1-\lambda)^t + \frac{8A^{3/2} \sqrt{\gamma \log(2/\delta)}}{(1-\gamma)^{5/2}} \cdot \sqrt{\eta}.
 \end{aligned}$$

Let $\varepsilon > 0$ be any number and Π_ε be the associated covering set of the policy class Π . Using a union bound over all choices of $\pi \in \Pi_\varepsilon$ and all timesteps $t \leq T$ we get that for any ε, δ , with probability at

least $1 - \delta$, we have that for any policy $\pi \in \Pi_\varepsilon$ and any time $t \leq T$,

$$|\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \leq \frac{1}{(1-\gamma)^2} (1-\lambda)^t + \frac{8A^{3/2} \sqrt{\gamma \log(2T|\Pi_\varepsilon|/\delta)}}{(1-\gamma)^{5/2}} \cdot \sqrt{\eta}. \quad (\text{D.3})$$

For the rest of the argument we will generate a crude bound over $\|v_t\|_{1,\infty}$. To this end will bound $\langle v_t, \pi \rangle$ for all $\pi \in \mathcal{B}_{\infty,1}$. By Lemma 20 since $\|\cdot\|_{\infty,1}$ and $\|\cdot\|_{1,\infty}$ are duals of each other this will imply a bound on $\|v_t\|_{1,\infty}$. Therefore consider any $\pi \in \mathcal{B}_{\infty,1}$. Using the definitions of v_t, ψ_t and ζ_t we get that

$$\begin{aligned} \langle v_t, \pi \rangle &\leq (1-\lambda)\langle v_{t-1}, \pi \rangle + \lambda|\psi_t(\pi)| + (1-\lambda)\zeta_t(\pi) \\ &\leq (1-\lambda)\langle v_{t-1}, \pi \rangle + \lambda \frac{A}{(1-\gamma)^2} + \eta \cdot \frac{4A^2\gamma}{(1-\gamma)^3} \\ &= (1-\lambda)\langle v_{t-1}, \pi \rangle + \eta \cdot \frac{8A^2\gamma}{(1-\gamma)^3}, \end{aligned}$$

where the last equality uses the choice of $\lambda = \frac{4\eta A\gamma}{(1-\gamma)}$. Now we will show by induction that for all t , $\langle v_t, \pi \rangle \leq \frac{2A}{(1-\gamma)^2}$. The base case is immediate. For the inductive case, it follows via the following computation using the choice of λ ,

$$\begin{aligned} \langle v_t, \pi \rangle &\leq (1-\lambda)\langle v_t, \pi \rangle + \eta \cdot \frac{8A^2\gamma}{(1-\gamma)^3} \\ &\leq (1-\lambda) \cdot \frac{2A}{(1-\gamma)^2} + \lambda \cdot \frac{2A}{(1-\gamma)^2} \leq \frac{2A}{(1-\gamma)^2}. \end{aligned}$$

This implies that for all t , $\|v_t\|_{1,\infty} \leq \frac{2A}{(1-\gamma)^2}$. It can be shown using Lemma 9 that for all t , $\|V_\mu^{\pi_t}\|_{1,\infty} \leq \frac{1}{(1-\gamma)^2}$. This implies that for all t , $\|v_t - \nabla V_\mu^{\pi_t}\|_{1,\infty} \leq \frac{4A}{(1-\gamma)^2}$.

Now consider any $\pi \in \Pi$ and any $\varepsilon > 0$, then by the covering property we have that there exists a $\pi' \in \Pi_\varepsilon$ such that for all t ,

$$|\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle| \leq |\langle v_t - \nabla V_\mu^{\pi_t}, \pi' \rangle| + \frac{4A\varepsilon}{(1-\gamma)^2}.$$

Combining the above with (D.3) completes the proof for all $\pi \in \Pi$.

To extend the statement to all π in the convex hull of Π , note that purely as a function of π , $f(\pi) \triangleq \max\{\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle, -\langle v_t - \nabla V_\mu^{\pi_t}, \pi \rangle\}$ is a convex function in π . Therefore one of its maxima over a convex set must lie at the boundary of the convex set. This implies that establishing the statement of any $\pi \in \Pi$ is sufficient to establish the statement for $\pi \in \text{CH}(\Pi)$. \blacksquare

Appendix E. Basic results

E.1. Duality of $\|\cdot\|_{\infty,1}$ and $\|\cdot\|_{1,\infty}$

Lemma 20 $\|\cdot\|_{1,\infty}$ and $\|\cdot\|_{\infty,1}$ are dual norms with respect to the matrix dot product, i.e.

$$\|Y\|_{1,\infty} = \max_{\|X\|_{\infty,1}=1} \langle X, Y \rangle.$$

Proof [Proof of Lemma 20] Let A_i denote the i^{th} row of a matrix A . Consider any X be such that $\|X\|_{\infty,1} = 1$. Then by Holder's inequality we have

$$\langle X, Y \rangle = \sum_{i=1}^M X_i^\top Y_i \leq \sum_{i=1}^M \|X_i\|_1 \|Y_i\|_\infty \leq \sum_{i=1}^M \|Y_i\|_\infty = \|Y\|_{1,\infty}$$

Now, construct a matrix X such that $X_{i,j} = \mathbb{I}_{j=\text{argmax}_{j'} |Y_{i,j'}|} \text{sign}(Y_{i,\text{argmax}_{j'} |Y_{i,j'}|})$ breaking ties arbitrarily for the cases where argmax is non-unique. Clearly, having one unit-sized entry per row, $\|X\|_{\infty,1} = 1$. Moreover, observe that the sequence of inequalities stated above is tight for such choice of X , because $X_i^\top Y_i = \|Y_i\|_\infty$ and $\|X_i\|_1 = 1$ hold for any $i \in [M]$ by definition of X . ■

E.2. Smoothness of V_μ^π in $\|\cdot\|_{\infty,1}$ (Proof of Lemma 18)

Proof [Proof of Lemma 18] Using the performance difference lemma [Kakade and Langford \(2002\)](#), we have

$$V_\mu^{\pi'} - V_\mu^\pi = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_\mu^{\pi'}} \left[Q^\pi(\cdot|s)^\top \pi'(s) - Q^\pi(\cdot|s)^\top \pi(s) \right]. \quad (\text{E.1})$$

Comparing this to the gradient characterization ([Lemma 9](#)), we have

$$\begin{aligned} & |V_\mu^{\pi'} - V_\mu^\pi - \langle \nabla V_\mu^\pi, \pi' - \pi \rangle| \\ & \leq \frac{1}{1-\gamma} \left| \mathbb{E}_{s \sim d_\mu^{\pi'}} \left[Q^\pi(s, \cdot)^\top \pi'(s) - Q^\pi(s, \cdot)^\top \pi(s) \right] - \mathbb{E}_{s \sim d_\mu^\pi} \left[Q^\pi(s, \cdot)^\top \pi'(s) - Q^\pi(s, \cdot)^\top \pi(s) \right] \right| \\ & \leq \frac{1}{1-\gamma} \|d_\mu^{\pi'} - d_\mu^\pi\|_1 \max_{s \in S} \left| Q^\pi(s, \cdot)^\top \pi'(s) - Q^\pi(s, \cdot)^\top \pi(s) \right| \\ & \leq \frac{1}{1-\gamma} \|d_\mu^{\pi'} - d_\mu^\pi\|_1 \max_{s \in S} \{ \|Q^\pi(s, \cdot)\|_\infty \|\pi'(s) - \pi(s)\|_1 \} \\ & \leq \frac{1}{(1-\gamma)^2} \|d_\mu^{\pi'} - d_\mu^\pi\|_1 \|\pi' - \pi\|_{\infty,1} \end{aligned}$$

where the last inequality follows from $\|Q^\pi(s, \cdot)\|_\infty \leq \frac{1}{1-\gamma}$. Now, define for any policy π , define $P^\pi(s'|s) = \sum_{a \in A} P(s'|s, a) \pi(a|s)$ as the associated Markov transition operator. First, for any distribution $d \in \Delta_S$, we have

$$\begin{aligned} \|(P^{\pi'} - P^\pi)d\|_1 &= \sum_{s' \in S} \left| \sum_{a \in A, s \in S} P(s'|s, a) (\pi'(a|s) - \pi(a|s)) d(s) \right| \\ &\leq \sum_{s' \in S, a \in A, s \in S} P(s'|s, a) |\pi'(a|s) - \pi(a|s)| d(s) \\ &= \sum_{s \in S, a \in A} |\pi'(a|s) - \pi(a|s)| d(s) \\ &\leq \max_{s \in S} \sum_{a \in A} |\pi'(a|s) - \pi(a|s)| = \|\pi' - \pi\|_{\infty,1} \end{aligned}$$

Generalizing this to t successive applications of the Markov operator, we have the following via an inductive argument. Suppose for $t-1$, we have that for all distributions $d \in \Delta_S$ we have that

$\|((P^{\pi'})^t - (P^\pi)^t)d\|_1 \leq (t-1)\|\pi' - \pi\|_{\infty,1}$. Now consider the case for t ,

$$\begin{aligned} \|((P^{\pi'})^t - (P^\pi)^t)d\|_1 &\leq \|((P^{\pi'})^t - (P^{\pi'})^{t-1}P^\pi)d\|_1 + \|((P^{\pi'})^{t-1}P^\pi - (P^\pi)^t)d\|_1 \\ &= \|(P^{\pi'})^{t-1}(P^{\pi'} - P^\pi)d\|_1 + \|((P^{\pi'})^{t-1} - (P^\pi)^{t-1})P^\pi d\|_1 \\ &\leq \|(P^{\pi'} - P^\pi)d\|_1 + \|((P^{\pi'})^{t-1} - (P^\pi)^{t-1})P^\pi d\|_1 \\ &\leq \|\pi' - \pi\|_{\infty,1} + (t-1)\|\pi' - \pi\|_{\infty,1} \\ &= t\|\pi' - \pi\|_{\infty,1} \end{aligned}$$

Using the definition of d_μ^π , we have

$$\begin{aligned} \|d_\mu^{\pi'} - d_\mu^\pi\|_1 &\leq (1-\gamma) \sum_{t \geq 0} \gamma^t \|((P^{\pi'})^t - (P^\pi)^t)\mu\|_1 \\ &\leq (1-\gamma) \sum_{t \geq 0} \gamma^t t \|\pi' - \pi\|_{\infty,1} \\ &= \frac{\gamma}{1-\gamma} \|\pi' - \pi\|_{\infty,1}. \end{aligned} \tag{E.2}$$

where the last line uses the identity $\sum_{t \geq 0} \gamma^t t = \frac{\gamma}{(1-\gamma)^2}$, completing the proof. For the second part of the lemma, consider the following. Using Lemma 9, we have that,

$$\begin{aligned} \langle \nabla V_\mu^{\pi'} - \nabla V_\mu^{\pi''}, \pi \rangle &= \frac{1}{1-\gamma} \left(\mathbb{E}_{s \sim d_\mu^{\pi'}} \left[Q^{\pi'}(s, \cdot)^\top \pi(s) \right] - \mathbb{E}_{s \sim d_\mu^{\pi''}} \left[Q^{\pi''}(s, \cdot)^\top \pi(s) \right] \right) \\ &= \frac{1}{1-\gamma} \left(\mathbb{E}_{s \sim d_\mu^{\pi'}} \left[(Q^{\pi'}(s, \cdot) - Q^{\pi''}(s, \cdot))^\top \pi(s) \right] \right) + \\ &\quad \frac{1}{1-\gamma} \left(\mathbb{E}_{s \sim d_\mu^{\pi'}} \left[Q^{\pi''}(s, \cdot)^\top \pi(s) \right] - \mathbb{E}_{s \sim d_\mu^{\pi''}} \left[Q^{\pi''}(s, \cdot)^\top \pi(s) \right] \right). \end{aligned}$$

We now bound the two terms above separately. Applying Cauchy-Schwartz repeatedly and noting the definition of $Q^\pi(s, a) = r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a)V^\pi(s)$ we have that

$$\begin{aligned} \mathbb{E}_{s \sim d_\mu^{\pi'}} \left[(Q^{\pi'}(s, a) - Q^{\pi''}(s, a))^\top \pi(s) \right] &\leq \max_{s, a \in \mathcal{S} \times \mathcal{A}} |(Q^{\pi'}(s, a) - Q^{\pi''}(s, a))| \\ &= \gamma \max_{s, a \in \mathcal{S} \times \mathcal{A}} \left| \sum_{s' \in \mathcal{S}} P(s'|s, a)(V^{\pi'}(s) - V^{\pi''}(s)) \right| \\ &\leq \gamma \max_{s \in \mathcal{S}} |V^{\pi'}(s) - V^{\pi''}(s)| \\ &= \frac{\gamma}{1-\gamma} \max_{s \in \mathcal{S}} \left| \mathbb{E}_{s' \sim d_s^{\pi'}} \left[Q^{\pi''}(\cdot|s')^\top \pi'(s') - Q^{\pi''}(\cdot|s')^\top \pi''(s') \right] \right| \\ &\leq \frac{\gamma}{(1-\gamma)^2} \|\pi' - \pi''\|_{\infty,1}. \end{aligned}$$

Here the second last inequality follows from the performance difference lemma (E.1). Furthermore for the second term we have that,

$$\begin{aligned} \mathbb{E}_{s \sim d_\mu^{\pi'}} \left[Q^{\pi''}(s, \cdot)^\top \pi(s) \right] - \mathbb{E}_{s \sim d_\mu^{\pi''}} \left[Q^{\pi''}(s, \cdot)^\top \pi(s) \right] &\leq \frac{1}{(1-\gamma)} \|d_\mu^{\pi'} - d_\mu^{\pi''}\|_1 \\ &\leq \frac{\gamma}{(1-\gamma)^2} \|\pi' - \pi''\|_{\infty,1}. \end{aligned}$$

where the inequality follows from (E.2). Putting the above statements together we easily see that

$$\langle \nabla V_{\mu}^{\pi'} - \nabla V_{\mu}^{\pi''}, \pi \rangle \leq \frac{2\gamma}{(1-\gamma)^3} \|\pi' - \pi''\|_{\infty,1}.$$

■

Appendix F. Characterization of the Functional gradient and Hessian of the Value Function (Proofs of Lemmas 9 and 14)

For the sake of this section, we introduce new notation that will help us state analytic derivatives cleanly. For any policy $\pi \in \mathbb{R}^{S \times A}$ define a function of π , T^π with the signature $T^\pi : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^{S \times A \times SA}$ such that for all $((s', a'), (s, a))$,

$$T^\pi((s', a'), (s, a)) = P(s'|s, a)\pi(a'|s').$$

Similarly define $P^\pi(s'|s) : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^{S \times S}$ such that for all (s', s) ,

$$P^\pi(s', s) \triangleq \sum_{a \in \mathcal{A}} P(s'|s, a)\pi(a|s).$$

Finally define $\mu^\pi : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^{SA}$ such that for all (s, a)

$$\mu^\pi(s, a) = \mu(s)\pi(a|s).$$

For any (\tilde{s}, \tilde{a}) , we define the following partial derivatives for $(\tilde{s}, \tilde{a})^{\text{th}}$ entry of the input π as $\frac{\partial T^\pi}{\partial \pi(\tilde{a}|\tilde{s})} \in \mathbb{R}^{SA \times SA}$ and $\frac{\partial P^\pi}{\partial \pi(\tilde{a}|\tilde{s})} \in \mathbb{R}^{S \times S}$ such that

$$\frac{\partial T^\pi}{\partial \pi(\tilde{a}|\tilde{s})}((s', a'), (s, a)) = P(\tilde{s}|s, a)\mathbb{I}_{s'=\tilde{s} \wedge a'=\tilde{a}}, \quad \frac{\partial P^\pi}{\partial \pi(\tilde{a}|\tilde{s})}(s', s) = P(s'|\tilde{s}, \tilde{a})\mathbb{I}_{s=\tilde{s}}.$$

It can be seen that for any $\pi \in (\Delta_A)^S$, T^π , P^π are stochastic matrices. Therefore since $\gamma < 1$, we have that $I - \gamma T^\pi$ and $I - \gamma P^\pi$ are invertible.

Next, consider any start state distribution $\mu \in \Delta_S$ and a reward vector $r \in \mathbb{R}^{S \times A}$. As defined before the steady-state distribution, $d_\mu^\pi : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^{SA}$, the Q-function $Q^\pi : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^{SA}$ and the value function $V^\pi : \mathbb{R}^{S \times A} \rightarrow \mathbb{R}^S$ are also functions of a policy π and thus similar partial derivatives for any \tilde{s}, \tilde{a} can be defined here as well. In this section to make the notation more explicit we define $d_{\mu, S}^\pi \in \Delta_S$ such that for all $s \in S$, $d_{\mu, S}^\pi(s) = \sum_{a \in \mathcal{A}} (d_{\mu, S}^\pi(s, a))$. Note that in other parts of the paper we have referred to $d_{\mu, S}^\pi(s)$ as just $d_\mu^\pi(s)$ but since we need to explicitly use the vector $d_{\mu, S}^\pi(s)$ in this section we make this notation explicit.

Now note by Bellman equations, we have that

$$\begin{aligned} d_\mu^\pi &= (1-\gamma)\mu^\pi + \gamma T^\pi d_\mu^\pi \implies d_\mu^\pi = (1-\gamma)(I - \gamma T^\pi)^{-1} \mu^\pi, \\ d_{\mu, S}^\pi &= (1-\gamma)\mu + \gamma P^\pi d_{\mu, S}^\pi \implies d_{\mu, S}^\pi = (1-\gamma)(I - \gamma P^\pi)^{-1} \mu, \\ Q^\pi &= r + \gamma (T^\pi)^\top Q^\pi \implies Q^\pi = (I - \gamma (T^\pi)^\top)^{-1} r. \end{aligned}$$

We will now use these notations repeatedly. As a warm-up, we provide a proof the policy gradient lemma (Lemma 9) first.

Proof [Proof of Lemma 9] First observe that $V_\mu^\pi = (\mu^\pi)^\top Q^\pi$. Now for all \tilde{s}, \tilde{a} we have that

$$\begin{aligned}
 \frac{\partial V_\mu^\pi}{\partial \pi(\tilde{a}|\tilde{s})} &= \left(\frac{\partial \mu^\pi}{\partial \pi(\tilde{a}|\tilde{s})} \right)^\top Q^\pi + (\mu^\pi)^\top \frac{\partial Q^\pi}{\partial \pi(\tilde{a}|\tilde{s})} \\
 &= \mu(\tilde{s})Q^\pi(\tilde{s}, \tilde{a}) + \gamma(\mu^\pi)^\top \left[(I - \gamma(T^\pi)^\top)^{-1} \frac{\partial (T^\pi)^\top}{\partial \pi(\tilde{a}|\tilde{s})} (I - \gamma(T^\pi)^\top)^{-1} r \right] \\
 &= \mu(\tilde{s})Q^\pi(\tilde{s}, \tilde{a}) + \gamma((I - \gamma(T^\pi))^{-1} \mu^\pi)^\top \left[\frac{\partial (T^\pi)^\top}{\partial \pi(\tilde{a}|\tilde{s})} (I - \gamma(T^\pi)^\top)^{-1} r \right] \\
 &= \mu(\tilde{s})Q^\pi(\tilde{s}, \tilde{a}) + \frac{\gamma}{1 - \gamma} (d_\mu^\pi)^\top \frac{\partial (T^\pi)^\top}{\partial \pi(\tilde{a}|\tilde{s})} Q^\pi \\
 &= \mu(\tilde{s})Q^\pi(\tilde{s}, \tilde{a}) + \frac{\gamma}{1 - \gamma} Q^\pi(\tilde{s}, \tilde{a}) \sum_{s \in S, a \in A} P(\tilde{s}|s, a) d_\mu^\pi(s, a) \\
 &= \mu(\tilde{s})Q^\pi(\tilde{s}, \tilde{a}) + \frac{\gamma}{1 - \gamma} Q^\pi(\tilde{s}, \tilde{a}) \sum_{s \in S, a \in A} P(\tilde{s}|s, a) \pi(a|s) d_\mu^\pi(s) \\
 &= \frac{d_\mu^\pi(\tilde{s})Q^\pi(\tilde{s}, \tilde{a})}{1 - \gamma}
 \end{aligned} \tag{F.1}$$

where we use that $d_\mu^\pi(s) = (1 - \gamma)\mu(s) + \gamma \sum_{s' \in S, a' \in A} P(s|s', a') \pi(a'|s') d_\mu^\pi(s')$. The statement of the lemma follows immediately now. \blacksquare

Proof [Proof of Lemma 14] Observe for any $s, a, \tilde{s}, \tilde{a}$ we have the following statement that follows by product rule,

$$\frac{\partial (d_\mu^\pi(s)Q^\pi(s, a))}{\partial \pi(\tilde{a}|\tilde{s})} = d_\mu^\pi(s) \frac{\partial Q^\pi(s, a)}{\partial \pi(\tilde{a}|\tilde{s})} + \frac{\partial d_\mu^\pi(s)}{\partial \pi(\tilde{a}|\tilde{s})} Q^\pi(s, a). \tag{F.2}$$

Further given any s, a we can define $d_{s,a}^\pi \in \mathbb{R}^{SA}$ to be the steady state distribution starting from state s and executing action a . In particular the following holds

$$d_{s,a}^\pi = (1 - \gamma)e_{s,a} + \gamma T^\pi d_\mu^\pi \implies d_{s,a}^\pi = (1 - \gamma)(I - \gamma T^\pi)^{-1} e_{s,a},$$

where $e_{s,a} \in \mathbb{R}^{SA}$ is the indicator vector of the $(s, a)^{th}$ coordinate. We now have the following for any $s, a, \tilde{s}, \tilde{a}$,

$$\begin{aligned}
 \frac{\partial Q^\pi(s, a)}{\partial \pi(\tilde{a}|\tilde{s})} &= e_{s,a}^\top \frac{\partial Q^\pi}{\partial \pi(\tilde{a}|\tilde{s})} \\
 &= \gamma e_{s,a}^\top (1 - \gamma(T^\pi)^\top)^{-1} \frac{\partial (T^\pi)^\top}{\partial \pi(\tilde{a}|\tilde{s})} (1 - \gamma(T^\pi)^\top)^{-1} r \\
 &= \frac{\gamma}{1 - \gamma} (d_{s,a}^\pi)^\top \frac{\partial (T^\pi)^\top}{\partial \pi(\tilde{a}|\tilde{s})} Q^\pi \\
 &= \frac{\gamma}{1 - \gamma} Q^\pi(\tilde{s}, \tilde{a}) \sum_{s' \in S, a' \in A} P(\tilde{s}|s', a') d_{s,a}^\pi(s', a').
 \end{aligned}$$

The above in particular implies that for any π' and any $s, a,$

$$\left\langle d_{\mu}^{\pi}(s) \frac{\partial Q^{\pi}(s, a)}{\partial \pi}, \pi' \right\rangle = \mathbb{E}_{s \sim d_{\mu}^{\pi}} \mathbb{E}_{s' \sim P(\cdot | s, a)} \mathbb{E}_{\tilde{s} \sim d_{s'}}^{\pi} \left[Q^{\pi}(\tilde{s}, \cdot)^{\top} \pi'(\tilde{s}) \right] = F^{\pi}(s, a | \pi'). \quad (\text{F.3})$$

Further, consider the following for any $s, \tilde{s}, \tilde{a},$

$$\begin{aligned} \frac{\partial d_{\mu}^{\pi}(s)}{\partial \pi(\tilde{a} | \tilde{s})} &= e_s^{\top} \frac{\partial d_{\mu, S}^{\pi}}{\partial \pi(\tilde{a} | \tilde{s})} \\ &= (1 - \gamma) \gamma e_s^{\top} (I - \gamma P^{\pi})^{-1} \frac{\partial P^{\pi}}{\partial \pi(\tilde{a} | \tilde{s})} (I - \gamma P^{\pi})^{-1} \mu \\ &= \gamma e_s^{\top} (I - \gamma P^{\pi})^{-1} \frac{\partial P^{\pi}}{\partial \pi(\tilde{a} | \tilde{s})} d_{\mu, S}^{\pi}. \end{aligned}$$

Let $\mu'_{\tilde{s}, \tilde{a}} \in \Delta_S$ such that for all $s',$

$$\mu'_{\tilde{s}, \tilde{a}}(s') = P(s' | \tilde{s}, \tilde{a}) d_{\mu}^{\pi}(\tilde{s}).$$

Therefore we have that,

$$\begin{aligned} \frac{\partial d_{\mu}^{\pi}(s)}{\partial \pi(\tilde{a} | \tilde{s})} &= \gamma e_s^{\top} (I - \gamma P^{\pi})^{-1} \mu'_{\tilde{s}, \tilde{a}} \\ &= \frac{\gamma}{1 - \gamma} e_s^{\top} d_{\mu'_{\tilde{s}, \tilde{a}}}^{\pi} \\ &= \frac{\gamma}{1 - \gamma} \sum_{s' \in S} d_{s'}^{\pi}(s) P(s' | \tilde{s}, \tilde{a}) d_{\mu}^{\pi}(\tilde{s}). \end{aligned}$$

The above in particular implies that for any π'' and any $\tilde{s}, \tilde{a},$

$$\left\langle Q^{\pi} \frac{\partial d_{\mu}^{\pi}}{\partial \pi(\tilde{a} | \tilde{s})}, \pi'' \right\rangle = \mathbb{E}_{\tilde{s} \sim d_{\mu}^{\pi}} \mathbb{E}_{s' \sim P(\cdot | \tilde{s}, \tilde{a})} \mathbb{E}_{s \sim d_{s'}}^{\pi} \left[Q^{\pi}(s, \cdot)^{\top} \pi''(s) \right] = F^{\pi}(\tilde{s}, \tilde{a} | \pi''). \quad (\text{F.4})$$

Combining (F.2), (F.3), (F.4) completes the proof. \blacksquare

Appendix G. Sampling subroutines

G.1. Properties of Q-sampler (Algorithm 2)

Proof [Proof of Lemma 10] Let E_t be the event that the first loop terminates at the t^{th} iteration. Then

$$\mathbb{P}(s_t = s) = \sum_{\tau=0}^{\infty} \mathbb{P}(E_{\tau}) \mathbb{E}[\mathbf{1}_{s_{\tau}=s} | \pi] = \sum_{\tau=0}^{\infty} (1 - \gamma) \gamma^{\tau} \mathbb{E}[\mathbf{1}_{s_{\tau}=s} | \pi, E_{\tau}] = d_{\pi}(s).$$

Let F_t be the event that the second loop terminates at the t^{th} iteration. Then

$$\begin{aligned}
 \mathbb{E}[R|s_t, a_t] &= \sum_{\tau=0}^{\infty} \mathbb{P}(F_\tau) \mathbb{E} \left[\frac{r(s_{t+\tau}, a_{t+\tau})}{1-\gamma} \middle| \pi, s_t, a_t, F_\tau \right] \\
 &= \sum_{\tau=0}^{\infty} (1-\gamma) \gamma^\tau \mathbb{E} \left[\frac{r(s_{t+\tau}, a_{t+\tau})}{1-\gamma} \middle| \pi, s_t, a_t, F_\tau \right] \\
 &= \mathbb{E} \left[\sum_{\tau=0}^{\infty} \gamma^\tau r(s_{t+\tau}, a_{t+\tau}) \middle| \pi, s_t, a_t, F_\tau \right] \\
 &= Q^\pi(s_t, a_t).
 \end{aligned}$$

Now, taking the marginal over the choice of a_t , we have for any $a \in A$

$$\mathbb{E} \left[\widehat{Q}(a) \middle| s_t, \pi \right] = A \mathbb{E} [R|s_t, a_t = a] \mathbb{P}(a_t = a) = Q^\pi(s_t, a).$$

Since E_t and F_t are geometric random variables with $(1-\gamma)$ probability of termination, the expected survival length of each is $\frac{1}{1-\gamma}$. \blacksquare

G.2. Properties of H-sampler (Algorithm 2)

Proof [Proof of Lemma 15] Since the sampling procedure of s coincides with the first phase of Algorithm 2, analogously we have that $s \sim d_\mu^\pi$. Similarly, by comparisons to the second phase of Algorithm 2, we have $\mathbb{E}[R|s_{t'}, a_{t'}] = Q^\pi(s_{t'}, a_{t'})$. Taking the marginal over the choice of $a_{t'}$, we have for any π' that

$$\mathbb{E} \left[\widehat{H}^\top \pi'(s') \middle| s, \pi \right] = F^\pi(s, \cdot | \pi').$$

Proof [Proof of Lemma 11] The first part of the claim is a synthesis of Lemma 15 and the characterization of the policy Hessian in Lemma 14. The episode length is bounded as each of the three rejection sampling phases has a survival length of $\frac{1}{1-\gamma}$. \blacksquare

Appendix H. Local-to-Global Lemmas (Proofs of Lemmas 4 and 7)

H.1. For an exploratory distribution

Proof [Proof of Lemma 4] Consider any $\pi \in \text{CH}(\Pi)$. Due to performance difference lemma (Agarwal et al., 2021), we have

$$\begin{aligned}
 V^* - V^\pi &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi^*}} \left[Q^\pi(s, \cdot)^\top \pi^*(s) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi^*}} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{1-\gamma} \left\| \frac{d^{\pi^*}}{d^\pi} \right\|_\infty \mathbb{E}_{s \sim d^\pi} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{(1-\gamma)^2} \left\| \frac{d^{\pi^*}}{\mu} \right\|_\infty \left(\min_{\pi' \in \Pi} \mathbb{E}_{s \sim d^\mu} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi'(s) \right] \right. \\
 &\quad \left. + \max_{\pi' \in \Pi} \mathbb{E}_{s \sim d^\mu} \left[Q^\pi(s, \cdot)^\top (\pi'(s) - \pi(s)) \right] \right) \\
 &\leq \frac{1}{(1-\gamma)^2} \left\| \frac{d^{\pi^*}}{\mu} \right\|_\infty \left(\epsilon_{\Pi, \mu} + (1-\gamma) \max_{\pi' \in \Pi} \langle \nabla V_\mu^\pi, \pi' - \pi \rangle \right),
 \end{aligned}$$

where the last line follows from Lemma 9 and the definition of $\epsilon_{\Pi, \mu}$. ■

H.2. For an exploratory policy class

Proof [Proof of Lemma 7] Consider any $\pi \in \text{CH}(\Pi)$. Due to performance difference lemma (E.1), we have

$$\begin{aligned}
 V^* - V^\pi &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi^*}} \left[Q^\pi(s, \cdot)^\top \pi^*(s) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi^*}} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{1-\gamma} \left\| \frac{d^{\pi^*}}{d^\pi} \right\|_\infty \mathbb{E}_{s \sim d^\pi} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi(s) \right] \\
 &\leq \frac{1}{1-\gamma} \left\| \frac{d^{\pi^*}}{d^\pi} \right\|_\infty \left(\min_{\pi' \in \Pi} \mathbb{E}_{s \sim d^\pi} \left[\max_{a \in \mathcal{A}} Q^\pi(s, a) - Q^\pi(s, \cdot)^\top \pi'(s) \right] \right. \\
 &\quad \left. + \max_{\pi' \in \Pi} \mathbb{E}_{s \sim d^\pi} \left[Q^\pi(s, \cdot)^\top (\pi'(s) - \pi(s)) \right] \right) \\
 &\leq \frac{1}{1-\gamma} \left\| \frac{d^{\pi^*}}{d^\pi} \right\|_\infty \left(\epsilon_\Pi + (1-\gamma) \max_{\pi' \in \Pi} \langle \nabla V^\pi, \pi' - \pi \rangle \right),
 \end{aligned}$$

where the last line follows from Lemma 9 and the definition of ϵ_Π . ■