Efficiently Escaping Saddle Points for Policy Optimization

Sadegh Khorasani[1]   **Saber Salehkaleybar**[2]   **Negar Kiyavash**[3]   **Niao He**[4]   **Matthias Grossglauser**[1]

[1]School of Computer and Communication Sciences, EPFL,, Lausanne, Switzerland
[2]Leiden Institute of Advanced Computer Science (LIACS), Leiden University, , Leiden, The Netherlands
[3]College of Management of Technology, EPFL, Lausanne, Switzerland
[4]Department of Computer Science ETH Zurich

Abstract

Policy gradient (PG) is widely used in reinforcement learning due to its scalability and good performance. In recent years, several variance-reduced PG methods have been proposed with a theoretical guarantee of converging to an approximate first-order stationary point (FOSP) with the sample complexity of $O(\epsilon^{-3})$. However, FOSPs could be bad local optima or saddle points. Moreover, these algorithms often use importance sampling (IS) weights which could impair the statistical effectiveness of variance reduction. In this paper, we propose a variance-reduced second-order method that uses second-order information in the form of Hessian vector products (HVP) and converges to an approximate second-order stationary point (SOSP) with sample complexity of $\tilde{O}(\epsilon^{-3})$. This rate improves the best-known sample complexity for achieving approximate SOSPs by a factor of $O(\epsilon^{-0.5})$. Moreover, the proposed variance reduction technique bypasses IS weights by using HVP terms. Our experimental results show that the proposed algorithm outperforms the state of the art and is more robust to changes in random seeds.

1 INTRODUCTION

Reinforcement Learning (RL) is an interactive learning approach where an agent learns how to choose the best action in an interactive environment based on received signals. In the past few years, RL has been applied with great success to many applications of interest such as autonomous driving [Shalev-Shwartz et al., 2016], games [Silver et al., 2017], and robot manipulation [Deisenroth et al., 2013]. RL can be formulated mathematically as a Markov Decision Process (MDP) where after taking an action, the state changes based on transition probabilities and the agent receives a reward according to the current state and the action taken. The agent takes actions according to a policy that maps the state space to action space, and the goal is to find a policy that maximizes the agent's average cumulative reward.

Policy gradient (PG) methods directly search for a policy that is parameterized by a parameter vector θ. PG methods can provide good policies in high-dimensional control tasks by first harnessing the power of deep neural networks for policy parameterization and subsequently optimizing the parameter vector θ. REINFORCE [Williams, 1992], PGT [Sutton et al., 2000], and GPOMDP [Baxter and Bartlett, 2001] are examples of classical PG methods that update the parameters using stochastic gradient ascent, as it is often infeasible to compute the gradient exactly in complex environments.

To manage the stochasticity in gradient update, in the RL literature, several approaches have been proposed. Sutton et al. [2000] presented the baseline technique to reduce the variance of gradient estimates. Konda and Tsitsiklis [2000] proposed an actor-critic algorithm that estimates the value function and utilizes it for the purpose of variance reduction. Schulman et al. [2015b] presented GAE to control both bias and variance by exploiting a temporal difference relation for the advantage function approximation. Schulman et al. [2015a] proposed TRPO, which considers a Kullback-Leibler (KL) divergence penalty term in order to ensure that the updated policy remains close to the current policy. Subsequently, Schulman et al. [2017] used a clipped surrogate objective function. It has been observed experimentally that the aforementioned algorithms have better performance compared to the vanilla PG method. However, no theoretical guarantees for the convergence rate of most of these algorithms are available.

In recent years, several methods have been proposed with theoretical guarantees for the convergence to an ϵ-approximate First-Order Stationary Point (ϵ-FOSP) of the objective function $J(\theta)$, i.e., $\|\nabla J(\theta)\| \leq \epsilon$. Most of these methods adopt recent variance reduction techniques pre-

sented originally in the context of stochastic optimization, and converge to ϵ-FOSP with the sample complexity of $O(\epsilon^{-3})$ (See Related Work section for more details). The two main drawbacks of these methods are as follows. First, most of these methods require importance sampling (IS) weights in the variance reduction part because the objective function in RL is non-oblivious in the sense that its trajectories depend on the policy that generates them. This degrades the effectiveness of variance reduction techniques, because the IS weights grow exponentially with the horizon length [Zhang et al., 2021]. Moreover, their analyses required strong assumptions such as the boundedness of variance of IS weights. The second drawback is the failure to provide guarantees beyond convergence to ϵ-FOSP. In many applications, the objective function $J(\theta)$ is nonconvex and FOSPs may include bad local optima and saddle points. This is one of the reasons why the aforementioned methods are too sensitive to parameter initialization and random seeds in practice. We argue that it is more desirable to obtain $(\epsilon, \sqrt{\rho\epsilon})$-approximate second-order stationary point $((\epsilon, \sqrt{\rho\epsilon})$-SOSP), i.e., $\|\nabla J(\theta)\| \leq \epsilon$ and $\lambda_{max}(\nabla^2 J(\theta)) \leq \sqrt{\rho\epsilon}$, where $\lambda_{max}(.)$ is the maximum eigenvalue and ρ is the Hessian Lipschitz constant[1].

In the context of optimization, Nesterov and Polyak [2006] showed that cubic regularized Newton (CRN) method escapes saddle points and converges to SOSP by incorporating second-order information, namely, the Hessian matrix. However, obtaining the Hessian is computationally intensive in high dimensions, which is the case in most RL applications where the policy is modeled by a (deep) neural network. Moreover, in the stochastic setting, only estimates of the gradient and Hessian are available. To address these challenges, Tripuraneni et al. [2018] proposed stochastic CRN (SCRN), which uses sub-sampled gradient and Hessian to converge to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-3.5})$[2]. Furthermore, they showed that using the result in [Carmon and Duchi, 2016], the update in each iteration can be obtained from Hessian vector products (HVP) instead of computing the full Hessian matrix. Later, Zhou and Gu [2020] proposed a variance-reduced version of SCRN based on SARAH [Nguyen et al., 2017] achieving sample complexity of $\tilde{O}(\epsilon^{-3})$.

In the context of RL, recently, Wang et al. [2022] proposed a second-order PG method that converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with a sample complexity of $\tilde{O}(\epsilon^{-3.5})$. The proposed method uses a variance reduction technique based on SARAH to estimate the gradients. This technique still requires IS weights and the customary strong assumptions on these weights (such as the boundedness of their vari-

ance). Concurrent to Wang et al. [2022], Maniyar et al. [2024], proposed a second-order PG method that converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with the best-known sample complexity rate of $O(\epsilon^{-3.5})$. A natural question to ask is whether there exists a second-order PG method that converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-3})$ but without using IS weights?

In this paper, we answer the above question in the affirmative by proposing Variance-Reduced Stochastic Cubic-regularized Policy gradient (VR-SCP) algorithm. The proposed algorithm updates the parameters in each iteration based on optimizing a stochastic second-order Taylor expansion of the objective function with a cubic penalty term where gradient estimates are obtained based on a novel variance reduction technique. In particular, our main contributions are as follows:

- VR-SCP converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-3})$, improving the best-known sample complexity [Wang et al., 2022, Maniyar et al., 2024] by a factor of $O(\epsilon^{-0.5})$.
- We propose a Hessian-aided variance reduction technique that incorporates HVP in estimating gradients, entirely bypassing IS weights. Our convergence analysis does not require strong assumptions on IS weights.
- To showcase the advantages of converging to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP, we define a new metric that incorporates both performance (the average return) and robustness (sensitivity to random seeds) of an RL algorithm, where the latter is crucial in terms of reproducibility of the results. Our experimental results show that not only VR-SCP outperforms state-of-the-art algorithms in terms of its theoretical guarantees, but also in terms of the aforementioned metric.

We should emphasize that our technique differs from HAPG [Shen et al., 2019] which also bypasses IS weights using HVP in three main aspects. First, in their analysis, it is required to update the parameters with a fixed step size of ϵ in order to bound the variance of gradient estimates which slows the training process in practice. Second, in HAPG, the number of computed HVPs per iteration is in the order of $O(1/\epsilon)$ while in our case, it depends on the norm of the update. Last but not least, HAPG only uses the second-order information in the form of HVP for variance reduction and hence achieves ϵ-FOSP, therefore, as we shall see in our experiments, it misses the main performance advantages of converging to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP.

The paper is organized as follows: In Section 2, we provide some definitions and background on variance-reduced methods in stochastic optimization and also the second-order method, SCRN. In Section 3, we describe the proposed algorithm and analyze its convergence rate for achieving $(\epsilon, \sqrt{\rho\epsilon})$-SOSP. In Section 4, we review related work in the RL literature. In Section 5, we define a new metric to evalu-

---

[1]The function J is said to have a Hessian Lipschitz constant ρ if for all $\theta_1, \theta_2 \in \mathbb{R}^d$: $\|\nabla^2 J(\theta_1) - \nabla^2 J(\theta_2)\| \leq \rho\|\theta_1 - \theta_2\|$.

[2]$\tilde{O}(.)$ is a variant of big-O notation, ignoring the logarithmic factors. In other words, $f(n) \in \tilde{O}(g(n))$ if there exists some positive constant k, such that $f(n) \in O(g(n)\log^k(g(n)))$.

ate RL algorithms in control tasks and compare the proposed algorithm with previous work based on this metric. Finally, we conclude the paper in Section 6.

2 PRELIMINARIES

2.1 NOTATIONS AND PROBLEM DEFINITION

An MDP can be represented as a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma, p_0 \rangle$, where $\mathcal{S}$ and $\mathcal{A}$ are state space and action space, respectively. The conditional probability of transition from state s to s' with action a is denoted by $P(s'|s,a)$. The probability distribution over the initial state s_0 is denoted by $p_0(s_0)$. The parameter $\gamma \in (0,1)$ denotes the discount factor. At each time step t, $r(s_t, a_t)$ returns the reward of taking action a_t in the state s_t. Actions are chosen according to the policy π where $\pi(a|s)$ is the probability of taking action a for a given state s. Here, we assume that the policy is parameterized with a vector $\theta \in \mathbb{R}^d$ and use shorthand notation π_θ for $\pi_\theta(a|s)$. For a given time horizon H, we define $\tau = (s_0, a_0, \cdots, s_{H-1}, a_{H-1})$ as a sequence of state-action pairs called a trajectory. $R(\tau)$ is a function that returns the discounted accumulated reward of each trajectory as follows: $R(\tau) := \sum_{h=0}^{H-1} \gamma^h r(s_h, a_h)$ where $\gamma \in (0,1)$ is the discount factor. For an arbitrary vector $v \in \mathbb{R}^d$, we denote Euclidean norm by $\|v\|_2$. For a matrix $A \in \mathbb{R}^{a \times b}$, $A[.] : \mathbb{R}^b \to \mathbb{R}^a$ takes a vector $v \in \mathbb{R}^b$ and returns the matrix-vector product Av. We show the maximum eigenvalue of a symmetric matrix $H \in \mathbb{R}^{d \times d}$ with $\lambda_{max}(H)$. For two vectors $v_1, v_2 \in \mathbb{R}^d$, we use $\langle v_1, v_2 \rangle$ to denote their vector product. The spectral norm ($\|A\|_2$) of a matrix A is defined as the square root of the largest eigenvalue of the symmetric matrix $A^T A$, where A^T is the transpose of matrix A. Throughout the rest of the paper, we may omit the subscript in the norm notation for the sake of brevity.

Variance-reduced methods for estimating the gradient vector were originally proposed for the stochastic optimization setting, i.e.,

$$\min_{\theta \in \mathbb{R}^d} F(\theta) = \mathbb{E}_{z \sim p(z)}[f(\theta, z)], \qquad (1)$$

where a sample z is drawn from the distribution $p(z)$ and $f(., z)$s are commonly assumed to be smooth and non-convex functions of θ. This setting is mainly considered in the supervised learning context where θ corresponds to the parameters of the training model and $z = (x, y)$ is the training sample, where x denotes the feature vector of the sample and y is the corresponding label. In this setting, the distribution $p(z)$ is invariant with respect to parameter θ.

In the RL setting, the goal is to maximize the expected cumulative reward:

$$\max_{\theta \in \mathbb{R}^d} J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta}[R(\tau)], \qquad (2)$$

where θ corresponds to the parameters of the policy and trajectory τ is drawn from distribution π_θ. The probability of observing a trajectory τ for a given policy π_θ is:

$$p(\tau|\pi_\theta) = p_0(s_0) \prod_{h=0}^{H-1} P(s_{h+1}|s_h, a_h)\pi_\theta(a_h|s_h). \quad (3)$$

Unlike supervised learning, the distribution of these trajectories depends on the parameters of policy π_θ. It can be shown that:

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta}\left[\sum_{h=0}^{H-1} \Psi_h(\tau)\nabla \log \pi_\theta(a_h|s_h) \right], \quad (4)$$

where $\Psi_h(\tau) = \sum_{t=h}^{H-1} \gamma^t r(s_t, a_t)$. Therefore, for any trajectory τ, $\hat{\nabla} J(\theta, \tau) := \sum_{h=0}^{H-1} \Psi_h(\tau)\nabla \log \pi_\theta(a_h|s_h)$ is an unbiased estimator of $\nabla J(\theta)$. The vanilla policy gradient updates θ as follows:

$$\theta \leftarrow \theta + \eta \hat{\nabla} J(\theta, \tau), \qquad (5)$$

where η is the learning rate.

The Hessian matrix of $J(\theta)$ can be written as follows [Shen et al., 2019]:

$$\nabla^2 J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta}[\nabla \Phi(\theta, \tau)\nabla \log p(\tau|\pi_\theta)^T + \nabla^2 \Phi(\theta, \tau)], \tag{6}$$

where $\Phi(\theta, \tau) = \sum_{h=0}^{H-1} \sum_{t=h}^{H-1} \gamma^t r(s_t, a_t) \log \pi_\theta(a_h|s_h)$ for a given trajectory τ. Hence, $\hat{\nabla}^2 J(\theta, \tau) := \nabla \Phi(\theta, \tau)\nabla \log p(\tau|\pi_\theta)^T + \nabla^2 \Phi(\theta, \tau)$ is an unbiased estimator of the Hessian matrix. Note that the computational complexity of the term $\nabla \Phi(\theta, \tau)\nabla \log p(\tau|\pi_\theta)^T$ is $O(Hd)$ where d is the dimension of the gradient vector. The second term is HVP which can also be computed in $O(Hd)$ using Pearlmutter's algorithm [Pearlmutter, 1994].

We can compute HVP using the above definition for the sample-based version of Hessian matrix $\hat{\nabla}^2 J(\theta, \tau)$ and an arbitrary vector $v \in \mathbb{R}^d$ as follows:

$$\hat{\nabla}^2 J(\theta, \tau)[v] = (\nabla \log p(\tau|\pi_\theta)^T v)\nabla \Phi(\theta, \tau) + \nabla^2 \Phi(\theta, \tau)[v]. \tag{7}$$

2.2 VARIANCE REDUCED METHODS FOR GRADIENT ESTIMATION

For any time $t \geq 1$ and a sequence of parameters $\{\theta_0, \theta_1, \cdots\}$, we can write the gradient at θ_t as follows:

$$\nabla J(\theta_t) = \nabla J(\theta_{t-1}) + \nabla J(\theta_t) - \nabla J(\theta_{t-1}). \quad (8)$$

Suppose that we have an unbiased estimate of $\nabla J(\theta_{t-1})$ at time $t-1$, denoted by v_{t-1}. If we have an unbiased estimate of $\nabla J(\theta_t) - \nabla J(\theta_{t-1})$ denoted by Δ_t, then we can add it to v_{t-1} in order to get an unbiased estimate of $\nabla J(\theta_t)$ at time t as follows:

$$v_t = v_{t-1} + \Delta_t. \qquad (9)$$

Let $\epsilon_t = v_t - \nabla J(\theta_t)$ and $\epsilon_{\Delta_t} = \Delta_t - (\nabla J(\theta_t) - \nabla J(\theta_{t-1}))$. Based on these definitions, we can rewrite the above equation as follows:

$$\epsilon_t = \epsilon_{t-1} + \epsilon_{\Delta_t}. \tag{10}$$

Thus, we have:

$$\mathbb{E}[\|\epsilon_t\|^2] = \mathbb{E}[\|\epsilon_{t-1}\|^2] + \mathbb{E}[\|\epsilon_{\Delta_t}\|^2] + 2\mathbb{E}[\langle\epsilon_{t-1}, \epsilon_{\Delta_t}\rangle]. \tag{11}$$

The above equation shows that if $\mathbb{E}[\langle\epsilon_{t-1}, \epsilon_{\Delta_t}\rangle]$ is sufficiently negative, then $\mathbb{E}[\|\epsilon_t\|^2]$ is decreasing in time. The updates in several previous variance reduction methods (such as in HAPG [Shen et al., 2019], MBPG [Huang et al., 2020], SRVR-PG [Xu et al., 2019], SVRG [Johnson and Zhang, 2013] and SARAH [Nguyen et al., 2017]) are consistent with (9) with different suggestions for Δ_t. Using the update in (9) may accumulate errors in the gradient estimates and some of the methods in the literature require checkpoints after some iterations and use a batch of stochastic gradients at these points to control the error.

2.3 STOCHASTIC CUBIC REGULARIZED NEWTON

In the context of optimization, Cubic Regularized Newton (CRN) [Nesterov and Polyak, 2006] obtains SOSP in general non-convex functions by updating the parameters using the cubic-regularized second-order expansion of the Taylor series at each iteration t as follows:

$$m_t(h) = \langle\nabla F(\theta_t), h\rangle + \frac{1}{2}\langle\nabla^2 F(\theta_t)h, h\rangle + \frac{M}{6}\|h\|^3, \tag{12}$$

$$h_t^* = \mathbf{argmin}_{h\in\mathbb{R}^d} m_t(h), \tag{13}$$

$$\theta_{t+1} = \theta_t + h_t^*.$$

In CRN, the sub-problem for obtaining the minimizer of $m_t(h)$, should be solved in each iteration. Although there is no closed-form solution for the sub-problem, Carmon and Duchi [2016] proposed a gradient descent-based algorithm to find an approximate solution.

As we often only have access to gradient and Hessian estimates of the objective function (herein, denoted by v_t and U_t, respectively), the aforementioned update rule in the context of RL becomes:

$$m_t(h) = \langle v_t, h\rangle + \frac{1}{2}\langle U_t h, h\rangle - \frac{M}{6}\|h\|^3, \tag{14}$$

$$h_t^* = \mathbf{argmax}_{h\in\mathbb{R}^d} m_t(h), \tag{15}$$

$$\theta_{t+1} = \theta_t + h_t^*. \tag{16}$$

3 VR-SCP ALGORITHM

In this section, we first describe our proposed algorithm, VR-SCP, which uses the second-order information in order to have a better estimate of the gradient vector and converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP. Next, we provide a convergence analysis of the proposed algorithm under some customary regularity assumptions in RL literature.

3.1 DESCRIPTION

The algorithm iterates in a loop starting from line 1. In line 2, the estimate of gradient, v_t, is computed. After every Q iterations, we have a checkpoint (when the condition $\mod(t, Q) = 0$ is satisfied) where v_t is set to the average of a batch of stochastic gradients. Otherwise, based on what we explained in (9), the following term will be added to the last gradient estimate v_{t-1}, as an estimate of $\nabla J(\theta_t) - \nabla J(\theta_{t-1})$:

$$\frac{1}{S_t}\sum_{s=1}^{S_t}\hat{\nabla}^2 J(\theta_{s,t}, \tau_s)(\theta_t - \theta_{t-1}), \tag{17}$$

where parameter $\theta_{s,t}$ is a point on the line between θ_{t-1} and θ_t and it is equal to $\left(1 - \frac{s}{S_t}\right)\theta_t + \frac{s}{S_t}\theta_{t-1}$. Furthermore, S_t is the number of points taken on the line, and τ_s is a trajectory drawn based on policy $\pi_{\theta_{s,t}}$. In line 3, VR-SCP defines a stochastic HVP function $U_t[.]$ which for every vector v, it computes $U_t[v] \leftarrow \frac{1}{|\mathcal{B}_h|}\sum_{\tau\in\mathcal{B}_h}\hat{\nabla}^2 J(\theta_t, \tau)[v]$. In line 4, Cubic-Subsolver runs a gradient descent-based algorithm to find an approximate solution of the sub-problem. In line 5, if there is a sufficient increase in $m_t(h)$, (i.e., $m_t(h_t) \geq \rho^{-1/2}\epsilon^{3/2}/6$), the parameter θ_t is updated based on the approximate solution h_t. Otherwise, Cubic-Finalsolver is called once and the algorithm terminates after updating the parameter θ_t. Cubic-Subsolver and Cubic-Finalsolver are based on [Tripuraneni et al., 2018] and the pseudo-codes of these two algorithms are given in Appendix C.

3.2 CONVERGENCE ANALYSIS

In this part, we analyze the convergence rate of the proposed algorithm under the bounded reward function and some common regularity assumptions on the policy.

Assumption 3.1 (Bounded reward). For $\forall s \in \mathcal{S}, \forall a \in \mathcal{A}$, $|R(s, a)| < R_0$ where $R_0 > 0$ is some constant.

Assumption 3.2 (Parameterization regularity). There exist constants $G, L_1, L_2 > 0$ such that for any $\theta_1, \theta_2, w \in \mathbb{R}^d$ and for any $s \in \mathcal{S}, a \in \mathcal{A}$:
(a) $\|\nabla\log\pi_{\theta_1}(a|s)\| \leq G$,
(b) $\|\nabla^2\log\pi_{\theta_1}(a|s)\| \leq L_1$,
(c) $\|(\nabla^2\log\pi_{\theta_1}(a|s) - \nabla^2\log\pi_{\theta_2}(a|s))w\| \leq L_2\|\theta_1 - \theta_2\|\|w\|$.

Assumptions 3.1 and 3.2 (a,b), commonly used in the RL literature to analyze the convergence of policy gradient meth-

Algorithm 1 Variance-reduced stochastic cubic regularized Newton-based policy gradient (VR-SCP)

Input: Batch $\mathcal{B}_{check}$ and $\mathcal{B}_h$, initial point θ_0, accuracy ϵ, cubic penalty parameter M, maximum number of iterations T, duration Q, and parameters L, ρ.

1: **for** $t = 0, \cdots, T - 1$ **do**
2:
$$v_t = \begin{cases} \dfrac{1}{|\mathcal{B}_{check}|} \displaystyle\sum_{\tau \in \mathcal{B}_{check}} \hat{\nabla} J(\theta_t, \tau), & \text{if } \mod(t, Q) = 0, \\[2ex] \dfrac{1}{S_t} \displaystyle\sum_{s=1}^{S_t} \hat{\nabla}^2 J(\theta_{s,t}, \tau_s)(\theta_t - \theta_{t-1}) + v_{t-1}, & \text{o.w.} \end{cases}$$

3: $\quad U_t[.] \leftarrow \frac{1}{|\mathcal{B}_h|} \sum_{\tau \in \mathcal{B}_h} \hat{\nabla}^2 J(\theta_t, \tau)[.]$
4: $\quad h_t \leftarrow$ **Cubic-Subsolver**$(U_t[.], v_t, M, L, \epsilon)$
5: $\quad$ **if** $m_t(h_t) > \rho^{-1/2} \epsilon^{3/2}/6$ **then**
6: $\qquad \theta_{t+1} \leftarrow \theta_t + h_t$
7: $\quad$ **else**
8: $\qquad h_t \leftarrow$ **Cubic-Finalsolver**$(U_t[.], v_t, M, L, \epsilon)$
9: $\qquad$ **return** $\theta_{t+1} \leftarrow \theta_t + h_t$
10: $\quad$ **end if**
11: $\quad t \leftarrow t + 1$
12: **end for**
13: **return** θ_T

---

ods, imply boundedness of the norm of the stochastic gradient as well as the individual smoothness of the stochastic gradient and Hessian.

Lemma 3.3. *[Shen et al., 2019, Wang et al., 2022] Under Assumptions 3.1 and 3.2, for any $\theta_1, \theta_2, w \in \mathbb{R}^d$ and for any trajectory τ, there exist constants W, L and ρ such that:*

$$\|\hat{\nabla} J(\theta_1, \tau)\|_2 \leq W,$$
$$\|\hat{\nabla} J(\theta_1, \tau) - \hat{\nabla} J(\theta_2, \tau)\|_2 \leq L\|\theta_1 - \theta_2\|_2,$$
$$\|\hat{\nabla}^2 J(\theta_1, \tau) - \hat{\nabla}^2 J(\theta_2, \tau)\| \leq \rho\|\theta_1 - \theta_2\|_2,$$

where $\rho = \frac{L_2 R_0 + 2R_0 GHL_1}{(1-\gamma)^2}, L = \frac{L_1 GR_0}{(1-\gamma)^2}, W = \frac{GR_0}{(1-\gamma)^2}$.

Remark 3.4. Assumption 3.2 (c) is also a common in analyzing second-order methods in RL. Examples of policies that satisfy this assumption are Gaussian [Pirotta et al., 2013] and soft-max policies [Masiha et al., 2022]. For instance, consider a Gaussian policy with standard deviation σ as follows:

$$\pi_\theta(a|s) = \mathcal{N}(\theta^T \mu(s), \sigma^2),$$

where $\mu(s) : \mathcal{S} \to \mathbb{R}^d$ is a feature map. It can be easily seen that $\nabla^2 \log \pi_\theta(a|s) = \mu(s)^T \mu(s)/\sigma^2$. Therefore, Gaussian policy satisfies Assumption 3.2 if $\mu(s)$, θ, and actions take values in a bounded domain.

Assumptions 3.1 and 3.2 also imply that the variance of gradient estimate $\hat{\nabla} J(\theta, \tau)$ and Hessian estimate $\hat{\nabla}^2 J(\theta, \tau)$ are bounded.

Lemma 3.5. *[Shen et al., 2019] Under Assumptions 3.1 and 3.2, for any point $\theta \in \mathbb{R}^d$ and trajectory τ, $\hat{\nabla} J(\theta, \tau)$ and $\hat{\nabla}^2 J(\theta, \tau)$, have bounded variances σ_1^2 and σ_2^2, respectively:*

$$\mathbb{E}[\|\hat{\nabla} J(\theta, \tau) - \nabla J(\theta)\|]^2 \leq \sigma_1^2$$
$$\mathbb{E}[\|\hat{\nabla}^2 J(\theta, \tau) - \nabla^2 J(\theta)\|^2] \leq \sigma_2^2, \quad (18)$$

where $\sigma_1^2 = \frac{G^2 R_0^2}{(1-\gamma)^4}$ and $\sigma_2^2 = \frac{H^2 G^4 R_0^2 + L_1^2 R_0^2}{(1-\gamma)^4}$.

Following the work [Nesterov and Polyak, 2006] on cubic regularized Newton method, we define the following quantity for showing convergence to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP.

Definition 3.6. For any $\theta \in \mathbb{R}^d$, we define $\mu(\theta)$ as follows:

$$\mu(\theta) = \max(\|\nabla J(\theta)\|^{3/2}, \lambda_{max}^3(\nabla^2 J(\theta))/\rho^{3/2}). \quad (19)$$

Based on the above definition, it can be easily shown that the point θ is $(\epsilon, \sqrt{\rho\epsilon})$-SOSP if and only if $\mu(\theta) \leq \epsilon^{3/2}$.

To show the convergence of the proposed algorithm, we use the following lemma which provides a bound on the norm of estimation error at each iteration t with high probability given history up to time $\lfloor t/Q \rfloor.Q$.

Lemma 3.7. *Let $\mathcal{F}_t$ be the history up to time t. Under Assumptions 3.1, 3.2 and for the values of S_t in Appendix B.1 and $|\mathcal{B}_{check}|$ in the statement of Theorem 3.9, conditioned on $\mathcal{F}_{\lfloor t/Q \rfloor.Q}$, with probability $1 - 2\delta(t - \lfloor t/Q \rfloor.Q)$, for all i between $\lfloor t/Q \rfloor.Q$ and t, we have:*

$$\|v_i - \nabla J(\theta_i)\|_2^2 \leq \frac{\epsilon^2}{30}. \quad (20)$$

All the proofs of lemmas and theorem appear in the Appendix.

Lemma 3.8. *For the value of $|\mathcal{B}_h|$ in the statement of Theorem 3.9, under Assumptions 3.1, 3.2, conditioned on $\mathcal{F}_t$, with probability $1 - \delta$, we have:*

$$\|U_t - \nabla^2 J(\theta_t)\|^2 \leq \frac{\epsilon\rho}{30}. \quad (21)$$

Theorem 3.9. *Under Assumptions 3.1, 3.2, for $|\mathcal{B}_{check}| \geq \frac{19440 W^2 \log^2(4T/\xi)}{\epsilon^2}$, $|\mathcal{B}_h| \geq \frac{1080 L^2 \log(4dT/\xi)}{\rho\epsilon}$, and S_t defined in Appendix B.1, cubic penalty parameter $M = 4\rho$, $\epsilon \leq 4L^2 \rho/M$ and maximum number of iterations $T \geq 25\Delta_J \rho^{1/2} \epsilon^{-3/2}$, Algorithm 1 guarantees that:*

$$\mu(\theta_{out}) \leq 1300 \epsilon^{3/2}, \quad (22)$$

with the probability $1 - \xi$ where $\Delta_J = J^* - J(\theta_0)$ and J^* is the optimal value of objective function.

Proof sketch. To prove Theorem 3.9, we establish bounds on $\|v_t - \nabla J(\theta_t)\|^2$ and $\|U_t - \nabla^2 J(\theta_t)\|^2$. A key step involves bounding the term

$$\frac{1}{S_t}\left\|\sum_{s=1}^{S_t}\hat{\nabla}^2 J(\theta_{s,t})(\theta_t - \theta_{t-1}) - \nabla J(\theta_t) + \nabla J(\theta_{t-1})\right\|,$$

which we show is a quadratic function of $\|\theta_t - \theta_{t-1}\|$ with high probability (see (53) in the appendix). Using this result, we adjust the batch sizes to derive high-probability bounds on $\|v_t - \nabla J(\theta_t)\|^2$ (Lemma 3.7) and $\|U_t - \nabla^2 J(\theta_t)\|^2$ (Lemma 3.8). Substituting these bounds into the recursive inequality in (77), we show that the proposed algorithm achieves the convergence rate stated in Theorem 3.9.

Corollary 3.10. *Under the Assumptions in the statement of Theorem 3.9, and $Q = \frac{\sqrt{\rho}M}{\sqrt{\epsilon}L}$, Algorithm 1 will return $(\epsilon, \sqrt{\rho\epsilon})$-SOSP after observing $\tilde{O}(\frac{1}{\epsilon^3})$ trajectories. Thus, it improves the best-known sample complexity in [Wang et al., 2022, Maniyar et al., 2024] by a factor of $O(\epsilon^{-0.5})$.*

4 RELATED WORK

Variance-reduced PG methods have been proposed in the RL literature to reduce the variance of the stochastic gradient and improve the training process [Papini et al., 2018, Xu et al., 2019, Zhang et al., 2021]. Variance-reduced methods such as SARAH [Nguyen et al., 2017], SAGA [Defazio et al., 2014], and SVRG [Johnson and Zhang, 2013] proposed in the context of stochastic optimization, are the basis of some more recent variance-reduced policy gradient methods in RL. These approaches use the difference $\nabla f(\theta', z) - \nabla f(\theta, z)$ for two consecutive points θ and θ' along iterations for the same randomness z to reduce the variance of gradient estimates. This difference can be easily computed in tasks such as supervised learning as the randomness z does not depend on the parameters to be optimized (e.g., parameters of the decision rule in supervised learning). However, in the RL setting, the distribution over trajectories depends on the parameters of the current policy. Thus, most variance-reduced methods in RL require IS weights in order to provide unbiased estimates of $\nabla J(\theta)$ for a given point θ based on a trajectory that is generated by a policy with another parameter θ'. As a result, the convergence analysis for these methods requires strong assumptions such as the boundedness of variance of IS weights. Examples of methods requiring such strong assumptions are SRVR-PG [Xu et al., 2020], ProxHSPGA [Pham et al., 2020], IS-MBPG [Huang et al., 2020], and PAGE-PG [Gargiani et al., 2022] which all of them converge to an ϵ-FOSP with the sample complexity of $O(\epsilon^{-3})$. HAPG is the first work achieving the same rate by using second-order information (in the form of HVP) to bypass IS weights. However, in their analysis, it is required to use a fixed step size of ϵ which slows the training process in practice. Moreover, the number of computed HVPs per iteration is in the order of $O(1/\epsilon)$. Very recently, two methods using mirror-descent algorithm based on Bregman divergence, called VR-MPO [Yang et al., 2022] and VR-BGPO [Huang et al., 2021] have been proposed. These methods achieve ϵ-FOSP if the mirror map in Bregman divergence is the l_2-norm.

In order to converge to SOSP, in the context of optimization, Nesterov and Polyak [2006] analyzed the cubic-regularized Newton (CRN) method in the deterministic case. In this case, a sub-problem formulated based on a second-order Taylor expansion of the objective function with a cubic penalty term is solved at each iteration. CRN uses the full Hessian matrix in the minimization of sub-problem which is not practical in large-scale applications. More recently, Carmon and Duchi [2016] proposed a gradient descent method to find ϵ-global optimum point of the sub-problem. The proposed method only requires computing HVPs (without requiring the full Hessian matrix) which can be computed with the same computational complexity of obtaining stochastic gradients [Pearlmutter, 1994]. In the stochastic setting, Tripuraneni et al. [2018] proposed stochastic CRN (SCRN) that uses sub-sampled gradients and Hessian vector products to solve the sub-problem at each iteration and converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-3.5})$. Later, Zhou and Gu [2020] introduced a variance-reduced version of SCRN with sample complexity of $\tilde{O}(\epsilon^{-3})$. In the context of RL, Yang et al. [2021] analyzed REINFORCE under some restrictive assumptions on the objective function and showed that it convergences to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-4.5})$. Recently, Wang et al. [2022] introduced a stochastic cubic-regularized policy gradient method that achieves $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with the sample complexity of $\tilde{O}(\epsilon^{-3.5})$. This method still requires IS weights and consequently strong assumptions about them. Concurrent to Wang et al. [2022], Maniyar et al. [2024] proposed ACR-PN, a second-order method which achieves $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with a sample complexity of $O(\epsilon^{-3.5})$. Our proposal, VR-SCP algorithm, does not require IS and achieves $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with sample complexity of $\tilde{O}(\epsilon^{-3})$. To the best of our knowledge, VR-SCP is the first algorithm in the literature to provide such a guarantee.

5 EXPERIMENTS

In this section, we evaluate the proposed algorithm and compare it with the related work in four control tasks in MuJoCo [Todorov et al., 2012] which is a physics simulator, with fast and accurate simulations in areas such as robotics, biomechanics, graphics, etc. We used Garage library [garage contributors, 2019] for our implementations as it allows for maintaining and integrating several RL algorithms[3].

---

[3]The code for all experiments is available at `https://github.com/sadegh16/VR-SCP`.

In the following, we briefly explain the control tasks in the four environments we consider. In Walker environment, a humanoid walker tries to move forward in a two-dimensional space. It can only fall forward or backward and the goal is to walk as long as possible without falling. In Reacher environment, there is an arm that tries to reach a specific point in a plane and the goal is to navigate the arm such that the tip of the arm gets as close as possible to that point. In Hopper environment, there is a two-dimensional one-legged robot and the goal is to make it hop in the forward (right) direction as long as possible. In Humanoid environment, there is a 3D bipedal robot like a human and the goal is to make it walk forward as fast as possible without falling over.

We compared our proposed algorithm with PG methods that provide theoretical guarantees: PAGE-PG [Gargiani et al., 2022], IS-MBPG [Huang et al., 2020] which is based on STROM, HAPG [Shen et al., 2019] which does not require IS weights, VR-BGPO [Huang et al., 2021] which is a mirror descent based algorithm, and ACR-PN [Maniyar et al., 2024] which is the state-of-the-art of second-order methods in the RL setting. These algorithms have guarantees on convergence to an approximate FOSP. We also considered REINFORCE [Sutton et al., 2000] as a baseline algorithm. There are some other approaches in the literature with theoretical guarantees: (SCR-PG [Wang et al., 2022], VRMPO [Yang et al., 2022], and STORM-PG [Ding et al., 2021]) but the official implementations are not publicly available and our request to access the code from the authors remained unanswered.

For each algorithm, we use the same set of Gaussian policies parameterized with neural networks consisting of two layers of 64 neurons each. Baselines and environment settings (such as maximum trajectory horizon, and reward scale) are considered the same for all algorithms. We chose a maximum horizon of 500 for Walker, Hopper, and Humanoid and 50 for Reacher. More details about experiments are given in Appendix D. It is worth mentioning that we used the official implementation of each algorithm. For REINFORCE, we used the implementation provided by Garage library [garage contributors, 2019]. For PAGE-PG, we found some issues in the official code and we decided to carefully implement it following the description of the original paper. The implementation of VR-SCP (our algorithm) is available as supplementary material.

There are two main challenges in evaluating PG methods experimentally. First, it has been observed that PG methods are often sensitive to parameter initialization and random seeds [Henderson et al., 2018]. Thus, it is quite difficult in some cases to reproduce previous results. Second, it is unclear how to compare algorithms in terms of performance (e.g., the average return over instances of an algorithm) and robustness (e.g., the standard deviation of return (STD return) over instances of an algorithm) simultaneously. For

Table 1: Comparison of VR-SCP with other variance-reduced methods in terms of PR. In each environment, the highest PR is in bold.

	Reacher	Walker	Humanoid	Hopper
VR-SCP (our algorithm)	**-10.88**	**486.08**	**484.19**	**891.48**
HAPG	-19.51	104.296	161.12	74.60
IS-MBPG	-21.76	204.32	201.01	312.21
PAGE-PG	-17.39	247.58	356.58	338.72
REINFORCE	-20.10	79.98	156.21	263.92
VR-BGPO	-15.15	320.51	409.67	861.92
ACR-PN	-17.19	200.84	260.09	74.05

instance, between an algorithm, with both a high average return and STD return and another algorithm with a lower average return but with also a lower STD return, which one is preferable?


Figure 1: Each configuration is evaluated with five different random seeds.

In this section, we aim to find a metric to simultaneously capture the average return as well as STD return of an algorithm. Additionally, we would like to assess an algorithm's sensitivity to random seeds, which is central for reproducibility. To address the first question, we define the following metric. For any algorithm A, after observing t number of state-actions pairs (called system probes), we compute the lower bound of the confidence interval of average return over n runs of the algorithm and denote it by $LCI_A(n, t)$. We define the performance-robustness (PR) metric by averaging $LCI_A(n, t)$ over all system probes $t = 1, \cdots, T$ as follows:

$$PR_A(n) = \frac{1}{T} \sum_{t=1}^{T} LCI_A(n, t), \qquad (23)$$

where T is the maximum number of system probes. Figure 1 illustrates why we take the average of $LCI_A(n, t)$ over all system probes $t = 1, \cdots, T$ in the definition of our PR metric. In this figure, four different configurations of

Figure 2: Comparison of VR-SCP with other variance reduction methods on four control tasks.

hyper-parameters for IS-MBPG algorithm [Huang et al., 2020] are considered in Walker environment. PR values for these configurations are $234.8, 246.9, 193.5$, and 199.1, respectively. Taking the average of $LCI_A(n, t)$ over the probes prevents us from choosing the hyper-parameters that aggressively improve the returns in the beginning but can degrade drastically by the end of the horizon.

We used grid search to tune the hyper-parameters of all the algorithms. For the algorithms except ours, the search space for each hyper-parameter was chosen based on the one from the original papers. For each configuration of the hyper-parameters, we ran each algorithm A, five times and computed $PR_A(5)$. We selected the configuration that maximized $PR_A(5)$ and then reported $PR_A(10)$ of each algorithm for the selected configuration based on 10 different runs in Table 1. Our proposed method achieved the highest PR in all environments compared to the other algorithms.

To address the second question, we consider the confidence interval of the performance to gauge the sensitivity of an algorithm to random seeds. In Figure 2, VR-SCP exhibits better performance and robustness compared to the other algorithms. It has relatively less variance as well as a higher average return. This could be explained by the fact that as VR-SCP converges to an $(\epsilon, \sqrt{\rho\epsilon})$-SOPS, it is less sensitive to random seeds which results in smaller confidence intervals. Additionally, as it avoids saddle points and possibly bad local optima it achieves a higher average return. This

is most evident in Walker environment, where the other algorithms get stuck in a saddle point or local maxima but VR-SCP escapes them. It is worth mentioning that, in our experiments, REINFORCE has comparable performance to some variance-reduced PG methods, while in previous work, the performance of REINFORCE was often reported much worse, which might have been due to poor tuning of its hyper-parameters.

6 CONCLUSION

We proposed a variance-reduced cubic regularized Newton PG method that uses second-order information in the form of HVP and converges to $(\epsilon, \sqrt{\rho\epsilon})$-SOSP with the sample complexity of $\tilde{O}(\epsilon^{-3})$. Such a guarantee ensures escaping saddle points and bad local optima. Our rate improves the best-known sample complexity by a factor of $O(\epsilon^{-0.5})$. Moreover, unlike most methods in the literature, we do not require IS weights in the variance reduction part. In the experiments, due to the sensitivity of the algorithms to random seeds, we defined a new measure to account simultaneously for the performance (average return) and robustness (standard deviation of return) of an algorithm and used it in our evaluations. Experimental results clearly showcase the advantage of VR-SCP both in terms of its performance and robustness in a variety of control tasks compared with other methods in the literature.

References

Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.

Yair Carmon and John C Duchi. Gradient descent efficiently finds the cubic-regularized non-convex newton step. *arXiv preprint arXiv:1612.00547*, 2016.

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.

Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2):388–403, 2013.

Yuhao Ding, Junzi Zhang, and Javad Lavaei. On the global convergence of momentum-based policy gradient. *arXiv preprint arXiv:2110.10116*, 2021.

The garage contributors. Garage: A toolkit for reproducible reinforcement learning research. `https://github.com/rlworkgroup/garage`, 2019.

Matilde Gargiani, Andrea Zanelli, Andrea Martinelli, Tyler Summers, and John Lygeros. Page-pg: A simple and loopless variance-reduced policy gradient method with probabilistic gradient estimation. *arXiv preprint arXiv:2202.00308*, 2022.

Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

Feihu Huang, Shangqian Gao, Jian Pei, and Heng Huang. Momentum-based policy gradient methods. In *International Conference on Machine Learning*, pages 4422–4433. PMLR, 2020.

Feihu Huang, Shangqian Gao, and Heng Huang. Bregman gradient policy optimization. In *International Conference on Learning Representations*, 2021.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26:315–323, 2013.

Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.

Mizhaan P Maniyar, LA Prashanth, Akash Mondal, and Shalabh Bhatnagar. A cubic-regularized policy newton algorithm for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 4708–4716. PMLR, 2024.

Saeed Masiha, Saber Salehkaleybar, Niao He, Negar Kiyavash, and Patrick Thiran. Stochastic second-order methods provably beat sgd for gradient-dominated functions. In *Advances in neural information processing systems*, 2022.

Yurii Nesterov and Boris T Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.

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 Pirotta, and Marcello Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pages 4026–4035. PMLR, 2018.

Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.

Nhan Pham, Lam Nguyen, Dzung Phan, Phuong Ha Nguyen, Marten Dijk, and Quoc Tran-Dinh. A hybrid stochastic policy gradient algorithm for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 374–385. PMLR, 2020.

Matteo Pirotta, Marcello Restelli, and Luca Bascetta. Adaptive step-size for policy gradient methods. *Advances in Neural Information Processing Systems*, 26, 2013.

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, 2015a.

John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015b.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.

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.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.

Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

Nilesh Tripuraneni, Mitchell Stern, Chi Jin, Jeffrey Regier, and Michael I Jordan. Stochastic cubic regularization for fast nonconvex optimization. *Advances in neural information processing systems*, 31, 2018.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12: 389–434, 2012.

Pengfei Wang, Hongyu Wang, and Nenggan Zheng. Stochastic cubic-regularized policy gradient method. *Knowledge-Based Systems*, 255:109687, 2022.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

Pan Xu, Felicia Gao, and Quanquan Gu. Sample efficient policy gradient methods with recursive variance reduction. *arXiv preprint arXiv:1909.08610*, 2019.

Pan Xu, Felicia Gao, and Quanquan Gu. An improved convergence analysis of stochastic variance-reduced policy gradient. In *Uncertainty in Artificial Intelligence*, pages 541–551. PMLR, 2020.

Long Yang, Qian Zheng, and Gang Pan. Sample complexity of policy gradient finding second-order stationary points. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10630–10638, 2021.

Long Yang, Yu Zhang, Gang Zheng, Qian Zheng, Pengfei Li, Jianhang Huang, and Gang Pan. Policy optimization with stochastic mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8823–8831, 2022.

Junyu Zhang, Chengzhuo Ni, Zheng Yu, Csaba Szepesvari, and Mengdi Wang. On the convergence and sample efficiency of variance-reduced policy gradient method. *arXiv preprint arXiv:2102.08607*, 2021.

Dongruo Zhou and Quanquan Gu. Stochastic recursive variance-reduced cubic regularization methods. In *International Conference on Artificial Intelligence and Statistics*, pages 3980–3990. PMLR, 2020.

Efficiently Escaping Saddle Points for Policy Optimization (Supplementary Material)

Sadegh Khorasani[1]   **Saber Salehkaleybar**[2]   **Negar Kiyavash**[3]   **Niao He**[4]   **Matthias Grossglauser**[1]

[1]School of Computer and Communication Sciences, EPFL,, Lausanne, Switzerland
[2]Leiden Institute of Advanced Computer Science (LIACS), Leiden University, , Leiden, The Netherlands
[3]College of Management of Technology, EPFL, Lausanne, Switzerland
[4]Department of Computer Science ETH Zurich

A   AUXILIARY LEMMAS

Lemma A.1. *(Vector Azuma-Hoeffding inequality) Let a_k be a vector martingale difference, that is,* $\mathbf{E}[a_k|\sigma(a_1, \cdots, a_{k-1})] = 0$ *where* $\sigma(a_1, \cdots, a_{k-1})$ *is* σ*-algebra of* $a_1, \cdots, a_{k-1}$ *and* $\|a_k\| \leq A_k$. *With probability at least* $1 - \delta$,

$$\left\| \sum_k a_k \right\| \leq 3\sqrt{\log(1/\delta) \sum_k A_k^2}. \tag{24}$$

Lemma A.2. *(Tropp [2012])(Matrix Azuma inequality) Consider a finite adapted sequence X_k of self-adjoint matrices in dimension d, and a fixed sequence $\{A_k\}$ of self-adjoint matrices that* $\mathbf{E}[X_k|\sigma(X_1, \cdots, X_{k-1})] = 0$ *where* $\sigma(X_1, \cdots, X_{k-1})$ *is* σ*-algebra of* $X_1, \cdots, X_{k-1}$ *and* $X_k^2 \preceq A_k^2$ *almost surly. Then with probability at least* $1 - \delta$,

$$\left\| \sum_k X_k \right\|_2 \leq 3\sqrt{\log(d/\delta) \sum_k \|A_k\|_2^2}. \tag{25}$$

Lemma A.3. *[Carmon and Duchi, 2016] Consider the following function:*

$$m(h) = \langle v, h \rangle + \frac{1}{2}\langle Uh, h \rangle - \frac{M}{6}\|h\|^3, \tag{26}$$

and its maximizer

$$h^* = \mathbf{argmax}_{h \in \mathbb{R}^d} m(h), \tag{27}$$

where $v \in \mathbb{R}^d, U \in \mathbb{R}^{d \times d}$ such that $\|U\|_2 \leq L$, and M is a positive constant.

If either of the following conditions holds:
1- $\sqrt{\epsilon/\rho} \leq \|h^*\|$
2- $\|v\| \geq \max\left(\frac{M\epsilon}{2\rho}, \sqrt{\frac{LM}{2}} \left(\frac{\epsilon}{\rho} \right)^{3/4} \right)$,

then, for $\epsilon \leq \frac{16L^2\rho}{M^2}$ and $\mathcal{T}(\epsilon) \geq C_s \frac{L}{M\sqrt{\epsilon/\rho}}$ (where C_s is a constant), with at least probability $1 - \delta$, the Cubic Sub-solver (see Algorithm 2) returns an $\hat{h}$ such that:

$$m(\hat{h}) \geq \frac{M\rho^{-3/2}\epsilon^{3/2}}{24}.$$

Lemma A.4. *[Carmon and Duchi, 2016] For any $v \in \mathbb{R}^d$, and positive scalars M and L, we define:*

$$R = \frac{L}{2M} + \sqrt{(\frac{L}{2M}) + \|v\|/M}. \tag{28}$$

For $m(h)$ and h^* defined in (26) and (27), respectively, 1) if $\|U\|_2 \leq L$, then $m(h)$ is $(L + 2MR)$-smooth, 2) at each iteration of Cubic-Finalsolver (Algorithm 3), $\|\Delta\| \leq \|h^*\|$ where Δ is the update vector defined in the Cubic-Finalsolver (see Algorithm 3).

Lemma A.5. *For any $v_t \in \mathbb{R}^d, U_t \in \mathbb{R}^{d \times d}, h \in \mathbb{R}^d, M \in \mathbb{R}$ such that $M/\rho \geq 2$, we have:*

$$\mu(\theta_t + h) \leq 9\left[M^3\rho^{-3/2}\|h\|^3 + M^{3/2}\rho^{-3/2}\left\|\nabla J(\theta_t) - v_t\right\|^{3/2} \right. \tag{29}$$

$$\left. + \rho^{-3/2}\left\|\nabla^2 J(\theta_t) - U_t\right\|^3 + M^{3/2}\rho^{-3/2}\|\nabla m_t(h)\|^{3/2} + M^{3/2}\rho^{-3/2}\left|\|h\| - \|h_t^*\|\right|^3 \right],$$

where $m_t(h)$ and h_t^* are defined in (14) and (15), respectively.

Proof: Lemma 3.3(c) shows the individual Lipschitz continuity of the Hessian estimator $\hat{\nabla}^2 J(\theta, \tau)$, i.e., for any $\theta_1, \theta_2 \in \mathbb{R}^d$,

$$\|\hat{\nabla}^2 J(\theta_1, \tau) - \hat{\nabla}^2 J(\theta_2, \tau)\| \leq \rho\|\theta_1 - \theta_2\|_2. \tag{30}$$

Based on above inequality, we can imply the Lipschitz continuity of the Hessian $\nabla^2 J(\theta)$ as follows:

$$\begin{aligned}
\|\nabla^2 J(\theta_1) - \nabla^2 J(\theta_2)\| &= \|\mathbb{E}_\tau[\hat{\nabla}^2 J(\theta_1, \tau)] - \mathbb{E}_\tau[\hat{\nabla}^2 J(\theta_2, \tau)]\| \\
&\leq \mathbb{E}_\tau[\|\hat{\nabla}^2 J(\theta_1, \tau) - \hat{\nabla}^2 J(\theta_2, \tau)\|] \\
&\leq \rho\|\theta_1 - \theta_2\|_2,
\end{aligned} \tag{31}$$

where the first inequality is due to Jensen's inequality (as the operator norm is convex) and the second inequality is according to (30).

Let $\{\lambda_1, \cdots, \lambda_d\}$ be the eigenvalues of $\nabla^2 J(\theta_1) - \nabla^2 J(\theta_2)$. By the definition of the operator norm: $\|\nabla^2 J(\theta_1) - \nabla^2 J(\theta_2)\| = \max_i |\lambda_i|$. According to (31), we have: $\rho\|\theta_1 - \theta_2\|_2 \geq |\lambda_i|$ for all eigenvalues λ_i's. Therefore, we can imply that $\rho\|\theta_1 - \theta_2\|\mathbf{I} - \nabla^2 J(\theta_1) + \nabla^2 J(\theta_2)$ is positive semi-definite, i.e.,

$$\nabla^2 J(\theta_1) - \nabla^2 J(\theta_2) \preceq \rho\|\theta_1 - \theta_2\|\mathbf{I}, \tag{32}$$

where the notation $\preceq$ denotes the Loewner order, that is, for two real symmetric matrices A and B, $A \preceq B$ if $B - A$ is positive semi-definite.

Now, we have:

$$\nabla^2 J(\theta_t + h) \preceq \nabla^2 J(\theta_t) + \rho\|h\|\mathbf{I} \tag{33}$$

$$\preceq \left\|\nabla^2 J(\theta_t) - U_t\right\|\mathbf{I} + U_t + \rho\|h\|\mathbf{I} \tag{34}$$

$$\preceq \left\|\nabla^2 J(\theta_t) - U_t\right\|\mathbf{I} + M/2\|h_t^*\|\mathbf{I} + \rho\|h\|\mathbf{I}, \tag{35}$$

where the first inequality is due to the Lipschitz continuity of the Hessian $\nabla^2 J(\theta)$ and the second inequality holds as $\nabla^2 J(\theta_t) - U_t \preceq \|\nabla^2 J(\theta_t) - U_t\|I$. Moreover, the last inequality is due to the fact that the second-order derivative of $m_t(h)$, is negative semi-definite at the optimal point h_t^*, i.e., $U_t - M/2\|h_t^*\| \preceq 0$. Having $\rho \leq M/2$, from the above inequality, we can imply that

$$\lambda_{max}(\nabla^2 J(\theta_t + h)) \leq \left\|\nabla^2 J(\theta_t) - U_t\right\| + M/2\|h_t^*\| + \rho\|h\| \tag{36}$$

$$\leq \left\|\nabla^2 J(\theta_t) - U_t\right\| + M\|h\| + M\left|\|h_t^*\| - \|h\|\right|, \tag{37}$$

where we applied triangle inequality in the last step, i.e., $M/2\|h_t^*\| \leq M/2\|h\| + M/2\left|\|h_t^*\| - \|h\|\right| \leq M/2\|h\| + M\left|\|h_t^*\| - \|h\|\right|$. Accordingly, we have:

$$\rho^{-3/2}\lambda_{max}^3(\nabla^2 J(\theta_t + h)) \leq 9\rho^{-3/2}\left[\left\|\nabla^2 J(\theta_t) - U_t\right\|^3 + M^3\|h\|^3 + M^3\left|\|h_t^*\| - \|h\|\right|^3\right], \tag{38}$$

where we used the inequality $(a + b + c)^3 \leq 9(a^3 + b^3 + c^3)$ for any $a, b, c \geq 0$.

Now, we obtain a bound on the norm of the gradient:

$$
\begin{aligned}
\left\| \nabla J(\theta_t + h) \right\| &\leq \left\| \nabla J(\theta_t + h) - \nabla J(\theta_t) - \nabla^2 J(\theta_t) h \right\| + \left\| \nabla J(\theta_t) - v_t \right\| \\
&\quad + \left\| \nabla^2 J(\theta_t) h - U_t h \right\| + \left\| v_t + U_t h - M/2 \|h\| h \right\| + M/2 \|h\|^2 \\
&\overset{(a)}{=} \left\| \nabla J(\theta_t + h) - \nabla J(\theta_t) - \nabla^2 J(\theta_t) h \right\| + \left\| \nabla J(\theta_t) - v_t \right\| \\
&\quad + \left\| \nabla^2 J(\theta_t) h - U_t h \right\| + \left\| \nabla m_t(h) \right\| + M/2 \|h\|^2 \\
&\overset{(b)}{\leq} \frac{M}{4} \|h\|^2 + \left\| \nabla J(\theta_t) - v_t \right\| + \left\| \nabla^2 J(\theta_t) h - U_t h \right\| \\
&\quad + \left\| \nabla m_t(h) \right\| + M/2 \|h\|^2 \\
&\overset{(c)}{\leq} \left\| \nabla J(\theta_t) - v_t \right\| + 1/M \left\| \nabla^2 J(\theta_t) - U_t \right\|^2 + \|\nabla m_t(h)\| + M \|h\|^2,
\end{aligned}
\tag{39}
$$

(a) Due to the definition of $m_t(h)$, we have: $\left\| v_t + U_t h - M/2 \|h\| h \right\| = \|\nabla m_t(h)\|$.

(b) According to Hessian Lipschitzness, we have: $\left\| \nabla J(\theta_t + h) - \nabla J(\theta_t) - \nabla^2 J(\theta_t) h \right\| \leq \rho/2 \|h\|^2 \leq M/4 \|h\|^2$ where we assumed that $M/\rho \geq 2$.

(c) We can bound the term $\|\nabla^2 J(\theta_t) h - U_t h\|$ using Cauchy-Schwarz and Young's inequalities as follows:

$$
\left\| (\nabla^2 J(\theta_t) - U_t) h \right\| \leq \left\| \nabla^2 J(\theta_t) - U_t \right\| \|h\| \leq 1/M \left\| \nabla^2 J(\theta_t) - U_t \right\|^2 + M/4 \|h\|^2.
\tag{40}
$$

Accordingly, we have:

$$
\|\nabla J(\theta_t + h)\|^{3/2} \leq 2 \left[\left\| \nabla J(\theta_t) - v_t \right\|^{3/2} + M^{-3/2} \left\| \nabla^2 J(\theta_t) - U_t \right\|^3 + \|\nabla m_t(h)\|^{3/2} + M^{3/2} \|h\|^3 \right],
\tag{41}
$$

where we used the inequality $(a + b + c + d)^{3/2} \leq 2(a^{3/2} + b^{3/2} + c^{3/2} + d^{3/2})$ for any $a, b, c, d \geq 0$.

Using (38) and (41) and according to Definition 3.6, we obtain an upper bound on $\mu(\theta_t + h)$:

$$
\mu(\theta_t + h) \leq 9 \left[M^3 \rho^{-3/2} \|h\|^3 + M^{3/2} \rho^{-3/2} \left\| \nabla J(\theta_t) - v_t \right\|^{3/2} \right.
\tag{42}
$$

$$
\left. + \rho^{-3/2} \left\| \nabla^2 J(\theta_t) - U_t \right\|^3 + M^{3/2} \rho^{-3/2} \|\nabla m_t(h)\|^{3/2} + M^3 \rho^{-3/2} \left| \|h\| - \|h_t^*\| \right|^3 \right],
$$

which is derived by multiplying the right-hand side of (41) with $1/8 M^{3/2} \rho^{-3/2}$ (note that $1/8 M^{3/2} \rho^{-3/2} \geq 1$ as $4\rho \leq M$) and then sum it with (38).

Lemma A.6. *For any $h \in \mathbb{R}^d$, we have:*

$$
-\frac{\rho}{8} \|h\|^3 - \frac{10}{\rho^2} \left\| \nabla^2 J(\theta_t) - U_t \right\|^3 \leq \left\langle (\nabla^2 J(\theta_t) - U_t) h, h \right\rangle,
\tag{43}
$$

$$
-\frac{\rho}{8} \|h\|^3 - \frac{6}{5\sqrt{\rho}} \left\| \nabla J(\theta_t) - v_t \right\|^{3/2} \leq \left\langle \nabla J(\theta_t) - v_t, h \right\rangle.
\tag{44}
$$

Proof: Using Cauchy-Schwarz inequality, we have:

$$
-\|h\|^2 \left\| \nabla^2 J(\theta_t) - U_t \right\| \leq \left\langle (\nabla^2 J(\theta_t) - U_t) h, h \right\rangle
\tag{45}
$$

On the other hand, using Young's inequality $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$, where $a = (\frac{3\rho}{2^4})^{2/3} \|h\|^2$, $b = (\frac{2^4}{3\rho})^{2/3} \left\|\nabla^2 J(\theta_t) - U_t\right\|$, $p = 3/2$ and $q = 3$, we have:

$$\|h\|^2 \left\|\nabla^2 J(\theta_t) - U_t\right\| \leq \frac{\rho}{8}\|h\|^3 + \frac{10}{\rho^2}\left\|\nabla^2 J(\theta_t) - U_t\right\|^3. \tag{46}$$

Putting (45) and (46) together, we derive (43). Similarly, we can use the following inequalities to derive (44),

$$-\|h\|\left\|\nabla J(\theta_t) - v_t\right\| \leq \left\langle \nabla J(\theta_t) - v_t, h \right\rangle \tag{47}$$

$$\|h\|\left\|\nabla J(\theta_t) - v_t\right\| \leq \frac{\rho}{8}\|h\|^3 + \frac{6}{5\sqrt{\rho}}\left\|\nabla J(\theta_t) - v_t\right\|^{3/2}, \tag{48}$$

where in the second inequality, we used Young's inequality $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$ where $a = (\frac{3\rho}{2^3})^{1/3}\|h\|$, $b = (\frac{2^3}{3\rho})^{1/3}\left\|\nabla^2 J(\theta_t) - U_t\right\|$, $p = 3$ and $q = 3/2$.

Lemma A.7. *[Zhang et al., 2021] For $\epsilon \leq 4L^2\rho/M$, Cubic-Finalsolver (Algorithm 3) will terminate within $C_F L/\sqrt{\rho\epsilon}$ iterations, where C_F is a constant.*

B   CONVERGENCE ANALYSIS

Consider $\xi = 4T\delta$. Let us define $\mathcal{F}_t$ as the history up to the point θ_t (i.e., σ-algebra of θ_0 to θ_t, $\mathcal{F}_t = \sigma(\theta_0, \theta_1, \cdots, \theta_t)$). We show the gradient estimate and sub-sampled Hessian of the objective function with v_t and U_t, respectively.

B.1   PROOF OF LEMMA 3.7

Lemma B.1. *Under Assumptions 3.1, 3.2, conditioned on $\mathcal{F}_{\lfloor t/Q \rfloor.Q}$, with probability at least $1 - 2\delta(t - \lfloor t/Q \rfloor.Q)$ for all i where $\lfloor t/Q \rfloor.Q \leq i \leq t$, we have:*

$$\left\|v_i - \nabla J(\theta_i)\right\|_2^2 \leq \frac{\epsilon^2}{30}, \tag{49}$$

where for any two consecutive points θ_{k-1} and θ_k in Algorithm 1, we set $S_k = C_2 Q \frac{\|\theta_k - \theta_{k-1}\|_2^2}{\epsilon^2}$, where $C_2 = 38880 L^2 \log^2(1/\delta) + 4^{5/4}\sqrt{1080\log(1/\delta)}\rho^2 L_1^3 M^{-7/4}$.

Proof: We have $v_t - \nabla J(\theta_t) = \sum_{k=\lfloor t/Q \rfloor.Q}^{t} u_k$ where:

$$u_k = \begin{cases} \frac{1}{S_k}\sum_{s=1}^{S_k} \hat{\nabla}^2 J(\theta_{s,k}, \tau_s)(\theta_k - \theta_{k-1}) - \nabla J(\theta_k) + \nabla J(\theta_{k-1}) & t \geq k > \lfloor t/Q \rfloor.Q \\ \hat{\nabla} J(\theta_k, \mathcal{B}_{check}) - \nabla J(\theta_k) & k = \lfloor t/Q \rfloor.Q, \end{cases} \tag{50}$$

where $\theta_{s,k} = (1 - \frac{s}{S_k})\theta_k + \frac{s}{S_k}\theta_{k-1}$. We now find a bound on the norm of u_k in both cases above. For $k > \lfloor t/Q \rfloor.Q$, we rewrite u_k as follows:

$$u_k = \underbrace{\frac{1}{S_k}\sum_{s=1}^{S_k} \hat{\nabla}^2 J(\theta_{s,k}, \tau_s)(\theta_k - \theta_{k-1}) - \nabla^2 J(\theta_{s,k})(\theta_k - \theta_{k-1})}_{u_k^a}$$

$$+ \underbrace{\frac{1}{S_k}\sum_{s=1}^{S_k} \nabla^2 J(\theta_{s,k})(\theta_k - \theta_{k-1}) - \nabla J(\theta_k) + \nabla J(\theta_{k-1})}_{u_k^b}.$$

To get the upper bound on $\|u_k\|$, we find an upper bound on the $\|u_k^a\|$ and $\|u_k^b\|$. For $\|u_k^a\|$, we define a_s as follows:

$$a_s = \hat{\nabla}^2 J(\theta_{s,k}, \tau_s)(\theta_k - \theta_{k-1}) - \nabla^2 J(\theta_{s,k})(\theta_k - \theta_{k-1}).$$

Due to individual gradient Lipschitzness (see Lemma 3.3), we have: $\|\hat{\nabla}^2 J(\theta_{s,k}, \tau_s)\| \leq L$. Moreover, using Jensen's inequality, $\|\nabla^2 J(\theta_{s,k})\| = \|\mathbb{E}[\hat{\nabla}^2 J(\theta_{s,k}, \tau_s)]\| \leq \mathbb{E}[\|\hat{\nabla}^2 J(\theta_{s,k}, \tau_s)\|] \leq L$. Therefore,

$$\|a_s\|_2 = \left\|\hat{\nabla}^2 J(\theta_{s,k}, \tau_s)(\theta_k - \theta_{k-1}) - \nabla^2 J(\theta_{s,k})(\theta_k - \theta_{k-1})\right\|_2 \leq 2L\|\theta_k - \theta_{k-1}\|_2.$$

Now, using vector Azuma-Hoeffding inequality from Lemma A.1, with at least probability $1 - \delta$, we have:

$$\begin{aligned}
\|u_k^a\|_2 &= \left\|\frac{1}{S_k}\sum_{s=0}^{S_k} a_s\right\|_2 \\
&\leq \frac{3}{S_k}\sqrt{\log(1/\delta)S_k(2L)^2\left\|\theta_k - \theta_{k-1}\right\|_2^2} \\
&\leq 6L\sqrt{\frac{\log(1/\delta)}{S_k}}\left\|\theta_k - \theta_{k-1}\right\|_2.
\end{aligned} \tag{51}$$

To bound $\|u_k^b\|$, we consider the following two term o_1 and o_2 in u_k^b:

$$\|u_k^b\|_2 = \left\|\underbrace{\frac{1}{S_k}\sum_{s=1}^{S_k}\nabla^2 J(\theta_{s,k})(\theta_k - \theta_{k-1})}_{o_1}\underbrace{-\nabla J(\theta_k) + \nabla J(\theta_{k-1})}_{o_2}\right\|. \tag{52}$$

We know that $\theta_{s,k} - \theta_{s-1,k} = \left[(1 - \frac{s}{S_k})\theta_k + \frac{s}{S_k}\theta_{k-1}\right] - \left[(1 - \frac{s-1}{S_k})\theta_k + \frac{s-1}{S_k}\theta_{k-1}\right] = \frac{1}{S_k}(\theta_{k-1} - \theta_k)$. Thus, we can rewrite o_1 as $\sum_{s=1}^{S_k}\nabla^2 J(\theta_{s,k})(\theta_{s,k} - \theta_{s-1,k})$ and o_2 as $\sum_{s=0}^{S_k}-\nabla J(\theta_{s,k}) + \nabla J(\theta_{s-1,k})$ using telescoping sum. Therefore:

$$\begin{aligned}
\|u_k^b\|_2 &\leq \sum_{s=1}^{S_k}\left\|-\nabla^2 J(\theta_{s,k})(\theta_{s,k} - \theta_{s-1,k}) + \nabla J(\theta_{s,k}) - \nabla J(\theta_{s-1,k})\right\|_2 \\
&\leq S_k \rho\frac{\|\theta_k - \theta_{k-1}\|_2^2}{S_k^2} = \rho\frac{\|\theta_k - \theta_{k-1}\|_2^2}{S_k},
\end{aligned} \tag{53}$$

where the second inequality is due to Hessian Lipschitzness of $J(\theta)$[1].

Summing (51) and (53), we have:

$$\|u_k\|_2 \leq 6L\sqrt{\frac{\log(1/\delta)}{S_k}}\left\|\theta_k - \theta_{k-1}\right\|_2 + \rho\frac{\left\|\theta_k - \theta_{k-1}\right\|_2^2}{S_k}. \tag{54}$$

Using inequality $\|A + B\|_2^2 \leq 2\|A\|_2^2 + 2\|B\|_2^2$,

$$\|u_k\|_2^2 \leq 72L^2\frac{\log(1/\delta)}{S_k}\left\|\theta_k - \theta_{k-1}\right\|_2^2 + 2\rho^2\frac{\left\|\theta_k - \theta_{k-1}\right\|_2^4}{S_k^2}. \tag{55}$$

From the above inequality, we can imply that the condition $\|u_k\|_2^2 \leq \frac{\epsilon^2}{540Q\log(1/\delta)}$ holds if S_k is greater than or equal to

$$S_k \geq C_{s_1}Q\frac{\left\|\theta_k - \theta_{k-1}\right\|_2^2}{\epsilon^2} + C_{s_2}\sqrt{Q}\frac{\left\|\theta_k - \theta_{k-1}\right\|_2^2}{\epsilon}, \tag{56}$$

---

[1] Hessian Lipschitz continuity of J implies that for a constant ρ and for all $\theta_1, \theta_2 \in \mathbb{R}^d$: $\|\nabla J(\theta_1) - \nabla J(\theta_2) - \nabla^2 J(\theta_2)(\theta_1 - \theta_2)\| \leq \rho\|\theta_1 - \theta_2\|^2$.

where $C_{s_1} = 38880L^2 \log^2(1/\delta)$ and $C_{s_2} = \sqrt{1080 \log(1/\delta)\rho^2}$. The above conditions has already been satisfied as we set S_k to $C_2 Q \frac{\|\theta_k - \theta_{k-1}\|_2^2}{\epsilon^2}$ in the statement of lemma where we assume that $\epsilon \leq 4L^2\rho/M$, and $Q = \frac{\sqrt{\rho M}}{\sqrt{\epsilon L}}$.

For $k = \lfloor t/Q \rfloor . Q$, using Azuma-Hoeffding inequality from Lemma A.1, with probability at least $1 - \delta$, we have

$$\|u_k\|_2 = \left\| \frac{1}{|\mathcal{B}_{check}|} \sum_{\tau \in \mathcal{B}_{check}} \hat{\nabla} J(\theta_k, \tau) - \nabla J(\theta_k) \right\|_2 \tag{57}$$

$$= \frac{1}{|\mathcal{B}_{check}|} \left\| \sum_{\tau \in \mathcal{B}_{check}} \hat{\nabla} J(\theta_k, \tau) - \nabla J(\theta_k) \right\|_2 \tag{58}$$

$$\leq \frac{3}{|\mathcal{B}_{check}|} \sqrt{4 \log(1/\delta) |\mathcal{B}_{check}| W^2} \tag{59}$$

$$\leq 3 \sqrt{4 \frac{\log(1/\delta)}{|\mathcal{B}_{check}|} W^2}, \tag{60}$$

where in the first inequality, we used the inequality $\|\hat{\nabla} J(\theta_k, \tau)\| \leq W$ according to Lemma 3.3.

Now, $\|u_k\|^2$ can be bounded as follows:

$$\|u_k\|_2^2 \leq 36 W^2 \frac{\log(1/\delta)}{|\mathcal{B}_{check}|} \tag{61}$$

$$\leq \frac{\epsilon^2}{540 \log(1/\delta)}, \tag{62}$$

where we set $|\mathcal{B}_{check}| = \frac{19440 W^2 \log^2(1/\delta)}{\epsilon^2}$ in the second inequality.

Based on what we proved above, with at least probability $1 - \delta(t - \lfloor t/Q \rfloor . Q)$, we have: $\|u_k\|_2^2 \leq \epsilon^2/(540 \log(1/\delta))$ for all $t \geq k > \lfloor t/Q \rfloor . Q$. Given that $\|u_k\|_2^2$'s are bounded by $\epsilon^2/(540 \log(1/\delta))$ for all $t \geq k > \lfloor t/Q \rfloor . Q$, using Azuma-Hoeffding inequality, with probability at least $1 - \delta$, we have:

$$\|v_i - \nabla J(\theta_i)\|_2^2 = \left\| \sum_{k=\lfloor t/Q \rfloor . Q}^{i} u_k \right\|_2^2 \tag{63}$$

$$\leq 9 \log(1/\delta) \left[(i - \lfloor t/Q \rfloor . Q) . \frac{\epsilon^2}{540 Q \log(1/\delta)} + \frac{\epsilon^2}{540 \log(1/\delta)} \right] \tag{64}$$

$$\leq 9 \log(1/\delta) \left[(t - \lfloor t/Q \rfloor . Q) . \frac{\epsilon^2}{540 Q \log(1/\delta)} + \frac{\epsilon^2}{540 \log(1/\delta)} \right] \tag{65}$$

$$\leq 9 \log(1/\delta) . \frac{\epsilon^2}{270 \log(1/\delta)} \tag{66}$$

$$\leq \frac{\epsilon^2}{30}, \tag{67}$$

for any $\lfloor t/Q \rfloor . Q \leq i \leq t$. Hence, with probability at least $1 - 2\delta(t - \lfloor t/Q \rfloor . Q)$, (67) holds for all i, where $\lfloor t/Q \rfloor . Q \leq i \leq t$.

B.2 PROOF OF LEMMA 3.8

Lemma B.2. *For $|\mathcal{B}_h|$ defined in the statement of Theorem 3.9, under Assumptions 3.1, 3.2, conditioned on $\mathcal{F}_t$, with probability at least $1 - \delta$, we have:*

$$\|U_t - \nabla^2 J(\theta_t)\|^2 \leq \frac{\epsilon \rho}{30}. \tag{68}$$

Proof: Using Azuma inequality from Lemma A.2, with probability at least $1 - \delta$, we have:

$$\|U_t - \nabla^2 J(\theta_t)\| = \left\| \frac{1}{|\mathcal{B}_h|} \sum_{\tau \in \mathcal{B}_h} \hat{\nabla}^2 J(\theta_t, \tau) - \nabla^2 J(\theta_t) \right\|_2 \tag{69}$$

$$= \frac{1}{|\mathcal{B}_h|} \left\| \sum_{\tau \in \mathcal{B}_h} \hat{\nabla}^2 J(\theta_t, \tau) - \nabla^2 J(\theta_t) \right\|_2 \tag{70}$$

$$\leq \frac{3}{|\mathcal{B}_h|} \sqrt{4 \log(d/\delta) |\mathcal{B}_h| L^2} \tag{71}$$

$$\leq 3\sqrt{4 \frac{\log(d/\delta)}{|\mathcal{B}_h|} L^2}, \tag{72}$$

where we used the the fact $\|\hat{\nabla}^2 J(\theta_t, \tau)\| \leq L$ and $\|\nabla^2 J(\theta_t)\| \leq L$ (as shown in the proof of Lemma B.1). Thus, $\left\| U_t - \nabla^2 J(\theta_t) \right\|^2$ can be bounded as follows:

$$\|U_t - \nabla^2 J(\theta_t)\|_2^2 \leq 36 L^2 \frac{\log(d/\delta)}{|\mathcal{B}_h|} \tag{73}$$

$$\leq \frac{\epsilon \rho}{30}, \tag{74}$$

where we set $|\mathcal{B}_h| = \dfrac{1080 L^2 \log(d/\delta)}{\rho \epsilon}$ in the second inequality.

B.3 PROOF OF THEOREM 3.9

Suppose that VR-SCP terminates at iteration $T^* - 1$. We claim that $T^* < T$, and accordingly Cubic-Finalsolver routine is executed. By contradiction, suppose that $T^* = T$. Then for all iteration $0 \leq t \leq T^* - 1$:

$$J(\theta_{t+1}) \geq J(\theta_t) + \langle \nabla J(\theta_t), h_t \rangle + \frac{1}{2} \langle \nabla^2 J(\theta_t) h_t, h_t \rangle - \frac{\rho}{6} \|h_t\|^3 \tag{75}$$

$$= J(\theta_t) + m_t(h_t) + \langle \nabla J(\theta_t) - v_t, h_t \rangle + \frac{1}{2} \langle (\nabla^2 J(\theta_t) - U_t) h_t, h_t \rangle + \frac{M - \rho}{6} \|h_t\|^3 \tag{76}$$

$$\geq J(\theta_t) + m_t(h_t) - \frac{6}{5\sqrt{\rho}} \left\| \nabla J(\theta_t) - v_t \right\|^{3/2} - \frac{10}{\rho^2} \left\| \nabla^2 J(\theta_t) - U_t \right\|^3 + \frac{\rho}{4} \|h_t\|^3, \tag{77}$$

where the equality is due to the definition of $m_t(h_t)$. Moreover, we use Lemma A.6 in the last inequality and consider $M = 4\rho$.

Since we assumed that $T^* = T$, for all $0 \leq t \leq T^* - 1$, we have:

$$m_t(h_t) \geq \rho^{-1/2} \epsilon^{3/2}/6. \tag{78}$$

By using Lemmas 3.7 and 3.8, for all $0 \leq t \leq T^* - 1$, with probability at least $1 - 3T\delta$:

$$\left\| \nabla J(\theta_t) - v_t \right\|^{3/2} \leq \epsilon^{3/2}/20, \tag{79}$$

$$\left\| \nabla^2 J(\theta_t) - U_t \right\|^3 \leq (\rho \epsilon)^{3/2}/160. \tag{80}$$

Using (78), (79) and (80) in (77),

$$J(\theta_{t+1}) - J(\theta_t) \geq \rho^{-1/2} \epsilon^{3/2}/6 + \frac{\rho}{4} \|h_t\|^3 - \rho^{-1/2} \epsilon^{3/2}/8. \tag{81}$$

Telescoping the last inequality from $t = 0$ to $t = T^* - 1$, we have:

$$\Delta_J \geq J(\theta_T) - J(\theta_0) \geq \sum_{t=0}^{T^*-1} \rho^{-1/2}\epsilon^{3/2}/24 + \frac{\rho}{4}\|h_t\|^3 \tag{82}$$

$$\geq T\rho^{-1/2}\epsilon^{3/2}/24 + \sum_{t=0}^{T^*-1} \frac{\rho}{4}\|h_t\|^3. \tag{83}$$

The last inequality contradicts with the assumption $T \geq 25\Delta_J\rho^{1/2}\epsilon^{-3/2}$. Thus, with probability at least $1 - 3T\delta$, Cubic-Finalsolver will be executed before ending the "for loop" of the Algorithm 1.

When Cubic-Finalsolver is executed, according to the Lemma A.3, with probability at least $1 - \delta$, none of the conditions in the statement of lemma hold. Thus we have $\sqrt{\epsilon/\rho} > \|h_t^*\|$ and $\|v_t\| < \max(M\epsilon/(2\rho), \sqrt{LM/2}(\frac{\epsilon}{\rho})^{3/4})$ with probability at least $1 - T\delta$. We use this observation to bound $\mu(\theta_{\tilde{T}} + h_{\tilde{T}})$ where $\tilde{T} = T^* - 1$ is the last iteration. $\mu(\theta_{\tilde{T}} + h_{\tilde{T}})$ can be bounded as follows:

$$\mu(\theta_{\tilde{T}} + h_{\tilde{T}}) \leq 9\left[M^3\rho^{-3/2}\|h_{\tilde{T}}\|^3 + M^{3/2}\rho^{-3/2}\left\|\nabla J(\theta_{\tilde{T}}) - v_{\tilde{T}}\right\|^{3/2} \right.$$

$$\left. + \rho^{-3/2}\left\|\nabla^2 J(\theta_{\tilde{T}}) - U_{\tilde{T}}\right\|^3 + M^{3/2}\rho^{-3/2}\|\nabla m_t(h_{\tilde{T}})\|^{3/2} + M^3\rho^{-3/2}\left|\|h_{\tilde{T}}\| - \|h_{\tilde{T}}^*\|\right|^3\right]$$

$$\leq 1300\epsilon^{3/2}, \tag{84}$$

where we know $\|h_{\tilde{T}}\| \leq \|h_{\tilde{T}}^*\|$ using Lemma A.4. Moreover, the output of Cubic-Finalsolver satisfies $\|\nabla m_t(h_{\tilde{T}})\| \leq \epsilon$. Furthermore, we used Lemma 3.7 and 3.8 to bound $M^{3/2}\rho^{-3/2}\|\nabla J(\theta_{\tilde{T}}) - v_{\tilde{T}}\|^{3/2}$ and $\rho^{-3/2}\|\nabla^2 J(\theta_{\tilde{T}}) - U_{\tilde{T}}\|^3$, respectively.

B.4 PROOF OF COROLLARY 3.10

We compute the number of the stochastic gradient and Hessian evaluations. First, we find a bound on the following term:

$$\sum_{t=0}^{\tilde{T}} \|h_t\|^2 \leq (\tilde{T} + 1)^{1/3}\left(\sum_{t=0}^{\tilde{T}} \|h_t\|^3\right)^{2/3} \leq (25\Delta_J\rho^{1/2}\epsilon^{-3/2})^{1/3} + (4\Delta_J/\rho)^{2/3} \leq \frac{\Delta_J}{8\rho^{1/2}\epsilon^{1/2}}, \tag{85}$$

where the first inequality is due to Hölder's inequality and the second one is due to the $\tilde{T} = T^* - 1 \leq 25\Delta_J\rho^{1/2}\epsilon^{-3/2}$ and $\Delta_J \geq \sum_{t=0}^{T^*-1} \frac{\rho}{4}\|h_t\|^3$ according to (82). Now we compute the whole stochastic gradient evaluations over $\tilde{T} = T^* - 1$:

$$\sum_{\mathrm{mod}(t,Q)=0}^{\tilde{T}} |\mathcal{B}_{check}| + \sum_{\mathrm{mod}(t,Q)\neq 0}^{\tilde{T}} S_t \overset{(i)}{\leq} \frac{C_1\tilde{T}}{Q\epsilon^2} + \sum_{\mathrm{mod}(t,Q)\neq 0}^{\tilde{T}} \frac{C_2 Q\|h_t\|^2}{\epsilon^2} \tag{86}$$

$$\leq \frac{C_1\tilde{T}}{Q\epsilon^2} + \frac{C_2 Q}{\epsilon^2}\sum_{t=0}^{\tilde{T}} \|h_t\|^2 \tag{87}$$

$$\leq \frac{C_1\tilde{T}}{Q\epsilon^2} + \frac{C_2 Q}{\epsilon^2}\sum_{t=0}^{\tilde{T}} \|h_t\|^2 \tag{88}$$

$$\overset{(ii)}{\leq} \frac{\sqrt{\epsilon L}}{\sqrt{\rho}M}\frac{C_1(25\Delta_J\rho^{1/2}\epsilon^{-3/2})}{\epsilon^2} + \frac{C_2\sqrt{\rho}M}{\epsilon^2\sqrt{\epsilon L}}\frac{\Delta_J}{8\rho^{1/2}\epsilon^{1/2}} \tag{89}$$

$$= \tilde{O}(\epsilon^{-3}). \tag{90}$$

(i) For the first term ($|\mathcal{B}_{check}|$), we consider $C_1 = 19440W^2\log^2(1/\delta)$. For the second term, we use S_k defined in Lemma B.1.

(ii) We set $Q = \frac{\sqrt{\rho}M}{\sqrt{\epsilon L}}$ and we know that $\tilde{T} = T^* - 1 \leq 25\Delta_J\rho^{1/2}\epsilon^{-3/2}$. Therefore we use (85) to bound $\sum_{t=0}^{\tilde{T}} \|h_t\|^2$.

Algorithm 2 Cubic-Subsolver

Input: $U[.], v, M, L, \epsilon$

1: **if** $||v|| \geq L^2/M$ **then**
2:     $R_c \leftarrow -\frac{v^T U[v]}{M||v||^2} + \sqrt{(\frac{v^T U[v]}{M||v||^2})^2 + 2||v||/M}$
3:     $\Delta \leftarrow \frac{v}{||v||} R_c$
4: **else**
5:     $\Delta \leftarrow 0, \sigma \leftarrow c'\frac{\sqrt{M\epsilon}}{L}, \eta \leftarrow \frac{1}{20L}$
6:     $\tilde{v} \leftarrow v + \sigma\mathcal{U}$ for $\mathcal{U} \sim Uniform(S^{d-1})$
7:     **for** $t = 0, \cdots, \mathcal{T}(\epsilon)$ **do**
8:         $\Delta \leftarrow \Delta + \eta(\tilde{v} + U[\Delta] - \frac{M}{2}||\Delta||\Delta)$
9:     **end for**
10: **end if**
11: $\Delta_m \leftarrow v^T\Delta + \frac{1}{2}\Delta^T U[\Delta] - \frac{M}{6}||\Delta||^3$
12: **return** Δ, Δ_m

Algorithm 3 Cubic-Finalsolver

Input: $U[.], v, M, L, \epsilon$

1: $\Delta \leftarrow 0, g_m \leftarrow v, \eta \leftarrow \frac{1}{20L}$
2: **while** $||v_m|| \geq \epsilon/2$ **do**
3:     $\Delta \leftarrow \Delta + \eta v_m$
4:     $v_m \leftarrow v + U[\Delta] - \frac{M}{2}||\Delta||\Delta$
5: **end while**
6: **return** Δ

Now we compute the Hessian vector evaluations in Cubic-Subslover and Cubic-Finalsolver. For each call of Cubic-Subslover, using Lemma A.3, we need $\mathcal{T}(\epsilon) \geq C_s\frac{L}{M\sqrt{\epsilon/\rho}}$ iterations and at each iteration we have $|\mathcal{B}_h| = \frac{1080L^2 \log(d/\delta)}{\rho\epsilon}$ samples. Hence in total, we have $T \times \mathcal{T}(\epsilon) \times |\mathcal{B}_h| = \frac{25\Delta_J\rho^{1/2}}{\epsilon^{3/2}} \times C_s\frac{L}{M\sqrt{\epsilon/\rho}} \times \frac{1080L^2 \log(d/\delta)}{\rho\epsilon} = C_3\epsilon^{-3}$ HVP evaluations where $C_3 = 25\Delta_J C_s\frac{1}{M}1080L^3 \log(d/\delta)$.

Using Lemma A.7, for Cubic-Finalsolver, the number of HVP evaluations is as follows:

$C_F L/\sqrt{\rho\epsilon} \times \frac{1080L^2 \log(d/\delta)}{\rho\epsilon} = C_4\epsilon^{-3/2}$ HVP evaluations where $C_4 = C_F\frac{1080L^3 \log(d/\delta)}{\rho^{3/2}}$. Thus in total, we have $\tilde{O}(\epsilon^{-3})$ stochastic gradient and HVP evaluations.

C   DESCRIPTIONS OF CUBIC-SUBSOLVER AND CUBIC-FINALSOLVER

The pseudo-code of Cubic-Subsolver is given in Algorithm 2 which uses gradient ascent to ensure that there is a sufficient increase in the objective of the sub-problem [Carmon and Duchi, 2016]. One difference compared with a simple gradient ascent is that at each call of Cubic-Subsolver, if the gradient norm is large enough, it uses the Cauchy point (line 3) which guarantees sufficient increase. In the gradient ascent part, the algorithm adds a small perturbation (line 6) (according to uniform distribution) to the gradient estimate v to escape the "hard" cases in the sub-problem and iterates for $\mathcal{T}(\epsilon) \geq C_s\frac{L}{M\sqrt{\epsilon/\rho}}$ iterations according to Lemma A.3.

Cubic-Subsolver and Cubic-Finalsolver have only access to v and HVP function $U[.]$. In Cubic-Finalsolver, in the WHILE loop, the gradient ascent update is applied till the WHILE condition fails (line 2).

D   DETAILS OF EXPERIMENTS

We used the default implementation of linear feature baseline and Gaussian MLP baseline from Garage library. The employed linear feature baseline is a linear regression model that takes observations for each trajectory and extracts new features such

as different powers of their lengths from the observations. These extracted features are concatenated to the observations and used to fit the parameters of the regression with the least square loss function.

We utilized a Linux server with Intel Xeon CPU E5-2680 v3 (24 cores) operating at 2.50GHz with 377 GB DDR4 of memory and Nvidia Titan X Pascal GPU. The computation was distributed over 48 threads to ensure a relatively efficient run time.

In the following table, we provide the fine-tuned parameters for each algorithm. Batch sizes are considered the same for all algorithms. The discount factor is also set to 0.99 for all the runs.

Table 2: Selected hyper-parameters for different algorithms.

	Reacher	Walker	Humanoid	Hopper
Max horizon	50	500	500	500
Neural network sizes	64×64	64×64	64×64	64×64
Activation functions	Tanh	Tanh	Tanh	Tanh
IS-MBPG lr	0.9	0.3	0.75	0.1
IS-MBPG c	100	12	5	50
IS-MBPG w	200	20	2	100
REINFORCE step-size	0.01	0.01	0.001	0.001
VR-SCP L	200	50	400	100
VR-SCP ρ	200	50	50	50
VR-SCP Q	10	2	5	2
PAGE-PG p_t	0.4	0.4	0.6	0.6
PAGE-PG step-size	0.01	0.001	0.0005	0.001
HAPG step-size	0.01	0.01	0.01	0.001
HAPG Q	5	10	10	10
VR-BGPO lr	0.8	0.75	0.8	0.75
VR-BGPO c	25	25	25	25
VR-BGPO w	1	1	1	1
VR-BGPO lam	0.0005	0.0005	0.0025	0.0005
ACR-PN $alpha$	100	10000	10000	10000