
On the Hidden Biases of Policy Mirror Ascent in Continuous Action Spaces

Amrit Singh Bedi¹ Souradip Chakraborty² Anjaly Parayil³ Brian Sadler⁴ Pratap Tokekar² Alec Koppel⁵

Abstract

We focus on parameterized policy search for reinforcement learning over continuous action spaces. Typically, one assumes the score function associated with a policy is bounded, which fails to hold even for Gaussian policies. To properly address this issue, one must introduce an exploration tolerance parameter to quantify the region in which it is bounded. Doing so incurs a persistent bias that appears in the attenuation rate of the expected policy gradient norm, which is inversely proportional to the radius of the action space. To mitigate this hidden bias, heavy-tailed policy parameterizations may be used, which exhibit a bounded score function, but doing so can cause instability in algorithmic updates. To address these issues, in this work, we study the convergence of policy gradient algorithms under heavy-tailed parameterizations, which we propose to stabilize with a combination of mirror ascent-type updates and gradient tracking. Our main theoretical contribution is the establishment that this scheme converges with constant batch sizes, whereas prior works require these parameters to respectively shrink to null or grow to infinity. Experimentally, this scheme under a heavy-tailed policy parameterization yields improved reward accumulation across a variety of settings as compared with standard benchmarks.

1. Introduction

In reinforcement learning (RL), an autonomous agent sequentially interacts with its environment and observes rewards incrementally across time (Sutton et al., 2017). This

framework has been successfully applied in continuous control (Schulman et al., 2015; Lillicrap et al., 2016), web services (Zou et al., 2019), personalized medicine (Kosorok & Moodie, 2015), among other contexts. Mathematically, RL may be defined by a Markov Decision Process (MDP) (Puterman, 2014), where an agent seeks to select actions to maximize the long-term accumulation of rewards, known as the value. The key distinguishing point of RL with classical optimal control is its ability to discern policies without a system dynamics model.

Algorithms for RL may be categorized as those which operate by approximately solving Bellman’s equations (Bellman, 1957; Watkins & Dayan, 1992) and policy gradient (PG) methods (Williams, 1992). While the former may be lower variance and converge faster (Even-Dar et al., 2003; Devraj & Meyn, 2017), typically they require representing a Q -function for every state-action pair, which is intractable for large spaces, the focus of this work. For this reason, we focus on PG methods.

Policy search hinges upon the Policy Gradient (PG) Theorem (Sutton et al., 2000), which expresses the gradient of the value function with respect to policy parameters as the expected value of the product of the score function of the policy and its associated Q function. Policy search has been classically studied from the perspective of dynamical systems (Kushner & Yin, 2003; Borkar, 2008), leading to an understanding of its asymptotic performance (Konda & Borkar, 1999; Konda & Tsitsiklis, 2000; Bhatnagar et al., 2009)). More recently, its non-asymptotic performance has come to the fore. Recent results (Bhandari & Russo, 2019; Zhang et al., 2020a; Agarwal et al., 2020) have established the global convergence of PG methods but require either softmax or direct policy parametrization. One may further refine the convergence constants via proximal regularization (Schulman et al., 2017; Tomar et al., 2020; Lan, 2021; Khodadadian et al., 2021). These results hold for finite state and action spaces. However, for continuous spaces the focus of this work) or general parameterizations, these results do not apply. Instead, stochastic gradient iteration for general non-convex objectives defines its performance in terms of convergence to stationarity (Bhatt et al., 2019; Zhang et al., 2020b), i.e., $\mathcal{O}(1/\sqrt{K})$ (or $\mathcal{O}(\epsilon^{-2})$) rate of decrease of the

¹Institute of Systems Research, University of Maryland, College Park, USA. ²Department of Computer Science, University of Maryland, College Park, USA. ³Microsoft Research, India ⁴DEVCOM Army Research Laboratory, Adelphi, MD, USA ⁵JP Morgan Chase AI Research, USA. . Correspondence to: Amrit Singh Bedi <amritbd@umd.edu>, Alec Koppel <aekoppel314@gmail.com >.

Algorithms	Hidden bias	Bregman term	SC
SVRPO (Xu et al., 2017)	Yes	No	N/A
SVRPG (Papini et al., 2018)	Yes	No	$\mathcal{O}(\epsilon^{-4})$
STORM-PG (Yuan et al., 2020)	Yes	No	$\mathcal{O}(\epsilon^{-4})$
RPG (Zhang et al., 2020b)	Yes	No	$\mathcal{O}(\epsilon^{-4})$
SRMA (This work)	No	Yes	$\mathcal{O}(\epsilon^{-4})$

Table 1. This table summarizes the existing sample complexity results for the policy gradient algorithm in literature for continuous state-action spaces. We note that all the results present in in the table are modified according to criteria $\mathbb{E}[\|\nabla F(\boldsymbol{\theta})\|] \leq \epsilon$ and constant batch size of $\mathcal{O}(1)$. True to our knowledge, this work provides the first optimal sample complexity results with finite batch size for the policy gradient algorithm with Bregman divergence as the regularization term without bounded score function assumption.

policy gradient norm¹.

Noticeably, these results require the score function to be uniformly bounded, which fails to encapsulate even standard Gaussian parameterizations with finite variance. This issue is identified in (Bedi et al., 2021), where in continuous space with an unbounded score function, instead the resultant rate is established as $\mathcal{O}(1/\sqrt{K}) + \mathcal{O}(\lambda)$. Here scalar λ is the exploration tolerance which quantifies the radius of action space. So, for Gaussian distribution and related parameterized families, a tradeoff between exploring the state space via extreme actions and converging to stationarity is present. Ideally, we would like λ to be very small, but doing so restricts the space over which action selection may take place (cf. (Bedi et al., 2021, Definition 4.1)). To achieve better space coverage, one must allow greater range of actions, but this yields large λ , which degrades the tightness of the radius of convergence to stationarity. Therefore, we pose the following question in this work:

“Which policy parameterizations achieve exact convergence to stationarity in continuous state and action space?”

This work presents an affirmative answer to this question through the identification that score function boundedness is a property of heavy-tailed distributions such as the Cauchy (Hutchinson, 1981; Focardi & Fabozzi, 2003; Barabási et al., 2003). This boundedness then mitigates the bias associated with the exploration tolerance parameter. Moreover, heavy-tailed policies in practice has been shown to yield improved state space coverage and escape spurious extrema in continuous space in (Chou et al., 2017; Papini et al., 2020; Garg et al., 2021). For these reason, we focus on policy search under heavy-tailed policy parameterizations². Substituting a

¹(ϵ -optimal solutions) In this work, we call $\boldsymbol{\theta} \in \mathbb{R}^d$ as ϵ -optimal if and only if $\mathbb{E}[\|\nabla F(\boldsymbol{\theta})\|] \leq \epsilon$. Sometimes, $\mathbb{E}[\|\nabla F(\boldsymbol{\theta})\|^2] \leq \epsilon$ is used in literature which means that our $\mathcal{O}(\epsilon^{-4})$ would be equivalent to $\mathcal{O}(\epsilon^{-2})$.

²As a side note, heavy-tailed policies induce heavy-tailed gradient noise, which has been associated with improved generalization

sub-Gaussian distribution with a heavy-tailed one, however, may cause instability in the parameter estimates due to increased likelihood of extreme action selection (Nemirovski et al., 2009). To ameliorate this issue, we prioritize proximal variants of PG (Schulman et al., 2017; Tomar et al., 2020), which employ Bregman divergence as a regularization. The idea of proximal algorithm was first introduced in (Rockafellar, 1976), and then generalized to Bregman divergences in (Nemirovski et al., 2009) for convex objectives. Later, mirror descent approach was extended to non-convex objectives in (Ghadimi et al., 2016) and shown to converge for increasing batch sizes. In the RL literature, the global convergence of policy mirror ascent has been established recently (Tomar et al., 2020; Khodadadian et al., 2021; Lan, 2021), but demands that one regularize the value function, and send the step-size or batch size to null or infinity (Ghadimi et al., 2016), respectively, in order to ensure convergence. In this work, we address this impracticality by introducing a gradient tracking mechanism that is able to applicable to operate with finite mini-batch sizes, which we call Stochastic Recursive Mirror Ascent Algorithm (SRMA). SRMA is inspired by (Cutkosky & Orabona, 2019), but differs in a key ways. In particular, we develop a gradient update that employs importance sampling to address distributional shift inherent to RL (Yuan et al., 2020); here, however, it is used in a novel context, which is to reduce persistent errors in the Bregman gradient, rather than for the purpose of variance reduction. This point of departure is underscored by the fact that both (Cutkosky & Orabona, 2019; Yuan et al., 2020) do not use proximal updates, and do not characterize how the convergence rate depend on the policy parameterization. Doing so then overlooks the hidden bias of PG in continuous space. A summary of related results in the literature is given in Table 1. Therefore, our main contributions are as follows:

- We use an often overlooked fundamental theoretical gap regarding the boundedness of the norm of score function in policy gradient method in continuous spaces to rigorously motivate the use of heavy-tailed parameterizations, such as the Cauchy policy, which are not afflicted by hidden bias.
- We develop a novel variant of policy mirror ascent that employs gradient tracking in its inner loop, which we call Stochastic Recursive Mirror Ascent (SRMA) (Algorithm 1). This scheme additionally employs importance sampling to correct for distributional shift that manifests when one tries to develop momentum techniques in RL.
- Our main theoretical result is the establishment of convergence to stationarity (Theorem 4.9) of SRMA,

in supervised learning recently (Simsekli et al., 2020a;b); however, we defer a rigorous exploration of this phenomenon in RL to future work.

which contrasts with prior results on stochastic mirror ascent that demand increasing batch sizes in order to obtain convergence, and closes a conspicuous gap in the literature for proximal methods applied to non-convex expected value objectives *without regularization under constant step-size* – see (Lan, 2021) for a study of the finite-state case. The convergence proof is done with a general set of assumptions which are explicitly proved to hold for the RL setting considered in this paper.

- We provide extensive simulations to show that policies learned with heavy-tailed distributions perform favorably in RL problems where biases associated with insufficiently covering the state spaces may be present (Sec. 5).

2. Markov Decision Problems

In reinforcement learning (RL), an autonomous agent traversing through a state space \mathcal{S} at state s , selects action $a \in \mathcal{A}$ and transitions to another state s' according to a Markov transition density $\mathbb{P}(s'|s, a)$. Upon reaching state s' , the environment reveals an instantaneous reward $r(s, a)$ which informs the merit of a given decision a starting from state s . Mathematically, this framework for interactive decision-making may be defined as a Markov Decision Process (MDP), whose components are $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, \gamma)$. The state \mathcal{S} and action space \mathcal{A} may either be finite or compact real vector space such that $\mathcal{S} \subseteq \mathbb{R}^q$ and $\mathcal{A} \subseteq \mathbb{R}^p$. Moreover, γ is a discount factor that determines how much future rewards are worth relative to the next step. As is well known in MDPs (Bertsekas & Shreve, 2004; Puterman, 2014), it suffices to hypothesize the decision-maker selects actions $a_t \sim \pi(\cdot|s)$ over a time-invariant distribution $\pi(a|s) := \Pr\{a_t = a | s_t = s\}$ called a policy, which denotes the probability of action a given the agent is in state s . The goal in RL is to determine the policy that accumulates the most long-term reward on average, i.e., the value:

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid s_0 = s, a_t = \pi(s_t) \right], \quad (1)$$

where s_0 denotes the initial state along a trajectory $\{s_t, a_t, r_t\}_{t=0}^{\infty}$ with short-hand notation $r_t = r(s_t, a_t)$. Here, the expectation in (1) is with respect to the product measure of randomized policy $a_t \sim \pi(\cdot|s_t)$ and state transition dynamics $s_{t+1} \sim \mathbb{P}(\cdot|s_t, a_t)$. For further reference, we define the action-value, i.e., Q -function as the value conditioned on an initially selected action as $Q^\pi(s, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid s_0 = s, a_0 = a, a_t = \pi(s_t) \right]$. Our focus is on policy search over parameterized families of policies, which hypothesizes that actions are selected according to a policy $\pi_\theta(\cdot|s_t)$ parameterized by vector $\theta \in \mathbb{R}^d$.

Then, we seek to estimate those parameters that maximize the cumulative return (Sutton et al., 2017) given by

$$\max_{\theta} J(\theta) \quad (2)$$

where, objective is given by $J(\theta) := V^{\pi_\theta}(s_0)$. Observe that (2) is non-convex in θ , and therefore, finding the optimal policy is challenging even in the deterministic setting. However, in RL, the search procedure necessarily interacts with the transition dynamics $\mathbb{P}(s'|s, a)$ as well. Before detailing how one may implement first-order stochastic search to solve (2), we introduce the widely used standard Gaussian policy parameterization, and clarify how its practice can lead to hidden bias.

Example 1 (Gaussian Parametrization). For continuous spaces, the Gaussian policy takes the form $\pi_\theta(a|s) = \mathcal{N}(a|\varphi(s)^\top \theta, \sigma^2)$, where the parameters θ determine the mean (centering) of a Gaussian distribution at $\varphi(s)$, and σ^2 is a fixed-variance hyper-parameter. Here, $\varphi(s)$ denotes the state space feature map, i.e., $\varphi : \mathcal{S} \rightarrow \mathbb{R}^d$ with $d \ll q$.

2.1. Unbounded Score Functions and Hidden Bias

Policy gradient (PG) method is an algorithm for RL which operates by implementing approximate gradient ascent in parameter space $\theta \in \mathbb{R}^d$ with respect to the value function (1). The key enabler of this method is the Policy Gradient Theorem (Sutton et al., 2017), which expresses search directions in parameter space as

$$\nabla J(\theta) = \frac{1}{1-\gamma} \cdot \mathbb{E}[\nabla \log \pi_\theta(a|s) \cdot Q^{\pi_\theta}(s, a)], \quad (3)$$

where the expectation is over $(s, a) \sim \rho_\theta(\cdot, \cdot)$ and $\rho_\theta(s, a) = \rho_{\pi_\theta}(s) \cdot \pi_\theta(a|s)$ is a probability distribution that denotes the *discounted state-action occupancy measure*, which is the product of the discounted state occupancy measure $\rho_{\pi_\theta}(s) = (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(s_t = s | s_0, \pi_\theta)$ and policy $\pi_\theta(a|s)$. In (Sutton et al., 2000), both $\rho_{\pi_\theta}(s)$ and $\rho_\theta(s, a)$ are established as valid probability distributions. To compute policy search directions, we consider an unbiased estimator of policy gradient via randomized horizon $T_k \sim \text{Geom}(1-\gamma^{1/2})$ with trajectory $\xi_k(\theta_k) = \{(s_0, a_0) \cdots (s_{T_k}, a_{T_k})\}$ given by

$$\begin{aligned} \nabla J(\theta_k, \xi_k(\theta_k)) \\ = \sum_{t=0}^{T_k} \gamma^{t/2} r_t \cdot \left(\sum_{\tau=0}^t \nabla \log \pi_{\theta_k}(a_\tau | s_\tau) \right), \end{aligned} \quad (4)$$

such that stochastic gradient is unbiased (proof is available in (Bedi et al., 2021, Lemma 1)) and $\xi_k(\theta_k)$ is the randomness in the stochastic gradient estimate at k . This estimator is a variant of the one proposed in (Baxter & Bartlett, 2001) but with a randomized horizon. This is in contrast to the existing literature where a fixed horizon length $T_k = H$

for all k is utilized (see (Papini et al., 2018; Yuan et al., 2020; Xu et al., 2017)). We remark here that a fixed horizon H actually leads to a bias-variance tradeoff in the gradient estimation rather than providing an unbiased estimate as discussed in (Baxter & Bartlett, 2001). Such tradeoff is not present for a randomized horizon based-estimator mentioned in (4) as it is unbiased, which is our motivation for using it in this work. Here, $\xi_k(\theta_k)$ is the trajectory collected by using θ_k as the policy parameter and hence it is function of θ_k . Note the summation in (4) over two index: index t denotes rollout trajectory information and τ is collecting the score function from starting to current index t . Using this scheme, stochastic policy gradient method iterates as

$$\theta_{k+1} = \theta_k + \eta \nabla J(\theta_k, \xi_k(\theta_k)), \quad (5)$$

where $\eta > 0$ denotes the step size. For general parametrized policy π_θ , by employing the iteration (5), one may obtain convergence to stationary points of (2) (see (Zhang et al., 2020b;a)). To do so, however, to date (see list in Table 1), most results require the score function to be deterministically bounded over the entire state space and action space. Unfortunately, this assumption is violated for the basic Gaussian policy (Example 1). To emphasize this point, note that score function, for scalar action space for simplicity, associated with the Gaussian policy (cf. Example 1) is

$$\nabla \log \pi_\theta(s, a) = \frac{(a - \varphi(s)^\top \theta) \varphi(s)}{\sigma^2}, \quad (6)$$

where $\varphi(s)$ is the state features such that $\|\varphi(s)\| \leq D$ for all s . (6) makes clear that the score function norm upper bound $|\nabla \log \pi_\theta(s, a)| \leq \mathcal{O}(D|a| + D^2\|\theta\|)$ is linear with respect to $|a|$ and $\|\theta\|$, which is unbounded unless the action space is compact. However, a Gaussian distribution is only valid over *infinite range*, leading to the following technical subtlety to address this issue. We note that identical logic applies in higher dimensions.

Exploration Tolerance and Hidden Bias (Bedi et al., 2021, Def. 4.1): Define $\mathcal{A}(\lambda)$ as the set of subsets of action space such that the score function has finite integral less than a scalar $\lambda > 0$ with respect to the policy

$$\mathcal{A}(\lambda) := \left\{ \mathcal{C} \subseteq \mathcal{A} : \int_{\mathcal{A} \setminus \mathcal{C}} \|\nabla \log \pi_\theta(a|s)\| \pi_\theta(a|s) da \leq \lambda \right\}, \quad (7)$$

for all s and θ . λ is the exploration tolerance parameter induced by a policy in an MDP whose score function is unbounded. In most existing analyses of PG algorithms, this quantity is ignored, leading to a misconception that PG methods converge exactly in expectation to a stationary point (Bhatt et al., 2019; Zhang et al., 2020b); however, when one properly accounts for λ , an $\mathcal{O}(\lambda)$ term appears as a persistent bias in the radius of convergence, i.e., PG in continuous space with a Gaussian policy yields an




Figure 1. Pathological Mountain Car example to demonstrate scenario when long and short-term incentives may be misaligned in continuous space: note there is a low reward state (red) and another high reward (red) state atop a higher hill. Policies that do not incentivize exploration get stuck at the spurious goal.

$\mathcal{O}(1/\sqrt{K} + \lambda)$ -suboptimal policy after K trajectories – see (Bedi et al., 2021)[Theorem 4.2]. For Gaussian policy, to make λ small, one must choose a large \mathcal{C} (set of possible actions), which is usually restricted by the practical physical limitations such as the bounded acceleration of a vehicle. By contrast, heavy-tailed distributions (Hutchinson, 1981; Focardi & Fabozzi, 2003; Barabási et al., 2003), for instance, the Cauchy, do not suffer this drawback. That is, their associated score functions are bounded. $r(s_t, a_t; \theta_t) =$

Example 2 (Cauchy Parametrization). Symmetric α stable, $\mathcal{S}\alpha\mathcal{S}$ distributions are a generalization of a centered Gaussian distribution with $\alpha \in (0, 2]$ as the tail index which determines the heaviness of the distribution’s tail (Nguyen et al., 2019). Denote random variable $\mathbf{X} \sim \mathcal{S}\alpha\mathcal{S}(\sigma)$ with associated characteristic function $\mathbb{E}[e^{i\omega\mathbf{X}}] = e^{-|\sigma|\omega^\alpha}$ and scale parameter $\sigma \in (0, \infty)$. Note that for $\alpha = 1$ we have a Cauchy distribution given by

$$\pi_\theta(a|s) = \frac{1}{\sigma\pi(1 + ((a - \varphi(s)^\top \theta)/\sigma)^2)}, \quad (8)$$

Interestingly, for the Cauchy distribution (cf. Example 2), the score function is given by

$$\nabla \log \pi_\theta(s, a) = \frac{2((a - \varphi(s)^\top \theta_1)/\sigma)}{1 + ((a - \varphi(s)^\top \theta_1)/\sigma)^2} \left(\frac{\varphi(s)}{\sigma} \right), \quad (9)$$

for any $\sigma > 0$. From (24), we can conclude that $\|\nabla \log \pi_\theta(s, a)\| \leq \frac{D}{\sigma}$ for all s, a , and θ .

With potential choice of policy parameterization detailed, we take a closer look at the relative merits and drawbacks. Intuitively, policies that select actions far from a learned mean parameter over actions may be beneficial in problems where state space coverage is essential. However, employing a Gaussian can result in persistent bias. To see how these issues manifest in practice, we develop the Pathological Mountain Car example next.

Pathological Mountain Car. Consider an environment with a car trapped between two mountains of different heights as

Algorithm 1 Policy Gradient with Stochastic Recursive Mirror Ascent (SRMA)

- 1: **Initialize** : Initial parameters, θ_0, β, γ , step-size η , $\mathbf{g}_0 = 0$
 - Repeat for** $k = 1, \dots$
 - 2: Sample trajectories $\xi_k(\theta_k)$ and $\xi_k(\theta_{k-1})$ of length $T_k \sim \text{Geom}(1 - \gamma^{1/2})$ using policies π_{θ_k} and $\pi_{\theta_{k-1}}$
 - 3: Estimate $\tilde{\nabla}J(\theta_{k-1}, \xi_k(\theta_k)), \nabla J(\theta_k, \xi_k(\theta_k))$ via (18) and (4), respectively
 - 4: $\hat{\mathbf{g}}_k = (1 - \beta)(\hat{\mathbf{g}}_{k-1} - \nabla J(\theta_{k-1}, \xi_k(\theta_k))) + \nabla J(\theta_k, \xi_k(\theta_k))$
 - 5: $\theta_{k+1} = \text{argmax}_{\theta} \{ \langle \hat{\mathbf{g}}_k, \theta \rangle - \frac{1}{\eta} D_{\psi}(\theta, \theta_k) \}$
 - 6: $k \leftarrow k + 1$
 - Until Convergence**
 - 7: **Return:** θ_k
-

shown in Fig. 1. Here, $s \in [-4.0, 3.709]$ denotes the state space, and the action a is a one-dimensional scalar representing the speed of the vehicle \dot{s} . The environment consists of two goal posts, a less-rewarding goal (red) at $s=2.667$ with a reward of 10 and a high reward (green) at $s=-4.0$ of 500. If we consider a Gaussian as a policy for this environment, then \mathcal{A} is \mathbb{R} . So there are two issues with using Gaussian policy for this environment. One major issue is in practice, we always define a set \mathcal{C} which denotes the practical range of actions around mean. For instance, we define $\mathcal{C} := [-5, 5]$ in the experiments (cf. 5) which comes from the practical limitations of the car. It prioritizes the actions around the mean, which can result in becoming stuck at the less rewarding states, never reaching the highest reward. We remark that this is the case with the experiments in existing literature as well (see Table 2). Therefore, there will always be a finite λ , and hence policy gradients converge to a policy that is not even a stationary point, but instead biased.

While encouraging action selection far from the mean may seem practically beneficial, and advantageous from the perspective of defining a policy whose score function is bounded, and hence mitigates the hidden bias of (7), doing so exhibits a downside. In particular, heavy-tailed policies result in search directions which may be unstable due to the high probability of taking actions far from mean. Surmounting this issue is the focus of Sec. 3, where we develop a mirror ascent-type algorithm for (2), discuss existing divergence issues with it due to the interaction of non-convexity and standard parameter selections. Consequently, we put forth a novel momentum variant, which is our main algorithmic innovation.

3. Stochastic Recursive Mirror Ascent

To conduct policy search when employing heavy-tailed parameterizations such as the Cauchy, we note that extreme action selection can cause numerical instabilities in the se-

quence of policy parameters in practice. To mitigate this issue, we take inspiration from proximal policy optimization (Schulman et al., 2017), which restricts movement of policy parameters through regularization of the gradient update. In the stochastic setting, such regularization can be rigorously substantiated through stochastic mirror ascent (Nemirovski et al., 2009; Ghadimi et al., 2016; Yang et al., 2019; Huang et al., 2020). Next, we discuss the technical limitations of existing stochastic mirror ascent approaches, which necessitate a modification that uses an additional recursive averaging step. The stochastic mirror ascent (SMA) update for (2) is given by

$$\theta_{k+1} = \text{argmax}_{\theta} \left\{ \langle \mathbf{g}_k, \theta \rangle - \frac{1}{\eta} D_{\psi}(\theta, \theta_k) \right\}, \quad (10)$$

where $\mathbf{g}_k := \nabla J(\theta_k, \xi_k(\theta_k))$ and D_{ψ} denotes a Bregman divergence defined with respect to the strongly convex function $\psi(\mathbf{x})$ with ζ as the strong convexity parameter. We remark that the update in (10) boils down to the standard stochastic gradient ascent (policy gradient in (18)) for $\psi(\theta) = \frac{1}{2} \|\theta\|^2$, and when $\psi(\theta) = \sum_j [\theta]_j \log[\theta]_j$ is the KL divergence, and the policy is tabular $\pi_{\theta} = \pi$, this reduces to Natural Policy Gradient. Extensions to parameterized settings are possible with Fischer information approximations, but we do not discuss this further. To analyze the update in (10), we define the Bregman gradient

$$\mathcal{G}_{\eta, \mathbf{g}_k}^{\psi}(\theta_k) = (\theta_k - \text{argmax}_{\theta} \left\{ \langle \mathbf{g}_k, \theta \rangle - \frac{1}{\eta} D_{\psi}(\theta, \theta_k) \right\}) / \eta \quad (11)$$

corresponding to the stochastic estimate of the gradient $\hat{\nabla}J(\theta_k)$ as an analogue of the fixed point of a gradient projection update, but instead with respect to the Bregman divergence – see (Ghadimi et al., 2016; Yang et al., 2019). Doing so allows us to rewrite (10) as

$$\theta_{k+1} = \theta_k + \eta \mathcal{G}_{\eta, \mathbf{g}_k}^{\psi}(\theta_k). \quad (12)$$

Optimality Criteria. To characterize convergence to first-order stationarity, we analyze the attenuation rate of the norm of the generalized gradient to a small constant ϵ as $\mathbb{E} [\|\mathcal{G}_{\eta, \mathbf{g}_k}^{\psi}(\theta)\|^2] \leq \epsilon$ which defines ϵ first-order stationarity, which is standard in the analysis of mirror ascent (Ghadimi et al., 2016). Importantly, in (Ghadimi et al., 2016), for batch size (number of stochastic gradient samples per iteration) of $B_k = \mathcal{O}(1)$, for any non-convex stochastic programming problem, under suitable conditions, we have that

$$\mathbb{E} [\|\mathcal{G}_{\eta, \mathbf{g}_k}^{\psi}(\theta)\|^2] \leq \frac{D_{\psi} + \frac{\sigma^2}{\zeta} \sum_{k=1}^K \eta_k}{\sum_{k=1}^K (\zeta \eta_k - L \eta_k^2)}, \quad (13)$$

where η_k is the step size used at each k . From the right hand side of (13), we can conclude that it is lower bounded by

$\frac{\sigma^2}{\zeta^2}$ independent of choice of step-size. To address this issue, one must allow the batch size B_k to increase with k to obtain convergence (Ghadimi et al., 2016; Yang et al., 2019). But obtaining a convergent algorithm for SMA algorithm with general non-convex *un-regularized* objective and fixed batch size per iteration remains a challenge. Moreover, increasing batch sizes at an unbounded rate is impractical in RL, as it is quite costly to sample ever-greater numbers of trajectories in between policy updates.

In this work, we address this issue via a recursive averaging step together with a difference of two gradient evaluations, i.e., gradient tracking, at each k given as

$$\hat{\mathbf{g}}_k = (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \nabla J(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k))) + \nabla J(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)), \quad (14)$$

$$\boldsymbol{\theta}_{k+1} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \left\{ \langle \hat{\mathbf{g}}_k, \boldsymbol{\theta} \rangle - \frac{1}{\eta} D_\psi(\boldsymbol{\theta}, \boldsymbol{\theta}_k) \right\}, \quad (15)$$

We introduce β as the step size for the gradient tracking update in (14). Different from the standard SMA update in (10), at each k , we need access to the stochastic gradients $\nabla J(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k))$ and $\nabla J(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k))$ mentioned in (14) evaluated at two different instances $\boldsymbol{\theta}_{k-1}$ and $\boldsymbol{\theta}_k$. Here, $\xi_k(\boldsymbol{\theta}_k)$ denotes the state action pair trajectory $\{(s_0, a_0, r_0) \cdots (s_{T_k}, a_{T_k}, r_{T_k}) \mid a_k \sim \pi_{\boldsymbol{\theta}_k}(\cdot \mid s_k)\}$ sampled using policy $\pi_{\boldsymbol{\theta}_k}$. The update in (14) is motivated from the momentum based tracking update proposed in (Cutkosky & Orabona, 2019), where two stochastic gradient evaluated using two different parameter instances at the same random variable are utilized for the gradient tracking. But using (14) in its current form results in a well known distribution shift issue (also appears in (Papini et al., 2018; Yuan et al., 2020)) because the trajectory sampling distribution (which we show explicitly by writing $\xi_k(\boldsymbol{\theta}_k)$ as function of $\boldsymbol{\theta}_k$) now depends upon parameter $\boldsymbol{\theta}$ which is not the case in (Cutkosky & Orabona, 2019). To resolve this issue, we employ an importance sampling (IS) based-modifications of the updates in (16)-(17) given as

$$\hat{\mathbf{g}}_k = (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \tilde{\nabla} J(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k))) + \nabla J(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)) \quad (16)$$

$$\boldsymbol{\theta}_{k+1} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \left\{ \langle \hat{\mathbf{g}}_k, \boldsymbol{\theta} \rangle - \frac{1}{\eta} D_\psi(\boldsymbol{\theta}, \boldsymbol{\theta}_k) \right\}, \quad (17)$$

where we define

$$\begin{aligned} & \tilde{\nabla} J(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \\ &= \sum_{\tau=0}^{T_k} \gamma^{t/2} \tau_t \cdot w_t(\tau \mid \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k) \cdot \sum_{s_\tau} \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau \mid s_\tau), \end{aligned} \quad (18)$$

where $w_t(\tau \mid \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)$ is the importance sampling weight given by $w_t(\tau \mid \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k) = \frac{\prod_{h=0}^t \pi_{\boldsymbol{\theta}_{k-1}}(a_h, s_h)}{\prod_{h=0}^t \pi_{\boldsymbol{\theta}_k}(a_h, s_h)}$. We call the updates in (16)-(17) as Stochastic Recursive Mirror Ascent

(SRMA). The overall proposed scheme for RL is summarized as Algorithm 1. Its important to note that this instantiation of gradient tracking is novel, in that previous uses were explicitly for variance reduction of non-proximal (Euclidean updates), whereas here they are specifically employed to reduce persistent bias [cf. (13)] associated with stochastic mirror ascent with constant batch size. Next, we shift towards establishing convergence of Algorithm 1.

4. Convergence Analysis

In this section, we establish that Algorithm 1 converges to stationarity in the sense of $\mathbb{E} [\|\mathcal{G}_{\eta, \mathbf{g}_k}^\psi(\boldsymbol{\theta})\|^2] \leq \epsilon$, with $\mathcal{G}_{\eta, \mathbf{g}_k}^\psi(\boldsymbol{\theta})$ as in (11). Without loss of generality, we reformulate the problem in the syntax of minimization, that is, we consider minimizing a function $F(\boldsymbol{\theta}) := -J(\boldsymbol{\theta})$, with $J(\boldsymbol{\theta})$ as the cumulative return under policy $\pi_{\boldsymbol{\theta}}$ in (2). Hence, the problem we consider for the analysis is given by

$$\min_{\boldsymbol{\theta}} F(\boldsymbol{\theta}). \quad (19)$$

Let $\nabla F(\boldsymbol{\theta}_k)$ denote the gradient of function $F(\boldsymbol{\theta})$ at $\boldsymbol{\theta}_k$ and $\xi_k(\boldsymbol{\theta}_k)$ is the randomness introduced at k to estimate the policy gradient. The associated stochastic unbiased gradient estimate is denoted as $\nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k))$. With this modified expression for the gradient and the reformulation in terms of minimization (19), the parameter update for $\boldsymbol{\theta}_k$ may be rewritten as

$$\begin{aligned} \hat{\mathbf{g}}_k &= (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k))) + \nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)), \\ \boldsymbol{\theta}_{k+1} &= \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \left\{ \langle \hat{\mathbf{g}}_k, \boldsymbol{\theta} \rangle + \frac{1}{\eta} D_\psi(\boldsymbol{\theta}, \boldsymbol{\theta}_k) \right\}. \end{aligned} \quad (20)$$

Next, we present the convergence rate results for the updates in (20) for any general non-convex function $F(\boldsymbol{\theta})$ which may be of independent interest. First, let us define the filtration $\mathcal{F}_k := \{\xi_u\}_{u < k}$, which collects randomness associated with trajectory realizations until step k . We use the definition of \mathcal{F}_k to state the following assumptions required for the following analysis.

Assumption 4.1. (Boundedness) The reward function $r(s, a)$ is bounded as $|r(s, a)| \leq U_R$ for all s and a . The state features $\varphi(s) \in \mathbb{R}^d$, are bounded, i.e., for some $D > 0$, $\|\varphi(s)\| \leq D$.

Assumption 4.2. (IS Variance) For any policy parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, let the importance sampling weights defined as $w(\tau \mid \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) := \frac{p(\tau \mid \boldsymbol{\theta}_1)}{p(\tau \mid \boldsymbol{\theta}_2)}$. Then it holds that $\operatorname{Var}(w(\tau \mid \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)) \leq C$, for some $C > 0$ where $\tau \sim p(\cdot \mid \boldsymbol{\theta}_2)$.

Note that when the non-convex objective $F(\boldsymbol{\theta}) = -J(\boldsymbol{\theta})$, then in the policy search setting, then above assumptions are sufficient for the SRMA algorithm to converge to the exact stationary point. But since, we are interested in doing the convergence analysis for general non-convex objective

$F(\boldsymbol{\theta})$, we mention next the general set of Assumptions 4.3-4.6 required for $F(\boldsymbol{\theta})$. For the RL setting considered in this paper where $F(\boldsymbol{\theta}) = -J(\boldsymbol{\theta})$, we prove in Appendix D that all the Assumptions 4.3-4.6 are satisfied.

Assumption 4.3. The stochastic estimate (4) is unbiased: $\mathbb{E}[\nabla F(\boldsymbol{\theta}_k, \xi_k) \mid \mathcal{F}_k] = \nabla F(\boldsymbol{\theta}_k)$ for all k .

Assumption 4.4. The variance of the stochastic gradient satisfies the growth condition

$$\mathbb{E}[\|\nabla F(\boldsymbol{\theta}_k, \xi_k) - \nabla F(\boldsymbol{\theta}_k)\|^2] \leq m_0 + m_1 \|\nabla F(\boldsymbol{\theta}_k)\|^2$$

for all k where $m_0 > 0$ and $m_1 \geq 0$ are scalars.

Assumption 4.5. The original gradient $\nabla F(\boldsymbol{\theta}_k)$ and the stochastic Bergman gradient $\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)$ satisfy the variance growth condition

$$\mathbb{E}[\|\nabla F(\boldsymbol{\theta}_k) - \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2] \leq m_2 + m_3 \mathbb{E}[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2]$$

where $m_2 > 0$, $m_3 \geq 0$ are scalars, and $\hat{\mathbf{g}}_k$ is as in (14).

Assumption 4.6. The objective $F(\cdot)$ is L -smooth.

Assumptions 4.3-4.6 are standard in the optimization literature. Assumption 4.5 imposes a bound between the original gradient estimate and the generalized gradient evaluated at the current biased estimate of the gradient $\hat{\mathbf{g}}_k$. Observe for the case when $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|^2$ and we utilize the stochastic unbiased gradient at each k , Assumption 4.5 simplifies to Assumption 4.4. Assumption 4.6 is related to the smoothness of the objective function F .

Before proceeding with the main result of this work, we recall two important properties of the generalized gradient from (Ghadimi & Lan, 2013, Lemma 1):

$$\langle \mathbf{g}_k, \mathcal{G}_{\eta, \mathbf{g}_k}^\psi(\boldsymbol{\theta}_k) \rangle \geq \zeta \|\mathcal{G}_{\eta, \mathbf{g}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \quad (21)$$

$$\|\mathcal{G}_{\eta, \mathbf{g}_1}^\psi(\boldsymbol{\theta}) - \mathcal{G}_{\eta, \mathbf{g}_2}^\psi(\boldsymbol{\theta})\| \leq \frac{1}{\zeta} \|\mathbf{g}_1 - \mathbf{g}_2\|. \quad (22)$$

The inequalities (21) and (22) will be used in the analysis. Next, we present an intermediate lemmas which bounds the stochastic errors associated with gradient estimation $\mathbb{E}[\|\hat{\mathbf{g}}_k - \nabla F(\boldsymbol{\theta}_k)\|^2]$.

Lemma 4.7. *The norm of the score function is bounded as $\|\nabla \log \pi_{\boldsymbol{\theta}}(a|s)\| \leq B$ for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$ for a Cauchy policy.*

See Appendix A for the proof of Lemma 4.7. We note that the statement of Lemma 4.7 explicitly provide a bound on the norm of score function which is usually assumed in the existing PG methods (Zhang et al., 2020b; Yuan et al., 2020; Xu et al., 2017; 2019). However, it cannot be assumed due to its severe dependence on the choice of policy parameterization. Next, we move to bound the error between the gradient $\hat{\mathbf{g}}_k$ used in Algorithm 1 and the original gradient $\nabla F(\boldsymbol{\theta}_k)$ as follows.

Refs.	Samples per Iteration	Sample Complexity
(Ghadimi et al., 2016)	$\mathcal{O}(1/\epsilon)$	$\mathcal{O}(1/\epsilon^2)$
(Yang et al., 2019)	$\mathcal{O}(1/\epsilon)$	$\mathcal{O}(1/\epsilon^2)$
This work	$\mathcal{O}(1)$	$\mathcal{O}(1/\epsilon^2)$

Table 2. Summary of related results in optimization literature to achieve $\mathbb{E}[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2] \leq \epsilon$.

Lemma 4.8. *Let $\epsilon_k := \mathbb{E}[\|\hat{\mathbf{g}}_k - \nabla F(\boldsymbol{\theta}_k)\|^2]$, then for all $k \geq 1$, it holds that*

$$\epsilon_k \leq (1 - \beta)^2 \epsilon_{k-1} + 2\eta^2 L_1 \mathbb{E}[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2] + 2m_0 \beta^2 + 2m_1 \beta^2 \|\nabla F(\boldsymbol{\theta}_k)\|^2, \quad (23)$$

where $L_1 := \left(2L^2 + \frac{2C_w U_R^2 B^2 (1 + \sqrt{\gamma})}{(1 - \gamma)(1 - \sqrt{\gamma})^2}\right)$.

See Appendix B in the Appendix for proof. The result in Lemma 4.8 bounds the per step expected value of the the norm of error for each k . Next we present the main theorem of this paper. We note that the tracking result in Lemma 4.8 incorporates random horizon T_k into derivation which is missing in the existing literature (Papini et al., 2018; Yuan et al., 2020; Xu et al., 2017).

Theorem 4.9. *Under Assumption 4.3-4.6, with step-size selections $\beta = C_1 \eta$ with $C_1 > 0$ and $\eta \leq \min\left\{\frac{\zeta}{10L'}, \frac{\zeta}{8\bar{m}_3 C_1^2}\right\}$, in order to achieve $\min_{1 \leq k \leq K} \mathbb{E}[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2] \leq \epsilon$ with $\epsilon \leq \min\left\{\frac{\zeta L'}{10}, \frac{\zeta}{8\bar{m}_3 C_1^2}\right\}$, the iterates in Algorithm 1 requires at least $K \geq \mathcal{O}\left(\frac{1}{\epsilon^2}\right)$ iterations with $\mathcal{O}(1)$ stochastic gradients samples at each k . Additionally, in order to achieve $\min_{1 \leq k \leq K} \mathbb{E}[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|] \leq \epsilon$, the iterates in Algorithm 1 requires at least $K \geq \mathcal{O}\left(\frac{1}{\epsilon^4}\right)$ with $\mathcal{O}(1)$ stochastic gradients samples at each k .*

We note that a related but simpler specification of step-size η also permitted: $\eta = \frac{\eta_0}{\sqrt{K}}$ in terms of final iteration index K with $\eta_0 = \min\left\{\frac{\zeta L'}{10}, \frac{\zeta}{8\bar{m}_3 C_1^2}\right\}$. See Appendix C for proof. Note that the use of recursive update for the stochastic gradient estimate in (16) permits us to achieve the $\mathcal{O}\left(\frac{1}{\epsilon^2}\right)$ oracle complexity with a We remark that the sample complexity result in Theorem 4.9 is general enough and holds for any non-convex objective $F(\boldsymbol{\theta})$ as well. constant batch size of gradients $\mathcal{O}(1)$ per iteration. In addition, we note that in the stochastic optimization literature, the related work on stochastic mirror ascent for non-convex objectives require an increasing batch size for convergence (Ghadimi & Lan, 2013; Yang et al., 2019). In contrast, the proposed algorithm is able to achieve the optimal convergence rate with finite batch size which is of great practical importance. We summarize the results in Table 2. Under RL setting considered in this paper, we prove in Appendix D that all the assumptions are satisfied. Further, in the PG analysis

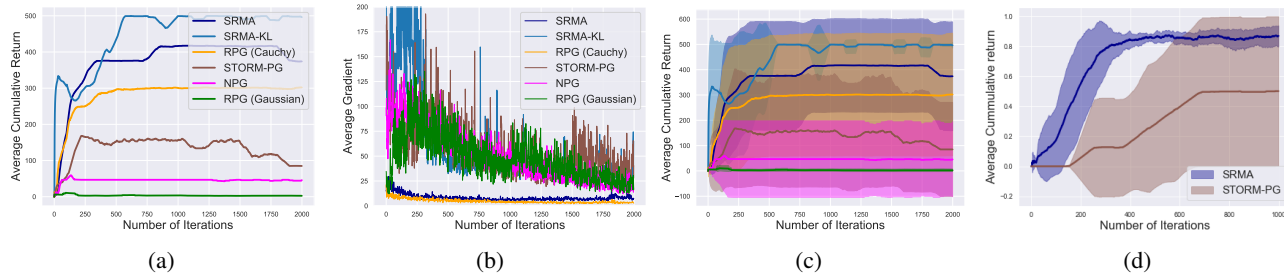


Figure 2. (a) This figure compares the average cumulative return obtained by the proposed Stochastic Recursive Mirror Ascent (SRMA) algorithm with the existing state of the art methods available to solve MDP in continuous state action space. We considered Pathological Mountain Car environment (cf. Fig. 1) for the experiments. We note that SRMA-KL is able to achieve the maximum reward as compared to other methods. (b) In this figure, we plot the average Bregman gradient (performance metric in Theorem 4.9) for the proposed algorithm and compares with others. We note that even if the gradient norm converges to zero for STORM-PG (which used Gaussian policy parametrization), the corresponding reward is low. This implies that Gaussian policy parametrization is not able to sufficiently explore the environment. (c) In this figure, we plot the average cumulative gradient shown in Fig. 2(a) but with confidence intervals. (d) Performance comparison of the proposed algorithm with STORM-PG on Mario environment mentioned in (Matheron et al., 2019).

in the existing literature, $\mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^{\psi}(\boldsymbol{\theta}_k)\| \right] \leq \epsilon$ is used as a performance metric, which requires $\mathcal{O} \left(\frac{1}{\epsilon^4} \right)$ (as mentioned in Table 1) sample complexity with batch size of $\mathcal{O}(1)$. Additionally, employing the Cauchy policy parametrization removes the bias present in the convergence of existing PG methods in continuous spaces (Bedi et al., 2021). In other words, λ (cf. (7)) is exactly zero for the proposed algorithm, which results in exact convergence to stationarity. We evaluate the experimental utility of SRMA next.

5. Experiments

This section validates the efficacy of proposed heavy tailed (Cauchy distribution) policy parameterizations in the continuous MDP of the Pathological Mountain Car (PMC) from Section 2. We consider an incentive structure in which the amount of energy expenditure, i.e., the action squared, is negatively penalized such that $r(s, a) = -a^2$ for $-4.0 < s < 3.7$, $s \neq 2.6$, $r(s, a) = 500 - a^2$ if $s = -4.0$, and $r(s, a) = 10 - a^2$ if $s = 2.6$. At each episode, the position of the car is initialized uniformly at random from $[1.15, 2.0]$. State s is constrained to an interval $[-4.0, 3.7]$ and action a lies in $[-6, 6]$. The discounted factor γ is 0.97 and we use a step size of 0.005. In Fig. 2, we compare performance of the proposed SRMA algorithm with the other existing state of the art techniques for continuous MDP such as NPG (Kakade, 2002), RPG (Zhang et al., 2020b), STORM-PG (Yuan et al., 2020). For experiments, we implemented two versions of the proposed SRMA algorithm called SRMA and SRMA-KL (see legends in Fig 2(a)). For the SRMA-KL version, we use KL divergence between the policies at k and $k - 1$ as the Bregman divergence term, which is also used for NPG implementation. In Fig 2(a), SRMA denotes implementation where we used standard Euclidean distance as the Bregman term. For SRMA and

SRMA-KL, we used Cauchy policy parametrization, and for RPG (Gaussian), NPG, and STORM-PG, we have used Gaussian policy parametrization.

Fig. 2(a) shows the cumulative return averaged over 15 runs of the experiments. It is clear that the SRMA algorithm which utilizes Cauchy policy parametrization yields better performance as relative to comparators which use Gaussian parametrization. Moreover, we note that simply replacing Gaussian parametrization by a Cauchy results in numerical instability, as may be observed by RPG (Cauchy), which is PG iteration with a Cauchy policy. This instability is corroborated in Fig. 2(c) where we observe the average rewards return along with standard deviation across runs, which is large for the orange curve. This justifies the use of policy mirror ascent-type updates. We also plot the average Bregman gradient in Fig. 2(b) to show the convergence of the proposed algorithm under constant step-size and batch-size as compared with the others methods available in literature.

6. Conclusion

In this work, we focused on policy gradient method for solving RL problems associated with infinite-horizon discounted returns. In some problems, the one-step reward may be very far from the value of a given state, which can cause policies to become mired at spurious behavior. Inspired by the relationship between score function boundedness and persistent bias that may arise when operating in continuous action spaces, we proposed to study policy search under heavy-tailed parameterizations. This parameterization introduced numerical challenges, namely, numerical instability in the agent’s trajectory, and potentially volatile changes in policy gradients. To address this issue, we studied mirror ascent-type updates, which we stabilized with recursive gradient tracking. The convergence of the resulting iterative

schemes was established under novel error bound conditions for the generalized Bregman gradient. Moreover, experimentally, we observed favorable performance of the proposed approach for escaping spurious stationary points.

7. Acknowledgements

This research was supported by Army Cooperative Agreement W911NF2120076.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. Optimality and approximation with policy gradient methods in markov decision processes. In *Conference on Learning Theory*, pp. 64–66. PMLR, 2020.
- Barabási, A.-L. et al. Emergence of scaling in complex networks. *Handbook of Graphs and Networks: From the Genome to the Internet*. Berlin: Wiley-VCH, 2003.
- Baxter, J. and Bartlett, P. L. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- Bedi, A. S., Parayil, A., Zhang, J., Wang, M., and Koppel, A. On the sample complexity and metastability of heavy-tailed policy search in continuous control. *arXiv preprint arXiv:2106.08414*, 2021.
- Bellman, R. E. *Dynamic Programming*. Courier Dover, 1957. ISBN 0486428095.
- Bertsekas, D. P. and Shreve, S. *Stochastic optimal control: the discrete-time case*. 2004.
- Bhandari, J. and Russo, D. Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*, 2019.
- Bhatnagar, S., Sutton, R., Ghavamzadeh, M., and Lee, M. Natural actor-critic algorithms. *Automatica*, 45(11):2471–2482, 2009.
- Bhatt, S., Koppel, A., and Krishnamurthy, V. Policy gradient using weak derivatives for reinforcement learning. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 5531–5537. IEEE, 2019.
- Borkar, V. S. *Stochastic approximation: A dynamical systems viewpoint*. Cambridge University Press, 2008.
- Chou, P.-W., Maturana, D., and Scherer, S. Improving stochastic policy gradients in continuous control with deep reinforcement learning using the beta distribution. In *International conference on machine learning*, pp. 834–843. PMLR, 2017.
- Cutkosky, A. and Orabona, F. Momentum-based variance reduction in non-convex sgd. *arXiv preprint arXiv:1905.10018*, 2019.
- Devraj, A. M. and Meyn, S. P. Zap q-learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 2232–2241, 2017.
- Even-Dar, E., Mansour, Y., and Bartlett, P. Learning rates for q-learning. *Journal of machine learning Research*, 5(1), 2003.
- Focardi, S. M. and Fabozzi, F. J. Fat tails, scaling, and stable laws: a critical look at modeling extremal events in financial phenomena. *The Journal of Risk Finance*, 2003.
- Garg, S., Zhanson, J., Parisotto, E., Prasad, A., Kolter, J. Z., Balakrishnan, S., Lipton, Z. C., Salakhutdinov, R., and Ravikumar, P. On proximal policy optimization’s heavy-tailed gradients. *arXiv preprint arXiv:2102.10264*, 2021.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIOPT*, 23(4):2341–2368, 2013.
- Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.
- Huang, F., Gao, S., Pei, J., and Huang, H. Momentum-based policy gradient methods. In *International Conference on Machine Learning*, pp. 4422–4433. PMLR, 2020.
- Hutchinson, J. E. Fractals and self similarity. *Indiana University Mathematics Journal*, 30(5):713–747, 1981.
- Kakade, S. M. A natural policy gradient. In *NeurIPS*, pp. 1531–1538, 2002.
- Khodadadian, S., Jhunjunwala, P. R., Varma, S. M., and Maguluri, S. T. On the linear convergence of natural policy gradient algorithm. *arXiv preprint arXiv:2105.01424*, 2021.
- Konda, V. R. and Borkar, V. S. Actor-critic-type learning algorithms for Markov decision processes. *SICON*, 38(1):94–123, 1999.
- Konda, V. R. and Tsitsiklis, J. N. Actor-critic algorithms. In *NeurIPS*, pp. 1008–1014, 2000.
- Kosorok, M. R. and Moodie, E. E. *Adaptive treatment strategies in practice: planning trials and analyzing data for personalized medicine*. SIAM, 2015.
- Kushner, H. J. and Yin, G. G. *Stochastic approximation and recursive algorithms and applications*. Springer, New York, NY, 2003.

- Lan, G. Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *arXiv preprint arXiv:2102.00135*, 2021.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations*, 2016.
- Matheron, G., Perrin, N., and Sigaud, O. The problem with ddpq: understanding failures in deterministic environments with sparse rewards. *arXiv preprint arXiv:1911.11679*, 2019.
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- Nguyen, T. H., Simsekli, U., Gürbüzbalaban, M., and Richard, G. First exit time analysis of stochastic gradient descent under heavy-tailed gradient noise. *arXiv preprint arXiv:1906.09069*, 2019.
- Papini, M., Binaghi, D., Canonaco, G., Pirota, M., and Restelli, M. Stochastic variance-reduced policy gradient. In *ICML*, pp. 4026–4035, 2018.
- Papini, M., Battistello, A., and Restelli, M. Balancing learning speed and stability in policy gradient via adaptive exploration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1188–1199. PMLR, 2020.
- Puterman, M. L. *Markov Decision Processes: Discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Rockafellar, R. T. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Simsekli, U., Sener, O., Deligiannidis, G., and Erdogdu, M. A. Hausdorff dimension, heavy tails, and generalization in neural networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 5138–5151. Curran Associates, Inc., 2020a.
- Simsekli, U., Zhu, L., Teh, Y. W., and Gurbuzbalaban, M. Fractional underdamped langevin dynamics: Retargeting sgd with momentum under heavy-tailed gradient noise. In *International Conference on Machine Learning*, pp. 8970–8980. PMLR, 2020b.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *NeurIPS*, pp. 1057–1063, 2000.
- Sutton, R. S., Barto, A. G., et al. *Reinforcement learning: An introduction*. 2 edition, 2017.
- Tomar, M., Shani, L., Efroni, Y., and Ghavamzadeh, M. Mirror descent policy optimization. *arXiv preprint arXiv:2005.09814*, 2020.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.
- Xu, P., Gao, F., and Gu, Q. Sample efficient policy gradient methods with recursive variance reduction. *arXiv preprint arXiv:1909.08610*, 2019.
- Xu, T., Liu, Q., and Peng, J. Stochastic variance reduction for policy gradient estimation. *arXiv preprint arXiv:1710.06034*, 2017.
- Yang, L., Zheng, G., Zhang, H., Zhang, Y., Zheng, Q., Wen, J., and Pan, G. Policy optimization with stochastic mirror descent. *arXiv preprint arXiv:1906.10462*, 2019.
- Yuan, H., Lian, X., Liu, J., and Zhou, Y. Stochastic recursive momentum for policy gradient methods. *arXiv preprint arXiv:2003.04302*, 2020.
- Zhang, J., Kim, J., O’Donoghue, B., and Boyd, S. Sample efficient reinforcement learning with reinforce. *arXiv preprint arXiv:2010.11364*, 2020a.
- Zhang, K., Koppel, A., Zhu, H., and Basar, T. Global convergence of policy gradient methods to (almost) locally optimal policies. *SIAM Journal on Control and Optimization*, 58(6):3586–3612, 2020b.
- Zou, L., Xia, L., Ding, Z., Song, J., Liu, W., and Yin, D. Reinforcement learning to optimize long-term user engagement in recommender systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2810–2818, 2019.

Appendix

A. Proof of Lemma 4.7

Since we consider Cauchy distribution, we note that the norm of the score function is given by

$$\|\nabla \log \pi_{\theta}(s, a)\| = \frac{2\|(\|a - \varphi(s)^\top \theta_1\|/\sigma)\|}{1 + (\|a - \varphi(s)^\top \theta_1\|/\sigma)^2} \left(\frac{\|\varphi(s)\|}{\sigma} \right), \quad (24)$$

which is of the form $\frac{2x}{1+x^2}$ and it holds that $\frac{2x}{1+x^2} \leq 1$ for all x . Hence we conclude that

$$\|\nabla \log \pi_{\theta}(s, a)\| \leq \left(\frac{\|\varphi(s)\|}{\sigma} \right) \leq \frac{D}{\sigma} := B.$$

Hence proved.

B. Proof of Lemma 4.8

Proof. The proof follows along the lines of (Cutkosky & Orabona, 2019, Lemma 2). Let us denote $\bar{\delta}_k = \nabla F(\theta_k)$, expand the expression $\hat{\mathbf{g}}_k - \delta_k$, and introducing $(1 - \beta)\bar{\delta}_{k-1}$, we obtain

$$\begin{aligned} \hat{\mathbf{g}}_k - \bar{\delta}_k &= (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k))) + \nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k \\ &= (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \bar{\delta}_{k-1} + \bar{\delta}_{k-1} - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) + \nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k) \\ &= (1 - \beta) (\hat{\mathbf{g}}_{k-1} - \bar{\delta}_{k-1}) - ((1 - \beta) [\tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) - \bar{\delta}_{k-1}] - \nabla F(\theta_k, \xi_k(\theta_k)) + \bar{\delta}_k). \end{aligned} \quad (25)$$

Note that in the above expression, $\xi_k(\theta_{k-1})$ denotes the trajectory collected at iteration k using the policy parameter θ_{k-1} and $\xi_k(\theta_k)$ denotes the trajectory collected at iteration k using θ_k as the policy parameter. Hence, this would imply that $\mathbb{E} [\bar{\delta}_{k-1} - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k))] = \mathbb{E} [\nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k] = 0$. Therefore, we have that

$$\begin{aligned} \mathbb{E} [\|\hat{\mathbf{g}}_k - \bar{\delta}_k\|^2] &= (1 - \beta)^2 \|\hat{\mathbf{g}}_{k-1} - \bar{\delta}_{k-1}\|^2 \\ &\quad + \mathbb{E} [\|(1 - \beta)(\tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) - \bar{\delta}_{k-1}) + \bar{\delta}_k - \nabla F(\theta_k, \xi_k(\theta_k))\|^2], \end{aligned} \quad (26)$$

where the cross term vanishes since $\hat{\mathbf{g}}_{k-1} - \bar{\delta}_{k-1}$ is independent of ξ_k and the second summand is zero mean. Let us define $\varepsilon_k := \mathbb{E} [\|\hat{\mathbf{g}}_k - \bar{\delta}_k\|^2]$, we could write (26) as

$$\varepsilon_k = (1 - \beta)^2 \varepsilon_{k-1} + \mathbb{E} [\|(1 - \beta)(\tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) - \bar{\delta}_{k-1}) + \bar{\delta}_k - \nabla F(\theta_k, \xi_k(\theta_k))\|^2]. \quad (27)$$

The second term on the right hand side of (26) can again be expanded as

$$\begin{aligned} &\mathbb{E} [\|(1 - \beta)(\tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) - \bar{\delta}_{k-1}) + \bar{\delta}_k - \nabla F(\theta_k, \xi_k(\theta_k))\|^2] \\ &= \mathbb{E} [\|(1 - \beta)(\bar{\delta}_{k-1} - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k)) + \nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k) + \beta(\nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k)\|^2] \end{aligned} \quad (28)$$

$$\leq 2(1 - \beta)^2 \mathbb{E} [\|\nabla F(\theta_k, \xi_k(\theta_k)) - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k))\|^2] + 2\beta^2 \mathbb{E} [\|\nabla F(\theta_k, \xi_k(\theta_k)) - \bar{\delta}_k\|^2], \quad (29)$$

where we have used the inequality $\mathbb{E} [\|X - \mathbb{E}[X] + Y\|^2] \leq 2\mathbb{E} [\|X\|^2] + 2\mathbb{E} [\|Y\|^2]$ for any random variables X and Y with bounded variances. Substitute Assumption 4.4 regarding the gradient estimation error for the second term to obtain:

$$\varepsilon_k \leq (1 - \beta)^2 \varepsilon_{k-1} + 2(1 - \beta)^2 \mathbb{E} [\|\nabla F(\theta_k, \xi_k(\theta_k)) - \tilde{\nabla} F(\theta_{k-1}, \xi_k(\theta_k))\|^2] + 2\beta^2 m_0 + 2m_1 \beta^2 \|\nabla F(\theta_k)\|^2. \quad (30)$$

Next, we focus on the second term on the right hand side of (46) as follows.

$$\begin{aligned}
 & \mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)) - \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) + \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \\
 &\leq 2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)) - \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] + 2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \\
 &\leq 2L^2\mathbb{E} \left[\left\| \boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1} \right\|^2 \right] + 2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right]. \tag{31}
 \end{aligned}$$

The second term in the right hand side of (31) is unique to the policy gradient settings. The importance weights based definition of gradient would help us to bound this term. Let us expand the second term on the right hand side of (31) as follows

$$\begin{aligned}
 & 2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \tag{32} \\
 &= 2\mathbb{E} \left[\left\| \sum_{t=0}^{T_k} \gamma^{t/2} r_t \cdot \left(\sum_{\tau=0}^t \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right) - \sum_{t=0}^{T_k} \gamma^{t/2} r_t \cdot w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k) \left(\sum_{\tau=0}^t \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right) \right\|^2 \right].
 \end{aligned}$$

After collecting the like terms together, we obtain

$$\begin{aligned}
 & 2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \tag{33} \\
 &= 2\mathbb{E} \left[\left\| \sum_{t=0}^{T_k} \gamma^{t/2} r_t \cdot (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \cdot \left(\sum_{\tau=0}^t \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right) \right\|^2 \right] \\
 &= 2\mathbb{E} \left[\sum_{t=0}^{T_k} \mathbb{E} \left[\left\| \gamma^{t/2} r_t \cdot (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \cdot \left(\sum_{\tau=0}^t \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right) \right\|^2 \right] \right], \tag{34}
 \end{aligned}$$

The second equality holds due to the fact that $\mathbb{E} [\nabla \log \pi_{\boldsymbol{\theta}}(a | s)] = 0$. Further, utilizing the upper bounds on reward and absolute bound on the score function norm, we can write

$$2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \tilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \tag{35}$$

$$\leq 2U_R^2 \mathbb{E} \left[\sum_{t=0}^{T_k} \mathbb{E} \left[\left\| \gamma^t \cdot (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \right\|^2 \cdot t \sum_{\tau=0}^t \left\| \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right\|^2 \right] \right] \tag{36}$$

$$\leq 2U_R^2 \mathbb{E} \left[\sum_{t=0}^{T_k} \mathbb{E} \left[\left\| \gamma^t \cdot (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \right\|^2 \cdot t \sum_{\tau=0}^t \left\| \nabla \log \pi_{\boldsymbol{\theta}_{k-1}}(a_\tau | s_\tau) \right\|^2 \right] \right] \tag{37}$$

$$\leq 2U_R^2 \mathbb{E} \left[\sum_{t=0}^{T_k} \mathbb{E} \left[\left\| \gamma^t \cdot (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \right\|^2 \cdot t^2 B^2 \right] \right] \tag{38}$$

$$\leq 2U_R^2 B^2 \mathbb{E} \left[\sum_{t=0}^{T_k} \gamma^t t^2 \mathbb{E} \left[\left\| (1 - w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \right\|^2 \right] \right] \tag{39}$$

$$= 2U_R^2 B^2 \mathbb{E} \left[\sum_{t=0}^{T_k} \gamma^t t^2 \text{Var}(w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)) \right], \tag{40}$$

which holds because $\mathbb{E}[w_t(\tau | \boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_k)] = 1$ (Xu et al., 2019, Lemma C.1). Next, utilizing Assumption 4.2 and (Xu et al., 2019, Lemma B.1), we can write

$$2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \widetilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \leq 2U_R^2 B^2 C_w \mathbb{E} \left[\sum_{t=0}^{T_k} \gamma^t t^2 \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2] \right] \quad (41)$$

$$\leq 2U_R^2 B^2 C_w \mathbb{E} \left[T_k^2 \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2] \sum_{t=0}^{T_k} \gamma^t \right] \quad (42)$$

$$\leq \frac{2U_R^2 B^2 C_w}{1-\gamma} \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2 \mathbb{E} [T_k^2]]. \quad (43)$$

Since $T_k \sim \text{Geom}(p)$ with $p = 1 - \gamma^{1/2}$, it holds that $\mathbb{E}[T_k^2] = (\mathbb{E}[T_k])^2 + \text{var}(T_k)$, which further implies that $\mathbb{E}[T_k^2] = \frac{1+\sqrt{\gamma}}{(1-\sqrt{\gamma})^2}$. Hence, we can write

$$2\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_{k-1})) - \widetilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \leq \frac{2U_R^2 B^2 C_w (1 + \sqrt{\gamma})}{(1-\gamma)((1-\sqrt{\gamma})^2)} \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2]. \quad (44)$$

Utilizing the upper bound in (44) into (31), we obtain

$$\mathbb{E} \left[\left\| \nabla F(\boldsymbol{\theta}_k, \xi_k(\boldsymbol{\theta}_k)) - \widetilde{\nabla} F(\boldsymbol{\theta}_{k-1}, \xi_k(\boldsymbol{\theta}_k)) \right\|^2 \right] \leq L_1 \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2], \quad (45)$$

where we define $L_1 := \left(2L^2 + \frac{2U_R^2 B^2 C_w (1 + \sqrt{\gamma})}{(1-\gamma)((1-\sqrt{\gamma})^2)} \right)$. Hence, using (45) into (30), we obtain

$$\varepsilon_k \leq (1-\beta)^2 \varepsilon_{k-1} + 2(1-\beta)^2 \eta^2 L_1 \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2] + 2\beta^2 m_0 + 2m_1 \beta^2 \|\nabla F(\boldsymbol{\theta}_k)\|^2.$$

From the definition of generalized Bregman gradient, we can write

$$\varepsilon_k \leq (1-\beta)^2 \varepsilon_{k-1} + 2(1-\beta)^2 L_1 \mathbb{E} [\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2] + 2\beta^2 m_0 + 2m_1 \beta^2 \|\nabla F(\boldsymbol{\theta}_k)\|^2 \quad (46)$$

$$\leq (1-\beta)^2 \varepsilon_{k-1} + 2\eta^2 L_1 \mathbb{E} [\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|^2] + 2\beta^2 m_0 + 2m_1 \beta^2 \|\nabla F(\boldsymbol{\theta}_k)\|^2, \quad (47)$$

where (47) follows from substituting in the difference of policy parameters by the generalized gradient in (12), and the second expression uses the fact that $\beta \leq 1$. \square

C. Proof of Theorem 4.9

Proof. From the smoothness of $F(\cdot)$ [cf. Assumption 4.6], it holds that

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) + \langle \nabla F(\boldsymbol{\theta}_k), \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \rangle + \frac{L}{2} \|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k\|^2. \quad (48)$$

From the update in (12) (note that the analysis is for descent update), it holds that

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k). \quad (49)$$

Let us utilize (49) into (48), we get

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \eta \langle \nabla F(\boldsymbol{\theta}_k), \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k) \rangle + \frac{L\eta^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2. \quad (50)$$

Next, we add subtract the gradient $\hat{\mathbf{g}}_k$ as follows

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \eta \langle \nabla F(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}_k + \hat{\mathbf{g}}_k, \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k) \rangle + \frac{L\eta_k^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \quad (51)$$

$$= F(\boldsymbol{\theta}_k) - \eta \langle \hat{\mathbf{g}}_k, \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k) \rangle + \frac{L\eta_k^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \eta \langle \mathbf{w}_k, \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k) \rangle \quad (52)$$

where we utilized the definition $\mathbf{w}_k := \hat{\mathbf{g}}_k - \nabla F(\boldsymbol{\theta}_k)$ is the stochastic error in the gradient in the second line. Next, using the lower bound from (21), we can substitute the second term by the square-norm of the Bregman gradient as

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \eta\zeta \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \frac{L\eta_k^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \eta \langle \mathbf{w}_k, \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k) \rangle. \quad (53)$$

Next, for the last inner product term on the right hand side of (53), we use $ab \leq \beta a^2 + \frac{1}{\beta} b^2$ for any $\beta > 0$ with $a = \mathbf{w}_k$ and $b = \eta \mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)$ to rewrite the preceding expression as

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \eta\zeta \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \beta \|\mathbf{w}_k\|^2 + \frac{\eta^2}{\beta} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \frac{L\eta^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2. \quad (54)$$

Next, after grouping like terms, we get

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \eta\zeta \left(1 - \frac{\eta}{\beta\zeta}\right) \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \frac{L\eta^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \beta \|\mathbf{w}_k\|^2. \quad (55)$$

By selecting $\beta \geq \frac{2\eta}{\zeta}$, it holds that,

$$F(\boldsymbol{\theta}_{k+1}) \leq F(\boldsymbol{\theta}_k) - \frac{\eta\zeta}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \frac{L\eta^2}{2} \|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 + \beta \|\mathbf{w}_k\|^2. \quad (56)$$

Taking the expectation on the both sides, we get

$$\mathbb{E}[F(\boldsymbol{\theta}_{k+1})] \leq \mathbb{E}[F(\boldsymbol{\theta}_k)] - \frac{\eta\zeta}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \frac{L\eta^2}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \beta \mathbb{E} \left[\|\mathbf{w}_k\|^2 \right]. \quad (57)$$

Before proceeding next, let us provide an upper bound on the term $\mathbb{E} \left[\|\mathbf{w}_k\|^2 \right]$ in the form of Lemma 4.8.

From the result in (57) and the statement of Lemma 4.8, we write

$$\varepsilon_{k+1} \leq (1 - \beta)^2 \varepsilon_k + 2\eta^2 L_{t_{\max}} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + 2\beta^2 m_0 + 2m_1 \beta^2 \|\nabla F(\boldsymbol{\theta}_k)\|^2 \quad (58)$$

$$\mathbb{E}[F(\boldsymbol{\theta}_{k+1})] \leq \mathbb{E}[F(\boldsymbol{\theta}_k)] - \frac{\eta\zeta}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \frac{L\eta^2}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \beta \varepsilon_k. \quad (59)$$

Adding (58) and (59) then yields

$$\mathbb{E}[F(\boldsymbol{\theta}_{k+1})] + \varepsilon_{k+1} \leq \mathbb{E}[F(\boldsymbol{\theta}_k)] + \varepsilon_k - \varepsilon_k - \frac{\eta\zeta}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \frac{L\eta^2}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + \beta \varepsilon_k \quad (60)$$

$$+ (1 - \beta)^2 \varepsilon_k + 2\eta^2 L_{t_{\max}} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + 2\beta^2 m_0 + 2m_1 \beta^2 \mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}_k)\|^2 \right]. \quad (61)$$

Define Lyapunov function $\Phi_k := \mathbb{E}[F(\boldsymbol{\theta}_k)] + \varepsilon_k$. With this definition, after rearranging the preceding expression, we then obtain

$$\Phi_{k+1} - \Phi_k + \frac{\eta\zeta}{2} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq A_k \varepsilon_k + \frac{5\eta^2 L' \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right]}{2} + 2\beta^2 m_0 + 2m_1 \beta^2 \mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}_k)\|^2 \right], \quad (62)$$

$L' = \max\{L, L_{t_{\max}}\}$ and $A_k = (1 - \beta)^2 + \beta - 1$. Next, our goal is to establish that Φ_k is decreasing. To do so, we study the coefficient A_k in front of ε_k in more detail

$$A_k = 1 + \beta^2 - 2\beta + \beta - 1 \quad (63)$$

$$= -\beta(1 - \beta). \quad (64)$$

Note that since $\beta \leq 1$, we have $A_k \leq 0$ for all t , which allows us to drop the first term on the right-hand side of (62) and hence write

$$\left(\frac{\eta\zeta}{2} - \frac{5\eta^2}{2L'} \right) \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\mathbf{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 2m_1 \beta^2 \mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}_k)\|^2 \right], \quad (65)$$

Next we use the error-bound condition on the second moment of the gradient $\mathbb{E} [\|\nabla F(\boldsymbol{\theta}_k)\|^2]$ in Assumption 4.5 as follows

$$\mathbb{E} [\|\nabla F(\boldsymbol{\theta}_k)\|^2] = \mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}_k) - \mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k) + \mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \quad (66)$$

$$\leq 2\mathbb{E} \left[\|\nabla F(\boldsymbol{\theta}_k) - \mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] + 2\mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \quad (67)$$

$$\leq 2m_2 + (2 + m_3) \|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2. \quad (68)$$

Utilizing the upper bound of (68) into the right hand side of (65), we get

$$\left(\frac{\eta\zeta}{2} - \frac{5\eta^2}{2L'} \right) \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 4m_1 m_2 \beta^2 + \tilde{m}_3 \beta^2 \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right], \quad (69)$$

where $\tilde{m}_3 = 2(2 + m_3)$.

By selecting $\eta \leq \frac{\zeta}{10L'}$ (first condition on η), we can lower bound the left hand side of (69), as

$$\frac{\eta\zeta}{4} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 4m_1 m_2 \beta^2 + \tilde{m}_3 \beta^2 \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right]. \quad (70)$$

After rearranging the terms, we get

$$\left(\frac{\eta\zeta}{4} - \tilde{m}_3 \beta^2 \right) \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 4m_1 m_2 \beta^2. \quad (71)$$

Let us select $\beta = C_1 \eta$ where C_1 is such that $C_1 > \frac{2}{\zeta}$ which would satisfy our requirement for β . This implies that

$$\frac{\eta\zeta}{8} \left(2 - \frac{8\tilde{m}_3 C_1^2 \eta}{\zeta} \right) \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 4m_1 m_2 \beta^2. \quad (72)$$

Again if we make sure that $\eta \leq \frac{\zeta}{8\tilde{m}_3 C_1^2}$ [second condition for η], we could lower bound the above expression as

$$\frac{\eta\zeta}{8} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \Phi_k - \Phi_{k+1} + 2\beta^2 m_0 + 4m_1 m_2 \beta^2. \quad (73)$$

Taking summation and lower bounding the left hand side with the minimum, we will obtain the similar rate as

$$\min_{1 \leq k \leq K} \mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right] \leq \mathcal{O} \left(\frac{1}{\eta K} + \eta \right), \quad (74)$$

where we used $\beta = C_1 \eta$ and absorb the constants in the \mathcal{O} notation. From the above expression, by selecting $\eta = \epsilon$, we will obtain $K = \mathcal{O}(\epsilon^{-2})$. The explicit values of step size by collecting the conditions for η and β together could be written as

$$\beta = C_1 \eta \quad \text{and} \quad \eta = \min \left\{ \frac{\zeta L'}{10}, \frac{\zeta}{8\tilde{m}_3 C_1^2} \right\}. \quad (75)$$

Similarly, if we choose, $\eta = \epsilon^2$, then from the Jensen's inequality, it would hold that $\mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\| \right] \leq \sqrt{\mathbb{E} \left[\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^\psi(\boldsymbol{\theta}_k)\|^2 \right]} \leq \epsilon$ with $K = \mathcal{O}(\frac{1}{\epsilon^4})$. \square

D. Verification of Assumptions 4.3-4.6 for the RL Setting

1. Utilizing the boundedness of the score function as mentioned in Lemma 4.7, combined with the proof of (Bedi et al., 2021, Lemma 1), it holds that Assumption 4.3 is satisfied.
2. Assumption 4.4 is regarding the variance of the stochastic gradient estimate which we assume to be bounded by some constant m_0 with $m_1 = 0$. This assumption is standard in the literature (see (Xu et al., 2017; 2019; Papini et al., 2018; Zhang et al., 2020b)).

3. To prove Assumption 4.5, we start by noting that $\nabla F(\boldsymbol{\theta}) = \mathbb{E}[\nabla F(\boldsymbol{\theta}, \xi)]$ because stochastic gradient estimator is unbiased. Now we take norm on both sides, we get $\|\nabla F(\boldsymbol{\theta})\| = \|\mathbb{E}[\nabla F(\boldsymbol{\theta}, \xi)]\|$. Since norm is convex, we know that $\|\nabla F(\boldsymbol{\theta})\| \leq \mathbb{E}\|\nabla F(\boldsymbol{\theta}, \xi)\|$. For the RL settings, from (4), we know that

$$\nabla F(\boldsymbol{\theta}, \xi) = \sum_{t=0}^T \gamma^{t/2} r_t \cdot \left(\sum_{\tau=0}^t \nabla \log \pi_{\boldsymbol{\theta}}(a_{\tau} | s_{\tau}) \right). \quad (76)$$

Taking norm on both sides, applying triangle inequality, and utilizing the bound on reward and norm of the score function, we note that

$$\|\nabla F(\boldsymbol{\theta}, \xi)\| \leq \frac{U_R B}{1 - \gamma} T. \quad (77)$$

Since $T \sim \text{Geom}(1 - \gamma^{1/2})$, after taking expectation, we can write that

$$\mathbb{E} [\|\nabla F(\boldsymbol{\theta}, \xi)\|^2] \leq \frac{U_R B(1 + \sqrt{\gamma})}{(1 - \gamma)(1 - \sqrt{\gamma})^2}. \quad (78)$$

From the state of Assumption 4.5, we note that term $\mathbb{E} [\|\nabla F(\boldsymbol{\theta}_k) - \mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^{\psi}(\boldsymbol{\theta}_k)\|^2]$ and use upper bound it as

$$\mathbb{E} [\|\nabla F(\boldsymbol{\theta}_k) - \mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^{\psi}(\boldsymbol{\theta}_k)\|^2] \leq 2\mathbb{E} [\|\nabla F(\boldsymbol{\theta}_k)\|^2] + 2\mathbb{E} [\|\mathcal{G}_{\eta, \hat{\boldsymbol{g}}_k}^{\psi}(\boldsymbol{\theta}_k)\|^2]. \quad (79)$$

Hence, Assumption 4.5 holds with $m_2 = \frac{2U_R B(1 + \sqrt{\gamma})}{(1 - \gamma)(1 - \sqrt{\gamma})^2}$ and $m_3 = 2$.

4. Since with the Cauchy policy parametrization, the score function is Lipschitz with parameter $L^{\pi} = \frac{2D^2}{\sigma^2} + \frac{7D}{\sigma^2} + 1$. Using the Lipschitz property of the score function, along with the boundedness of the score function $\|\nabla \log \pi_{\boldsymbol{\theta}}(a|s)\| \leq B$, it holds that the objective F is smooth with constant L (proof is provided in (Zhang et al., 2020b, Lemma 3.2)) and L is given by

$$L = \frac{U_R L^{\pi}}{(1 - \gamma)^2} + \frac{(1 + \gamma)U_R B^2}{(1 - \gamma)^3}. \quad (80)$$