

Global Convergence of Policy Gradient Methods in Reinforcement Learning, Games and Control

Shicong Cen
Carnegie Mellon University

Yuejie Chi
Carnegie Mellon University*

Abstract

Policy gradient methods, where one searches for the policy of interest by maximizing the value functions using first-order information, become increasingly popular for sequential decision making in reinforcement learning, games, and control. Guaranteeing the global optimality of policy gradient methods, however, is highly nontrivial due to nonconcavity of the value functions. In this exposition, we highlight recent progresses in understanding and developing policy gradient methods with global convergence guarantees, putting an emphasis on their finite-time convergence rates with regard to salient problem parameters.

1 Introduction

Sequential decision making is a canonical task that lies at the heart of a wide spectrum of disciplines such as reinforcement learning (RL), games and control, finding numerous applications in autonomous driving, robotics, supply chain management, resource scheduling, and so on. While classical approaches such as dynamic programming (Bertsekas, 2017) require model knowledge, modern practices advocate for the model-free approach as an alternative: one seeks the policy of interest directly based on data collected through interactions with the environments, without estimating the model. The advantage is that the model-free approach is often more memory efficient, as well as more agile to changes.

One prevalent class of model-free approaches is policy gradient (PG) methods, which follow an optimizer’s perspective by formulating a value maximization problem with regard to parameterized policies, and performing gradient updates to improve the policy iteratively — often based on noisy feedbacks received from the environments. PG methods and their variants have become the de facto standard practices in an increasing number of domains, due to their seamless integration with neural network parameterization and adaptivity to various problem setups involving discrete, continuous or mixed action and state spaces.

Despite the great empirical success of PG methods, little is known about their theoretical convergence properties — especially when it comes to finite-time global convergence —

due to the notorious nonconcavity of the value functions. Until very recently, such understandings begin to emerge, leveraging the fact that the PG methods are typically operated on highly structured model classes, whose induced optimization landscapes turn out to be much more benign and tractable than previously thought. The execution of problem-dependent tailored analyses, rather than relying on black-box optimization theory, fuels recent breakthroughs pioneered by Fazel et al. (2018); Agarwal et al. (2021); Bhandari and Russo (2019), to name just a few. The purpose of this article is to survey the latest efforts in understanding the global convergence of PG methods in the fields of RL, game theory and control, as well as highlight algorithmic ideas that enable fast global convergence, especially with regard to salient problem parameters that are crucial in practice.

Organization. The rest of this article is organized as follows. Section 2 reviews PG methods in single-agent RL, focusing on solving tabular Markov decision processes (MDP). Section 3 reviews PG methods in the game and multi-agent RL setting, using the two-player zero-sum matrix game and two-player zero-sum Markov game as illustrative examples. Section 4 moves onto PG methods in control, focusing on solving the linear quadratic regulator (LQR). We conclude in Section 5.

Notation. We use $\|A\|$, $\|A\|_F$ and $\sigma_{\min}(A)$ to represent the spectral norm, the Frobenius norm, and the smallest singular value of a matrix A , respectively. For two vectors a and b , we use $\frac{a}{b}$ to denote their entrywise division when exists,

*S. Cen and Y. Chi are with the Department of Electrical and Computer Engineering, Carnegie Mellon University; emails: {shicongc,yuejiec}@andrew.cmu.edu.

and $\|a\|_\infty$ denotes the entrywise maximum absolute value of a vector a . Given a set \mathcal{S} , we let $\Delta(\mathcal{S})$ represent the probability simplex over \mathcal{S} . Given two distributions p and q , $\text{KL}(p\|q)$ denotes the Kullback-Leibler (KL) divergence from q to p . Last but not least, let $\mathcal{P}_{\mathcal{C}}(\cdot)$ be the projection operator onto the set \mathcal{C} . To characterize the iteration complexity, we write $T(x) = \mathcal{O}(f(x))$ when $T(x) \leq Cf(x), \forall x \geq X$ for some $C, X > 0$. Here x is typically set to $1/\epsilon$ or other salient problem parameters. We use $\tilde{\mathcal{O}}(\cdot)$ to suppress logarithmic factors from the standard order notation $\mathcal{O}(\cdot)$.

2 Global convergence of policy gradient methods in RL

In this section, we review recent progresses in developing policy gradient methods for single-agent RL, focusing on the basic model of tabular Markov decision processes (MDPs).

2.1 Problem settings

Markov decision processes. An infinite-horizon discounted MDP models the sequential decision making problem as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with the state space \mathcal{S} , the action space \mathcal{A} , the transition kernel $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, the reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and the discount factor $\gamma \in (0, 1)$. Upon the agent choosing action $a \in \mathcal{A}$ at state $s \in \mathcal{S}$, the environment will move to a new state s' according to the transition probability $P(s', s, a)$ and assign a reward $r(s, a)$ to the agent. The action selection rule is implemented by a randomized policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$, where $\pi(a|s)$ specifies the probability of choosing action a in state s . We shall focus exclusively on the case where both \mathcal{S} and \mathcal{A} are finite throughout this article.

Value functions and Q-functions. The value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ is defined as the expected discounted cumulative reward starting at state s :

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

where $a_t \sim \pi(\cdot|s_t)$ is obtained by executing policy π and $s_{t+1} \sim P(\cdot|s_t, a_t)$ is generated by the MDP. The Q-function (or action-value function) $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined in a similar manner with an initial state-action pair (s, a) :

$$Q^\pi(s, a) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right].$$

In addition, the advantage function of policy π is defined as

$$A^\pi(s, a) := Q^\pi(s, a) - V^\pi(s).$$

It is well-known that there exists an optimal policy π^* that maximizes the value function $V^\pi(s)$ for all $s \in \mathcal{S}$ simultaneously (Bellman, 1952; Puterman, 2014). We denote the resulting optimal value function and Q-function by V^* and Q^* , which satisfy the well-known Bellman optimality equations. The optimal policy π^* can be implied from the optimal Q-function in a greedy fashion as

$$\pi^*(a|s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a), \quad \forall s \in \mathcal{S}. \quad (1)$$

Policy parameterizations. Given some prescribed initial distribution ρ over \mathcal{S} , policy optimization methods seek to maximize the value function

$$V^\pi(\rho) := V^{\pi_\theta}(\rho)$$

over the policy π , where $\pi := \pi_\theta$ is often parameterized via some parameter θ , and we overload the notation to denote by $V^\pi(\rho)$ the expected value function $\mathbb{E}_{s \sim \rho}[V^\pi(s)]$. Note that it is straightforward to observe that $V^*(\phi) - V^\pi(\phi) \leq \|\phi/\rho\|_\infty (V^*(\rho) - V^\pi(\rho))$ for general choices of $\phi \in \Delta(\mathcal{S})$, which characterizes the effect of possible discrepancy between the initial distributions in training and deploying. Common choices of policy parameterization are as follows.

- *Direct parameterization:* the policy is directly parameterized by

$$\pi_\theta(a|s) = \theta(s, a),$$

where $\theta \in \{\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|} : \theta(s, a) \geq 0, \sum_{a \in \mathcal{A}} \theta(s, a) = 1\}$.

- *Tabular softmax parameterization:* For $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, the policy π_θ is generated through softmax transform

$$\pi_\theta(a|s) = \frac{\exp(\theta(s, a))}{\sum_{a' \in \mathcal{A}} \exp(\theta(s, a'))},$$

leading to an unconstrained optimization over θ . Although beyond the scope of this exposition, softmax parameterization is more popular with function approximation, by replacing $\theta(s, a)$ with $f_\theta(s, a)$ typically implemented by neural networks.

Policy gradients. The gradient $\nabla_\theta V^{\pi_\theta}(\rho)$ plays an instrumental role in developing first-order policy optimization methods. To facilitate presentation, we shall first introduce the discounted state visitation distributions of a policy π :

$$d_{s_0}^\pi(s) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}(s_t = s | s_0),$$

where the expectation is with respect to the trajectory $(s_0, a_0, s_1, a_1, \dots)$ generated by the MDP under policy π . We

further denote by d_ρ^π the discounted state visitation distribution when s_0 is randomly drawn from distribution ρ , i.e.,

$$d_\rho^\pi(s) := \mathbb{E}_{s_0 \sim \rho} [d_{s_0}^\pi(s)].$$

The policy gradient of parameterized policy π_θ is then given by (Williams, 1992)

$$\begin{aligned} \nabla_\theta V^{\pi_\theta}(\rho) &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} [\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s, a)] \\ &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} [\nabla_\theta \log \pi_\theta(a|s) A^{\pi_\theta}(s, a)], \end{aligned}$$

which can be evaluated, for example, via REINFORCE (Williams, 1992). The use of the advantage function $A^{\pi_\theta}(s, a)$, rather than the Q-function $Q^{\pi_\theta}(s, a)$, often helps to reduce the variance of the estimated policy gradient.

For notational simplicity, we shall denote by $\theta^{(t)}$ and $\pi^{(t)}$ the parameter and the policy at the t -th iteration, and use $V^{(t)}$, $Q^{(t)}$, $A^{(t)}$, $d_\rho^{(t)}$ to denote $V^{\pi^{(t)}}$, $Q^{\pi^{(t)}}$, $A^{\pi^{(t)}}$, $d_\rho^{\pi^{(t)}}$, respectively. In addition, we assume the policy gradients and the value functions are exactly evaluated throughout this article, which enables us to focus on the optimization aspect of PG methods.

2.2 Projected policy gradient method

The most straightforward first-order policy optimization method is to adopt direct parameterization and perform projected gradient ascent updates:

$$\theta^{(t+1)} = \mathcal{P}_{\Delta(\mathcal{A})|S}(\theta^{(t)} + \eta \nabla_\theta V^{(t)}(\rho)), \quad (2)$$

or equivalently

$$\pi^{(t+1)} = \mathcal{P}_{\Delta(\mathcal{A})|S}(\pi^{(t)} + \eta \nabla_\pi V^{(t)}(\rho)),$$

where $\eta > 0$ is the learning rate, and

$$\nabla_{\theta(s,a)} V^{(t)}(\rho) = \nabla_{\pi(s,a)} V^{(t)}(\rho) = \frac{1}{1-\gamma} d_\rho^{(t)}(s) Q^{(t)}(s, a).$$

As the value function $V^{\pi_\theta}(\rho)$ is $\frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3}$ -smooth (Agarwal et al., 2021), setting the learning rate to $0 < \eta \leq \frac{(1-\gamma)^3}{2\gamma|\mathcal{A}|}$ ensures monotonicity of $V^{(t)}(\rho)$ in t . On the other end, it is critically established in Agarwal et al. (2021) that the value function satisfies the following gradient domination condition.

Lemma 1 (Variational gradient domination). *For any policy π , we have*

$$V^*(\rho) - V^\pi(\rho) \leq \frac{1}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\rho} \right\|_\infty \max_{\pi' \in \Delta(\mathcal{A})|S} (\pi' - \pi)^\top \nabla_\pi V^\pi(\rho).$$

The above lemma associates the optimality gap $V^*(\rho) - V^\pi(\rho)$ with a variational gradient term, allowing the iterates to converge globally as stated below.

Theorem 1 (Agarwal et al. (2021)). *With $0 < \eta \leq \frac{(1-\gamma)^3}{2\gamma|\mathcal{A}|}$, the iterates of the projected PG method (2) satisfies*

$$\min_{0 \leq t \leq T} V^*(\rho) - V^{(t)}(\rho) \leq \frac{4\sqrt{|\mathcal{S}|}}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\rho} \right\|_\infty \sqrt{\frac{2(V^*(\rho) - V^{(0)}(\rho))}{\eta T}}.$$

Theorem 1 establishes an iteration complexity of $\mathcal{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^6 \epsilon^2} \left\| \frac{d_\rho^{\pi^*}}{\rho} \right\|_\infty^2\right)$ for finding an ϵ -optimal policy, which is later improved to $\mathcal{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^5 \epsilon} \left\| \frac{d_\rho^{\pi^*}}{\rho} \right\|_\infty^2\right)$ (Xiao, 2022). However, the projection operator introduces $\mathcal{O}(\log |\mathcal{A}|)$ computational overhead every iteration and is less compatible with function approximation. This motivates the study of PG methods that are compatible with unconstrained optimization, e.g., by using softmax parameterization.

2.3 Softmax policy gradient method

With softmax parameterization, the policy gradient method writes

$$\theta^{(t+1)} = \theta^{(t)} + \eta \nabla_\theta V^{(t)}(\rho), \quad (3)$$

where

$$\nabla_{\theta(s,a)} V^{(t)}(\rho) = \frac{\eta}{1-\gamma} d_\rho^{(t)}(s) \pi^{(t)}(a|s) A^{(t)}(s, a).$$

Remarkably, Agarwal et al. (2021) established the asymptotic global convergence of the softmax PG method as follows.

Theorem 2 (Agarwal et al. (2021)). *With constant learning rate $0 < \eta \leq (1-\gamma)^3/8$, the softmax PG method converges to the optimal policy, i.e., $V^{(t)}(s) \rightarrow V^*(s)$ as $t \rightarrow \infty$ for all $s \in \mathcal{S}$.*

Mei et al. (2020) later demonstrated an iteration complexity of $\mathcal{O}\left(\frac{1}{c(\mathcal{M})^2 \epsilon}\right)$ for achieving an ϵ -optimal policy, where $c(\mathcal{M})$ is a trajectory-dependent quantity depending on salient problem parameters such as the number of states $|\mathcal{S}|$ and the effective horizon $(1-\gamma)^{-1}$. Unfortunately, this quantity $c(\mathcal{M})$ can be rather small and does not exclude the possibility of incurring excessively large iteration complexity, as demonstrated by the following hardness result (Li et al., 2023).

Theorem 3 (Li et al. (2023)). *There exist universal constants $c_1, c_2, c_3 > 0$ such that for any $\gamma \in (0.96, 1)$ and $|\mathcal{S}| \geq c_3(1-\gamma)^{-6}$, one can find a γ -discounted MDP such that the softmax PG method takes at least*

$$\frac{c_1}{\eta} |\mathcal{S}|^2 \frac{c_2}{1-\gamma}$$

iterations to reach $\|V^ - V^{(t)}\|_\infty \leq 0.15$.*

Therefore, though guaranteed to converge globally, softmax PG can take (super-)exponential time to even reduce the optimality gap within a constant level. Intuitively speaking, softmax PG method fails to achieve a reasonable convergence rate when the probability $\pi^{(t)}(a^*(s)|s)$ assigned to the optimal action $a^*(s)$ is close to zero. Agarwal et al. (2021) proposed to penalize the policy for getting too close to the border, by imposing a log barrier regularization

$$V_\omega^{\pi_\theta}(\rho) = V^{\pi_\theta}(\rho) + \frac{\omega}{|\mathcal{S}||\mathcal{A}|} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \log \pi_\theta(a|s).$$

With an appropriate choice of regularization parameter ω , the regularized softmax PG can achieve an ϵ -optimal policy within $\mathcal{O}\left(\frac{|\mathcal{S}|^2|\mathcal{A}|^2}{(1-\gamma)^6\epsilon^2} \left\| \frac{d_\rho^*}{\rho} \right\|_\infty^2\right)$ iterations (Agarwal et al., 2021). Nonetheless, this regularization scheme is not as popular in practice, compared to the entropy regularization scheme that will be discussed momentarily.

2.4 Natural policy gradient method

Both projected PG and softmax PG fall short of attaining an iteration complexity that is independent of salient problem parameters, especially with respect to the size of the state space $|\mathcal{S}|$. This ambitious goal can be achieved, somewhat surprisingly, by adopting Fisher information matrix as a preconditioner, which leads to the natural policy gradient (NPG) method (Kakade, 2002):

$$\theta^{(t+1)} = \theta^{(t)} + \eta(\mathcal{F}_\rho^{\theta^{(t)}})^\dagger \nabla_\theta V^{(t)}(\rho), \quad (4)$$

where

$$\mathcal{F}_\rho^\theta := \mathbb{E}_{s \sim d_\rho^\pi, a \sim \pi_\theta(\cdot|s)} \left[(\nabla_\theta \log \pi_\theta(a|s)) (\nabla_\theta \log \pi_\theta(a|s))^\top \right]$$

is the Fisher information matrix, and \dagger denotes the Moore-Penrose pseudoinverse. With softmax parameterization, the NPG updates take the form

$$\theta^{(t+1)} = \theta^{(t)} + \frac{\eta}{1-\gamma} A^{(t)},$$

or equivalently,

$$\pi^{(t+1)}(a|s) \propto \pi^{(t)}(a|s) \exp\left(\frac{\eta Q^{(t)}(s, a)}{1-\gamma}\right),$$

It is noted that the (softmax) NPG update rule coincides with the multiplicative weights update (MWU) method (Cesa-Bianchi and Lugosi, 2006), and that the update rule does not depend on the initial state distribution ρ . Shani et al. (2020) first established a global convergence rate of $\mathcal{O}\left(\frac{1}{(1-\gamma)^2\sqrt{T}}\right)$ using decaying learning rate $\eta_t = \mathcal{O}\left(\frac{1-\gamma}{\sqrt{t}}\right)$, which was improved by Agarwal et al. (2021) using a constant learning rate η , stated as follows.

Theorem 4 (Agarwal et al. (2021)). *With uniform initialization $\theta^{(0)} = 0$ and constant learning rate $\eta > 0$, the iterates of NPG satisfy*

$$V^*(\rho) - V^{(T)}(\rho) \leq \frac{1}{T} \left(\frac{\log |\mathcal{A}|}{\eta} + \frac{1}{(1-\gamma)^2} \right).$$

Encouragingly, as long as $\eta \geq \frac{(1-\gamma)^2}{\log |\mathcal{A}|}$, the iteration complexity of NPG methods becomes $\mathcal{O}\left(\frac{1}{(1-\gamma)^2 T}\right)$, which is independent of the size of the state-action space. On the complementary side, the iteration complexity of NPG is lower bounded by $\frac{\Delta}{(1-\gamma)|\mathcal{A}|} \exp(-\eta\Delta T)$ — established in Khodadadian et al. (2021) — where the optimal advantage function gap $\Delta = \min_s \min_{a \neq a^*(s)} |A^*(s, a)| \geq 0$ is determined by the MDP instance. As the lower bound attains its maximum $\frac{1}{(1-\gamma)e\eta T}$ when $\Delta = \frac{1}{\eta T}$, it is immediate that the sublinear rate in Theorem 4 cannot be improved in T . Nonetheless, two strategies to achieve even faster linear convergence with NPG updates include (i) adopting increasing/adaptive learning rates (Khodadadian et al., 2021; Bhandari and Russo, 2021; Lan, 2023; Xiao, 2022), or (ii) introducing entropy regularization (Cen et al., 2022; Lan, 2023; Zhan et al., 2023), which we shall elaborate on next.

2.5 Entropy regularization

Introducing entropy regularization is a popular technique in practice to promote exploration (Haarnoja et al., 2017). Specifically, one seeks to optimize the entropy-regularized value function defined as

$$V_\tau^\pi(s) = V^\pi(s) + \frac{\tau}{1-\gamma} \mathbb{E}_{s' \sim d_s^\pi} [\mathcal{H}(\pi(\cdot|s'))],$$

where $\mathcal{H}(\pi(\cdot|s)) = -\sum_{a \in \mathcal{A}} \pi(a|s) \log \pi(a|s)$ is the entropy of policy $\pi(\cdot|s)$, and $\tau > 0$ serves as the regularization parameter known as the temperature. The entropy-regularized Q -function is defined as

$$Q_\tau^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a)} [V_\tau^\pi(s')].$$

The resulting optimal value function, Q -function and optimal policy are denoted by V_τ^* , Q_τ^* , and π_τ^* . From an optimization perspective, the entropy term adds curvature to the value function and ensures that the optimal policy π_τ^* is unique. Interestingly, in contrast to the greedy optimal policy for the unregularized problem in (1), the optimal policy of the entropy-regularized problem reflects “bounded rationality” in decision making, namely

$$\pi_\tau^*(\cdot|s) \propto \exp(Q_\tau^*(s, \cdot)/\tau).$$

It should be noted, however, adding the entropy regularization generally does not make $V_\tau^\pi(\rho)$ concave unless τ is unreasonably large. As the entropy function is bounded by $\log |\mathcal{A}|$,

the optimal entropy-regularized policy is also guaranteed to be approximately optimal for the unregularized RL problem in the following sense:

$$V^{\pi_\tau^*}(\rho) \geq V^*(\rho) - \frac{\tau \log |\mathcal{A}|}{1 - \gamma}.$$

Motivated by its benign convergence, we consider NPG for the entropy-regularized problem:

$$\theta^{(t+1)} = \theta^{(t)} + \eta(\mathcal{F}_\rho^{\theta^{(t)}})^\dagger \nabla_\theta V_\tau^{(t)}(\rho),$$

which can be equivalently written as

$$\pi^{(t+1)}(a|s) \propto \pi^{(t)}(a|s)^{1 - \frac{\eta\tau}{1-\gamma}} \exp\left(\frac{\eta Q_\tau^{(t)}(s, a)}{1 - \gamma}\right). \quad (5)$$

The following theorem shows that with appropriate choices of constant learning rate η , entropy-regularized NPG converges to the unique optimal policy π_τ^* at a linear rate.

Theorem 5. *For constant learning rate $0 < \eta \leq (1 - \gamma)/\tau$ and uniform initialization, the entropy-regularized NPG updates (5) satisfy*

$$\|V_\tau^* - V_\tau^{(T)}\|_\infty \leq \frac{15(1 + \tau \log |\mathcal{A}|)}{1 - \gamma} (1 - \eta\tau)^{T-1}$$

and

$$V_\tau^*(\rho) - V_\tau^{(T)}(\rho) \leq \left\| \frac{\rho}{\nu_\tau^*} \right\|_\infty \left(\frac{1 + \tau \log |\mathcal{A}|}{1 - \gamma} + \frac{(1 - \gamma) \log |\mathcal{A}|}{\eta} \right) \cdot \max \left\{ \gamma, 1 - \frac{\eta\tau}{1 - \gamma} \right\}^T.$$

Here, ν_τ^* is the stationary state distribution of policy π_τ^* .

The first and the second bounds are due to [Cen et al. \(2022\)](#) and [Lan \(2023\)](#)¹ respectively, where they lead to slightly different iteration complexities. Taken collectively, entropy-regularized NPG takes no more than

$$\tilde{\mathcal{O}}\left(\min\left\{\frac{1}{\eta\tau} \log \frac{1}{\epsilon}, \max\left\{\frac{1}{1 - \gamma}, \frac{1 - \gamma}{\eta\tau}\right\} \log \frac{\|\rho/\nu_\tau^*\|_\infty}{\epsilon}\right\}\right)$$

iterations to find a policy satisfying $V_\tau^*(\rho) - V_\tau^\pi(\rho) \leq \epsilon$. We make note that the difference stems from different analysis approaches: [Cen et al. \(2022\)](#) built their analysis upon the contraction property of the soft Bellman operator (the entropy-regularized counterpart of the original Bellman operator), while [Lan \(2023\)](#) made use of the connection between regularized NPG and regularized mirror descent. This can be

¹We discard some of the simplification steps therein and state the convergence result for a wider range of learning rate η .




Figure 1: Comparison of PG and NPG methods with entropy regularization for a bandit problem ($\gamma = 0$) with 3 actions associated with rewards 1.0, 0.9 and 0.1. The regularization parameter is set to $\tau = 0.1$.

observed from the following equivalence: the update rule (5) can be equivalently expressed as

$$\begin{aligned} \pi^{(t+1)}(\cdot|s) = \arg \min_{p \in \Delta(\mathcal{A})} & \langle p, -Q_\tau^{(t)}(s, \cdot) \rangle - \tau \mathcal{H}(p) \\ & + \frac{1}{\eta_{\text{MD}}} \text{KL}(p \| \pi^{(t)}(\cdot|s)), \end{aligned}$$

with $\eta_{\text{MD}} = \frac{\eta}{1 - \gamma - \eta\tau}$. The analysis of regularized RL can be further generalized to adopt non-strongly convex regularizers ([Lan, 2023](#)), non-smooth regularizers ([Zhan et al., 2023](#)), state-wise policy updates ([Lan et al., 2023](#)), and so on.

3 Global convergence of policy gradient methods in games

In this section, we review recent progresses in developing policy gradient methods for games and multi-agent RL, focusing on the basic models of two-player zero-sum matrix games and two-player zero-sum Markov games.

3.1 Problem settings

Two-player zero-sum matrix games. Given two players each taking actions from action spaces \mathcal{A} and \mathcal{B} , the matrix game aims to solve the following saddle-point optimization problem

$$\max_{\mu \in \Delta(\mathcal{A})} \min_{\nu \in \Delta(\mathcal{B})} V^{\mu, \nu} := \mu^\top A \nu, \quad (6)$$

where $A \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{B}|}$ denotes the payoff matrix with $\|A\|_\infty \leq 1$, $\mu \in \Delta(\mathcal{A})$ and $\nu \in \Delta(\mathcal{B})$ stand for the mixed/randomized policies of each player, defined respectively as distributions over the probability simplex $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$. Here, one player seeks to maximize the value function (i.e., the max player) while the other player wants to minimize it (i.e., the min player). It is well-known since [Neumann \(1928\)](#) that the max

and min operators in (6) can be exchanged without affecting the solution. A pair of policies (μ^*, ν^*) is said to be a *Nash equilibrium (NE)* of the matrix game if

$$V^{\mu^*, \nu} \geq V^{\mu^*, \nu^*} \geq V^{\mu, \nu^*} \quad \forall (\mu, \nu) \in \Delta(\mathcal{A}) \times \Delta(\mathcal{B}). \quad (7)$$

In words, the NE corresponds to when both players play their best-response strategies against their opponents.

Two-player zero-sum Markov games. Moving onto sequential decision-making, we consider an infinite-horizon discounted Markov game which is defined as $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{B}, P, r, \gamma\}$, with discrete state space \mathcal{S} , action spaces of two players \mathcal{A} and \mathcal{B} , transition probability P , reward function $r : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1]$ and discount factor $\gamma \in [0, 1)$. A policy $\mu : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ (resp. $\nu : \mathcal{S} \rightarrow \Delta(\mathcal{B})$) defines how the max player (resp. the min player) reacts to a given state s , where the probability of taking action $a \in \mathcal{A}$ (resp. $b \in \mathcal{B}$) is $\mu(a|s)$ (resp. $\nu(b|s)$). The transition probability kernel $P : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \Delta(\mathcal{S})$ defines the dynamics of the Markov game, where $P(s'|s, a, b)$ specifies the probability of transiting to state s' from state s when the players take actions a and b respectively. The value function and Q -function of a given policy pair (μ, ν) is defined in a similar way as in single-agent RL:

$$V^{\mu, \nu}(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, b_t) \mid s_0 = s \right],$$

$$Q^{\mu, \nu}(s, a, b) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, b_t) \mid s_0 = s, a_0 = a, b_0 = b \right].$$

The minimax game value on state s is defined by

$$V^*(s) = \max_{\mu} \min_{\nu} V^{\mu, \nu}(s) = \min_{\nu} \max_{\mu} V^{\mu, \nu}(s).$$

Similarly, the minimax Q -function $Q^*(s, a, b)$ is defined by

$$Q^*(s, a, b) = r(s, a, b) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a, b)} V^*(s'). \quad (8)$$

It is established by [Shapley \(1953\)](#) that there exists a pair of stationary policy (μ^*, ν^*) attaining the minimax value on all states ([Filar and Vrieze, 2012](#)), and is called the NE of the Markov game. We seek to obtain a pair of ϵ -optimal NE policy pair — denoted by ϵ -NE — $(\hat{\mu}^*, \hat{\nu}^*)$ that satisfies

$$V^{\mu, \hat{\nu}^*}(s) - \epsilon \leq V^{\hat{\mu}^*, \hat{\nu}^*}(s) \leq V^{\hat{\mu}^*, \mu}(s) + \epsilon$$

for any $\mu \in \Delta(\mathcal{A})^{\mathcal{S}}$, $\nu \in \Delta(\mathcal{B})^{\mathcal{S}}$ and $s \in \mathcal{S}$.

With the success of the NPG method in single-agent RL, it is attempting to apply it in the context of two-player zero-sum games, where each player executes the NPG updates independently by treating the other player as part of the environment. Notably, the NPG dynamics coincide with the




Figure 2: Cycles of MWU trajectories in a rock-paper-scissors game from a benign initialization close to the equilibrium.

well-studied MWU method, or Hedge, that stems from online optimization and game theory, whose *average-iterate* policy (i.e., $(\frac{1}{T} \sum_{t=1}^T \mu^{(t)}, \frac{1}{T} \sum_{t=1}^T \nu^{(t)})$) is shown to achieve a converging NE-gap at the rate of $\mathcal{O}(1/\sqrt{T})$ for two-player zero-sum matrix games. However, two technical challenges remain: last-iterate convergence and generalization to Markov games, which we shall discuss separately in the sequel.

3.2 Global last-iterate convergence

Average-iterate convergence guarantees fall short of determining if the learning trajectory converges towards NE (referred to as *last-iterate* convergence that applies to $(\mu^{(T)}, \nu^{(T)})$ alone) or enters recurrent cycles instead. In addition, for large-scale applications that involve the use of neural networks, average-iterate convergence is also unsatisfactory as averaging neural networks is intractable. [Mertikopoulos et al. \(2018\)](#) demonstrated that MWU, when adopted by both agents in a two-player zero-sum matrix game, suffers from the Poincaré recurrence phenomenon that forbids the method from converging. This holds for other methods in the family of “Follow the Regularized Leader” (FTRL), which necessitates algorithmic modifications to the original NPG/MWU updates.

Optimism. [Rakhlin and Sridharan \(2013\)](#) proposed the use of *optimistic updates*, which extrapolate the gradient information with that from the previous iteration. In the context of two-player zero-sum matrix games, Optimistic MWU (OMWU) updates at the t -th iteration can be written as

$$\begin{cases} \mu^{(t+1)}(a) & \propto \mu^{(t)}(a) \exp(\eta(2A\nu^{(t)} - A\nu^{(t-1)})) \\ \nu^{(t+1)}(b) & \propto \nu^{(t)}(b) \exp(-\eta(2A^\top \mu^{(t)} - A^\top \mu^{(t-1)})) \end{cases}$$

Compared with the original MWU method

$$\begin{cases} \mu^{(t+1)}(a) & \propto \mu^{(t)}(a) \exp(\eta A \nu^{(t)}) \\ \nu^{(t+1)}(b) & \propto \nu^{(t)}(b) \exp(-\eta A^\top \mu^{(t)}) \end{cases},$$

OMWU estimates the reward for the next iteration by appending the terms $A\nu^{(t)} - A\nu^{(t-1)}$ and $A^\top \mu^{(t)} - A^\top \mu^{(t-1)}$ to the current payoff vectors $A\nu^{(t)}$ and $A^\top \mu^{(t)}$. The updates can be equivalently formalized as Algorithm 1, where $\{\bar{\mu}^{(t)}, \bar{\nu}^{(t)}\}_{t=0}^\infty$ are auxiliary variables that can be viewed as some “predictive” sequence for facilitating the analysis. [Rakhlin and Sridharan \(2013\)](#) first proved that OMWU yields

Algorithm 1: OMWU for two-player zero-sum matrix games

- 1 **Input:** Learning rate η , (optional) regularization parameter τ .
 - 2 **Initialization:** Set $\mu^{(0)}, \nu^{(0)}, \bar{\mu}^{(0)}$ and $\bar{\nu}^{(0)}$ as uniform policies. Set $\tau = 0$ when not using regularization.
 - 3 **for** $t = 0, \dots, \infty$ **do**
 - 4 When $t \geq 1$, update $\bar{\mu}^{(t)}$ and $\bar{\nu}^{(t)}$ as

$$\begin{cases} \bar{\mu}^{(t)}(a) & \propto \bar{\mu}^{(t-1)}(a)^{1-\eta\tau} \exp(\eta A \nu^{(t)}) \\ \bar{\nu}^{(t)}(b) & \propto \bar{\nu}^{(t-1)}(b)^{1-\eta\tau} \exp(-\eta A^\top \mu^{(t)}). \end{cases}$$

Update $\mu^{(t+1)}$ and $\nu^{(t+1)}$ as

$$\begin{cases} \mu^{(t+1)}(a) & \propto \bar{\mu}^{(t)}(a)^{1-\eta\tau} \exp(\eta A \nu^{(t)}) \\ \nu^{(t+1)}(b) & \propto \bar{\nu}^{(t)}(b)^{1-\eta\tau} \exp(-\eta A^\top \mu^{(t)}) \end{cases}.$$
-

$\mathcal{O}(\log T)$ regret in two-player zero-sum matrix games, thus achieving a faster $\tilde{\mathcal{O}}(1/T)$ average-iterate convergence to NE. [Daskalakis et al. \(2021\)](#) demonstrated that OMWU achieves near-optimal $\mathcal{O}(\text{poly}(\log T))$ regret for multi-player general-sum games as well. [Daskalakis and Panageas \(2019\)](#) established asymptotic last-iterate convergence of OMWU assuming that the NE is unique. [Wei et al. \(2021b\)](#) demonstrated OMWU converges linearly to the NE under the same uniqueness assumption.

Theorem 6 (([Wei et al., 2021b](#), informal)). *For a two-player zero-sum matrix game with unique NE $\zeta^* = (\mu^*, \nu^*)$, the last iterate of OMWU $\zeta^{(T)} = (\mu^{(T)}, \nu^{(T)})$ converges to ζ^* linearly with constant learning rate $\eta \leq 1/8$.*

[Wei et al. \(2021b\)](#) also investigated optimistic gradient descent ascent (OGDA), another variant of optimistic update rules by focusing on projected gradient updates (2), and derived global linear convergence guarantees without placing

assumptions on NE uniqueness. A concrete iteration complexity, however, remains elusive as both convergence rates depend on unspecified problem-dependent parameters.

Regularization. Regularization has been proven instrumental in enabling faster convergence for single-agent RL. In the context of game theory, entropy regularization is closely related to quantal response equilibrium (QRE) ([McKelvey and Palfrey, 1995](#)), an extension to NE with bounded rationality. Formally speaking, when the two agents seek to maximize their own entropy-regularized payoffs

$$\max_{\mu \in \Delta(\mathcal{A})} \min_{\nu \in \Delta(\mathcal{B})} V_\tau^{\mu, \nu} := \mu^\top A \nu + \tau \mathcal{H}(\mu) - \tau \mathcal{H}(\nu),$$

the resulting equilibrium $\zeta_\tau^* = (\mu_\tau^*, \nu_\tau^*)$ is referred to as a QRE and satisfies

$$\begin{cases} \mu_\tau^*(a) & \propto \exp([A \nu_\tau^* / \tau]_a) \\ \nu_\tau^*(b) & \propto \exp([A^\top \mu_\tau^* / \tau]_b) \end{cases}, \quad \forall a \in \mathcal{A}, b \in \mathcal{B}.$$

An approximate ϵ -QRE is defined in a similar way to ϵ -NE, by replacing $V^{\mu, \nu}$ in (7) with its regularized counterpart $V_\tau^{\mu, \nu}$, and $\epsilon/2$ -QRE is guaranteed to be an ϵ -NE by setting $\tau = \tilde{\mathcal{O}}(\epsilon)$. [Cen et al. \(2021\)](#) proposed entropy-regularized OMWU (summarized in Algorithm 1) to combine the ideas of regularization and optimism, and proved the method converges to the QRE at a linear rate without assumption on the uniqueness of NE.

Theorem 7 (([Cen et al., 2021](#), informal)). *With constant learning rate $0 < \eta \leq \min\{1/(2\tau + 2), 1/4\}$, the last iterate $\zeta^{(T)} = (\mu^{(T)}, \nu^{(T)})$ generated by entropy-regularized OMWU converges to the QRE ζ_τ^* at a linear rate $1 - \eta\tau$.*

This gives an iteration complexity of $\mathcal{O}((1 + 1/\tau) \log 1/\epsilon)$ for finding an ϵ -QRE in a last-iterate manner, or $\tilde{\mathcal{O}}(1/\epsilon)$ for finding an ϵ -NE by setting $\tau = \tilde{\mathcal{O}}(\epsilon)$. One might wonder if using regularization alone can also ensure the last-iterate convergence, which amounts to studying the update rule

$$\begin{cases} \mu^{(t)}(a) & \propto \mu^{(t-1)}(a)^{1-\eta\tau} \exp(\eta A \nu^{(t)}) \\ \nu^{(t)}(b) & \propto \nu^{(t-1)}(b)^{1-\eta\tau} \exp(-\eta A^\top \mu^{(t)}) \end{cases}.$$

The answer turns out to be yes, as investigated recently in [Sokota et al. \(2023\)](#); [Pattathil et al. \(2023\)](#). The same contraction rate $1 - \eta\tau$ can be obtained albeit with a more restrictive choice of learning rate $\eta = \mathcal{O}(\tau)$, which translates to an iteration complexity of $\tilde{\mathcal{O}}(1/\epsilon^2)$ for finding an ϵ -NE, slower than the entropy-regularized OMWU by a factor of $\tilde{\mathcal{O}}(1/\epsilon)$.

3.3 Extension to Markov games

Given some prescribed initial state distribution ρ , one can focus on solving the following saddle-point optimization problem

$$\max_{\mu \in \Delta(\mathcal{A})^{|\mathcal{S}|}} \min_{\nu \in \Delta(\mathcal{B})^{|\mathcal{S}|}} V^{\mu, \nu}(\rho)$$

for two-player zero-sum Markov games. A key property for two-player zero-sum matrix games is that the value function $V^{\mu, \nu}$ is bilinear in the policy space, which unfortunately no longer holds in the Markov setting. Two strategies have been successful in establishing provable policy gradient methods for two-player zero-sum Markov games, by either leveraging tools from saddle-point optimization theory for nonconvex-nonconcave functions or exploiting recursive structures of the value function.

Conventional wisdom in nonconvex-nonconcave saddle-point optimization suggests adopting two-timescale rules, which enforces a much smaller learning rate on one of the players. [Daskalakis et al. \(2020\)](#) demonstrated that when the two players adopt projected gradient descent ascent (GDA) updates with two-timescale learning rates, the method achieves an average-iterate convergence to ϵ -NE within $\tilde{\mathcal{O}}(\text{poly}(\epsilon^{-1}, |\mathcal{S}|, |\mathcal{A}|, |\mathcal{B}|))$ iterations (omitting additional instance-dependent parameters). [Zeng et al. \(2022\)](#) showed that softmax policy GDA updates with entropy regularization yields a last-iterate convergence with an improved iteration complexity with regard to ϵ^{-1} .

However, two-timescale rules give asymmetric convergence guarantees, i.e., only the slow learner is guaranteed to find an approximate NE policy, while the fast learner approximates the best response to the slow learner throughout the learning process. A natural question arises: is it possible to design symmetric algorithms with an improved iteration complexity?

Smooth value updates. Instead of sticking to policy updates using vanilla gradient information $\nabla_{\mu} V^{\mu, \nu}(\rho)$ and $\nabla_{\nu} V^{\mu, \nu}(\rho)$, another line of works seek to divide the updates into two parts, where the policy updates are provided by two-player zero-sum matrix game algorithms with $Q^{(t)}(s) = [Q^{(t)}(s, a, b)]_{a \in \mathcal{A}, b \in \mathcal{B}}$, where

$$Q^{(t)}(s, a, b) = r(s, a, b) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a, b)} [V^{(t-1)}(s')], \quad (9a)$$

presuming the role of payoff matrices for all $s \in \mathcal{S}$, and the value updates are given by

$$V^{(t)}(s) = (1 - \alpha_t)V^{(t-1)}(s) + \alpha_t f^{(t)}(s). \quad (9b)$$

Here, $\alpha_t > 0$ is the learning rate for the value function and $f^{(t)}(s)$ is a one-step look-ahead value estimator for state s , typically defined as

$$f^{(t)}(s) = \mu^{(t)}(s)^{\top} Q^{(t)}(s) \nu^{(t)}(s), \quad (10)$$

or

$$f^{(t)}(s) = \mu^{(t)}(s)^{\top} Q^{(t)}(s) \nu^{(t)}(s) + \tau \mathcal{H}(\mu^{(t)}(s)) - \tau \mathcal{H}(\nu^{(t)}(s)) \quad (11)$$

when we incorporate entropy regularization. These methods are akin to the prevalent actor-critic type algorithms. The algorithm procedure is summarized in [Algorithm 2](#). Note that when setting $\alpha_t = 1$ and $f^{(t)}(s)$ to the one-step minimax game value as

$$f^{(t)}(s) = \max_{\mu(s) \in \Delta(\mathcal{A})} \min_{\nu(s) \in \Delta(\mathcal{B})} \mu(s)^{\top} Q^{(t)}(s) \nu(s),$$

we recover the classical value iteration for two-player zero-sum Markov games.

Algorithm 2: Actor-critic for two-player zero-sum Markov games

- 1 **Input:** Learning rate for Q -value function $\{\alpha_t\}_{t=0}^{\infty}$, learning rate for policies η , policy optimization method for two-player zero-sum matrix game `matrix_alg`, (optional) regularization parameter τ .
 - 2 **Initialization:** Set $Q^{(0)} = 0$ and $\mu^{(0)}, \nu^{(0)}$ as uniform policies.
 - 3 **for** $t = 0, 1, \dots$ **do**
 - 4 **for all** $s \in \mathcal{S}$ **do in parallel**
 - 5 Invoke `matrix_alg` with payoff matrix $Q^{(t)}(s)$ and learning rate η to update $\mu^{(t+1)}(s), \nu^{(t+1)}(s)$.
 - 6 Update $Q^{(t+1)}(s)$ and $V^{(t+1)}(s)$ according to (9) with learning rate α_{t+1} .
-

[Wei et al. \(2021a\)](#) first demonstrated that [Algorithm 2](#) with `matrix_alg` OGDA and decaying learning rate $\alpha_t = \frac{2/(1-\gamma)+1}{2/(1-\gamma)+t}$ yields both average-iterate and last-iterate convergences to NE. [Cen et al. \(2023\)](#) achieved an improved convergence rate by adopting entropy regularization and the OMWU method (e.g. [Algorithm 1](#)).

Theorem 8 (([Cen et al., 2023](#), Theorem 1)). *Algorithm 2 with `matrix_alg` entropy-regularized OMWU, entropy-regularized value updates (11) and constant learning rates $\eta = \mathcal{O}((1-\gamma)^3/|\mathcal{S}|)$, $\alpha_t = \eta\tau$ guarantees last-iterate convergence to the QRE at a linear rate $1 - \eta\tau$.*

The above theorem demonstrates an iteration complexity of $\tilde{\mathcal{O}}\left(\frac{|\mathcal{S}|}{(1-\gamma)^4\tau} \log 1/\epsilon\right)$ for finding an ϵ -QRE, or $\tilde{\mathcal{O}}\left(\frac{|\mathcal{S}|}{(1-\gamma)^5\epsilon}\right)$ for finding an ϵ -NE. It remains an open problem to achieve an iteration complexity with better dependency on $|\mathcal{S}|$ and $(1-\gamma)^{-1}$.

4 Global convergence of policy gradient methods in control

In this section, we briefly review policy gradient methods for control, focusing on a standard control problem called linear quadratic regulators (LQRs), based primarily on the excellent work of [Fazel et al. \(2018\)](#). We refer interested readers to [Hu et al. \(2023\)](#) for a recent comprehensive survey on the developments of policy optimization for general control problems including but not limited to \mathcal{H}_∞ control, risk-sensitive control, to name just a few.

4.1 Problem settings

Linear quadratic regulator (LQR). Consider a discrete-time linear dynamic system

$$x^{(t+1)} = Ax^{(t)} + Bu^{(t)},$$

where $x^{(t)} \in \mathbb{R}^d$ and $u^{(t)} \in \mathbb{R}^k$ are the state and the input at time t , $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times k}$ specify system transition matrices. The linear quadratic regulator (LQR) problem in the infinite horizon is defined as

$$\min_{u^{(t)}} \mathbb{E}_{x^{(0)} \sim \mathcal{D}} \left[\sum_{t=0}^{\infty} \left(x^{(t)\top} Q x^{(t)} + u^{(t)\top} R u^{(t)} \right) \right], \quad (12)$$

where \mathcal{D} determines the distribution of the initial state $x^{(0)}$, $Q \in \mathbb{R}^{d \times d}$ and $R \in \mathbb{R}^{k \times k}$ are positive definite matrices parametrizing the costs. Classical optimal control theory ([Anderson and Moore, 2007](#)) tells us that under certain stability conditions (e.g., controllability), it is ensured that the optimal cost is finite and can be achieved by a linear controller

$$u^{(t)} = -K^* x^{(t)},$$

where $K^* \in \mathbb{R}^{k \times d}$ is the optimal control gain matrix. From an optimization perspective, it is therefore natural to cast the LQR problem as optimizing over all linear controllers $u^{(t)} = -Kx^{(t)}$ with $K \in \mathbb{R}^{k \times d}$ to minimize the cost:

$$\min_{K \in \mathcal{K}} C(K) := \mathbb{E}_{x^{(0)} \sim \mathcal{D}} \left[x^{(0)\top} P_K x^{(0)} \right],$$

where

$$P_K = \sum_{t=0}^{\infty} ((A - BK)^\top)^t (Q + K^\top R K) (A - BK)^t.$$

The feasible set $\mathcal{K} = \{K : \|A - BK\|_2 < 1\}$ ensures P_K is well defined for all $K \in \mathcal{K}$.

4.2 Policy gradient method

The policy gradient method for the LQR problem is simply defined as

$$K^{(t+1)} = K^{(t)} - \eta \nabla_K C(K^{(t)}), \quad (13)$$

where $\eta > 0$ is the learning rate. In addition, the policy gradient $\nabla_K C(K)$ can be written as

$$\nabla_K C(K) = 2((R + B^\top P_K B)K - B^\top P_K A) \Sigma_K,$$

where $\Sigma_K = \mathbb{E}_{x^{(0)} \sim \mathcal{D}} \left[\sum_{t=0}^{\infty} x^{(t)} x^{(t)\top} \right]$ denotes the state correlation matrix.

Like the RL problem, the objective function $C(K)$ is non-convex in general, which makes it challenging to claim a global convergence guarantee. Fortunately, the LQR problem satisfies the following gradient dominance condition ([Fazel et al., 2018](#)).

Lemma 2 (Gradient dominance). *Suppose that $\lambda = \sigma_{\min}(\mathbb{E}_{x^{(0)} \sim \mathcal{D}} [x^{(0)} x^{(0)\top}]) > 0$. It holds that*

$$C(K) - C(K^*) \leq \frac{\|\Sigma_{K^*}\|_2}{\lambda^2 \sigma_{\min}(R)} \|\nabla_K C(K)\|_F^2.$$

The gradient dominance property provides hope for attaining global linear convergence of the policy gradient method to the optimal policy, yet to complete the puzzle, another desirable ingredient is the smoothness of the cost function $C(K)$. While this cannot be established in its full generality as $C(K)$ becomes infinity when K moves beyond \mathcal{K} , fortunately, smoothness can be established for any given sublevel set $\mathcal{K}_{\bar{\gamma}} = \{K \in \mathcal{K} : C(K) \leq \bar{\gamma}\}$, which suffices to establish the desired convergence results as follows.

Theorem 9 ([Fazel et al. \(2018\)](#)). *Assume that $C(K^{(0)})$ is finite. With an appropriate constant learning rate*

$$\eta = \text{poly} \left(\frac{\lambda \sigma_{\min}(Q)}{C(K^{(0)})}, \|A\|_2^{-1}, \|B\|_2^{-1}, \|R\|_2^{-1}, \sigma_{\min}(R) \right),$$

the policy gradient method (13) converges to the optimal policy at a linear rate:

$$C(K^{(t+1)}) - C(K^*) \leq \left(1 - \frac{\lambda^2 \sigma_{\min}(R) \eta}{\|\Sigma_{K^*}\|_2} \right) (C(K^{(t)}) - C(K^*)).$$

To find an ϵ -optimal control policy $K^{(T)}$ satisfying $C(K^{(T)}) - C(K^*) \leq \epsilon$, the above theorem ensures that the policy gradient method takes no more than

$$\frac{\|\Sigma_{K^*}\|_2}{\lambda^2 \sigma_{\min}(R) \eta} \log \frac{C(K^{(0)}) - C(K^*)}{\epsilon}$$

iterations.

4.3 Natural policy gradient method

To facilitate the development of NPG, we consider a linear policy with additive Gaussian noise, specified as

$$u^{(t)} \sim \pi(\cdot|x^{(t)}) = \mathcal{N}(-Kx^{(t)}, \sigma^2 I).$$

The NPG update rule (Kakade, 2002) then reads like

$$\text{vec}(K^{(t+1)}) = \text{vec}(K^{(t)}) - \eta(\mathcal{F}_{\mathcal{D}}^{K^{(t)}})^{\dagger} \text{vec}(\nabla_K C(K^{(t)})), \quad (14)$$

where $\text{vec}(K)$ flattens $K \in \mathbb{R}^{k \times d}$ into a vector in row-major order, and the Fisher information matrix $\mathcal{F}_{\mathcal{D}}^K \in \mathbb{R}^{kd \times kd}$ is given by

$$\begin{aligned} \mathcal{F}_{\mathcal{D}}^K &= \mathbb{E} \left[\sum_{t=0}^{\infty} \text{vec}(\nabla_K \log \pi(u^{(t)}|x^{(t)})) \text{vec}(\nabla_K \log \pi(u^{(t)}|x^{(t)}))^{\top} \right] \\ &= \sigma^{-2} \mathbb{E} \left[\sum_{t=0}^{\infty} \text{diag}(x^{(t)} x^{(t)\top}, \dots, x^{(t)} x^{(t)\top}) \right] \\ &= \sigma^{-2} \text{diag}(\Sigma_K, \dots, \Sigma_K). \end{aligned}$$

Merging the dummy variance σ^2 into the learning rate η , and reshaping back into the matrix form, the NPG update rule (14) can be equivalently rewritten as

$$K^{(t+1)} = K^{(t)} - \eta \nabla_K C(K^{(t)}) \Sigma_{K^{(t)}}^{-1},$$

which modifies the update direction using the state correlation matrix. This allows for an improved progress following a single update, as demonstrated by the following lemma.

Lemma 3 (Fazel et al. (2018)). *Assume that $C(K^{(t)})$ is finite and the learning rate satisfies $\eta \leq \frac{1}{\|R+B^{\top}P_{K^{(t)}}B\|_2}$, the NPG update satisfies*

$$C(K^{(t+1)}) - C(K^*) \leq \left(1 - \frac{\lambda \sigma_{\min}(R)\eta}{\|\Sigma_{K^*}\|_2}\right) (C(K^{(t)}) - C(K^*)).$$

The improvement is twofold: the convergence rate is improved by a factor of λ ; in addition, it allows for a larger learning rate that would not be possible under a smoothness-based analysis. By adopting $\eta = \frac{1}{\|R\|_2 + \|B\|_2^2 C(K^{(0)}) \lambda^{-1}}$, one can show that the learning rate requirement is met throughout the trajectory, which leads to an iteration complexity of

$$\frac{\|\Sigma_{K^*}\|_2}{\lambda \sigma_{\min}(R)\eta} \log \frac{C(K^{(0)}) - C(K^*)}{\epsilon}$$

for finding an ϵ -optimal control policy. Last but not least, it is possible to achieve even-faster convergence rate by assuming access to more complex oracle, i.e., the Gauss-Newton method. The update rule is given by

$$K^{(t+1)} = K^{(t)} - \eta(R + B^{\top}P_{K^{(t)}}B)^{-1} \nabla_K C(K^{(t)}) \Sigma_{K^{(t)}}^{-1},$$

which allows a constant learning rate as large as $\eta = 1$ and an iteration complexity of

$$\frac{\|\Sigma_{K^*}\|_2}{\lambda} \log \frac{C(K^{(0)}) - C(K^*)}{\epsilon}.$$

5 Conclusions

Policy gradient methods remain to be at the forefront of data-driven sequential decision making, due to its simplicity and flexibility in integrating with other advances in computation from adjacent fields such as high performance computing and deep learning. Due to space limits, our focus is constrained on the optimization aspects of policy gradient methods, assuming access to exact gradient or policy evaluations. In reality, these information need to be estimated by samples collected via various mechanisms and thus noisy, leading to deep interplays between statistics and optimization. We hope this exposition provides a teaser to invite more interest from the optimization community to work in this area.

Acknowledgement S. Cen and Y. Chi are supported in part by the grants ONR N00014-19-1-2404, NSF CCF-1901199, CCF-2106778 and CNS-2148212. S. Cen is also gratefully supported by Wei Shen and Xuehong Zhang Presidential Fellowship, Boeing Scholarship, and JP Morgan Chase PhD Fellowship.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *The Journal of Machine Learning Research*, 22(1):4431–4506.
- Anderson, B. D. and Moore, J. B. (2007). *Optimal control: linear quadratic methods*. Courier Corporation.
- Bellman, R. (1952). On the theory of dynamic programming. *Proceedings of the National Academy of Sciences of the United States of America*, 38(8):716.
- Bertsekas, D. P. (2017). *Dynamic programming and optimal control (4th edition)*. Athena Scientific.
- Bhandari, J. and Russo, D. (2019). Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*.
- Bhandari, J. and Russo, D. (2021). On the linear convergence of policy gradient methods for finite mdps. In *International Conference on Artificial Intelligence and Statistics*, pages 2386–2394. PMLR.
- Cen, S., Cheng, C., Chen, Y., Wei, Y., and Chi, Y. (2022). Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 70(4):2563–2578.

- Cen, S., Chi, Y., Du, S. S., and Xiao, L. (2023). Faster last-iterate convergence of policy optimization in zero-sum markov games. In *International Conference on Learning Representations*.
- Cen, S., Wei, Y., and Chi, Y. (2021). Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems*, 34:27952–27964.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Daskalakis, C., Fishelson, M., and Golowich, N. (2021). Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems*, 34:27604–27616.
- Daskalakis, C., Foster, D. J., and Golowich, N. (2020). Independent policy gradient methods for competitive reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 5527–5540.
- Daskalakis, C. and Panageas, I. (2019). Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *Innovations in Theoretical Computer Science*.
- Fazel, M., Ge, R., Kakade, S., and Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476.
- Filar, J. and Vrieze, K. (2012). *Competitive Markov decision processes*. Springer Science & Business Media.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. (2017). Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*, pages 1352–1361.
- Hu, B., Zhang, K., Li, N., Mesbahi, M., Fazel, M., and Başar, T. (2023). Toward a theoretical foundation of policy optimization for learning control policies. *Annual Review of Control, Robotics, and Autonomous Systems*, 6:123–158.
- Kakade, S. M. (2002). A natural policy gradient. In *Advances in neural information processing systems*, pages 1531–1538.
- Khodadadian, S., Jhunjhunwala, P. R., Varma, S. M., and Maguluri, S. T. (2021). On the linear convergence of natural policy gradient algorithm. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 3794–3799. IEEE.
- Lan, G. (2023). Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *Mathematical programming*, 198(1):1059–1106.
- Lan, G., Li, Y., and Zhao, T. (2023). Block policy mirror descent. *SIAM Journal on Optimization*, 33(3):2341–2378.
- Li, G., Wei, Y., Chi, Y., and Chen, Y. (2023). Softmax policy gradient methods can take exponential time to converge. *Mathematical Programming*, pages 1–96.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38.
- Mei, J., Xiao, C., Szepesvari, C., and Schuurmans, D. (2020). On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR.
- Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. (2018). Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2703–2717. SIAM.
- Neumann, J. V. (1928). Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320.
- Pattathil, S., Zhang, K., and Ozdaglar, A. (2023). Symmetric (optimistic) natural policy gradient for multi-agent learning with parameter convergence. In *International Conference on Artificial Intelligence and Statistics*, pages 5641–5685. PMLR.
- Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Rakhlín, S. and Sridharan, K. (2013). Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems*, 26.
- Shani, L., Efroni, Y., and Mannor, S. (2020). Adaptive trust region policy optimization: Global convergence and faster rates for regularized mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5668–5675.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100.
- Sokota, S., D’Orazio, R., Kolter, J. Z., Loizou, N., Lanctot, M., Mitliagkas, I., Brown, N., and Kroer, C. (2023). A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. In *International Conference on Learning Representations (ICLR)*.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2021a). Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive Markov games. In *Conference on learning theory*, pages 4259–4299. PMLR.

- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2021b). Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations (ICLR)*.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Xiao, L. (2022). On the convergence rates of policy gradient methods. *The Journal of Machine Learning Research*, 23(1):12887–12922.
- Zeng, S., Doan, T. T., and Romberg, J. (2022). Regularized gradient descent ascent for two-player zero-sum Markov games. *Advances in Neural Information Processing Systems*, 35.
- Zhan, W., Cen, S., Huang, B., Chen, Y., Lee, J. D., and Chi, Y. (2023). Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *SIAM Journal on Optimization*, 33(2):1061–1091.