

Independent Natural Policy Gradient Methods for Potential Games: Finite-time Global Convergence with Entropy Regularization

Shicong Cen, Fan Chen, Yuejie Chi

Abstract—A major challenge in multi-agent systems is that the system complexity grows dramatically with the number of agents as well as the size of their action spaces, which is typical in real world scenarios such as autonomous vehicles, robotic teams, network routing, etc. It is hence in imminent need to design decentralized or independent algorithms where the update of each agent is only based on their local observations without the need of introducing complex communication/coordination mechanisms. In this work, we study the finite-time convergence of independent entropy-regularized natural policy gradient (NPG) methods for potential games, where the difference in an agent’s utility function due to unilateral deviation matches exactly that of a common potential function. The proposed entropy-regularized NPG method enables each agent to deploy symmetric, decentralized, and multiplicative updates according to its own payoff. We show that the proposed method converges to the quantal response equilibrium (QRE)—the equilibrium to the entropy-regularized game—at a sublinear rate, which is independent of the size of the action space and grows at most sublinearly with the number of agents. Appealingly, the convergence rate further becomes *independent* with the number of agents for the important special case of identical-interest games, leading to the first method that converges at a *dimension-free* rate. Our approach can be used as a smoothing technique to find an approximate Nash equilibrium (NE) of the unregularized problem without assuming that stationary policies are isolated.

I. INTRODUCTION

Reinforcement learning (RL) has garnered a growing amount of interest in recent years, due to its excellent empirical performance in a wide variety of applications, such as Go [44], motor control [25], chip design [38], and so on. There is a growing interest to applying RL techniques such as Q-learning and policy gradient methods to the setting with multiple agents, i.e., multi-agent reinforcement learning (MARL) problems.

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

F. Chen is with the Department of Mathematics at Peking University; email: chern@pku.edu.cn.

The work of S. Cen and Y. Chi is supported in part by the grants ONR N00014-19-1-2404, ARO W911NF-18-1-0303, NSF CCF-1901199, CCF-2007911, CCF-2106778 and CNS-2148212. S. Cen is also gratefully supported by Wei Shen and Xuehong Zhang Presidential Fellowship, and Nicholas Minnici Dean’s Graduate Fellowship in Electrical and Computer Engineering at Carnegie Mellon University.

While it seems appealing to apply single-agent RL methods to each agent in a multi-agent system in a straightforward fashion, this approach neglects non-stationarity of the environment due to the presence of other agents, and thus lacks theoretical support in general. The complication has given rise to the paradigm of *centralized training with decentralized execution* (CTDE) [29], where the policies are first obtained through training with a centralized controller with access to all agents’ observations and then disseminated to each agent for execution. However, this approach falls short of adapting to changes in the environment without retraining and raises privacy concerns as well. It is hence of great interest to understand and design more versatile *independent learning* algorithms that only depend on the agents’ local observations, require minimal coordination between agents, and provably converge.

In this work, we focus on independent learning algorithms for potential games [40], an important class of games that admit a potential function to capture the differences in each agent’s utility function induced by unilateral deviations. In particular, the analysis established in this work is tailored to potential games in their most basic setting, i.e., static potential games, an important stepping stone to the more general Markov setting. Despite its simple formulation and decades-long research, however, the computational underpinnings of such problems are still far from mature, especially when it comes to finding the Nash equilibrium (NE) of potential games in a decentralized manner. While several recent works have made significant breakthroughs by achieving logarithmic regrets with independent learning dynamics [12], [2], these results only guarantee convergence to coarse correlated equilibrium or correlated equilibrium, which are arguably much weaker equilibrium concepts than NE and hence do not lead to an approximate NE solution.

A. Our contributions

We seek to find the quantal response equilibrium (QRE) [34], the prototypical extension of NE for games with bounded rationality [43], where each agent runs independent natural policy gradient (NPG) methods [22] involving symmetric, decentralized, and multiplicative updates according to its own payoff. This amounts to solving a potential game with entropy regularization, whose algorithmic role has been studied in the setting of single-agent RL [36], [7] as well as two-player zero-sum games [8], but yet to be explored in more general

settings. Our contributions are summarized below.

- **Finite-time global convergence of independent entropy-regularized NPG methods.** We show that independent entropy-regularized NPG methods provably converge to the QRE of a potential game, and it takes no more than

$$\mathcal{O}\left(\frac{\min\{\sqrt{N}, \Phi_{\max}\}\Phi_{\max}}{\tau^2\epsilon^2}\right)$$

iterations to find an ϵ -optimal QRE (to be defined precisely). Here, N stands for the number of agents, $\tau > 0$ the entropy regularization parameter, and $\Phi_{\max} > 0$ the maximum value of the potential function.

- **Finite-time global convergence to ϵ -NE without isolation assumption.** By setting the entropy regularization parameter τ sufficiently small, the result translates to finding an approximate NE with non-asymptotic convergence guarantees, thereby obviating the additional assumption in prior literature [16], [42], [52] that requires the set of stationary policies to be isolated. Specifically, it takes no more than

$$\tilde{\mathcal{O}}\left(\frac{\min\{\sqrt{N}, \Phi_{\max}\}\Phi_{\max}}{\epsilon^4}\right)$$

iterations to find an ϵ -NE for the unregularized potential game, where $\tilde{\mathcal{O}}$ hides logarithmic dependencies.

These rates give the first set of iteration complexities—to the best of our knowledge—that are independent of the size of the action spaces, up to logarithmic factors. In addition, the iteration complexities exhibit a sublinear dependency with the number of agents, outperforming existing NE-finding algorithms whose complexities depend at least linearly with the number of agents. Even more appealingly, when interpreting our convergence rates for the important special case of identical-interest games with bounded payoffs [39], they further become independent with the number of agents, leading to the *first* method that achieves a *dimension-free* convergence rate of $\tilde{\mathcal{O}}(1/\epsilon^4)$ to find an ϵ -NE.

B. Related works

We review some related works, focusing on the theoretical advances on policy gradient methods and independent learning in games.

a) Global convergence of policy gradient methods:

Only recently theoretical understandings on the global convergence of policy gradient (PG) methods have emerged, mostly in the single-agent setting, including but not limited to [15], [1], [36], [7], [5], [28], [26], [35], [46], [48]. In addition, several works developed finite-time guarantees of independent PG methods for zero-sum two-player Markov games [13], [47], [54], [8] in the competitive MARL setting. A recent line work has been

successful in extending PG methods with direct parameterization to Markov potential games [53], [24], [14], [31]. In addition, [52] studies the finite-time convergence rate of PG methods with softmax parameterization for Markov potential games. Given that NPG methods often have better finite-time convergence rates than vanilla PG methods in the single-agent setting, our work focuses on the understanding of NPG methods for potential games.

b) *Fast convergence of natural policy gradient methods with entropy regularization:* Entropy regularization as a de facto trick to promote exploration in RL [18] and has been shown to provably accelerate convergence of policy gradient methods for single-agent RL [36], [7], [8], [51], [23]. In particular, combining entropy regularization with NPG methods leads to fast linear convergence at a desirable dimension-free rate [7], [51], [23], which continues to hold in two-player zero-sum games [8]. Extending such results to potential games, however, is non-trivial, due to the non-uniqueness of NE even with regularization. [16], [42], [52] established asymptotic convergence of independent NPG methods for Markov potential games with an additional assumption that requires the set of stationary policies to be isolated. [19] demonstrated asymptotic convergence of NPG with diminishing step sizes for potential games in the bandit feedback setting. In addition, [52] proposed to use a log-barrier regularization along with NPG to sidestep the isolation assumption and achieved the same iteration complexity as that of PG methods with direct parameterization. In contrast, we consider NPG with entropy regularization, which achieves a convergence rate that has better dependencies with the size of the action spaces and the number of agents.

c) *Independent learning in general-sum games:* Considerable progress has been made towards understanding independent learning dynamics in general-sum games [12], [2] and general-sum Markov games (also known as stochastic games) [45], [21], [30] by establishing non-asymptotic convergence to correlated equilibrium and coarse correlated equilibrium. However, such successes do not directly extend to potential games where NE is of interest. Specialized analysis for potential games is thus needed as finding approximate NE in a two-player game can be PPAD-hard even with full information [11]. Notably, there have been attempts to establish asymptotic convergence with independent learning dynamics [32], [33], [49] for weakly acyclic games [50], which includes potential games as a special case.

C. Notation and paper organization

We use $\Delta(\mathcal{A})$ to denote the probability simplex over the set \mathcal{A} . For a vector $\mathbf{a} \in \mathcal{A}^N$, we use $a_i \in \mathcal{A}$ and $a_{-i} \in \mathcal{A}^{N-1}$ to denote the entry with index i and all the rest entries as a vector, respectively. The application of scalar functions such as \exp and \log to vectors are defined in an entry-wise fashion. Let $\mathbf{1}$ be the all-one vector, and $[N] = \{1, \dots, N\}$. Given two distributions π_1 and π'_1 over \mathcal{A} , the

Kullback-Leibler (KL) divergence from π'_1 to π_1 is defined by $\text{KL}(\pi_1 \| \pi'_1) = \sum_{a \in \mathcal{A}} \pi_1(a) (\log \pi_1(a) - \log \pi'_1(a))$. Note that KL divergence is additive for product distributions in the sense that $\text{KL}(\pi \| \pi') = \sum_{i \in [N]} \text{KL}(\pi_i \| \pi'_i)$ for $\pi = \pi_1 \times \dots \times \pi_N \in \Delta(\mathcal{A})^N$ and $\pi' = \pi'_1 \times \dots \times \pi'_N \in \Delta(\mathcal{A})^N$. We denote Jeffrey divergence [20] by $J(\pi, \pi') = \text{KL}(\pi \| \pi') + \text{KL}(\pi' \| \pi)$, which is the symmetric version of the KL divergence.

The rest of this paper is organized as follows. Section II presents the backgrounds of the potential game setup. Section III introduces independent NPG methods and presents the finite-time global convergence guarantees. Detailed proof and numerical results can be found in [6]. Finally, we conclude in Section IV.

II. POTENTIAL GAMES WITH ENTROPY REGULARIZATION

In this section, we introduce the basics of potential games, as well as the incorporation of entropy regularization into its formulation.

A. Potential games

A strategic game $\mathcal{G} = \{N, \mathcal{A}, \{u_i\}_{i \in [N]}\}$ consists of N agents each with an individual utility or payoff function $u_i : \mathcal{A}^N \rightarrow [0, 1]$, $i \in [N]$, where \mathcal{A} is, without loss of generality, a finite action space shared by all agents. The policy or mixed strategy of agent i is denoted by $\pi_i \in \Delta(\mathcal{A})$, which is a distribution over the action space \mathcal{A} . By an abuse of notation, let $u_i(\pi)$ denote agent i 's expected utility function under the joint policy $\pi = \pi_1 \times \dots \times \pi_N \in \Delta(\mathcal{A})^N$, i.e., $u_i(\pi) = \mathbb{E}_{a_i \sim \pi_i, \forall i \in [N]} [u_i(\mathbf{a})]$, where we denote the action profile (a_1, \dots, a_N) of all agents by $\mathbf{a} \in \mathcal{A}^N$. We shall often instead write $\mathbf{a} = (a_i, a_{-i})$ where $a_{-i} = \{a_j\}_{j \neq i}$ collects the actions of all agents but i ; similarly, we write $\pi = (\pi_i, \pi_{-i})$, where $\pi_{-i} = \{\pi_j\}_{j \neq i}$ collects the policies of all agents but i .

The game \mathcal{G} is said to be a potential game if there exists a potential function $\Phi : \mathcal{A}^N \rightarrow \mathbb{R}$ such that

$$u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i}) = \Phi(a_i, a_{-i}) - \Phi(a'_i, a_{-i})$$

for any $a_i, a'_i \in \mathcal{A}$, $a_{-i} \in \mathcal{A}^{N-1}$ and $i \in [N]$. We assume that

$$0 \leq \Phi(\mathbf{a}) \leq \Phi_{\max}, \quad \forall \mathbf{a} \in \mathcal{A}^N, \quad (1)$$

where Φ_{\max} upper bounds the potential function. An important special case of the potential game is when all the agents share the same utility function, known as the identical-interest game [39]. It is straightforward to see that for an identical-interest game, we can set $\Phi = u_i$ for all $i \in [N]$, and therefore $\Phi_{\max} = 1$ due to the fact that the individual payoff is bounded in $[0, 1]$.

By linearity of expectation, we have

$$u_i(\pi_i, \pi_{-i}) - u_i(\pi'_i, \pi_{-i}) = \Phi(\pi_i, \pi_{-i}) - \Phi(\pi'_i, \pi_{-i}),$$

where, again with slight abuse of notation, we denote

$$\Phi(\pi) = \mathbb{E}_{\mathbf{a} \sim \pi} [\Phi(\mathbf{a})] = \mathbb{E}_{a_i \sim \pi_i, \forall i \in [N]} [\Phi(\mathbf{a})],$$

for any $\pi_i, \pi'_i \in \Delta(\mathcal{A})$, $\pi_{-i} \in \Delta(\mathcal{A})^{N-1}$ and $i \in [N]$.

a) *Nash equilibrium*: We now introduce the important notion of *Nash equilibrium* in a potential game.

Definition 2.1 (Nash equilibrium): A joint policy π^* is called a *Nash equilibrium* (NE) when it holds that

$$u_i(\pi_i, \pi_{-i}) \geq u_i(\pi'_i, \pi_{-i}), \quad \forall \pi'_i \in \Delta(\mathcal{A}), \forall i \in [N].$$

In other words, every agent cannot improve its utility function by deviating from the current policy. It is known that there exists at least one NE in a strategic game with finite agents and actions [41]. It follows immediately that the policy or strategy profile maximizing Φ in a potential game is an NE.

b) *Marginalized utility*: Before continuing, let us introduce an important quantity called the marginalized utility $r_i^\pi : \mathcal{A} \rightarrow \mathbb{R}$:

$$r_i^\pi(a) = \mathbb{E}_{a_{-i} \sim \pi_{-i}} [u_i(a, a_{-i})], \quad (2)$$

which can be viewed as the “single-agent” payoff or reward function when the policies of other agents are fixed. It is immediate to see that the utility function u_i can be written as

$$u_i(\pi) = \mathbb{E}_{\mathbf{a} \sim \pi} [u_i(\mathbf{a})] = \mathbb{E}_{a \sim \pi_i} [r_i^\pi(a)] = \langle r_i^\pi, \pi_i \rangle.$$

Here and throughout this paper, we shall often abuse the notation to treat π , π_i and $r_i^{(t)}$ as vectors.

B. Entropy-regularized potential games

The *quantal response equilibrium* (QRE) is proposed by McKelvey and Palfrey [34] as a seminal extension to the Nash equilibrium, which enables players to combat randomness in payoffs. A QRE or logit equilibrium $\pi_\tau^* = \pi_{\tau,1}^* \times \dots \times \pi_{\tau,N}^*$ necessitates every agent to maximize its own utility function with entropy regularization [37], i.e.,

$$u_{i,\tau}(\pi_{\tau,i}^*, \pi_{\tau,-i}^*) \geq u_{i,\tau}(\pi'_i, \pi_{\tau,-i}^*), \quad \forall \pi'_i \in \Delta(\mathcal{A}),$$

where the entropy-regularized individual utility function is given by

$$u_{i,\tau}(\pi) = u_i(\pi) + \tau \mathcal{H}(\pi_i).$$

Here, $\pi = \pi_1 \times \dots \times \pi_N$, $\tau > 0$ is the regularization parameter, and $\mathcal{H}(\pi_i) = -\sum_{a \in \mathcal{A}} \pi_i(a|s) \log \pi_i(a|s)$ is the Shannon entropy of the policy π employed by agent i . By introducing the regularized potential function

$$\Phi_\tau(\pi) = \Phi(\pi) + \tau \mathcal{H}(\pi) := \Phi(\pi) + \tau \sum_{i \in [N]} \mathcal{H}(\pi_i),$$

it is easy to verify

$$u_{i,\tau}(\pi_i, \pi_{-i}) - u_{i,\tau}(\pi'_i, \pi_{-i}) = \Phi_\tau(\pi_i, \pi_{-i}) - \Phi_\tau(\pi'_i, \pi_{-i}).$$

for any $\pi_i, \pi'_i \in \Delta(\mathcal{A})$, $\pi_{-i} \in \Delta(\mathcal{A})^{N-1}$ and $i \in [N]$, as long as the unregularized game is a potential game.

a) *Fixed-point characterization of QRE*: An equivalent interpretation of QRE is to let each agent assign the probability mass in its policy according to every action's utility in a bounded rationality fashion [43]:

$$\pi_{\tau,i}^*(a) \propto \exp\left(r_i^{\pi^*}(\tau)/\tau\right), \quad \forall i \in [N], \quad (3)$$

where $r_i^{\pi^*}$ is the marginalized utility of π^* defined in (2). Note that the above relation defines a fixed-point equation of π^* .

III. FINITE-TIME GLOBAL CONVERGENCE OF INDEPENDENT NATURAL POLICY GRADIENT METHODS

A popular approach in the game theory literature to find an NE of a potential game is for each agent to switch to the best or better response policy, one at a time, and is generally referred to as *best-response dynamics*. This approach converges to an NE in finite iterations [40] and underlies the algorithm design of a considerable number of works on, e.g., cut games [10], congestion games [9], weakly acyclic games [49], and, more recently, their extensions in the Markovian setting [45], [4]. It is noted, however, that this approach isolates itself from the independent learning paradigm as the update sequence needs to be scheduled in a centralized manner that is not often possible. Therefore, it is greatly desirable to design independent update rules, where each agent updates simultaneously without observing the payoffs of other agents, that achieves faster convergence. In this section, we answer this call by developing the independent natural policy gradient method to solve (entropy-regularized) potential games with finite-time global convergence guarantees.

A. Independent natural policy gradient method

In policy optimization, it is common practice to parameterize the policy class in a way that obviates the need for tackling probability simplex constraint. We consider the standard softmax parameterization, where every agent i generates its own policy π_{θ_i} parameterized with $\theta_i \in \mathbb{R}^{|\mathcal{A}|}$ through the softmax transform:

$$\pi_{\theta_i}(a) = \frac{\exp(\theta_i(a))}{\sum_{a \in \mathcal{A}} \exp(\theta_i(a))}.$$

Every agent i evaluates and updates its policy independently using the *natural policy gradient* (NPG) method [22]:

$$\theta_i \leftarrow \theta_i + \eta(\mathcal{F}^{\theta_i})^\dagger \nabla_{\theta_i} u_{i,\tau}(\pi), \quad (4)$$

where $(\mathcal{F}^{\theta_i})^\dagger$ denotes the Moore-Penrose pseudo-inverse of the Fisher information matrix \mathcal{F}^{θ_i} , which is defined as

$$\mathcal{F}^{\theta_i} = \mathbb{E}_{a \sim \pi_{\theta_i}(\cdot)} \left[(\nabla_{\theta_i} \log \pi_{\theta_i}(a)) (\nabla_{\theta_i} \log \pi_{\theta_i}(a))^\top \right],$$

and $\eta > 0$ is the learning rate. Moreover, the gradient $\nabla_{\theta_i} u_{i,\tau}(\pi)$ can be expressed as

$$\nabla_{\theta_i} u_{i,\tau}(\pi) = r_i^\pi - \tau \log \pi_i - \tau \mathbf{1}.$$

It turns out that with some algebra, the NPG update rule (4) can be equivalently rewritten with respect to the policies in use [7]:

$$\pi_i^{(t+1)}(a) \propto \pi_i^{(t)}(a)^{1-\eta\tau} \exp(\eta r_i^{(t)}(a)), \quad (5)$$

where $\pi_i^{(t)}$ denotes agent i 's policy in the t -th iteration, and $r_i^{(t)} := r_i^{\pi^{(t)}}$ denotes the marginalized utility of $\pi^{(t)}$ (cf. (2)). The complete procedure is summarized in Algorithm 1.

Algorithm 1: Independent NPG for Entropy-regularized Potential Games

```

1 Input: Regularization parameter  $\tau > 0$ , step size
   for policy update  $\eta > 0$ .
2 Initialization: Set  $\pi_i^{(0)}$  as uniform policy for all
    $i \in [N]$ .
3 for  $t = 0, 1, \dots$  do
4   for all agent  $i \in [N]$  do in parallel
5     Observe agent  $i$ 's marginalized utility  $r_i^{(t)}$ .
6     Perform policy update
7     
$$\pi_i^{(t+1)}(a) \propto \pi_i^{(t)}(a)^{1-\eta\tau} \exp(\eta r_i^{(t)}(a)).$$


```

To better understand the update rule (5) as well as prepare for follow-up analysis, we introduce $\pi_i^{\star(t)}$ to denote agent i 's best-response policy in the t -th iteration, which is the policy that obeys

$$u_{i,\tau}(\pi_i^{\star(t)}, \pi_{-i}^{(t)}) = \max_{\pi'_i} u_{i,\tau}(\pi'_i, \pi_{-i}^{(t)}). \quad (6)$$

It is easily seen that

$$\pi_i^{\star(t)}(a) \propto \exp(r_i^{(t)}(a)/\tau). \quad (7)$$

Therefore, the updated policy in (5) can be regarded as a multiplicative combination of the current policy $\pi_i^{(t)}$ and the best-response policy $\pi_i^{\star(t)}$, where the weight is controlled by the learning rate η . Note that the unregularized counterpart of the method is equivalent to Multiplicative Weights Update method (MWU) [27], [3] or Hedge [17].

B. Finite-time global convergence

We are now ready to present our main theorem concerning the finite-time global convergence of independent NPG for solving entropy-regularized potential games. We introduce

$$\text{NE-gap}(\pi) = \max_{i \in [N], \pi'_i \in \Delta(\mathcal{A})} [u_i(\pi'_i, \pi_{-i}) - u_i(\pi_i, \pi_{-i})]$$

and

$$\text{QRE-gap}_\tau(\pi) = \max_{i \in [N], \pi'_i \in \Delta(\mathcal{A})} [u_{i,\tau}(\pi'_i, \pi_{-i}) - u_{i,\tau}(\pi_i, \pi_{-i})]$$

to characterize how close the joint policy π is to an equilibrium. A joint policy π is said to be an ϵ -QRE

(resp. ϵ -NE) when $\text{QRE-gap}(\pi) \leq \epsilon$ (resp. $\text{NE-gap}(\pi) \leq \epsilon$). For notational simplicity, we denote

$$\Phi_\tau^{(t)} := \Phi_\tau(\pi^{(t)}), \quad \text{QRE-gap}_\tau^{(t)} := \text{QRE-gap}_\tau(\pi^{(t)}),$$

and $\text{NE-gap}^{(t)} := \text{NE-gap}(\pi^{(t)})$.

Our main theorem is as follows, whose proof is deferred to [6].

Theorem 3.1: Suppose that the learning rate η satisfies $\eta \leq \frac{1}{2(\min\{\sqrt{N}, 2\Phi_{\max}\} + \tau)}$, then for independent NPG updates (5), it holds that

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T \text{QRE-gap}_\tau^{(t)} \\ & \leq \frac{2}{\eta\tau T} \left(\tau \left\| \log \pi^{(0)} - \log \pi^{*(0)} \right\|_\infty + \sqrt{2\eta T(\Phi_\tau^{(T)} - \Phi_\tau^{(0)})} \right). \end{aligned}$$

Theorem 3.1 suggests that the average iterate of independent NPG converges to an ϵ -QRE at a sublinear rate when we initialize it via uniform policies, as indicated in the following corollary. The proof can be found in [6].

Corollary 3.2: Assume the independent NPG method is initialized with uniform policies at all agents. Setting the learning rate $\eta = 1/(2(\min\{\sqrt{N}, 2\Phi_{\max}\} + \tau))$ and $\tau = \mathcal{O}(1)$, then independent NPG updates ensure that $\frac{1}{T} \sum_{t=1}^T \text{QRE-gap}_\tau^{(t)} \leq \epsilon$ with at most

$$T = \mathcal{O} \left(\frac{\min\{\sqrt{N}, \Phi_{\max}\} \Phi_{\max}}{\tau^2 \epsilon^2} \right)$$

iterations.

a) *Finding approximate NEs:* It is possible to leverage the entropy-regularized potential game to find an approximate NE by setting the regularization parameter sufficiently small. Note that

$$\begin{aligned} \text{NE-gap}(\pi) &= \max_{i \in [N], \pi'_i \in \Delta(\mathcal{A})} [u_i(\pi'_i, \pi_{-i}) - u_i(\pi_i, \pi_{-i})] \\ &\leq \max_{i \in [N], \pi'_i \in \Delta(\mathcal{A})} [u_{i,\tau}(\pi'_i, \pi_{-i}) - u_{i,\tau}(\pi_i, \pi_{-i})] \\ &\quad + \max_{i \in [N], \pi'_i \in \Delta(\mathcal{A})} [-\tau \mathcal{H}(\pi'_i) + \tau \mathcal{H}(\pi_i)] \\ &\leq \text{QRE-gap}_\tau(\pi) + \tau \log |\mathcal{A}|. \end{aligned}$$

Therefore, by setting the entropy regularization at $\tau = \frac{\epsilon}{2 \log |\mathcal{A}|}$, with at most

$$T = \tilde{\mathcal{O}} \left(\frac{\min\{\sqrt{N}, \Phi_{\max}\} \Phi_{\max}}{\epsilon^4} \right)$$

iterations, we can ensure $\frac{1}{T} \sum_{t=1}^T \text{NE-gap}^{(t)} \leq \epsilon$.

b) *Comparisons with prior art:* Importantly, our iteration complexities do not depend on the size of the action space (up to logarithmic factors), which is in sharp contrast to existing analyses of potential games using other policy gradient approaches, such as direct PG [53], [24], [14], [31] and NPG with log-barrier regularization [52], where the iteration complexity scales as $\tilde{\mathcal{O}}(N|\mathcal{A}|\Phi_{\max}/\epsilon^2)$ to find an ϵ -approximate NE. In comparison, while our rate $\tilde{\mathcal{O}}(\min\{\sqrt{N}, \Phi_{\max}\} \Phi_{\max}/\epsilon^4)$ is

worse in terms of ϵ , it is almost independent of the size $|\mathcal{A}|$ of the action space, as well as exhibits only a sublinear dependency with the number of agents N , thus can be beneficial for problems with large action spaces and a large number of agents. Furthermore, for the special case of identical-interest games [39] where $\Phi_{\max} = 1$, the convergence rate of our method simplifies to

$$\tilde{\mathcal{O}} \left(\frac{1}{\epsilon^4} \right),$$

which leads to the first method that achieves a dimension-free iteration complexity (up to a logarithmic factor) for finding an ϵ -NE without imposing any isolation assumptions.

IV. CONCLUSIONS AND DISCUSSIONS

This paper studies independent NPG methods for entropy-regularized potential games and develops a sublinear rate of convergence to quantum response equilibrium, which is independent of the size of the action spaces up to logarithmic factors and grows only sublinearly with respect to the number of agents. In addition, the method achieves the first *dimension-free* convergence rate for the important special case of identical-interest games, where the rate is independent of both the size of the action space and the number of agents. The approach can also be used as a smoothing technique to find Nash equilibria by setting the regularization parameter sufficiently small, without imposing the isolation assumption as often required in prior works. This work leaves open a number of interesting questions:

- Can we tighten the convergence rate in terms of the dependencies on ϵ ?
- Can we extend the analysis to establish finite-time global convergence for Markov potential games?

We leave the answers to future work.

REFERENCES

- [1] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. Optimality and approximation with policy gradient methods in Markov decision processes. In *Conference on Learning Theory*, pages 64–66. PMLR, 2020.
- [2] I. Anagnostides, C. Daskalakis, G. Farina, M. Fishelson, N. Golowich, and T. Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. *arXiv preprint arXiv:2111.06008*, 2021.
- [3] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [4] G. Arslan and S. Yüksel. Decentralized q-learning for stochastic teams and games. *IEEE Transactions on Automatic Control*, 62(4):1545–1558, 2016.
- [5] J. Bhandari and D. Russo. Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*, 2019.
- [6] S. Cen, F. Chen, and Y. Chi. Independent natural policy gradient methods for potential games: Finite-time global convergence with entropy regularization. *arXiv preprint arXiv:2204.05466*, 2022.
- [7] S. Cen, C. Cheng, Y. Chen, Y. Wei, and Y. Chi. Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 2021.

[8] S. Cen, Y. Wei, and Y. Chi. Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems*, 34, 2021.

[9] S. Chien and A. Sinclair. Convergence to approximate nash equilibria in congestion games. *Games and Economic Behavior*, 71(2):315–327, 2011.

[10] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 349–360. Springer, 2006.

[11] C. Daskalakis. On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms (TALG)*, 9(3):1–35, 2013.

[12] C. Daskalakis, M. Fishelson, and N. Golowich. Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems*, 34, 2021.

[13] C. Daskalakis, D. J. Foster, and N. Golowich. Independent policy gradient methods for competitive reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 5527–5540, 2020.

[14] D. Ding, C.-Y. Wei, K. Zhang, and M. R. Jovanović. Independent policy gradient for large-scale markov potential games: Sharper rates, function approximation, and game-agnostic convergence. *arXiv preprint arXiv:2202.04129*, 2022.

[15] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476, 2018.

[16] R. Fox, S. McAleer, W. Overman, and I. Panageas. Independent natural policy gradient always converges in Markov potential games. *arXiv preprint arXiv:2110.10614*, 2021.

[17] Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1–2):79–103, 1999.

[18] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

[19] A. Heliou, J. Cohen, and P. Mertikopoulos. Learning with bandit feedback in potential games. *Advances in Neural Information Processing Systems*, 30, 2017.

[20] H. Jeffreys. *The theory of probability*. OUP Oxford, 1998.

[21] C. Jin, Q. Liu, Y. Wang, and T. Yu. V-learning—a simple, efficient, decentralized algorithm for multiagent RL. *arXiv preprint arXiv:2110.14555*, 2021.

[22] S. M. Kakade. A natural policy gradient. *Advances in neural information processing systems*, 14, 2001.

[23] G. Lan. Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *arXiv preprint arXiv:2102.00135*, 2021.

[24] S. Leonardos, W. Overman, I. Panageas, and G. Piliouras. Global convergence of multi-agent policy gradient in Markov potential games. *arXiv preprint arXiv:2106.01969*, 2021.

[25] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.

[26] G. Li, Y. Wei, Y. Chi, Y. Gu, and Y. Chen. Softmax policy gradient methods can take exponential time to converge. *arXiv preprint arXiv:2102.11270*, 2021.

[27] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.

[28] Y. Liu, K. Zhang, T. Basar, and W. Yin. An improved analysis of (variance-reduced) policy gradient and natural policy gradient methods. *Advances in Neural Information Processing Systems*, 33, 2020.

[29] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.

[30] W. Mao and T. Başar. Provably efficient reinforcement learning in decentralized general-sum markov games. *Dynamic Games and Applications*, pages 1–22, 2022.

[31] W. Mao, L. Yang, K. Zhang, and T. Basar. On improving model-free algorithms for decentralized multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 15007–15049. PMLR, 2022.

[32] J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, 2007.

[33] J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. Payoff-based dynamics for multiplayer weakly acyclic games. *SIAM Journal on Control and Optimization*, 48(1):373–396, 2009.

[34] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

[35] J. Mei, C. Xiao, B. Dai, L. Li, C. Szepesvári, and D. Schuurmans. Escaping the gravitational pull of softmax. *Advances in Neural Information Processing Systems*, 33, 2020.

[36] J. Mei, C. Xiao, C. Szepesvari, and D. Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR, 2020.

[37] P. Mertikopoulos and W. H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016.

[38] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, A. Nazi, et al. A graph placement methodology for fast chip design. *Nature*, 594(7862):207–212, 2021.

[39] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of economic theory*, 68(1):258–265, 1996.

[40] D. Monderer and L. S. Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.

[41] J. Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

[42] G. Palaiopanos, I. Panageas, and G. Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. *Advances in Neural Information Processing Systems*, 30, 2017.

[43] R. Selten. Evolution, learning and economic behavior. 1989.

[44] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7877):484–489, 2016.

[45] Z. Song, S. Mei, and Y. Bai. When can we learn general-sum Markov games with a large number of players sample-efficiently? *arXiv preprint arXiv:2110.04184*, 2021.

[46] L. Wang, Q. Cai, Z. Yang, and Z. Wang. Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*, 2019.

[47] C.-Y. Wei, C.-W. Lee, M. Zhang, and H. Luo. Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive Markov games. *arXiv preprint arXiv:2102.04540*, 2021.

[48] L. Xiao. On the convergence rates of policy gradient methods. *arXiv preprint arXiv:2201.07443*, 2022.

[49] H. P. Young. *Strategic learning and its limits*. OUP Oxford, 2004.

[50] H. P. Young. Individual strategy and social structure. In *Individual Strategy and Social Structure*. Princeton University Press, 2020.

[51] W. Zhan, S. Cen, B. Huang, Y. Chen, J. D. Lee, and Y. Chi. Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *arXiv preprint arXiv:2105.11066*, 2021.

[52] R. Zhang, J. Mei, B. Dai, D. Schuurmans, and N. Li. On the effect of log-barrier regularization in decentralized softmax gradient play in multiagent systems. *arXiv preprint arXiv:2202.00872*, 2022.

[53] R. Zhang, Z. Ren, and N. Li. Gradient play in stochastic games: stationary points, convergence, and sample complexity. *arXiv preprint arXiv:2106.00198*, 2021.

[54] Y. Zhao, Y. Tian, J. D. Lee, and S. S. Du. Provably efficient policy gradient methods for two-player zero-sum Markov games. *arXiv preprint arXiv:2102.08903*, 2021.