

# $\alpha^\alpha$ -RANK: SCALABLE MULTI-AGENT EVALUATION THROUGH EVOLUTION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Although challenging, strategy profile evaluation in large connected learner networks is crucial for enabling the next wave of machine learning applications. Recently,  $\alpha$ -Rank, an evolutionary algorithm, has been proposed as a solution for ranking joint policy profiles in multi-agent systems.  $\alpha$ -Rank claimed scalability through a polynomial time implementation with respect to the total number of pure strategy profiles. In this paper, we formally prove that such a claim is not grounded. In fact, we show that  $\alpha$ -Rank exhibits an exponential complexity in number of agents, hindering its application beyond a small finite number of joint profiles. Realizing such a limitation, we contribute by proposing a scalable evaluation protocol that we title  $\alpha^\alpha$ -Rank. Our method combines evolutionary dynamics with stochastic optimization and double oracles for *truly* scalable ranking with linear (in number of agents) time and memory complexities. Our contributions allow us, for the first time, to conduct large-scale evaluation experiments of multi-agent systems, where we show successful results on large joint strategy profiles with sizes in the order of  $\mathcal{O}(2^{25})$  (i.e.,  $\approx 33$  million strategies) – a setting not evaluable using current techniques.

## 1 INTRODUCTION

Scalable policy evaluation and learning have been long-standing challenges in multi-agent reinforcement learning (MAREL) with two difficulties obstructing progress. First, joint-strategy spaces exponentially explode when a large number of strategic decision-makers is considered, and second, the underlying game dynamics may exhibit cyclic behavior (e.g. the game of Rock-Paper-Scissor) rendering an appropriate evaluation criteria non-trivial.

Focusing on the second challenge, much work in multi-agent systems followed a game-theoretic treatment proposing fixed-points, e.g., Nash (Nash et al., 1950) equilibrium, as potentially valid evaluation metrics. Though appealing, such measures are normative only when prescribing behaviors of perfectly rational agents – an assumption rarely met in reality Grau-Moya et al. (2018); Wen et al. (2019). In fact, many game dynamics have been proven not converge to any fixed-point equilibria (Hart & Mas-Colell, 2003; Viossat, 2007), but rather to limit cycles (Palaiopoulos et al., 2017; Bowling & Veloso, 2001). Apart from these aforementioned inconsistencies, solving for a Nash equilibrium even for “simple” settings, e.g. two-player games is known to be PPAD-complete (Chen & Deng, 2005) – a demanding complexity class when it comes to computational requirements.

To address some of the above limitations, Omidshafiei et al. (2019) recently proposed  $\alpha$ -Rank as a graph-based game-theoretic solution to multi-agent evaluation.  $\alpha$ -Rank adopts Markov Conley Chains to highlight the presence of cycles in game dynamics, and attempts to compute stationary distributions as a mean for strategy profile ranking. Though successful in small-scale applications,  $\alpha$ -Rank severely suffers in scalability *contrary to* polynomial time claims made in Omidshafiei et al. (2019). In fact, we show that  $\alpha$ -Rank exhibits exponential time and memory complexities shedding light on the small-scale empirical study conducted in Omidshafiei et al. (2019), whereby the largest reported game included only four agents with four available strategies each.

In this work, we put forward  $\alpha^\alpha$ -Rank as a scalable alternative for multi-agent evaluation with *linear time and memory demands*. Our method combines numerical optimization with evolutionary game theory for a scalable solver capable of handling large joint spaces with millions of strategy profiles. To handle even larger profiles, e.g., tens to hundreds of millions, we further introduce an oracle




Figure 1: Example of population based evaluation on  $N = 3$  learners each with 3 strategies and 5 copies. a) Each population obtains a fitness value  $\mathcal{P}_i$  depending on the strategies chosen, b) mutation strategy (red star), and c) population either selecting original strategy, or adopting the novel strategy.

(McMahan et al., 2003) mechanism transforming joint evaluation into a sequence of incremental sub-games with varying sizes. Given our algorithmic advancements, we justify our claims in a large-scale empirical study involving systems with  $\mathcal{O}(2^{25})$  possible strategy profiles. We first demonstrate the computation advantages of  $\alpha$ -Rank on varying size stochastic matrices against other implementations in Numpy, PyTorch, and OpenSpiel (Lanctot et al., 2019). With these successes, we then consider experiments unsolvable by current techniques. Precisely, we evaluate multi-agent systems in self-driving and Ising model scenarios each exhibiting a prohibitively-large strategy space (i.e., order of thousands for the former, and tens of millions for the latter). Here, we again show that  $\alpha$ -Rank is capable of recovering correct strategy ranking in such complex domains.

## 2 $\alpha$ -RANK & ITS LIMITATIONS

In  $\alpha$ -Rank, strategy profiles of  $N$  agents are evaluated through an evolutionary process of mutation and selection. Initially, agent populations are constructed by creating multiple copies of each learner  $i \in \{1, \dots, N\}$  assuming that all agents (in one population) execute the same unified policy. With this,  $\alpha$ -Rank then simulates a multi-agent game played by randomly sampled learners from each population. Upon game termination, each participating agent receives a payoff to be used in policy mutation and selection after its return to the population. Here, the agent is faced with a probabilistic choice between switching to the mutation policy, continuing to follow its current policy, or randomly selecting a novel policy (other than the previous two) from the pool. This process repeats with the goal of determining an evolutionary strong profile that spreads across the population of agents. Each of the above three phases is demonstrated in Fig. 1 on a simple example of three agents – depicted by different symbols – each equipped with three strategies – depicted by the colors.

**Mathematical Formalisation, Notation, and Definitions:** We next formalize the process posed by  $\alpha$ -Rank, which will lead to its limitations, and also pave the way for our own proposed solution. We consider  $N$  agents with each agent  $i$  having access to a set of strategies of size  $s_i$ . At round  $k$  of the evaluation process, we denote the strategy profile for agent  $i$  by  $\mathcal{S}_i^{[k]} = \{\pi_{i,1}^{[k]}, \dots, \pi_{i,s_i}^{[k]}\}$ , with  $\pi_{i,j}^{[k]} : \mathcal{X} \times \mathcal{A}_i \rightarrow [0, 1]$  representing the  $j^{th}$  allowed policy of the learner.  $\mathcal{X}$  represents the set of states and  $\mathcal{A}_i$  is the set of actions for agent  $i$ . With this, we define a joint strategy profile for all participating agents as policies belonging to the joint strategy pool,  $\mathcal{S}_{\text{joint}} = \mathcal{S}_1^{[k]} \times \mathcal{S}_2^{[k]} \dots \times \mathcal{S}_N^{[k]}$ :  $\pi_{\text{joint}}^{[k]} = \{\pi_{1,j_1}^{[k]}, \dots, \pi_{N,j_N}^{[k]}\}$  with  $\pi_{i,j_i} \in \mathcal{S}_i^{[k]}$  and  $j_i \in \{1, \dots, s_i\}$ .

To evaluate performance, we assume each agent is additionally equipped with a payoff (reward) function  $\mathcal{P}_i^{[k]} : \mathcal{S}_{\text{joint}} \rightarrow \mathbb{R}_+$ . Crucially, the domain of  $\mathcal{P}_i^{[k]}$  is the pool of *joint strategies* so to accommodate the effect of other learners on the  $i^{th}$  player performance further complicating the evaluation process. Finally, given a joint profile  $\pi_{\text{joint}}^{[k]}$ , we define the corresponding joint payoff to be the collection of all individual payoff functions, i.e.,  $\mathcal{P}_{\text{joint}}^{[k]} = \{\mathcal{P}_1^{[k]}(\pi_{\text{joint}}^{[k]}), \dots, \mathcal{P}_N^{[k]}(\pi_{\text{joint}}^{[k]})\}$ .

After attaining rewards from the environment, each agent returns to its population and faces a choice between switching to a mutation policy, exploring a novel policy, or sticking to the current one Such

a choice is probabilistic and defined proportional to rewards. Precisely, agent  $i$  adopts

$$\begin{aligned} \mathbb{P}(\pi_{i,a}^{[k]} \rightarrow \pi_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}) &= \frac{e^{\alpha \mathcal{P}_i^{[k]}(\pi_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})}}{e^{\alpha \mathcal{P}_i^{[k]}(\pi_{i,a}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})} + e^{\alpha \mathcal{P}_i^{[k]}(\pi_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})}} - \frac{\mu}{2} \text{ for } (\pi_{i,a}^{[k]}, \pi_{i,b}^{[k]}) \in \mathcal{S}_i^{[k]} \times \mathcal{S}_i^{[k]}, \\ \mathbb{P}(\pi_{i,a}^{[k]} \rightarrow \pi_{i,c}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}) &= \frac{\mu}{K_i - 2}, \quad \forall \pi_{i,c}^{[k]} \in \mathcal{S}_i^{[k]} \setminus \{\pi_{i,a}^{[k]}, \pi_{i,b}^{[k]}\}, \end{aligned}$$

with  $\mu \in \mathbb{R}_+$  denoting an exploration parameter<sup>1</sup>,  $\boldsymbol{\pi}_{-i}^{[k]} = \boldsymbol{\pi}^{[k]} \setminus \pi_i^{[k]}$  representing policies followed by other agents at round  $k$ , and  $\alpha \in \mathbb{R}_+$  an intensity ranking parameter. As noted in Omidshafiei et al. (2019), one can relate the above switching process to a random walk on a Markov chain with states defined as elements in  $\mathcal{S}_{\text{joint}}^{[k]}$  and transition probabilities through payoff functions. In particular,

each entry in the transition probability matrix  $\mathbf{T} \in \mathbb{R}^{|\mathcal{S}_{\text{joint}}^{[k]}| \times |\mathcal{S}_{\text{joint}}^{[k]}|}$  refers to the probability of one agent switching from one policy in a relation to attained payoffs. Precisely, consider any two joint strategy profiles  $\boldsymbol{\pi}_{\text{joint}}^{[k]}$  and  $\hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}$  that differ in only *one* individual strategy for the  $i^{\text{th}}$  agent, i.e., there exists a unique agent such that  $\boldsymbol{\pi}_{\text{joint}}^{[k]} = \{\pi_{i,a}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}\}$  and  $\hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]} = \{\hat{\pi}_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}\}$  with  $\pi_{i,a}^{[k]} \neq \hat{\pi}_{i,b}^{[k]}$ , we set  $[\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} = 1/\sum_{i=1}^N (s_i - 1) \rho_{\pi_{i,a}^{[k]}, \hat{\pi}_{i,b}^{[k]}}(\boldsymbol{\pi}_{-i}^{[k]})$ , with  $\rho_{\pi_{i,a}^{[k]}, \hat{\pi}_{i,b}^{[k]}}(\boldsymbol{\pi}_{-i}^{[k]})$  defining the probability that one copy of agent  $i$  with strategy  $\pi_{i,a}^{[k]}$  invades the population with all other agents (in that population) playing  $\hat{\pi}_{i,b}^{[k]}$ . Following Pinsky & Karlin (2010) for  $\mathcal{P}_i^{[k]}(\pi_{i,a}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}) \neq \mathcal{P}_i^{[k]}(\hat{\pi}_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})$ , such a probability is formalized as:

$$\rho_{\pi_{i,a}^{[k]}, \hat{\pi}_{i,b}^{[k]}}(\boldsymbol{\pi}_{-i}^{[k]}) = \frac{1 - e^{-\alpha(\mathcal{P}_i^{[k]}(\pi_{i,a}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}) - \mathcal{P}_i^{[k]}(\hat{\pi}_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})}}{1 - e^{-m\alpha(\mathcal{P}_i^{[k]}(\pi_{i,a}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}) - \mathcal{P}_i^{[k]}(\hat{\pi}_{i,b}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]})}}, \quad \text{and } 1/m \text{ otherwise,} \quad (1)$$

with  $m$  being the size of the population. So far, we presented relevant derivations for the  $(\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]})$  entry of the state transition matrix when exactly the  $i^{\text{th}}$  agent differs in exactly one strategy. Having one policy change, however, only represents a subset of allowed variations, where two more cases need to be considered. Now we restrict our attention to variations in joint policies involving more than two individual strategies, i.e.,  $|\boldsymbol{\pi}_{\text{joint}}^{[k]} \setminus \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}| \geq 2$ . Here, we set<sup>2</sup>  $[\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} = 0$ . Consequently, the remaining event of self-transitions can be thus written as  $[\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} = 1 - \sum_{\hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} [\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}}$ . Summarising the above three cases, we can then write the  $(\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]})$ 's entry of the Markov chain's transition matrix as:

$$[\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} = \begin{cases} 1/\sum_{i=1}^N (s_i - 1) \rho_{\pi_{i,a}^{[k]}, \hat{\pi}_{i,b}^{[k]}}(\boldsymbol{\pi}_{-i}^{[k]}) & \text{if } |\boldsymbol{\pi}_{\text{joint}}^{[k]} \setminus \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}| = 1, \\ 1 - \sum_{\hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} [\mathbf{T}^{[k]}]_{\boldsymbol{\pi}_{\text{joint}}^{[k]}, \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}} & \text{if } \boldsymbol{\pi}_{\text{joint}}^{[k]} = \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]} \text{ self-loops,} \\ 0 & \text{if } |\boldsymbol{\pi}_{\text{joint}}^{[k]} \setminus \hat{\boldsymbol{\pi}}_{\text{joint}}^{[k]}| \geq 2, \end{cases} \quad (2)$$

The goal in  $\alpha$ -Rank is to establish an ordering in policy profiles dependent on evolutionary stability of each joint strategy. In other words, higher ranked strategies are these that are prevalent in populations with higher average times. Formally, such a notion can be easily derived as the limiting vector  $\mathbf{v}^{[k]} = \lim_{t \rightarrow \infty} \left[ [\mathbf{T}^{[k]}]^T \right]^t \mathbf{v}_0$  of our Markov chain when evolving from an initial distribution  $\mathbf{v}_0$ . Knowing that the limiting vector is a stationary distribution, one can calculate strategy rankings as the solution to the following eigenvector problem:

$$[\mathbf{T}^{[k]}]^T \mathbf{v}^{[k]} = \mathbf{v}^{[k]}. \quad (3)$$

**Limitations of  $\alpha$ -Rank:** Though the work in Omidshafiei et al. (2019) seeks to determine a solution to the above problem, it is worth mentioning that  $\alpha$ -Rank suffers from one major drawback—the

<sup>1</sup>Please note that in the original paper  $\mu$  is heuristically set to a small positive constant to ensure at maximum two varying policies per-each population. Theoretical justification can be found in Fudenberg & Imhof (2006)

<sup>2</sup>This assumption significantly reduces the analysis complexity as detailed in Fudenberg & Imhof (2006).

scalability—that we remedy in this paper. We note that the solution methodology in  $\alpha$ -Rank is in fact unscalable to settings involving more than a hand-full of agents. Particularly, authors claim polynomial complexities of their solution to the problem in Eqn. 3. Though polynomial, such a complexity, however, is polynomial in an exponential search space, i.e., the space of joint strategy profiles. As such, the polynomial complexity claim is not grounded, and need to be investigated. In short,  $\alpha$ -Rank exhibits an exponential (in terms of the number of agents) complexity for determining a ranking, thus rendering it inapplicable to settings involving more than a small amount of agents.

In what comes next, we first discuss traditional approaches that could help solve the Eqn. 3; soon we realize an off-the-shelf solution is unavailable. Hence, we commence to propose an efficient evaluation algorithm, i.e.,  $\alpha^\alpha$ -Rank, based on stochastic optimization with suitable complexities and rigorous theoretical guarantees. At the end, we propose a search heuristic to further scale up our method by introducing oracles and we name it by  $\alpha^\alpha$ -Oracle.

### 3 SCALABLE EVALUATION FOR MULTI-AGENT SYSTEMS

The problem of computing stationary distributions is a long-standing classical problem from linear algebra. Various techniques including *power method*, *PageRank*, *eigenvalue decomposition*, and *mirror descent* can be utilized for solving the problem in Eqn. 3. As we demonstrate next, any such implementation scales exponentially in the number of learners, as we summarize in Table 1.

#### 3.1 TRADITIONAL APPROACHES.

**Power Method.** One of the most common approaches to computing the solution in Eqn. 3 is the power method. Power method computes the stationary vector  $\mathbf{v}_k$  by constructing a sequence  $\{\mathbf{v}_j\}_{j \geq 0}$  from a non-zero initial vector  $\mathbf{v}_0$  by applying  $\mathbf{v}_{j+1} = 1/\|\mathbf{T}^{[k],\top} \mathbf{v}_j\| \mathbf{T}^{[k],\top} \mathbf{v}_j$ . Though viable, we first note that the power method exhibits an exponential memory complexity in terms of the number of agents. To formally derive the bound, define  $n$  to represent the total number of joint strategy profiles (i.e.,  $n = |\mathcal{S}_{\text{joint}}^{[k]}| = \prod_{i=1}^N s_i$ ), and  $m_{\text{size}}$  the total number of transitions between the states of the Markov chain in Section 2. By the construction, one can easily see that  $m_{\text{size}} = n(\sum_{i=1}^N s_i - N + 1)$  as each row and column in  $\mathbf{T}^{[k]}$  contains  $\sum_{i=1}^N s_i - N + 1$  non-zero elements. Hence, memory complexity of such an implementation are in the order of  $\mathcal{O}(m_{\text{size}}) = \mathcal{O}(n[\sum_{i=1}^N s_i - N + 1]) \approx \mathcal{O}(\prod_{i=1}^N s_i \sum_{i=1}^N s_i) = \mathcal{O}(n \sum_{i=1}^N s_i)$ .

Analyzing its time complexity, on the other hand, requires a careful consideration that links convergence rates with the resulting graph topology of the Markov chain. Precisely, the convergence rate of the power method is dictated by the second-smallest eigenvalue of the normalized Laplacian,  $\mathcal{L}_{\mathbb{G}}$ , of the graph,  $\mathbb{G}$ , associated to the Markov chain in Section 2, i.e.,  $\|\mathbf{v}^{[k]} - \mathbf{v}_j\|_2 \leq C(\mu_2(\mathcal{L}_{\mathbb{G}}))^j$ , with  $\mu_2(\mathcal{L}_{\mathbb{G}})$  being the second-smallest eigenvalue of  $\mathcal{L}_{\mathbb{G}}$ , and  $C > 0$ . Hence, as long as the second-smallest eigenvalue of the normalized Laplacian is well-behaved, one would expect suitable time complexity guarantees. To this end, we prove the following lemma

**Lemma:** [Second-Smallest Eigenvalue] *Consider the Markov chain defined in Section 2 with states in  $\mathcal{S}_{\text{joint}}^{[k]}$  and transition probability matrix  $\mathbf{T}^{[k]}$ . The second-smallest eigenvalue of the normalized Laplacian of the graph associated with the Markov chain is given by:  $\mu_2(\mathcal{L}_{\mathbb{G}}) = \frac{\min_i s_i - 1}{\sum_{i=1}^N s_i - N + 1}$ , with  $s_i$  denoting the number of strategies of agent  $i$ .*

Due to space constraints, the full proof of the above lemma is refrained to Appendix A.1. The importance of Lemma A.1 is that the resultant time complexity of the power method is also exponential of the form  $\mathcal{O}(n \times \log[(\sum_{i=1}^N s_i - N + 1)/(\min_i s_i - 1)]) = \mathcal{O}(n \log n) = \mathcal{O}(nN)$ . The above results is true as  $(\min_i s_i - 1)/(\sum_{i=1}^N s_i - N + 1) \leq (\min_i s_i - 1)/\log n = \Omega(1/\log n)$ . Hence,  $(\sum_{i=1}^N s_i - N + 1)/(\min_i s_i - 1) = \mathcal{O}(\log n)$ .

**PageRank.** Inspired by ranking web-pages on the internet, one can consider PageRank (Page et al., 1999) for computing the solution to the eigenvalue problem in Eqn. 3. Applied to our setting, we first realize that the memory is analogous to the power method that is  $\mathcal{O}(m_{\text{size}}) = \mathcal{O}(n \sum_{i=1}^N s_i)$ , and the time complexity are in the order of  $\mathcal{O}(m_{\text{size}} + n) = \mathcal{O}(n + n[\sum_{i=1}^N s_i - N + 1]) \approx \mathcal{O}(n \sum_{i=1}^N s_i)$ .

Table 1: Time and space complexity comparison ( $\epsilon$ : precision parameter).

Method	Time	Memory
Power Method	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N \log s_i\right)$	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N s_i\right)$
PageRank	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N \log s_i\right)$	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N s_i\right)$
Eig. Decomp.	$\mathcal{O}\left(\prod_{i=1}^N s_i^w\right)$	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N s_i\right)$
OSMD	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N \log s_i\right)$	$\mathcal{O}\left(\prod_{i=1}^N s_i \sum_{i=1}^N s_i\right)$
<b>Our Method</b>	$\mathcal{O}\left(\frac{1}{\epsilon^4} \sum_{i=1}^N s_i\right)$	$\mathcal{O}\left(\frac{1}{\epsilon^2} \sum_{i=1}^N s_i\right)$

**Eigenvalue Decomposition.** Apart from the above, we can also consider the problem as a standard eigenvalue decomposition task (also what the original  $\alpha$ -Rank is implemented according to Lanctot et al. (2019)) and adopt the method in Coppersmith & Winograd (1990) to compute the stationary distribution. Unfortunately, state-of-the-art techniques for eigenvalue decomposition also require exponential memory and exhibit a time complexity of the form  $\mathcal{O}(n^\omega) = \mathcal{O}\left(\prod_{i=1}^N s_i^w\right)$  with  $\omega \in [2, 2.376]$ . Clearly, these bounds restrict  $\alpha$ -Rank to small number of agents  $N$ .

**Mirror Descent.** The ordered subsets mirror descent (Ben-Tal et al., 2001) requires at each iteration a projection on standard  $n$ -dimensional simplex:  $\Xi_n = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\top \mathbf{1} = 1 \ \& \ \mathbf{x} \succeq \mathbf{0}\}$ . As stated in the paper, the computing of this projection requires  $\mathcal{O}(n \log n)$  time. In our setting,  $n = \prod_{i=1}^N s_i$  is the total number of joint strategy profiles. Hence, the projection step is exponential in the number of agents  $N$ . This makes mirror descent inapplicable for  $\alpha$ -Rank when  $N$  is large.

### 3.2 OUR PROPOSAL: AN OPTIMIZATION-BASED SOLUTION.

Rather than seeking an exact solution to the problem in Eqn. 3, one can consider approximate solvers by defining a constraint optimization objective:

$$\min_{\mathbf{x}} \frac{1}{n} \left\| \mathbf{T}^{[k], \top} \mathbf{x} - \mathbf{x} \right\|_2^2 \quad \text{s.t. } \mathbf{x}^\top \mathbf{1} = 1, \text{ and } \mathbf{x} \succeq \mathbf{0}. \quad (4)$$

The constrained objective in Eqn. 4 simply seeks a vector  $\mathbf{x}$  minimizing the distance between  $\mathbf{x}$ , itself, and  $\mathbf{T}^{[k], \top} \mathbf{x}$  (i.e., attempting to solve  $\mathbf{T}^{[k], \top} \mathbf{x} = \mathbf{x}$ ) while ensuring that  $\mathbf{x}$  lies on an  $n$ -dimensional simplex (i.e.,  $\mathbf{x}^\top \mathbf{1} = 1$ , and  $\mathbf{x} \succeq \mathbf{0}$ ). Due to time and memory complexities required for computing exact solutions, we focus on determining an approximate vector  $\tilde{\mathbf{x}}$  defined to be the solution to the following relaxed problem of Eqn. 4:

$$\min_{\mathbf{x}} \frac{1}{n} \left\| \mathbf{T}^{[k], \top} \mathbf{x} - \mathbf{x} \right\|_2^2 \quad \text{s.t. } |\mathbf{x}^\top \mathbf{1} - 1| \leq \delta \text{ for } 0 < \delta < 1, \text{ and } \mathbf{x} \succeq \mathbf{0}. \quad (5)$$

The optimization problem in Eqn. 5 can be solved using a barrier-like technique that we detail below. Before that, it is instructive to clarify the connection between the original and the relaxed problems

**Proposition:** [Connections to Markov Chain] *Let  $\tilde{\mathbf{x}}$  be a solution to the relaxed optimization problem in Eqn. 5. Then,  $\tilde{\mathbf{x}} / \|\tilde{\mathbf{x}}\|_1 = \mathbf{v}^{[k]}$  is the stationary distribution to the Markov chain in Section 2.*

Importantly, the above proposition allows us to focus on solving the problem in Eqn. 5 which only exhibits inequality constraints. Problems of this nature can be solved by considering a barrier function leading to an unconstrained finite sum minimization problem. To do so, denoting  $\mathbf{b}_i^{[k]}$  to be the  $i^{\text{th}}$  row of  $\mathbf{T}^{[k], \top} - \mathbf{I}$ , we can write  $\frac{1}{n} \left\| \mathbf{T}^{[k], \top} \mathbf{x} - \mathbf{x} \right\|_2^2 = \frac{1}{n} \sum_{i=1}^n \left( \mathbf{x}^\top \mathbf{b}_i^{[k]} \right)^2$ . Introducing logarithmic barrier-functions, with  $\lambda > 0$  being a penalty parameter, we arrive at

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{n} \sum_{i=1}^n \left( \mathbf{x}^\top \mathbf{b}_i^{[k]} \right)^2 - \lambda \log \left( \delta^2 - [\mathbf{x}^\top \mathbf{1} - 1]^2 \right) - \frac{\lambda}{n} \sum_{i=1}^n \log(x_i), \quad (6)$$

Eqn. 6 is a standard finite minimization problem that can be solved using any off-the-shelf stochastic optimization algorithm, e.g., stochastic gradients, ADAM (Kingma & Ba, 2014) among others. A stochastic gradient execution involves sampling a strategy profile  $i_t \sim [1, \dots, n]$  at iteration  $t$ ,

and then executing a descent step:  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t)$ , with  $\nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t)$  being a sub-sampled gradient of Eqn. 6, and  $\lambda$  being a scheduled penalty parameter with  $\lambda_{t+1} = \lambda_t/\gamma$  for some  $\gamma > 1$ ,

$$\nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t) = 2 \left( \mathbf{b}_{i_t}^{[k], \top} \mathbf{1} \right) \mathbf{b}_{i_t}^{[k]} + \frac{2\lambda_t (\mathbf{x}_t^\top \mathbf{1} - 1)}{\delta^2 - (\mathbf{x}_t^\top \mathbf{1} - 1)^2} - \lambda_t/n \left[ \frac{1}{[\mathbf{x}_t]_1}, \dots, \frac{1}{[\mathbf{x}_t]_n} \right]^\top. \quad (7)$$

See Phase I in Algorithm 1 for the pseudo-code. We can further derive a convergence theorem of:

**Theorem:** [Convergence of Barrier Method] *Let  $\tilde{\mathbf{x}}_\lambda$  be the output of a gradient algorithm descending in the objective in Eqn. 6, after  $\mathcal{T}$  iterations, then*

$$\mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k], \top} \right] \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2^2 \right] \leq \mathcal{O} \left( \lambda + \frac{1}{\sqrt{\mathcal{T}}} + \frac{n^2}{\gamma^{\mathcal{T}} \lambda} \right),$$

where expectation is taken w.r.t. all randomness of a stochastic gradient implementation, and  $\gamma > 1$  is a decay-rate for  $\lambda$ , i.e.,  $\lambda^{[t+1]} = \lambda^{[t]}/\gamma$ .

---

### Algorithm 1 $\alpha^\alpha$ -Rank: Scalable Multi-Agent Evaluation

---

- 1: **Inputs: Evaluation Parameters:** Initial a vector  $\mathbf{x}_0 = 1/n\mathbf{1}$ , total number of iterations  $T = \mathcal{O}(1/\epsilon^2)$ , decaying learning rate  $\{\eta_t\}_{t=1}^T$ , penalty parameter  $\lambda \approx \mathcal{O}(\epsilon)$ ,  $\lambda$  decay rate  $\gamma > 1$ , total number of joint strategy profiles  $n$ , and a constraint relaxation term  $\delta$ . **Oracle Parameters:** initialize a subset of strategy pools for all agents  $\{\mathcal{S}_i^{[0]}\}$  by randomly sampling from  $\{\mathcal{S}_i\}$
  - 2: Set outer iteration count  $k = 0$
  - 3: **while** stopping criteria **do:**
  - 4:     **Phase I: Scalable Policy Evaluation (Section 3.2):**
  - 5:     **for**  $t = 0 \rightarrow T - 1$  **do:**
  - 6:         Uniformly sample one strategy profile  $i_t^{[k]} \sim \{1, \dots, n\}$
  - 7:         Construct  $b_{i_t}^{[k]}$  as the  $i_t^{[k]}$  row of  $\mathbf{T}^{[k], \top} - \mathbf{I}$  with  $\mathbf{T}^{[k]}$  entries defined Eqn. 2
  - 8:         Update solution  $\mathbf{x}_{t+1}^{[k]} = \mathbf{x}_t^{[k]} - \eta_t \nabla_{\mathbf{x}} f_{i_t^{[k]}}(\mathbf{x}_t^{[k]})$  as in Eqn. 7 & set  $\lambda_{t+1} = \lambda_t/\gamma$
  - 9:         Extract  $\pi_{\text{joint}}^{[k], \text{top}}$  from approximate stationary distribution  $\mathbf{v}^{[k]} = \mathbf{x}_T^{[k]} / \|\mathbf{x}_T^{[k]}\|_1$
  - 10:     **Phase II (if turned on): Scalable Policy Evaluation with Oracle (Section 3.3):**
  - 11:     **for** each agent  $i$  **do:**
  - 12:         Compute the best-response strategy  $\pi_i^*$  by solving Eqn. 8.
  - 13:         Update strategy pools for each agent  $i$  as  $\mathcal{S}_i^{[k]} = \mathcal{S}_i^{[k-1]} \cup \pi_i^*$
  - 14:     Set  $k = k + 1$
  - 15: **Return:** Best performing strategy profile  $\pi_{\text{joint}, \star}$  across all agents.
- 

The proof of the above theorem (see the full proof in Appendix A.2) is interesting by itself, a more important aspect is the memory and time complexity implications posed by our algorithm. Theorem A.2 implies that after  $T = \mathcal{O}(1/\epsilon^2)$  iterations with  $\epsilon$  being a precision parameter, our algorithm outputs a vector  $\tilde{\mathbf{x}}_\lambda \succeq \mathbf{0}$  such that

$$\|\tilde{\mathbf{x}}_\lambda^\top \mathbf{1} - 1\| \leq \delta \quad \text{and} \quad \mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k], \top} \right] \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2 \right] \approx \mathcal{O}(\epsilon).$$

Moreover, one can easily see<sup>3</sup> that after  $T$  steps, the overall time and memory complexities of our update rules are given by  $\mathcal{O}(T^2 \sum_{i=1}^N s_i)$  and  $\mathcal{O}(T \sum_{i=1}^N s_i)$ , respectively. Using  $T = \mathcal{O}(1/\epsilon^2)$  eventually leads to a memory complexity of  $\mathcal{O}(1/\epsilon^2 \sum_{i=1}^N s_i)$  and  $\mathcal{O}(1/\epsilon^4 \sum_{i=1}^N s_i)$  for time (see the comparison in Table. 1). Hence, our algorithm is able to achieve an **exponential reduction**, in terms of number of agents, in both memory and time complexities.

### 3.3 HEURISTIC SEARCH BY INTRODUCING ORACLES

So far, we have presented scalable multi-agent evaluations through stochastic optimization. We can further boost scalability (to tens of millions of joint profiles) of our method by introducing an oracle mechanism. The heuristic of oracles was first introduced in solving large-scale zero-sum matrix

<sup>3</sup>More details on these derivations can be found in the Appendix A.3

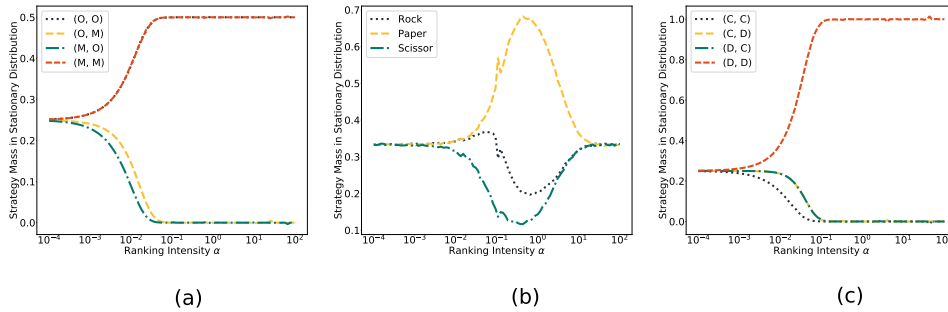


Figure 2: Ranking intensity sweep on (a) Battle of Sexes (b) Biased RPS (c) Prisoner’s Dilemma.




Figure 3: Comparisons of time and memory complexities on varying sizes of random matrices.

games (McMahan et al., 2003). The idea is to first create a restricted sub-game in which all players are only allowed to play a restricted number of strategies, which are then expanded by adding incorporating each of the players’ best-responses to opponents; the sub-game will be replayed with agents’ augmented strategy pools before a new round of best responses is found. The worse-case scenario of introducing oracles would be to solve the original evaluation problem in full size. The best response is assumed to be given by an *oracle* that can be simply implemented by a grid search. Precisely, given the top-rank profile  $\pi_{-i}^{[k],\text{top}}$  at iteration  $k$ , the goal for agent  $i$  is to select<sup>4</sup> the optimal  $\pi_i^*$  from the pre-defined strategy pool  $\mathcal{S}_i$  to maximize the reward

$$\pi_i^* = \arg \max_{\pi_i \in \mathcal{S}_i} \mathbb{E}_{\pi_i, \pi_{-i}^{[k],\text{top}}} \left[ \sum_{h \geq 0} \gamma_i^{[k],h} \mathcal{P}_i^{[k]}(\mathbf{x}_h^{[k]}, \mathbf{u}_{i,h}^{[k]}, \mathbf{u}_{-i,h}^{[k]}) \right], \tag{8}$$

with  $\mathbf{x}_h^{[k]}$  denoting the state,  $\mathbf{u}_{i,h}^{[k]} \sim \pi_i(\mathbf{u}_{i,h}^{[k]} | \mathbf{x}_{i,h}^{[k]})$ ,  $\mathbf{u}_{-i,h}^{[k]} \sim \pi_{-i}^{[k],\text{high}}(\mathbf{u}_{-i,h}^{[k]} | \mathbf{x}_{-i,h}^{[k]})$  denoting the actions from agent  $i$  and the opponents, respectively. The heuristic of solving the full game from restricted sub-games is crucial especially when it is prohibitively expensive to list all joint-strategy profiles, e.g., in scenarios involving tens-of-millions of joint profiles.

For a complete exposition, we summarize the pseudo-code in Algorithm 1. In the first phase, vanilla  $\alpha^\alpha$ -Rank is executed (lines 4-9), while in the second (lines 11 - 13),  $\alpha^\alpha$ -Rank with Oracle (if turned on) is computed. To avoid any confusion, we refer to the latter as  $\alpha^\alpha$ -Oracle. Note that even though in the two-player zero-sum games, the oracle algorithm (McMahan et al., 2003) is guaranteed to converge to the minimax equilibrium. Providing valid convergence guarantees for  $\alpha^\alpha$ -Oracle is an interesting direction for future work. In this paper, we rather demonstrate the effectiveness of such an approach in a large-scale empirical study as shown in Section 4.

### 4 EXPERIMENTS

In this section, we evaluate the scalability properties of  $\alpha^\alpha$ -Rank<sup>5</sup>. Precisely, we demonstrate that our method is capable of successfully recovering optimal policies in self-driving car simulations and in the Ising model where strategy spaces are in the order of up to tens-of-millions of possible strategies. We note that these sizes are well beyond the capabilities of state-of-the-art methods, e.g.,

<sup>4</sup>Note that the oracle can also be used to find an improved policy that is not in the original strategy pool, similar to PSRO (Lanctot et al., 2017), which we leave to the future work.

<sup>5</sup>All of the experiments are run by a single machine with 64 GB memory, and 10-core Intel I9-9900X CPU.

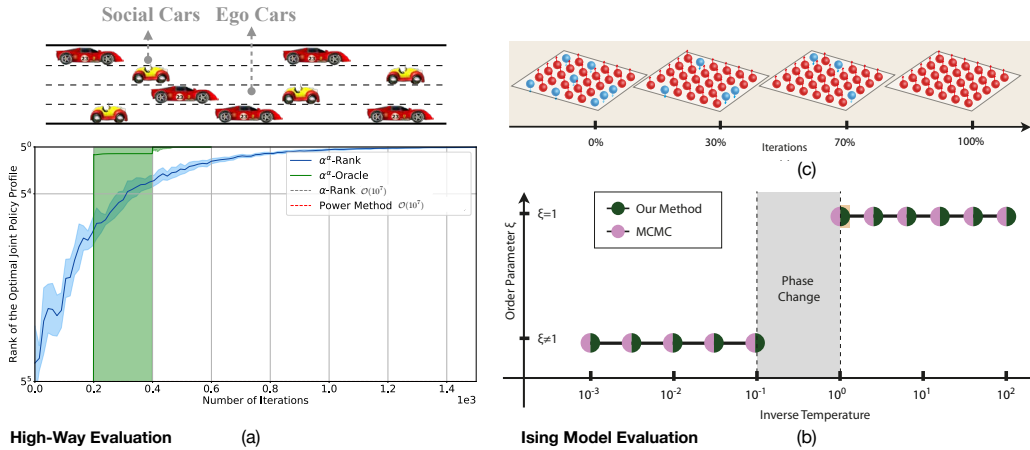


Figure 4: Large-scale multi-agent evaluations. (a) Convergence of the optimal joint-strategy profile of self-driving simulation on the highway. (b) Status of the Ising-model equilibrium measured by  $\xi = \frac{|N_{\uparrow} - N_{\downarrow}|}{|N|}$ . (c) Convergence of the top-rank profile from  $\alpha^{\alpha}$ -Oracle under  $\tau = 1, \alpha = 1$ .

$\alpha$ -Rank (Omidshafiei et al., 2019) that considers at maximum four agents with four strategies, or AlphaStar which handles about 600 strategies as detailed in Vinyals et al. (2019).

**Sparsity Data Structures.** During the implementation phase, we realised that the transition probability,  $T^{[k]}$ , of the Markov chain induces a sparsity pattern (each row and column in  $T^{[k]}$  contains  $\sum_{i=1}^N s_i - N + 1$  non-zero elements, check Section 3.2) that if exploited can lead to significant speed-ups. To fully leverage such sparsity, we tailored a novel data structure for sparse storage and computations needed by Algorithm 1. More details can be found in Appendix B.1.

**Correctness of Ranking Results.** Before conducting large-scale sophisticated experiments, it is instructive to validate the correctness of our results on the simple cases especially those reported by Omidshafiei et al. (2019). We therefore test on three normal-form games. Due to space constraints, we refrain the full description of these tasks to Appendix B.2. Fig. 2 shows that, in fact, results generated by  $\alpha^{\alpha}$ -Rank, the Phase I of Algorithm 1, are consistent with  $\alpha$ -Rank’s results.

**Complexity Results on Random Matrices.** We measured the time and memory needed by our method for computing the stationary distribution with varying sizes of simulated random matrices. Baselines includes eigenvalue decomposition from Numpy, optimization tools in PyTorch, and  $\alpha$ -Rank from OpenSpiel (Lanctot et al., 2019). For our algorithm we terminated execution with gradient norms being below a predefined threshold of 0.01. According to Fig. 3,  $\alpha^{\alpha}$ -Rank can achieve three orders of magnitude reduction compared to eigenvalue decomposition in terms of time. Most importantly, the performance gap keeps developing with the increasing matrix size.

**Autonomous Driving on Highway:** High-way (Leurent, 2018) provides an environment for simulating self-driving scenarios with social vehicles designed to mimic real-world traffic flow as strategy pools. We conducted a ranking experiment involving 5 agents each with 5 strategies, i.e. a strategy space in the order of  $\mathcal{O}(5^5)$  (3125 possible strategy profiles). Agent strategies varied between “rational” and “dangerous” drivers, which we encoded using different reward functions during training (complete details of defining reward functions can be found in Appendix C.2). Under this setting, we know, upfront, that optimal profile corresponds to all agents is five rational drivers. Cars trained using value-iteration and rewards averaged from 200 test trails were reported. Due to the size of the strategy space, we considered both  $\alpha^{\alpha}$ -Rank and  $\alpha^{\alpha}$ -Oracle. We set  $\alpha^{\alpha}$ -Oracle to run 200 iterations of gradient updates in solving the top-rank strategy profile (Phase I in Algorithm 1). Results depicted in Fig. 4(a) clearly demonstrate that both our implementations are capable of recovering the correct highest ranking strategy profile. We also note that though such sizes are feasible using  $\alpha$ -Rank and the power-method, our results achieve 4 orders of magnitude reduction in total number of iterations.

**Ising Model Experiment:** The Ising model (Ising, 1925) is the model for describing ferromagnetism in statistical mechanics. It assumes a system of magnetic spins, where each spin  $a^j$  is either an up-spin,  $\uparrow$ , or down-spin,  $\downarrow$ . The system energy is defined by  $E(\mathbf{a}, h) = -\sum_j (h^j a^j + \frac{\lambda}{2} \sum_{k \neq j} a^j a^k)$  with  $h^j$  and  $\lambda$  being constant coefficients. The probability of one set of spin configuration is  $P(\mathbf{a}) = \frac{\exp(-E(\mathbf{a}, h)/\tau)}{\sum_{\mathbf{a}} \exp(-E(\mathbf{a}, h)/\tau)}$  where  $\tau$  is the environmental temperature. Finding the equilibrium of



the system is notoriously hard because it is needed to enumerate all possible configurations in computing  $P(\mathbf{a})$ . Traditional approaches include Markov Chain Monte Carlo (MCMC). An interesting phenomenon is the *phase change*, i.e., the spins will reach an equilibrium in the low temperatures, with the increasing  $\tau$ , such equilibrium will suddenly break and the system becomes chaotic.

Here we try to observe the phase change through multi-agent evaluation methods. We assume each spins as an agent, and the reward to be  $r^j = h^j a^j + \frac{\lambda}{2} \sum_{k \neq j} a^j a^k$ , and set  $\alpha = \frac{1}{\tau}$  to build the link between Eqn. 1 and  $P(\mathbf{a})$ . We consider the top-rank strategy profile from  $\alpha^\alpha$ -Oracle as the system equilibrium and compare it against the ground truth from MCMC. We consider a five-by-five 2D model which induces a prohibitively-large strategy space of size  $2^{25}$  (tens of millions) to which the existing baselines, including  $\alpha^\alpha$ -Rank on the single machine, are inapplicable. Fig. 4(b) illustrates that our method identifies the same phase change as what MCMC suggests. We show an example of how  $\alpha^\alpha$ -Oracle’s top-ranked profile finds the system equilibrium in Fig. 4(c) at  $\tau = 1$ .

## 5 CONCLUSION

In this paper, we demonstrated that the approach in Omidshafiei et al. (2019) exhibits exponential time and memory complexities. We then proposed  $\alpha^\alpha$ -Rank as a scalable solution for multi-agent evaluation with linear time and memory demands. In a set of experiments, we demonstrated that our method is truly scalable capable of handling large strategy spaces.

There are a lot of interesting avenues for future research. First, we plan to theoretically analyze convergence properties of the resulting oracle algorithm, and further introduce policy learning through oracles. Second, we plan take our method to the real-world by conducting multi-robot experiments.

## REFERENCES

- Sasmita Barik, Ravindra Bapat, and Sukanta Pati. On the laplacian spectra of product graphs. *Applicable Analysis and Discrete Mathematics*, 9, 04 2015. doi: 10.2298/AADM150218006B.
- Aharon Ben-Tal, Tamar Margalit, and Arkadi Nemirovski. The ordered subsets mirror descent optimization method with applications to tomography. *SIAM Journal on Optimization*, 12(1):79–108, 2001.
- Michael Bowling and Manuela Veloso. Convergence of gradient dynamics with a variable learning rate. In *ICML*, pp. 27–34, 2001.
- X Chen and X Deng. Settling the complexity of 2-player nash equilibrium. eccc. Technical report, Report, 2005.
- Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of symbolic computation*, 9(3):251–280, 1990.
- Drew Fudenberg and Lorens A Imhof. Imitation processes with small mutations. *Journal of Economic Theory*, 131(1):251–262, 2006.
- Jordi Grau-Moya, Felix Leibfried, and Haitham Bou-Ammar. Balancing two-player stochastic games with soft q-learning. *arXiv preprint arXiv:1802.03216*, 2018.
- Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. A unified game-theoretic approach to multiagent reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 4190–4203, 2017.

- Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. Openspiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.
- Edouard Leurent. An environment for autonomous driving decision-making. <https://github.com/eleurent/highway-env>, 2018.
- H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 536–543, 2003.
- John F Nash et al. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, second edition, 2006.
- Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos.  $\alpha$ -rank: Multi-agent evaluation by evolution. *Scientific Reports, Nature*, 2019.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems*, pp. 5872–5882, 2017.
- Mark Pinsky and Samuel Karlin. *An introduction to stochastic modeling*. Academic press, 2010.
- Oriol Vinyals, Igor Babuschkin, Junyoung Chung, Michael Mathieu, Max Jaderberg, Wojciech M Czarnecki, Andrew Dudzik, Aja Huang, Petko Georgiev, Richard Powell, et al. Alphastar: Mastering the real-time strategy game starcraft ii. *DeepMind Blog*, 2019.
- Yannick Viossat. The replicator dynamics does not lead to correlated equilibria. *Games and Economic Behavior*, 59(2):397–407, 2007.
- Ying Wen, Yaodong Yang, Rui Luo, Jun Wang, and Wei Pan. Probabilistic recursive reasoning for multi-agent reinforcement learning. *arXiv preprint arXiv:1901.09207*, 2019.

## APPENDIX

## A COMPREHENSIVE PROOFS

## A.1 LEMMA OF SECOND-SMALLEST EIGENVALUE

**Lemma:** [Second-Smallest Eigenvalue] *Consider the Markov chain defined in Section 2 with states in  $\mathcal{S}_{joint}^{[k]}$  and transition probability matrix  $\mathbf{T}^{[k]}$ . The second-smallest eigenvalue of the normalized Laplacian of the graph associated with the Markov chain is given by:  $\mu_2(\mathcal{L}_{\mathbb{G}}) = \frac{\min_i s_i - 1}{\sum_{i=1}^N s_i - N + 1}$ , with  $s_i$  denoting the number of strategies of agent  $i$ .*

*Proof:* For simplicity we drop round index  $k$  in the below derivation. Notice, the underlying graph for the constructed Markov Chain can be represented as a Cartesian product of  $N$  complete graphs  $\mathcal{K}_{s_i}$ <sup>6</sup>:

$$\mathbb{G} = \mathcal{K}_{s_1} \times \mathcal{K}_{s_2} \times \dots \times \mathcal{K}_{s_N} \quad (9)$$

Indeed, two vertices  $\boldsymbol{\pi}^{[k]}, \hat{\boldsymbol{\pi}}^{[k]} \in \mathbb{G}$  are connected by the edge if and only these joint strategy profiles differ in at most one individual strategy, i.e  $\exists! i \in \{1, \dots, N\} : \boldsymbol{\pi}^{[k]} = \{\pi_{ia}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}\}, \hat{\boldsymbol{\pi}}^{[k]} = \{\hat{\pi}_{ib}^{[k]}, \boldsymbol{\pi}_{-i}^{[k]}\}$ . Hence, the spectral properties of  $\mathbb{G}$  can be described in terms of spectral properties of  $\mathcal{K}_{s_i}$  as follows (Barik et al., 2015):

$$\begin{aligned} \text{Spectr}(\mathbb{G}) &= \\ &\left\{ \mu_{i_1}(\mathbb{L}_{\mathcal{K}_{s_1}}^{(un)}) + \dots + \mu_{i_N}(\mathbb{L}_{\mathcal{K}_{s_N}}^{(un)}) \right\}_{i_1=1, i_2=1, \dots, i_N=1}^{s_1, s_2, \dots, s_N} \\ \text{Eigenvectors of } \mathbb{G} &= \\ &\left\{ \boldsymbol{\vartheta}_{i_1, 1} \otimes \boldsymbol{\vartheta}_{i_2, 2} \otimes \dots \otimes \boldsymbol{\vartheta}_{i_N, N} \right\}_{i_1=1, i_2=1, \dots, i_N=1}^{s_1, s_2, \dots, s_N} \end{aligned}$$

where  $\mu_i(\mathbb{L}_{\mathcal{K}_{s_j}}^{(un)})$  is the  $i^{th}$  eigenvalue of the unnormalized Laplacian of the complete graph  $\mathcal{K}_{s_j}$  and  $\boldsymbol{\vartheta}_{i,j}$  is the corresponding eigenvector<sup>7</sup>. The spectrum of unnormalized Laplacian of the complete graph  $\mathcal{K}_{s_i}$  is given by  $\text{Spectr}(\mathcal{K}_{s_i}) = \{0, s_i - 1\}$  and the only eigenvector corresponding to zero eigenvalue is  $\mathbf{1} \in \mathbb{R}^{s_i}$ . Therefore, the minimum non-zero eigenvalue of unnormalized Laplacian of  $\mathbb{G}$  is given by  $\min_i s_i - 1$ . Finally, due to the fact that  $\mathbb{G}$  is a regular graph (with degree of each node is equal to  $\sum_{i=1}^N s_i - N + 1$ ), the smallest non-zero eigenvalue of the normalized Laplacian of  $\mathbb{G}$  is given by  $\frac{\min_i s_i - 1}{\sum_{i=1}^N s_i - N + 1}$ .

Giving this result, the overall time complexity of Power Method is bounded by  $\mathcal{O}\left(n \times \log \left[ \frac{\sum_{i=1}^N s_i - N + 1}{\min_i s_i - 1} \right]\right) = \mathcal{O}(n \log n)$ . Indeed, notice that  $\frac{\min_i s_i - 1}{\sum_{i=1}^N s_i - N + 1} \leq \frac{\min_i s_i - 1}{\log n} = \Omega\left(\frac{1}{\log n}\right)$ , hence,  $\frac{\sum_{i=1}^N s_i - N + 1}{\min_i s_i - 1} = \mathcal{O}(\log n)$ . As for the memory complexity, Power Method requires has the same requirements as PageRank algorithm.<sup>8</sup> These results imply that Power Method scales exponentially with number of agents  $N$ , and therefore, inapplicable when  $N$  is large.

## A.2 THEOREM OF CONVERGENCE OF BARRIER METHOD

**Theorem:** [Convergence of Barrier Method] *Let  $\tilde{\boldsymbol{x}}_\lambda$  be the output of a gradient algorithm descending in the objective in Eqn. 6, after  $\mathcal{T}$  iterations, then*

$$\mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k, \mathcal{T}]} \right] \tilde{\boldsymbol{x}}_\lambda - \tilde{\boldsymbol{x}}_\lambda \right\|_2^2 \right] \leq \mathcal{O} \left( \lambda + \frac{1}{\sqrt{\mathcal{T}}} + \frac{n^2}{\gamma \mathcal{T} \lambda} \right),$$

where expectation is taken w.r.t. all randomness of a stochastic gradient implementation, and  $\gamma > 1$  is a decay-rate for  $\lambda$ , i.e.,  $\lambda^{[t+1]} = \lambda^{[t]} / \gamma$ . See the Algorithm 2.

<sup>6</sup>Here,  $\mathcal{K}_p$  denotes a complete graph with  $p$  nodes.

<sup>7</sup>In other words,  $\mathbb{L}_{\mathcal{K}_{s_j}} \boldsymbol{\vartheta}_{i,j} = \mu_i(\mathbb{L}_{\mathcal{K}_{s_j}}) \boldsymbol{\vartheta}_{i,j}$  for all  $i = 1, \dots, s_j$  and  $j = 1, \dots, N$ .

<sup>8</sup>Due to necessity to store matrix  $\mathbf{T}^{[k]}$

**Algorithm 2** Log-Barrier Stochastic Gradient Descent

- 1: **Input:**  $\mathbf{x}_0 = \frac{1}{n}\mathbf{1}$ ,  $\mathcal{T}$ ,  $\lambda$ ,  $\delta$ ,  $\epsilon$ ,  $\{\eta_t\}_{t=1}^{\mathcal{T}}$ ,  $\gamma > 1$ .
- 2: **Output:**  $\hat{\nu}$
- 3: Set  $\delta_0 = \delta$ ,  $\lambda_0 = \lambda$ .
- 4: **for**  $t=0$  *to*  $\mathcal{T}-1$  **do:**
- 5:   Sample  $i_t \sim [1, \dots, n]$  and compute:

$$\nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t) = 2(\mathbf{b}_{i_t}^{[k],\top} \mathbf{1}) \mathbf{b}_{i_t}^{[k]} + \frac{2\lambda_t(\mathbf{x}_t^\top \mathbf{1} - 1)}{\delta^2 - (\mathbf{x}_t^\top \mathbf{1} - 1)^2} \mathbf{1} - \frac{\lambda_t}{n} \begin{bmatrix} \frac{1}{[\mathbf{x}_t]_1} \\ \vdots \\ \frac{1}{[\mathbf{x}_t]_n} \end{bmatrix}$$

- 6:   Update  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t)$ .
- 7:   Update  $\lambda_{t+1} = \frac{\lambda_t}{\gamma}$ .
- 8: **end for**
- 9: Set  $\hat{\nu} = \mathbf{x}_{\mathcal{T}}$

*Proof:* Let  $\tilde{\mathbf{x}}$  and  $\tilde{\mathbf{x}}_\lambda^*$  be the solutions of Eqn. (5) and Eqn. (6) respectively. Convergence guarantees for logarithmic barrier method (Nocedal & Wright, 2006) with penalty parameter  $\lambda$  and barrier parameter  $\gamma$  gives:

$$\left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda^* - \tilde{\mathbf{x}}_\lambda^* \right\|_2^2 - \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}} - \tilde{\mathbf{x}} \right\|_2^2 = \mathcal{O} \left( \frac{n^2}{\gamma^{\mathcal{T}} \lambda} \right) \quad (10)$$

and using  $\left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}} = \tilde{\mathbf{x}}$  in (10) gives:

$$\left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda^* - \tilde{\mathbf{x}}_\lambda^* \right\|_2^2 = \mathcal{O} \left( \frac{n^2}{\gamma^{\mathcal{T}} \lambda} \right) \quad (11)$$

Applying the convergence guarantees of stochastic gradient descent method to convex function  $F(\mathbf{x})$  gives:

$$\mathbb{E} [F(\tilde{\mathbf{x}}_\lambda) - F(\tilde{\mathbf{x}}_\lambda^*)] \leq \mathcal{O} \left( \frac{1}{\sqrt{\mathcal{T}}} \right)$$

Using the definition of function  $F(\mathbf{x})$ :

$$n \mathbb{E} [F(\tilde{\mathbf{x}}_\lambda) - F(\tilde{\mathbf{x}}_\lambda^*)] \geq \mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2^2 \right] - \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda^* - \tilde{\mathbf{x}}_\lambda^* \right\|_2^2 - \mathcal{O}(\lambda)$$

Combining the last two results:

$$\mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2^2 \right] - \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda^* - \tilde{\mathbf{x}}_\lambda^* \right\|_2^2 \leq \mathcal{O}(\lambda) + \mathcal{O} \left( \frac{1}{\sqrt{\mathcal{T}}} \right)$$

Finally, using (11) gives:

$$\mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2^2 \right] \leq \mathcal{O} \left( \lambda + \frac{1}{\sqrt{\mathcal{T}}} + \frac{n^2}{\gamma^{\mathcal{T}} \lambda} \right)$$

By choosing  $\lambda = \frac{\epsilon}{3}$  and  $\mathcal{T} = \max \left\{ \frac{9}{\epsilon^2}, \frac{2 \log \frac{3n}{\epsilon}}{\log \gamma} \right\}$ , we have  $\lambda + \frac{1}{\sqrt{\mathcal{T}}} + \frac{n^2}{\gamma^{\mathcal{T}} \lambda} \leq \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon$ . Hence,

$$\mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2^2 \right] \leq \mathcal{O}(\epsilon)$$

### A.3 TIME AND MEMORY COMPLEXITY OF $\alpha^\alpha$ -RANK

The above result implies given precision parameter  $\epsilon > 0$ , after  $\mathcal{T} = \mathcal{O}\left(\frac{1}{\epsilon^2}\right)$  iterations, Algorithm 2 outputs vector  $\tilde{\mathbf{x}}_\lambda \succ \mathbf{0}$  such that:

$$|\tilde{\mathbf{x}}_\lambda^\top \mathbf{1} - 1| \leq \delta, \text{ and } \mathbb{E} \left[ \left\| \left[ \mathbf{T}^{[k]} \right]^\top \tilde{\mathbf{x}}_\lambda - \tilde{\mathbf{x}}_\lambda \right\|_2 \right] \leq \mathcal{O}(\epsilon)$$

Hence, by tuning parameters  $\delta$  and  $\epsilon$  one can approximate a stationary distribution vector  $\boldsymbol{\nu}^{[k]}$ .

Algorithm 2 starts with uniform distribution vector  $\mathbf{x}_0 = \frac{1}{n}\mathbf{1}$  and at step  $t$  it updates the previous iterative  $\mathbf{x}_t$  by a rule given in line 6. Let  $\mathcal{N}_t = \{j \in [n] : \exists r \in [t] \text{ with } [\mathbf{b}_{i_r}^{[k]}]_j \neq 0\}$ <sup>9</sup>. Then, for  $j \notin \mathcal{N}_t$  all entries  $[\mathbf{x}_{t+1}]_j$  are equal to each other and updated as:

$$[\mathbf{x}_{t+1}]_j = [\mathbf{x}_t]_j + \frac{2\eta_t \lambda_t (\mathbf{x}_t^\top \mathbf{1} - 1)}{\delta^2 - (\mathbf{x}_t^\top \mathbf{1} - 1)^2} - \frac{\eta_t \lambda_t}{n[\mathbf{x}_t]_j}$$

Given value  $\mathbf{x}_t^\top \mathbf{1}$ , the above computation takes  $\mathcal{O}(1)$  time and space. For entries  $j \in \mathcal{N}_t$ , all entries  $[\mathbf{x}_{t+1}]_j$  might be different (worst case) and, therefore, update

$$[\mathbf{x}_{t+1}]_j = [\mathbf{x}_t]_j - 2\mathbf{1}^\top \mathbf{b}_{i_t}^{[k]} [\mathbf{b}_{i_t}^{[k]}]_j + \frac{2\eta_t \lambda_t (\mathbf{x}_t^\top \mathbf{1} - 1)}{\delta^2 - (\mathbf{x}_t^\top \mathbf{1} - 1)^2} - \frac{\eta_t \lambda_t}{n[\mathbf{x}_t]_j}$$

need to be performed at most  $\mathcal{O}(|\mathcal{N}_t|)$  times. Notice, only  $\mathcal{O}(\sum_{i=1}^N s_i - N + 1)$  of these entries require the same computation of  $\mathbf{1}^\top \mathbf{b}_{i_t}^{[k]}$  and all other entries can be updated in constant time and space. Hence, the computation of all these entries can be done in  $\mathcal{O}(|\mathcal{N}_t|)$  time and space. Finally, term  $\mathbf{x}_{t+1}^\top \mathbf{1}$  can be computed in  $\mathcal{O}(|\mathcal{N}_t|)$  time and space. Hence, storing and updating of all entries  $[\mathbf{x}_{t+1}]_j$  require  $\mathcal{O}\left(t \sum_{i=1}^N s_i\right)$  time and space. Therefore, after  $\mathcal{O}(\mathcal{T})$  steps, the overall time and memory complexity of Algorithm 2 are given by  $\mathcal{O}\left(\mathcal{T}^2 \sum_{i=1}^N s_i\right)$  and  $\mathcal{O}\left(\mathcal{T} \sum_{i=1}^N s_i\right)$  respectively. Using  $\mathcal{T} = \mathcal{O}\left(\frac{1}{\epsilon^2}\right)$  eventually gives  $\mathcal{O}\left(\frac{1}{\epsilon^4} \sum_{i=1}^N s_i\right)$  for time and  $\mathcal{O}\left(\frac{1}{\epsilon^2} \sum_{i=1}^N s_i\right)$  for space complexity. Our result is summarized in Table 1.

## B IMPLEMENTATION OF $\alpha^\alpha$ -RANK

### B.1 THE DATA STRUCTURE FOR SPARSITY

The transitional probability matrix  $\mathbf{T}$  in  $\alpha^\alpha$ -Rank is sparse; each row and column in  $\mathbf{T}^{[k]}$  contains  $\sum_{i=1}^N s_i - N + 1$  non-zero elements (see Section 3.2). To fully leverage such sparsity, we design a new data structure (see Fig. 5) for the storage and computation. Compared to standard techniques (e.g., COO, CSR, and CRS<sup>10</sup>) that store (row, column, value) of a sparse vector, our data structure adopts a more efficient protocol that stores (defaults, positions, biases) leading to improvements in computational efficiency, which gives us additional advantages in computational efficiency. We reload the operations for such data structure including *addition*, *scalar multiplication*, *dot product*, *element-wise square root*, *L1 norm*. We show the example of addition in Fig. 5.

<sup>9</sup>Here  $[n] = \{1, \dots, n\}$  and  $[t] = \{0, \dots, t\}$ .

<sup>10</sup><https://docs.scipy.org/doc/scipy/reference/sparse.html>

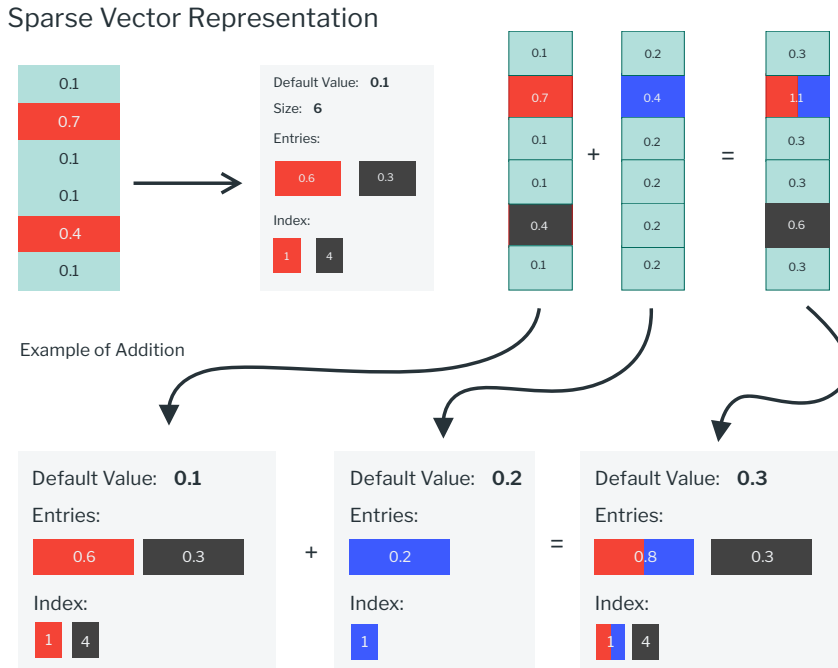


Figure 5: Sparse vector representation in  $\alpha^\alpha$ -Rank.

## B.2 VALIDITY CHECK ON NORMAL-FORM GAMES.

Our algorithm provides the expected ranking in all three normal-form games shown in Fig. 6, which is consistent with the results in  $\alpha$ -Rank (Omidshafiei et al., 2019).

**Battle of sexes.** Battle of sexes is an asymmetric game  $R_{OM} = \begin{bmatrix} 3, 2 & 0, 0 \\ 0, 0 & 2, 3 \end{bmatrix}$ .  $\alpha^\alpha$ -Rank suggests that populations would spend an equal amount of time on the profile (O,O) and (M,M) during the evolution. The distribution mass of (M,O) drops to 0 faster than that of (O,M), this is because deviating from (M,O) for either player has a larger gain (from 0 to 3) than deviating from (O,M) (from 0 to 2).

**Biased Rock-Paper-Scissor.** We consider the biased RPS game  $R_{RPS} = \begin{bmatrix} 0 & -0.5 & 1 \\ 0.5 & 0 & -0.1 \\ -1 & 0.1 & 0 \end{bmatrix}$ . As it is a single-population game, we adopt the transitional probability matrix of Eqn. 11 in Omidshafiei et al. (2019). Such game has the inherent structure that Rock/Paper/Scissor is equally likely to be invaded by a mutant, e.g., the scissor population will always be fixated by the rock population, therefore, our method suggests the long-term survival rate for all three strategies are the same  $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ . Note this is different from the Nash equilibrium solution that is  $(\frac{1}{16}, \frac{5}{8}, \frac{5}{16})$ .

**Prison’s Dilemma.** In prison’s dilemma  $R_{CD} = \begin{bmatrix} -1, -1 & -3, 0 \\ 0, -3 & -2, -2 \end{bmatrix}$ , cooperation is an evolutionary transient strategy since the cooperation player can always be exploited by the defection player. Our method thus yields (D, D) as the only strategy profile that could survive in the long-term evolution.

## C ADDITIONAL DETAILS FOR EXPERIMENTS

### C.1 HYPER-PARAMETERS SETTINGS

For all of our experiments, the gradient updates include two phases: warm-up phase and Adam (Kingma & Ba, 2014) phase. In the warm-up phase, we used standard stochastic gradient descent; after that, we replace SGD with Adam till the convergence. In practice, we find this yields faster convergence than normal stochastic gradient descent. As our algorithm does column sampling for the stochastic matrix (i.e. batch size equals to one), adding momentum term intuitively help stabilize the learning. The warm-up step is 100 for all experiments




Figure 6: Ranking intensity sweep on (a) Battle of Sexes (b) Biased RPS (c) Prisoner’s Dilemma

We also implement infinite  $\alpha$  Lanctot et al. (2019), when calculating transition matrix (or its column), where our noise term is set to be 0.01.

For most of our experiments that involve  $\alpha^\alpha$ -rank, we set the terminating condition to be, when the gradient norm is less than  $10^{-9}$ . However, for Random Matrix experiment, we set the terminating gradient norm to be  $10^{-2}$

- Learning rate to be in between 15 - 17
- Alpha (ranking intensity) to be in between 1 - 2.5
- Number of Population to be between 25 - 55 (in integer)

For all of the Adam experiments, after the warmup-step we chooses to decay  $\delta$  and  $\lambda$  by 0.999 for each time steps, where we have  $\delta$  to always be 0.1. Similarly,  $\lambda$  starts at the value 0.5. However, in speed and memory experiment, we chooses the decay to be 0.9

Experiments	$m$	$\alpha$	$\eta$	Max. Iteration	$\beta_1$	$\lambda$	$\delta$
NFG (without self-play)	50	$10^{-4} \sim 10^2$	1.0	1000	0.9	0.5	0.1
NFG (self-play)	50	$10^{-4} \sim 10^2$	0.03	1000	0.9	0.5	0.1
Random Matrix	n/a	n/a	0.01	1000	n/a	0.1	0.01
Car Experiment (SGD)	40	1.0	15.0	2000	0.999	0.5	0.1
Car Experiment (Oracle)	40	1.0	1.0	200	0.999	0.5	0.1
Ising Model	40	90.0	0.01	4000	0.999	0.5	0.1

List of symbols and names

- Population size:  $m$
- Ranking Intensity:  $\alpha$
- Learning rate:  $\eta$

### C.2 SELF-DRIVING CAR EXPERIMENT

The environmental reward given to each agent is calculated by

$$\text{Collision Reward} + \frac{\text{Current Velocity}}{\text{Default Velocity}} \times \text{Proportional Speed Reward}$$

Collision Reward is calculated when agent collided with either social car or other agents.

All of our value iteration agents are based on Leurent (2018) environment discretization, which represents the environment in terms of time to collision MDP, taking into account that the other agents are moving in constant speed. For all experiments, we run value-iteration for 200 steps with the discounting factor of 0.99.

For each controllable cars, the default speed is randomized to be between 10 to 25, while the social cars, the speed are randomized to be between 23 to 25.

We define five types of driving behaviors (one rational + four dangerous) by letting each controlled car have a different ego reward function during training (though the reward we report is the environmental reward which cannot be changed). By setting this, we can make sure, at upfront, the best joint-strategy strategy should be all cars to drive rationally.

Table 2: Reward settings.

	<b>Collision Reward</b>	<b>Proportional Speed Reward</b>
Environment reward	-2.0	0.4
Rational driver	-2.0	0.4
Dangerous driver 1	10.0	10.0
Dangerous driver 2	20.0	10.0
Dangerous driver 3	30.0	10.0
Dangerous driver 4	40.0	10.0