

α^α -Rank: Practically Scaling α -Rank through Stochastic Optimisation

Yaodong Yang*
Huawei Technologies R&D U.K.
University College London
yaodong.yang@huawei.com

Phu Sakulwongtana
Huawei Technologies R&D U.K.
phu.sakulwongtana@huawei.com

Rasul Tutunov*
Huawei Technologies R&D U.K.
rasul.tutunov@huawei.com

Haitham Bou Ammar
Huawei Technologies R&D U.K.
University College London
haitham.ammar@huawei.com

ABSTRACT

Recently, α -Rank, a graph-based algorithm, has been proposed as a solution to ranking joint policy profiles in large scale multi-agent systems. α -Rank claimed tractability through a polynomial time implementation with respect to the total number of pure strategy profiles. Here, we note that inputs to the algorithm were not clearly specified in the original presentation; as such, we deem complexity claims as not grounded, and conjecture solving α -Rank is NP-hard.

The authors of α -Rank suggested that the input to α -Rank can be an exponentially-sized payoff matrix; a claim promised to be clarified later. Even though α -Rank exhibits a polynomial-time solution with respect to such an input, we further reflect additional critical problems. We demonstrate that due to the need of constructing an exponentially large Markov chain, α -Rank is infeasible beyond a small finite number of agents. We ground these claims by adopting amount of dollars spent as a non-refutable evaluation metric. Realising such scalability issue, we present a stochastic implementation of α -Rank with a double oracle mechanism allowing for reductions in joint strategy spaces. Our method, α^α -Rank, does not need to save exponentially-large transition matrix, and can terminate early under required precision. Although theoretically our method exhibits similar worst-case complexity guarantees compared to α -Rank, it allows us, for the first time, to practically conduct large-scale multi-agent evaluations. On $10^4 \times 10^4$ random matrices, we achieve $1000\times$ speed reduction. Furthermore, we also show successful results on large joint strategy profiles with a maximum size in the order of $O(2^{25})$ (≈ 33 million joint strategies) – a setting not evaluable using α -Rank with reasonable computational budget.

KEYWORDS

α -Rank; complexity analysis; multi-agent evaluation; tractability

ACM Reference Format:

Yaodong Yang, Rasul Tutunov*, Phu Sakulwongtana, and Haitham Bou Ammar. 2020. α^α -Rank: Practically Scaling α -Rank through Stochastic Optimisation. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

*Equal contribution.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Scalable policy evaluation and learning have been long-standing challenges in multi-agent reinforcement learning (MARL) 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 behaviour (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 [21] equilibrium, as potentially valid evaluation metrics [35, 37]. Though appealing, such measures are normative only when prescribing behaviours of perfectly rational agents – an assumption rarely met in reality [9, 16, 33, 34]. In fact, many game dynamics have been proven not converge to any fixed-point equilibria [10, 32, 36], but rather to limit cycles [2, 24]. Apart from these challenges, solving for a Nash equilibrium even for “simple” settings, e.g. two-player games is known to be PPAD-complete [3] – a demanding complexity class when it comes to computational requirements.

To address some of the above limitations, [22] recently proposed α -Rank as a graph-based game-theoretic solution to multi-agent evaluation. α -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. In a novel attempt, the authors reduce multi-agent evaluation to computing a stationary distribution of a Markov chain. Namely, consider a set of N agents each having a strategy pool of size k , a Markov chain is, first, defined over the graph of joint strategy profiles with a transition matrix $T \in \mathbb{R}^{k^N \times k^N}$, and then a stationary distribution $\mathbf{v} \in \mathbb{R}^{k^N}$ is computed solving: $T\mathbf{v} = \mathbf{v}$. The probability mass in \mathbf{v} then represents the ranking of joint-strategy profile.

Extensions of α -Rank have been developed on various instances. [27] adapted α -Rank to model games with incomplete information. [20] combined α -Rank with the policy search space oracle (PSRO) [15] and claimed their method to be a *generalised* training approach for multi-agent learning. Unsurprisingly, these work inherit the same claim of tractability from α -Rank. For example, the abstract in [20] reads “ α -Rank, which is unique (thus faces no equilibrium selection issues, unlike Nash) and *tractable* to compute in general-sum, *many-player* settings.”

In this work, we contribute to refine the claims made in α -Rank dependent on its input type. We thoroughly argue that α -Rank exhibits a prohibitive computational and memory bottleneck that is hard to remedy even if pay-off matrices were provided as inputs. We measure such a restriction using money spent as a non-refutable metric to assess α -Rank’s validity scale. With this in mind, we then present a stochastic solver that we title α^α -Rank as a scalable and memory efficient alternative. Our method reduces memory constraints, and makes use of the oracle mechanism for reductions in joint strategy spaces. This, in turn, allows us to run large-scale multi-player experiments, including evaluation on self-driving cars and Ising models where the maximum size involves tens of millions of joint strategies.

2 A REVIEW OF α -RANK

In α -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, α -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 dominant profile that spreads across the population of agents. Fig. 1 demonstrates a simple example of a three-player game, each playing three strategies.

Mathematical Formulation: To formalise α -Rank, we consider N agents each, denoted by i , having access to a set of strategies of size k_i . We refer to the strategy set for agent i by $\mathcal{S}_i = \{\pi_{i,1}, \dots, \pi_{i,k_i}\}$, $k_i = |\mathcal{S}_i|$, with $\pi_{i,j} : \mathcal{X} \times \mathcal{A}_i \rightarrow [0, 1]$ 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 . A joint strategy profile is a set of policies for all participating agents in the joint strategy set, i.e., $\mathcal{S}_{\text{joint}} = \mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_N$: $\pi_{\text{joint}} = \{\pi_{1,j_1}, \dots, \pi_{N,j_N}\}$, with $\pi_{i,j_i} \in \mathcal{S}_i$ and $j_i \in \{1, \dots, k_i\}$. We assume $k = k_1 = \dots = k_N$ hereafter.

To evaluate performance, we assume each agent is additionally equipped with a payoff (reward) function $\mathcal{P}_i : \mathcal{S}_{\text{joint}} \rightarrow \mathbb{R}_+$. Crucially, the domain of \mathcal{P}_i is the pool of *joint strategies* so as to accommodate the effect of other learners on the i^{th} player’s performance. Finally, given a joint profile π_{joint} , we define the corresponding joint payoff to be the collection of all individual payoff functions, i.e., $\mathcal{P}_{\text{joint}} = \{\mathcal{P}_1(\pi_{\text{joint}}), \dots, \mathcal{P}_N(\pi_{\text{joint}})\}$. After attaining payoffs from the environment, each agent returns to its population and faces a choice between switching the whole population to a mutation policy, exploring a novel policy, or sticking to the current one. Such a choice is probabilistic and defined proportional to rewards by

$$\mathbb{P}(\pi_{i,a} \rightarrow \pi_{i,b}, \boldsymbol{\pi}_{-i}) = \frac{e^{\alpha \mathcal{P}_i(\pi_{i,b}, \boldsymbol{\pi}_{-i})}}{e^{\alpha \mathcal{P}_i(\pi_{i,a}, \boldsymbol{\pi}_{-i})} + e^{\alpha \mathcal{P}_i(\pi_{i,b}, \boldsymbol{\pi}_{-i})}} - \frac{\mu}{2}$$

$$\text{for } (\pi_{i,a}, \pi_{i,b}) \in \mathcal{S}_i \times \mathcal{S}_i,$$

$$\mathbb{P}(\pi_{i,a} \rightarrow \pi_{i,c}, \boldsymbol{\pi}_{-i}) = \frac{\mu}{K_i - 2}, \quad \forall \pi_{i,c} \in \mathcal{S}_i \setminus \{\pi_{i,a}, \pi_{i,b}\},$$

with $\mu \in \mathbb{R}_+$ being an exploration parameter¹, $\boldsymbol{\pi}_{-i} = \boldsymbol{\pi} \setminus \pi_i$ representing policies followed by other agents, and $\alpha \in \mathbb{R}_+$ an ranking-intensity parameter. Large α ensures that the probability that a sub-optimal strategy overtakes a better strategy is close to zero.

As noted in [22], one can relate the above switching process to a random walk on a Markov chain with states defined as elements in $\mathcal{S}_{\text{joint}}$. Essentially, the Markov chain models the *sink strongly connected components* (SSCC) of the *response graph* associated with the game. The response graph of a game is a directed graph where each node corresponds to each joint strategy profile, and directed edges if the deviating player’s new strategy is a better response to that player, and the SSCC of a directed graph are the (group of) nodes with no out-going edges.

Each entry in the transition probability matrix $T \in \mathbb{R}^{|\mathcal{S}_{\text{joint}}| \times |\mathcal{S}_{\text{joint}}|}$ of Markov chain refers to the probability of one agent switching from one policy in a relation to attained payoffs. Consider any two joint strategy profiles π_{joint} and $\hat{\pi}_{\text{joint}}$ that differ in only *one* individual strategy for the i^{th} agent, i.e., there exists a unique agent such that $\pi_{\text{joint}} = \{\pi_{i,a}, \boldsymbol{\pi}_{-i}\}$ and $\hat{\pi}_{\text{joint}} = \{\hat{\pi}_{i,b}, \boldsymbol{\pi}_{-i}\}$ with $\pi_{i,a} \neq \hat{\pi}_{i,b}$, we set $[T]_{\pi_{\text{joint}}, \hat{\pi}_{\text{joint}}} = \frac{1}{\sum_{l=1}^N (k_l - 1)} \rho_{\pi_{i,a}, \hat{\pi}_{i,b}}(\boldsymbol{\pi}_{-i})$, with $\rho_{\pi_{i,a}, \hat{\pi}_{i,b}}(\boldsymbol{\pi}_{-i})$ defining the probability that one copy of agent i with strategy $\pi_{i,a}$ invades the population with all other agents (in that population) playing $\hat{\pi}_{i,b}$. Following [26], for $\mathcal{P}_i(\pi_{i,a}, \boldsymbol{\pi}_{-i}) \neq \mathcal{P}_i(\hat{\pi}_{i,b}, \boldsymbol{\pi}_{-i})$, such a probability is formalised as

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

and $\frac{1}{m}$ otherwise, with m being the size of the population. So far, we presented relevant derivations for the $(\pi_{\text{joint}}, \hat{\pi}_{\text{joint}})$ entry of the state transition matrix when exactly the i^{th} agent differs in exactly one strategy. Having one policy change, however, only represents a subset of allowed variations; 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., $|\pi_{\text{joint}} \setminus \hat{\pi}_{\text{joint}}| \geq 2$. Here, we set² $[T]_{\pi_{\text{joint}}, \hat{\pi}_{\text{joint}}} = 0$. Consequently, the remaining event of self-transitions can be thus written as $[T]_{\pi_{\text{joint}}, \hat{\pi}_{\text{joint}}} = 1 - \sum_{\hat{\pi} \neq \pi_{\text{joint}}} [T]_{\pi_{\text{joint}}, \hat{\pi}}$. Summarising the above three cases, we can write the $(\pi_{\text{joint}}, \hat{\pi}_{\text{joint}})$ ’s entry of the Markov chain’s transition matrix as:

$$[T]_{\pi_{\text{joint}}, \hat{\pi}_{\text{joint}}} = \begin{cases} \frac{1}{\sum_{l=1}^N (k_l - 1)} \rho_{\pi_{i,a}, \hat{\pi}_{i,b}}(\boldsymbol{\pi}_{-i}), & \text{if } |\pi_{\text{joint}} \setminus \hat{\pi}_{\text{joint}}| = 1, \\ 1 - \sum_{\hat{\pi} \neq \pi_{\text{joint}}} [T]_{\pi_{\text{joint}}, \hat{\pi}}, & \text{if } \pi_{\text{joint}} = \hat{\pi}_{\text{joint}}, \\ 0, & \text{if } |\pi_{\text{joint}} \setminus \hat{\pi}_{\text{joint}}| \geq 2, \end{cases} \quad (2)$$

The goal in α -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 time of survival. Formally, such a notion can be easily derived as the limiting vector $\boldsymbol{v} = \lim_{t \rightarrow \infty} [[T]^T]^t \boldsymbol{v}_0$ of our Markov chain when evolving from an initial distribution \boldsymbol{v}_0 .

¹In α -Rank, μ is heuristically set to a small positive constant to ensure at maximum two varying policies per-each population. Theoretical justification can be found in [6].
²This assumption significantly reduces the analysis complexity as detailed in [6].

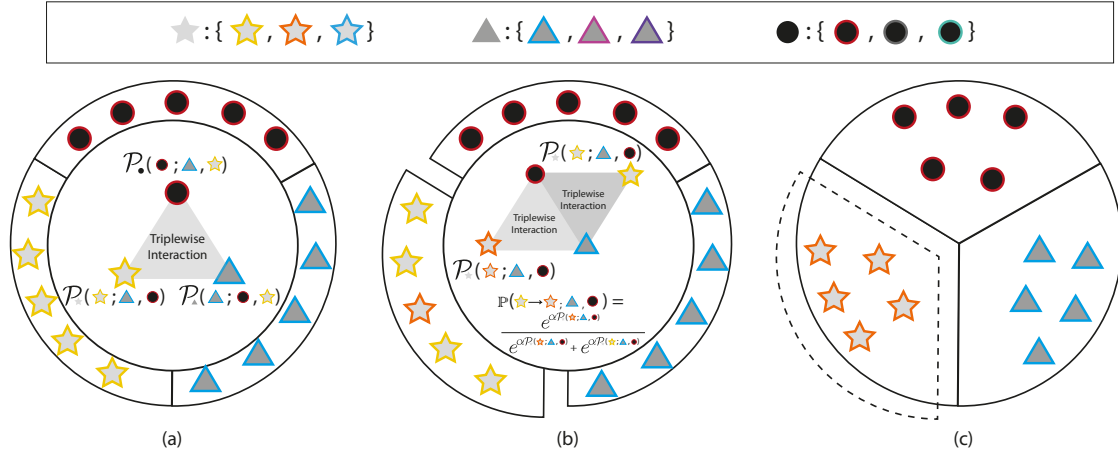


Figure 1: Example of α -Rank evaluation on $N = 3$ players (star, triangle, circle) each with $|s| = 3$ strategies (denoted by the colours) and $m = 5$ copies. a) Each population obtains a fitness value \mathcal{P}_i depending on the strategies chosen, b) one mutation strategy (red star) occurs, and c) the population either selects the original strategy, or being fixated by the mutation strategy.

Algorithm 1 α -Rank (see Section 3.1.1 in [22])

- 1: **(Unspecified) Inputs:** $\mathcal{S}_{\text{joint}}$, Multi-agent Simulator
- 2: Listing all possible joint-strategy profiles, for each profile, run the multi-agent simulator to get the payoff values for all players $\mathcal{P}_{\text{joint}} = \{\mathcal{P}_1(\pi_{\text{joint}}), \dots, \mathcal{P}_N(\pi_{\text{joint}})\}, \forall \pi_{\text{joint}} \in \mathcal{S}_{\text{joint}}$.
- 3: Construct Markov chain's transition matrix T by Eqn. 2.
- 4: Compute the stationary distribution \mathbf{v} by Eqn. 3.
- 5: Rank all π_{joint} in \mathbf{v} based on their probability masses.
- 6: **Outputs:** The ranked list of \mathbf{v} (each element refers to the time that players spend in playing that π_{joint} during evolution).

Knowing that the limiting vector is a stationary distribution, one can calculate the solution to the following eigenvector problem:

$$[T]^T \mathbf{v} = \mathbf{v}. \tag{3}$$

We summarised the pseudo-code of α -Rank in Algorithm 1. As the input to α -Rank is unclear and turns out to be controversial later, we point the readers to the original description in Section 3.1.1 of [22], and the practical implementation of α -Rank from [14] for self-judgement. In what comes next, we demonstrate that the tractability claim of α -Rank needs to be relaxed as the algorithm exhibits exponential time and memory complexities in number of players dependent on the input type considered. This, consequently, renders α -Rank inapplicable to large-scale multi-agent systems contrary to the original presentation.

3 CLAIMS & REFINEMENTS

Original presentation of α -Rank claims to be tractable in the sense that it runs in polynomial time with respect to the total number of joint-strategy profiles. Unfortunately, such a claim is not clear without a formal specification of the inputs to Algorithm 3.1.1 in [22]. In fact, we, next, demonstrate that α -Rank's can easily exhibit exponential complexity under the input of $N \times k$ table, rendering it inapplicable beyond finite small number of players. We also present a conjecture stating that determining the top-rank joint strategy profile in α -Rank is in fact NP-hard.

3.1 On α -Rank's Computational Complexity

Before diving into the details of our arguments, it is first instructive to note that *tractable* algorithms are these that exhibit a worst-case polynomial running time in the size of their input [25]. Mathematically, for a size \mathcal{I} input, a polynomial time algorithm adheres to an $O(\mathcal{I}^d)$ complexity for some constant d independent of \mathcal{I} .

Following the presentation in Section 3.1.1 in [22], α -Rank assumes availability of a game simulator to construct a payoff matrix quantifying performance of joint strategy profiles. As such, we deem that necessary inputs for such a construction is of the size $\mathcal{I} = N \times k$, where N is the total number of agents and k is the total number of strategies per agent, where we assumed $k = k_i = k_j$ for simplicity.

Following the definition above, if α -Rank possesses polynomial complexity then it should attain a time proportional to $O((N \times k)^d)$ with d being a constant independent of N and k . As the algorithm requires to compute a stationary distribution of a Markov chain described by a transition matrix T with k^N rows and columns, the time complexity of α -Rank amounts to $O(k^N)$. Clearly, this result demonstrates exponential, thus intractable, complexity in the number of agent N . In fact, we conjecture that determining top rank joint strategy profile using α -Rank with an $N \times k$ input is NP-hard.

CONJECTURE 3.1 (α -RANK IS NP-HARD). Consider N agents each with k strategies. Computing top-rank joint strategy profile with respect to the stationary distribution of the Markov chain's transition matrix, T , is **NP-hard**.

Reasoning: To illustrate the point of the conjecture above, imagine N agents each with k strategies. Following the certificate argument for determining complexity classes, we ask the question:

"Assume we are given a joint strategy profile π_{joint} , is π_{joint} top rank w.r.t the stationary distribution of the Markov chain?"

To determine an answer to the above question, one requires an evaluation mechanism of some sort. If the time complexity of this mechanism is polynomial with respect to the input size, i.e., $N \times k$, then one can claim that the problem belongs to the NP complexity class. However, if the aforementioned mechanism exhibits an exponential time complexity, then the problem belongs to the NP-hard

Table 1: Time and space complexity comparison given N (number of agents) $\times k$ (number of strategies) as inputs.

Method	Time	Memory
Power Method	$O(k^{N+1}N)$	$O(k^{N+1}N)$
PageRank	$O(k^{N+1}N)$	$O(k^{N+1}N)$
Eig. Decomp.	$O(k^{N\omega})$	$O(k^{N+1}N)$
Mirror Descent	$O(k^{N+1} \log k)$	$O(k^{N+1}N)$

class. When it comes to α -Rank, we believe a mechanism answering the above question would require computing a holistic solution of the problem, which, unfortunately, is exponential (i.e., $O(k^N)$). Crucially, if our conjecture proves correct, we do not see how α -Rank can handle more than a finite small number of agents. \square

3.2 On Optimisation-Based Techniques

Given exponential complexity as derived above, we can resort to approximations of stationary distributions that aim at determining ϵ -close solution for some precision parameter $\epsilon > 0$. Here, we note that a problem of this type is a long-standing classical problem from linear algebra. Various techniques including Power method, PageRank, eigenvalue decomposition, and mirror descent can be utilised. Briefly surveying this literature, we demonstrate that any such implementation (unfortunately) scales exponentially in the number of players. For a quick summary, please consult Table 1.

Power Method. One of the most common approaches to computing a stationary distribution is the power method that computes the stationary vector \mathbf{v} by constructing a sequence $\{\mathbf{v}_j\}_{j \geq 0}$ from a non-zero initialisation \mathbf{v}_0 by applying $\mathbf{v}_{j+1} = \frac{1}{\|\mathbf{r}^T \mathbf{v}_j\|} \mathbf{r}^T \mathbf{v}_j$. 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., $n = k^N$, and m the total number of transitions between the states of the Markov chain. By construction, one can easily see that $m = n(kN - N + 1)$ as each row and column in T contains $kN - N + 1$ non-zero elements. Hence, memory complexity of such implementation is in the order of

$$O(m) = O(n[kN - N + 1]) \approx O(k^N kN).$$

The time complexity of a power method, furthermore, is given by $O(m \times \mathcal{T})$, where \mathcal{T} is the total number of iterations. Since m is of the order $n \log n$, the total complexity of such an implementation is also exponential.

PageRank. Inspired by ranking web-pages on the internet, one can consider PageRank [23] for computing the solution to the eigenvalue problem presented above. Applied to our setting, we first realise that the memory is analogous to the power method that is $O(m) = O(k^{N+1}N)$, and the time complexity are in the order of $O(m + n) \approx O(k^{N+1}N)$.

Eigenvalue Decomposition. Apart from the above, we can also consider the problem as a standard eigenvalue decomposition task (also what is used to implement α -Rank in [14]) and adopt the method in [4] to compute the stationary distribution. Unfortunately, state-of-the-art techniques for eigenvalue decomposition also require exponential memory ($O(k^{N+1}N)$) and exhibit a time

complexity of the form $O(n^\omega) = O(k^{N\omega})$ with $\omega \in [2, 2.376]$ [4]. Clearly, these bounds restrict α -Rank to small agent number N .

Mirror Descent. Another optimisation-based alternative is the ordered subsets mirror descent algorithm [1]. This is an iterative procedure requiring projection step on the standard n -dimensional simplex on every iteration: $\Delta_n = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \mathbf{1} = 1 \ \& \ \mathbf{x} \geq \mathbf{0}\}$. As mentioned in [1], computing this projection requires $O(n \log n)$ time. Hence, the projection step is exponential in the number of agents N . This makes mirror descent inapplicable to α -Rank when N is large.

Apart from the methods listed above, we are aware of other approaches that could solve the leading eigenvector for big matrices, for example the online learning approach [7], the sketching methods [30], and the subspace iteration with Rayleigh-Ritz acceleration [8]. The trade-off of these methods is that they usually assume special structure of the matrix, such as being Hermitian or at least positive semi-definite, which α -Rank however does not fit. Importantly, they can not offer any advantages on the time complexity either.

4 RECONSIDERING α -RANK'S INPUTS

Having discussed our results with the authors, we were suggested that “inputs” to α -Rank are **exponentially-sized** payoff matrices, i.e., assuming line 2 in Algorithm 1 as an input. Though polynomial in an exponentially-sized input, this consideration does not resolve problems mentioned above. In this section, we further demonstrate additional theoretical and practical problems when considering the advised “input” by the authors.

4.1 On the Definition of Agents

α -Rank redefines a strategy to correspond to the agents under evaluation differentiating them from players in the game (see line 4 in Section 3.1.1 and also Fig. 2a in [22]). Complexity results are then given in terms of these “agents”, where tractability is claimed. We would like to clarify that such definitions do not necessarily reflect the true underlying time complexity, whereby without formal input definitions, it is difficult to claim tractability.

To illustrate, consider solving a travelling salesman problem in which a traveller needs to visit a set of cities while returning to the origin following the shortest route. Although it is well-known that the travelling salesman problem is NP-hard, following the line of thought presented in α -Rank, one can show that such a problem reduces to a polynomial time (linear, i.e., tractable) problem in the size of “meta-cities”, which is not a valid claim.

So what are the “meta-cities”, and what is wrong with the above argument?

A strategy in the travelling salesman problem corresponds to a permutation in the order of cities. Rather than operating with number of cities, following α -Rank, we can construct the space of all permutation calling each a “meta-city” (or agent)³. Having enumerated all permutations, somehow, searching for the shortest route can be performed in polynomial time. Even though, one can state that solving the travelling salesman problem is polynomial in the size of permutations, it is incorrect to claim that any such algorithm is tractable.

³How to enumerate all these permutations of cities is analogous to enumerating an exponentially sized matrix in α -Rank if $N \times k$ was not the input to α -Rank.




Figure 2: Money cost of constructing the transition matrix T in computing α -Rank (line 3 in Algorithm 1). Note that one trillion dollar is the world’s total hardware budget [5]. The projected contours shows that due to the exponentially-growing size of α -Rank’s “input”, under reasonable budget, α -Rank is infeasible to handle more than ten agents.

The same exact argument can be made for α -Rank, whereby having a polynomial time algorithm in an exponentially-sized space does not at all imply tractability⁴. It is for this reason, that reporting complexity results needs to be done with respect to the size of the input without any redefinition (i.e. agents in multi-agent systems, and cities in the travelling salesman problem).

As is clear so-far, inputs to α -Rank lack clarity. Confused on the form of the input, we realise that the we are left with two choices: 1) list of all joint strategy profiles, or 2) a table of the size $N \times k$ – collection of all of the players’ strategy pools. If we are to follow the first direction, the claims made in the paper are of course correct; however, this by no means resolves the problem as it is not clear how one would construct such an input in a tractable manner. Precisely, given an $N \times k$ table (collection of all of the players’ strategy pools) as input, constructing the aforementioned list requires exponential time (k^N). In other words, providing α -Rank with such a list only hides the exponential complexity burden in a pre-processing step. Analogously, applying this idea to the travelling salesman problem described above would hide the exponential complexity under a pre-processing step used to construct all possible permutations. Provided as inputs, the travelling salesman problem can now be solved in linear time, i.e., transforming an intractable problem to a tractable one by a mere redefinition.

4.2 Dollars Spent: A Non-Refutable Metric

Admittedly, our arguments have been mostly theoretical and can become controversial dependent on the setting one considers. To abolish any doubts, we followed the advice given by the authors and considered the input of α -Rank to be exponentially-sized payoff

⁴Note that this claim does not apply on the complexity of solving Nash equilibrium. For example, in solving zero-sum games, polynomial tractability is never claimed on the number of players, whereas α -Rank claims tractable in the number of players.

Table 2: Cost of getting the payoff table $\mathcal{P}_{\text{joint}}$ (line 2 in Algorithm 1) for the experiments conducted in Omidshafiei et al. [22]. We list the numbers by the cost of running one joint-strategy profile \times the number of joint-strategy profiles considered. Detailed computation can be found here.

Game Env.	PetaFlop/s-days	Cost (\$)	Time (days)
AlphaZero Go [29]	$1,413 \times 7$	207M	1.9M
AlphaGo Zero [28]	$1,181 \times 7$	172M	1.6M
AlphaZero Chess [29]	17×1	352K	3.2K
MuJoCo Soccer [18]	0.053×10	4.1K	72
Leduc Poker [15]	0.006×9	420	7
Kuhn Poker [11]	$< 10^{-4} \times 256$	< 1	–
AlphaStar [31]	52,425	244M	1.3M

matrices. We then conducted an experiment measuring dollars spent to evaluate scalability of running just line 3 in Algorithm 1, while considering the tasks reported in [22].

Assuming $\mathcal{P}_{\text{joint}} = \{\mathcal{P}_1(\pi_{\text{joint}}), \dots, \mathcal{P}_N(\pi_{\text{joint}})\}$, $\forall \pi_{\text{joint}} \in \mathcal{S}_{\text{joint}}$ is given at no cost, the total amount of floating point operations (FLOPS) needed for constructing T given in Eqn. 2 is $9k^N N(k-1) + k^N N(k-1) + 0 = 10k^N N(k-1)$. In terms of money cost needed for just building T , we plot the dollar amount in Fig. 2 considering the Nvidia Tesla K80 GPU⁵ which can process under single precision at maximum 5.6 TFlop/s at a price of 0.9 \$/hour⁶. Clearly, Fig. 2 shows that due to the fact that α -Rank needs to construct a Markov chain with an exponential size in the number of agents, it is only “money” feasible on tasks with at most tens of agents. It is also worth noting that our analysis is optimistic in the sense that we have not considered costs of storing T nor computing Eqn. 3.

Conclusion I: *Given exponentially-sized payoff matrices, constructing transition matrices in α -Rank for about 20 agents each with 8 strategies requires about one trillion dollars in budget.*

Though assumed given, in reality, the payoff values $\mathcal{P}_{\text{joint}}$ come at a non-trivial cost themselves, which is particularly true in reinforcement learning tasks [28]. Here, we take a closer look at the amount of money it takes to attain payoff matrices for the experiments listed in [22] that we present in Table 2. Following the methodology in here, we first count the total FLOPS each model uses under the unit of PetaFlop/s-day that consists of performing 10^{20} operations per second in one day. For each experiment, if the answer to “how many GPUs were trained and for how long” was not available, we then traced back to the neural architecture used and counted the operations needed for both forward and backward propagation. The cost in time was then transformed from PetaFlop/s-day using Tesla K80 as discussed above. In addition, we also list the cost of attaining payoff values from the most recent AlphaStar model [31]. It is obvious that although α -Rank could take the payoff values as “input” at a hefty price, the cost of acquiring such values is not negligible, e.g., payoff values from GO cost about 207M \$, and require a single GPU to run for more than five thousand years⁷!

Conclusion II: *Acquiring necessary inputs to α -Rank easily becomes intractable giving credence to our arguments in Section 4.1.*

⁵https://en.wikipedia.org/wiki/Nvidia_Tesla

⁶<https://aws.amazon.com/ec2/instance-types/p2/>

⁷In practice, the game outcomes are noisy, multiple runs are often needed (check Theorem 3.2 in [27]), which will turn the numbers in Table 2 to an even larger scale.

5 A PRACTICAL SOLUTION TO α -RANK

One can consider approximate solutions to the problem in Eqn. 3. As briefly surveyed in Section 3, most current methods, unfortunately, require exponential *time and memory* complexities. We believe achieving a solution that aims at reducing time complexity is an interesting and open question in linear algebra in general, and leave such a study to future work. Here, we rather contribute by a stochastic optimisation method that can attain a solution through random sampling of payoff matrices without the need to store exponential-size input. Contrary to memory requirements reported in Table 1, our method requires a linear (in number of agents) per-iteration complexity of the form $\mathcal{O}(Nk)$. It is worth noting that most other techniques need to store exponentially-sized matrices before commencing with any numerical instructions. Though we do not theoretically contribute to reductions in time complexities, we do, however, augment our algorithm with a double-oracle heuristic for joint strategy space reduction. In fact, our experiments reveal that α^α -Rank can converge to the correct top-rank strategies in hundreds of iterations in large strategy spaces, i.e., spaces with ≈ 33 million profiles.

Optimisation Problem Formulation: Computing the stationary distribution can be rewritten as an optimisation problem:

$$\min_{\mathbf{x}} \frac{1}{n} \|T^T \mathbf{x} - \mathbf{x}\|_2^2 \quad \text{s.t. } \mathbf{x}^T \mathbf{1} = 1, \text{ and } \mathbf{x} \geq 0, \quad (4)$$

where the constrained objective in Eqn. 4 simply seeks a vector \mathbf{x} minimising the distance between \mathbf{x} , itself, and $T^T \mathbf{x}$ while ensuring that \mathbf{x} lies on an n -dimensional. To handle exponential complexities needed for acquiring exact solutions, we pose a relaxation the problem in Eqn. 4 and focus on computing an approximate solution vector $\tilde{\mathbf{x}}$ instead, where $\tilde{\mathbf{x}}$ solves:

$$\min_{\mathbf{x}} \frac{1}{n} \|T^T \mathbf{x} - \mathbf{x}\|_2^2 \quad \text{s.t. } |\mathbf{x}^T \mathbf{1} - 1| \leq \delta \text{ for } 0 < \delta < 1, \text{ and } \mathbf{x} \geq 0. \quad (5)$$

Before proceeding, however, it is worth investigating the relation between the solutions of the original (Eqn. 4) and relaxed (Eqn. 5) problems. We summarise such a relation in the following proposition that shows that determining $\tilde{\mathbf{x}}$ suffices for computing a stationary distribution of α -Rank's Markov chain:

PROPOSITION 5.1 (CONNECTIONS TO MARKOV CHAIN). Let $\tilde{\mathbf{x}}$ be a solution to the relaxed optimisation problem in Eqn. 5. Then, $\tilde{\mathbf{x}} / \|\tilde{\mathbf{x}}\|_1 = \mathbf{v}$ is the stationary distribution of Eqn. 3 in Section 2.

Importantly, the above proposition, additionally, allows us to focus on solving the problem in Eqn. 5 that only exhibits inequality constraints. Problems of this nature can be solved by considering a barrier function leading to an unconstrained finite sum minimisation problem. By denoting \mathbf{b}_i to be the i^{th} row of $T^T - I$, we can, thus, write: $\frac{1}{n} \|T^T \mathbf{x} - \mathbf{x}\|_2^2 = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}^T \mathbf{b}_i)^2$. Introducing logarithmic barrier-functions, with $\lambda > 0$ being a penalty parameter, we have:

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

Eqn. 6 represents a standard finite minimisation problem, which can be solved using any off-the-shelf stochastic optimisation methods, e.g., stochastic gradients, ADAM [13]. A stochastic gradient execution involves sampling a strategy profile $i_t \sim [1, \dots, n]$ at iteration t , and then executing a descent step: $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \nabla_{\mathbf{x}} f_{i_t}(\mathbf{x}_t)$,

Algorithm 2 α^α -Oracle: Practical Multi-Agent Evaluation

```

1: Inputs: Number of trails  $\mathcal{N}$ , total number of iterations  $T$ , decaying learning rate  $\{\eta_t\}_{t=1}^T$ , penalty parameter  $\lambda$ ,  $\lambda$  decay rate  $\gamma > 1$ , and a constraint relaxation term  $\delta$ , initialise  $p = 0$ .
2: while  $p \leq \mathcal{N}$  do:
3:   Set the counter of running oracles,  $k = 0$ 
4:   Initialise the strategy set  $\{\mathcal{S}_i^{[0]}\}$  by sub-sampling from  $\{\mathcal{S}_i\}$ 
5:   while  $\{\mathcal{S}_i^{[k]}\} \neq \{\mathcal{S}_i^{[k-1]}\}$  AND  $k \geq 1$  do:
6:     Compute total number of joint profiles  $n = \prod_{i=1}^N |\mathcal{S}_i^{[k]}|$ 
7:     Initial a vector  $\mathbf{x}_0 = \mathbf{1}/n$ 
8:     for  $t = 0 \rightarrow T - 1$  do: //  $\alpha^\alpha$ -Rank update
9:       Uniformly sample one strategy profile  $i_t^{[k]} \sim \{1, \dots, n\}$ 
10:      Construct  $\mathbf{b}_{i_t}^{[k]}$  as the  $i_t^{[k]}$  row of  $T^{[k], T} - I$ 
11:      Update  $\mathbf{x}_{t+1}^{[k]} = \mathbf{x}_t^{[k]} - \eta_t \nabla_{\mathbf{x}} f_{i_t^{[k]}}(\mathbf{x}_t^{[k]})$  by Eqn. 7
12:      Set  $\lambda_{t+1} = \lambda_t / \gamma$ 
13:      Get  $\pi_{\text{joint}}^{[p], \text{top}}$  by ranking the prob. mass of  $\mathbf{v}^{[k]} = \frac{\mathbf{x}_t^{[k]}}{\|\mathbf{x}_t^{[k]}\|_1}$ 
14:      Set  $k = k + 1$ 
15:      for each agent  $i$  do: // The oracles (Section 5.1)
16:        Compute the best response  $\pi_i^*$  to  $\pi_{\text{joint}}^{[p], \text{top}}$  by Eqn. 8
17:        Update the strategy set by  $\mathcal{S}_i^{[k]} = \mathcal{S}_i^{[k-1]} \cup \pi_i^*$ 
18:      Set  $p = p + 1$ 
19: Return: The best performing joint-strategy profile  $\pi_{\text{joint}, \star}$  among  $\{\pi_{\text{joint}}^{[1: \mathcal{N}], \text{top}}\}$ .
```

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

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

To avoid any confusion, we name the above stochastic approach of solving α -Rank via Eqn. 6 & 7 as α^α -Rank and present its pseudo-code in Algorithm 2. When comparing our algorithm to these reported in Table 1, it is worth highlighting that computing updates using Eqn. 7 requires no storage of the full transition or payoff matrices as updates are performed only using sub-sampled columns as shown in line 11 in Algorithm 2.

5.1 Efficient Exploration via Oracles

Stochastic sampling enables to solve α -Rank with no need to store the transition matrix T ; however, the size of the column \mathbf{b}_i (i.e., $\prod_{i=1}^N k_i$) can still be prohibitively large. Here we further boost scalability of our method by introducing an *oracle* mechanism. The heuristic of oracles was first proposed in solving large-scale zero-sum matrix games [19]. The idea is to first create a sub-game in which all players are only allowed to play a restricted number of strategies, which are then expanded by adding each of the players' best-responses to their opponents; the sub-game will be replayed with agents' augmented strategy sets before a new round of best responses is computed.

The best response is assumed to be given by an *oracle* that can be simply implemented by a grid search, where given the top-rank profile π_{-i}^{top} at iteration k , the goal for agent i is to select the optimal

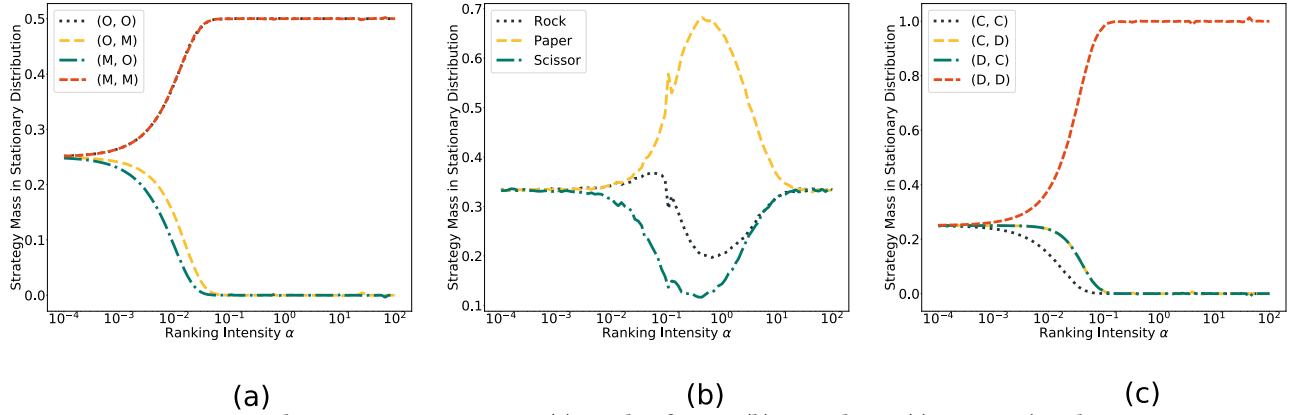


Figure 3: Ranking intensity sweep on (a) Battle of Sexes (b) Biased RPS (c) Prisoner’s Dilemma.




Figure 4: Comparisons of time and memory complexities on varying sizes of random matrices.

π_i^* from a pre-defined strategy set \mathcal{S}_i to maximise its reward:

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

with \mathbf{x}_h denoting the state, $\mathbf{u}_{i,h} \sim \pi_i(\cdot | \mathbf{x}_{i,h})$, $\mathbf{u}_{-i,h} \sim \pi_{-i}^{\text{top}}(\cdot | \mathbf{x}_{-i,h})$ denoting the actions from agent i and the opponents, respectively. Though worse-case scenario of introducing oracles would require solving the original evaluation problem, our experimental results on large-scale systems demonstrate efficiency by converging early.

For a complete exposition, we summarise the pseudo-code of our proposed method, named as α^α -Oracle, in Algorithm 1. α^α -Oracle degenerates to α^α -Rank (lines 6 – 13) if one initialises strategy sets of agents by the full size at the beginning, i.e., $\{\mathcal{S}_i^{[0]}\} \triangleq \{\mathcal{S}_i\}$.

Providing valid convergence guarantee for α^α -Oracle is an interesting direction for future work. In fact, recently [20] proposed a close idea of adopting an oracle mechanism into α -Rank without any stochastic solver however. Interestingly, it is reported that bad initialisation can lead to failures in recovering top-rank strategies. Contrary to the results reported in [20], we rather demonstrate the effectiveness of our approach through running multiple trails of initialisation for $\{\mathcal{S}_i^{[0]}\}$. In addition, we also believe the stochastic nature of α^α -Oracle potentially prevents from being trapped by the local minimal from sub-games.

6 EXPERIMENTS

In this section, we demonstrate the scalability of α^α -Rank in successfully recovering optimal policies in self-driving car simulations

and in the Ising model—a setting with tens-of-millions of possible strategies. We note that these sizes are far beyond the capability of state-of-the-art methods; α -Rank [22] considers at maximum 4 agents with 4 strategies. All of our experiments were run only on a single machine with 64 GB memory and 10-core Intel i9 CPU.

Sparsity Data Structure: 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 $\sum_{i=1}^N k_i^{[k]} - N + 1$ non-zero elements, check Section 5) that if exploited can lead to significant speed-up. To fully leverage such sparsity, we tailored a novel data structure for sparse storage and computations needed by Algorithm 2. More details are in Appendix.

Correctness of Ranking Results: As Algorithm 2 is a generalisation (in terms of scalability) of α -Rank, it is instructive to validate the correctness of our results on three simple matrix games. Due to space constraints, we refrain the full description of these tasks to Appendix. Fig. 3, however, shows that, in fact, results generated by α^α -Rank are consistent with these reported in [22].

Complexity Comparisons on Random Matrices: To further assess scalability, we measured the time and memory needed by our method for computing stationary distributions of varying sizes of simulated random matrices. Baselines included eigenvalue decomposition from Numpy, optimisation tools from PyTorch, and α -Rank from OpenSpiel [14]. We terminated execution of α^α -Rank when gradient norms fell-short a predefined threshold of **0.01**. According to Fig. 4, α^α -Rank can achieve three orders of magnitude reduction in time (i.e. 1000x faster) compared to default α -Rank

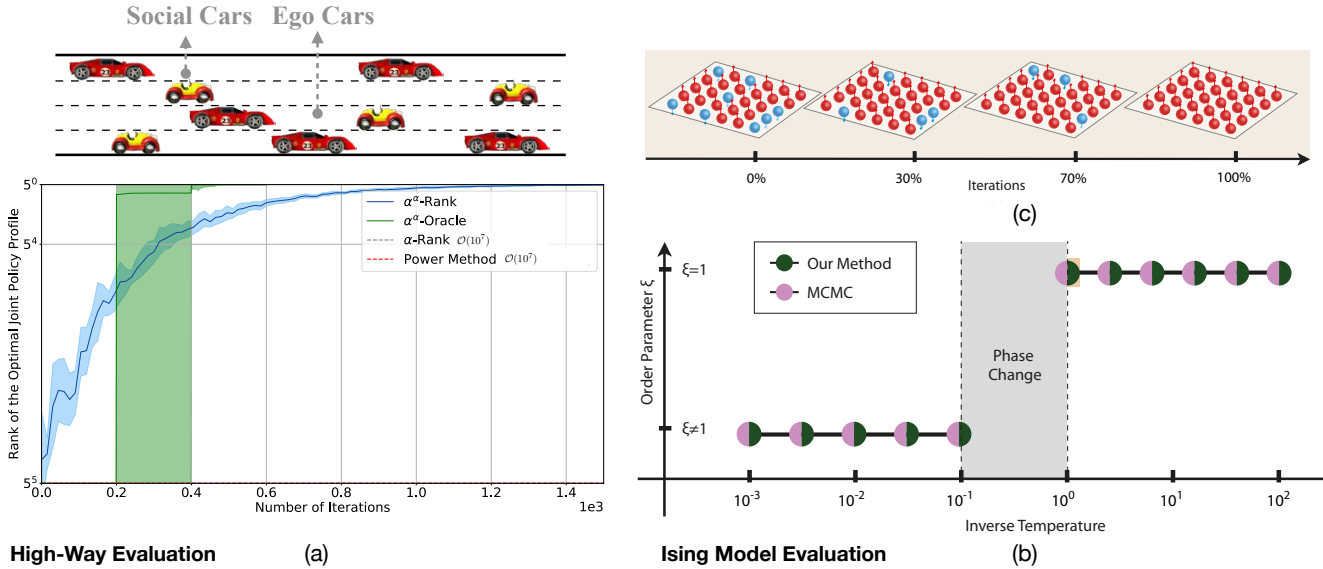


Figure 5: Large-scale multi-agent evaluations. (a) Convergence of the optimal joint-strategy profile in self-driving simulation. (b) Status of the Ising-model equilibrium measured by $\xi = \frac{|N_1 - N_{\downarrow}|}{|N|}$. (c) Change of the top-rank profile from α^α -Oracle.

implementation from [14]. Memory-wise, our method uses only half of the space when considering, for instance, $10^4 \times 10^4$ matrices.

Autonomous Driving on Highway: Having assessed correctness and scalability, we now present novel application domains on large-scale *multi-agent/multi-player* systems. For that we made used of high-way [17]; an environment for simulating self-driving scenarios with social vehicles designed to mimic real-world traffic flow. We conducted a ranking experiment involving 5 agents each with 5 strategies, i.e., a strategy space in the order of $O(5^5)$ (3125 possible strategy profiles). Agent strategies varied between “rational” and “dangerous” drivers, which we encoded using different reward functions during training (complete details of reward functions are in Appendix). Under this setting, we knew, upfront, that optimal profile corresponds to all agents being five rational drivers.

We considered both α^α -Rank and α^α -Oracle, and reported the results by running 1000 random seeds. We set α^α -Oracle to run 200 iterations of gradient updates in solving the top-rank strategy profile (lines 8 – 12 in Algorithm 2). Results depicted in Fig. 5(a) clearly demonstrate that both our proposed methods are capable of recovering the correct highest ranking strategy profile. α^α -Oracle converges faster than α^α -Rank, which we believe is due to the oracle mechanism saving time in inefficiently exploring “dangerous” drivers upon one observation. We also note that although such size of problem are feasible using α -Rank and the Power Method, our results achieve 4 orders of reduction in number of iterations.

Ising Model Experiment: We repeated the above experiment on the Ising model [12] that is typically used 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 $E(\mathbf{a}, \mathbf{h}) = -\sum_j (h^j a^j + \frac{\lambda}{2} \sum_{k \neq j} a^j a^k)$ with h^j and λ being constant coefficients. The probability of one spin configuration is $P(\mathbf{a}) = \frac{\exp(-E(\mathbf{a}, \mathbf{h})/\tau)}{\sum_{\mathbf{a}} \exp(-E(\mathbf{a}, \mathbf{h})/\tau)}$ where τ 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 τ , 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 $r^j = h^j a^j + \frac{\lambda}{2} \sum_{k \neq j} a^j a^k$. We consider the top-rank strategy profile from α^α -Oracle as the system equilibrium and compare it against the ground truth from MCMC. We consider a 5×5 2D model which induces a prohibitively-large strategy space of the size 2^{25} (≈ 33 million strategies) to which existing methods are inapplicable. Fig. 5(b) illustrates that our method identifies the same phase change as that of MCMC. We also show an example of how α^α -Oracle’s top-ranked profile finds the system’s equilibrium when $\tau = \alpha = 1$ in Fig. 5(c). Note that the problem of 25 agent with 2 strategies goes far beyond the capability of α -Rank on one single machine (billions of elements in T).

7 CONCLUSIONS & FUTURE WORK

In this paper, we presented major bottlenecks prohibiting α -Rank from scaling beyond tens of agents. Dependent on the type of input, α -Rank’s time and memory complexities can easily become exponential. We further argued that notions introduced in α -Rank can lead to confusing tractability results on notoriously difficult NP-hard problems. To eradicate any doubts, we empirically validated our claims by presenting dollars spent as a non-refutable metric.

Realising these problems, we proposed a scalable alternative for multi-agent evaluation based on stochastic optimisation and double oracles, along with rigorous scalability results on a variety of benchmarks. For future work, we plan to understand the relation between α -Rank’s solution and that of a Nash equilibrium. Second, we will attempt to conduct a theoretical study on the convergence of our proposed α^α -Oracle algorithm.

REFERENCES

- [1] Aharon Ben-Tal, Tamar Margalit, and Arkadi Nemirovski. 2001. The ordered subsets mirror descent optimization method with applications to tomography. *SIAM Journal on Optimization* 12, 1 (2001), 79–108.
- [2] Michael Bowling and Manuela Veloso. 2001. Convergence of gradient dynamics with a variable learning rate. In *ICML*. 27–34.
- [3] X Chen and X Deng. 2005. *Settling the complexity of 2-player Nash equilibrium*. ECCS. Technical Report. Report.
- [4] Don Coppersmith and Shmuel Winograd. 1990. Matrix multiplication via arithmetic progressions. *Journal of symbolic computation* 9, 3 (1990), 251–280.
- [5] Statista Research Department. [n. d.]. Forecast information technology (IT) hardware spending worldwide from 2013 to 2019 (in billion U.S. dollars). ([n. d.]). <https://www.statista.com/statistics/422802/hardware-spending-forecast-worldwide/>
- [6] Drew Fudenberg and Lorenz A Imhof. 2006. Imitation processes with small mutations. *Journal of Economic Theory* 131, 1 (2006), 251–262.
- [7] Dan Garber, Elad Hazan, and Tengyu Ma. 2015. Online Learning of Eigenvectors.. In *ICML*. 560–568.
- [8] Gene H Golub and Henk A Van der Vorst. 2000. Eigenvalue computation in the 20th century. *J. Comput. Appl. Math.* 123, 1-2 (2000), 35–65.
- [9] Jordi Grau-Moya, Felix Leibfried, and Haitham Bou-Ammar. 2018. Balancing two-player stochastic games with soft q-learning. *arXiv preprint arXiv:1802.03216* (2018).
- [10] Sergiu Hart and Andreu Mas-Colell. 2003. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review* 93, 5 (2003), 1830–1836.
- [11] Johannes Heinrich, Marc Lanctot, and David Silver. 2015. Fictitious self-play in extensive-form games. In *International Conference on Machine Learning*. 805–813.
- [12] Ernst Ising. 1925. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei* 31, 1 (1925), 253–258.
- [13] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [14] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. 2019. OpenSpiel: A Framework for Reinforcement Learning in Games. *arXiv preprint arXiv:1908.09453* (2019).
- [15] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *Advances in Neural Information Processing Systems*. 4190–4203.
- [16] Felix Leibfried, Jordi Grau-Moya, and Haitham Bou-Ammar. 2017. An Information-Theoretic Optimality Principle for Deep Reinforcement Learning. *CoRR* abs/1708.01867 (2017). [arXiv:1708.01867](http://arxiv.org/abs/1708.01867) <http://arxiv.org/abs/1708.01867>
- [17] Edouard Leurent. 2018. An Environment for Autonomous Driving Decision-Making. <https://github.com/eleurent/highway-env>. (2018).
- [18] Siqi Liu, Guy Lever, Josh Merel, Saran Tunyasuvunakool, Nicolas Heess, and Thore Graepel. 2019. Emergent coordination through competition. *arXiv preprint arXiv:1902.07151* (2019).
- [19] H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. 2003. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*. 536–543.
- [20] Paul Muller, Shayegan Omidshafiei, Mark Rowland, Karl Tuyls, Julien Perolat, Siqi Liu, Daniel Hennes, Luke Marris, Marc Lanctot, Edward Hughes, Zhe Wang, Guy Lever, Nicolas Heess, Thore Graepel, and Remi Munos. 2019. A Generalized Training Approach for Multiagent Learning. (2019). [arXiv:cs.MA/1909.12823](https://arxiv.org/abs/1909.12823)
- [21] John F Nash et al. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences* 36, 1 (1950), 48–49.
- [22] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. 2019. α -Rank: Multi-Agent Evaluation by Evolution. *Scientific Reports, Nature* (2019).
- [23] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [24] Gerasimos Palaioupanos, Ioannis Panageas, and Georgios Piliouras. 2017. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems*. 5872–5882.
- [25] Christos H Papadimitriou. 2003. *Computational complexity*. John Wiley and Sons Ltd.
- [26] Mark Pinsky and Samuel Karlin. 2010. *An introduction to stochastic modeling*. Academic press.
- [27] Mark Rowland, Shayegan Omidshafiei, Karl Tuyls, Julien Perolat, Michal Valko, Georgios Piliouras, and Remi Munos. 2019. Multiagent Evaluation under Incomplete Information. (2019). [arXiv:cs.MA/1909.09849](https://arxiv.org/abs/1909.09849)
- [28] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484.
- [29] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, et al. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815* (2017).
- [30] Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. 2017. Practical sketching algorithms for low-rank matrix approximation. *SIAM J. Matrix Anal. Appl.* 38, 4 (2017), 1454–1485.
- [31] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* (2019), 1–5.
- [32] Yannick Viossat. 2007. The replicator dynamics does not lead to correlated equilibria. *Games and Economic Behavior* 59, 2 (2007), 397–407.
- [33] Ying Wen, Yaodong Yang, Rui Lu, and Jun Wang. 2019. Multi-Agent Generalized Recursive Reasoning. *arXiv preprint arXiv:1901.09216* (2019).
- [34] Ying Wen, Yaodong Yang, Rui Luo, Jun Wang, and Wei Pan. 2019. Probabilistic recursive reasoning for multi-agent reinforcement learning. *arXiv preprint arXiv:1901.09207* (2019).
- [35] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean Field Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*. 5571–5580.
- [36] Yaodong Yang, Lantao Yu, Yiwei Bai, Ying Wen, Weinan Zhang, and Jun Wang. 2018. A study of ai population dynamics with million-agent reinforcement learning. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2133–2135.
- [37] Haifeng Zhang, Weizhe Chen, Zeren Huang, Minne Li, Yaodong Yang, Weinan Zhang, and Jun Wang. 2019. Bi-level Actor-Critic for Multi-agent Coordination. *arXiv preprint arXiv:1909.03510* (2019).