

# Learning to Identify Top Elo Ratings: A Dueling Bandits Approach

Xue Yan<sup>1,2</sup>, Yali Du<sup>\* 3</sup>, Binxin Ru<sup>4</sup>, Jun Wang<sup>5</sup>, Haifeng Zhang<sup>1,2</sup>, Xu Chen<sup>6</sup>

<sup>1</sup> Institute of Automation, Chinese Academy of Sciences, China

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences, China

<sup>3</sup> Department of Informatics, King's College London, UK

<sup>4</sup> Machine Learning Research Group, University of Oxford, UK

<sup>5</sup> Department of Computer Science, University College London, UK

<sup>6</sup> Gaoling School of Artificial Intelligence, Renmin University of China, China

yanxue2021@ia.ac.cn, yali.du@kcl.ac.uk, robin@robots.ox.ac.uk,  
jun.wang@cs.ucl.ac.uk, haifeng.zhang@ia.ac.cn, xu.chen@ruc.edu.cn

## Abstract

The Elo rating system is widely adopted to evaluate the skills of (chess) game and sports players. Recently it has been also integrated into machine learning algorithms in evaluating the performance of computerised AI agents. However, an accurate estimation of the Elo rating (for the top players) often requires many rounds of competitions, which can be expensive to carry out. In this paper, to improve the sample efficiency of the Elo evaluation (for top players), we propose an efficient online match scheduling algorithm. Specifically, we identify and match the top players through a dueling bandits framework and tailor the bandit algorithm to the gradient-based update of Elo. We show that it reduces the per-step memory and time complexity to constant, compared to the traditional likelihood maximization approaches requiring  $O(t)$  time. Our algorithm has a regret guarantee of  $\tilde{O}(\sqrt{T})$ , sublinear in the number of competition rounds and has been extended to the multidimensional Elo ratings for handling intransitive games. We empirically demonstrate that our method achieves superior convergence speed and time efficiency on a variety of gaming tasks.

## Introduction

In this paper, we investigate the selection of best multi-agent strategies under the Elo rating systems. The evaluation of the competition outcome has received lots of attention, especially in view of the successful usage of reinforcement learning in StarCraft (Vinyals et al. 2019; Han et al. 2019; Du et al. 2019), Game of Go (Silver et al. 2017) and video games (Mnih et al. 2015). The Elo rating system (Elo 1978) is a predominant and valuable algorithm for evaluating and ranking agents. In the widely adopted Bradley-Terry model (Hunter et al. 2004) for Elo, each player is assigned a numerical rating which is updated with competition outcomes via online stochastic gradient descent. Further, for dealing with non-transitive relations between interacting agents such as the game of *Rock-Paper-Scissors*, Balduzzi et al. (2018)

proposes multidimensional Elo (mElo), which decomposes a game into transitive and cyclic parts to handle intransitive skills and evaluates different strategies by computing Nash-averaging.

In practical settings when a competition is expensive to conduct, updating Elo rating in a sample efficient way is highly valuable. To achieve such sample efficiency, we need a way to select the most informative pairs for evaluation. Two popular sampling approaches are Round-robin (Rasmussen and Trick 2008) and Elimination tournament (Groh et al. 2012); The Round-robin (Rasmussen and Trick 2008) is widely used in sport scheduling to balance the total time, venue usage and fairness of tournaments. It would arrange each team to play against all the others in as few as possible days while satisfying some constraints such as each team not playing twice in the same day to promote game fairness. By contrast, the Elimination tournament (Groh et al. 2012) only allows the winners at each round to proceed to the next round, so the stronger team will have the chance to play more times. A recent approach, RG-UCB, Rowland et al. (2019) introduces an adaptive sampling scheme to estimate the accurate ranking among all agents. RG-UCB considers sampling of agent match-ups as a collection of pure exploration bandit problems (Bubeck, Munos, and Stoltz 2011) and requires enough pairwise comparison for estimating each pair of strategies.

However, these tournament matching/sampling methods suffer from two major limitations which prohibit their wide usage in the modern large scale evaluations. Firstly, both the Round-robin and the Elimination tournament organise competitions following a pre-designed schedule, and the Elimination tournament scheduling may need some prior knowledge on the players' skill. Also each pair of players only compete once in both schemes so the results can be highly noisy. Secondly, the main idea behind the matching schemes of the RG-UCB and the Round-robin is random sampling, which fails to pay more attention to more promising players and/or pairs with higher uncertainty in competition outcome. Thus, they are less sample efficient in identifying the best players.

\*Corresponding to Yali Du (yali.du@kcl.ac.uk).

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this work, we propose two sampling algorithms, named MaxIn-Elo and MaxIn-mElo, for the update of Elo and mElo rating systems respectively. Specifically, we maintain a candidate set with promising players using UCB-based (Upper Confidence Bound) dueling bandits and then select the pair with the highest uncertainty in competition outcome at each round. On the one hand, our algorithms are adapted to the gradient-based update of Elo rating systems, thus more memory and time efficient compared to a prior work, MaxInP (Saha and Gopalan 2020), which relies on maximum likelihood estimation (MLE). On the other hand, we extend our method to update mElo (multidimensional Elo) ratings to handle intransitive games, while Saha and Gopalan (2020) is based on a generalized linear model and can only fit to the transitive games. To the best of our knowledge, this is the first work that enables online gradient-based update for dueling bandits, and a theoretical guarantee on the cumulative regret is provided. Also compared to a previous dueling bandit method descending through randomly sampled gradients at each time step (Yue and Joachims 2009), our method, by selecting the pairs with higher information gains from a set of top candidates, is more sample efficient and are guaranteed to converge at  $\tilde{O}(\sqrt{T})^1$ .

In summary, our contributions are three-fold: Firstly, we are the first to propose two online active sampling algorithms MaxIn-Elo and MaxIn-mElo that select maximum informative pairs with dueling bandits to update Elo and mElo ratings. Secondly, we give the regret analysis of our proposed MaxIn-Elo and show that the regret converge at  $\tilde{O}(\sqrt{T})$ . Thirdly, we demonstrate empirically on synthetic and real-world games that our algorithms achieve significantly lower cumulative regret than all baselines. Notably, our methods outperform MaxInP which uses maximum likelihood for more accurate estimation while with lower time and memory complexity.

## Related Work

Multi-agent evaluation has attracted wide attention in ranking of players (Silver et al. 2017; Lai 2015; Arneson, Hayward, and Henderson 2010; Gruslys et al. 2018) and in selecting stronger strategies in meta games (Muller et al. 2020; Czarnecki et al. 2020). There are many methods used for multi-agent evaluation problem. The Elo rating system is widely used for two-player games such as chess and tennis. It increases (decreases) player’s rating according to player wins (loss) a competition, and updates ratings by online stochastic gradient descent (SGD), which is computationally efficient and simple to implement. While the Elo rating system cannot handle intransitive games such as rock-paper-scissors, multidimensional Elo (Melo) (Balduzzi et al. 2018) was introduced. It decomposes the win-loss matrix of an intransitive game into the transitive component and cyclic component baked in Hodge decomposition theory (Jiang et al. 2011).  $\alpha$ -rank (Omidshafiei et al. 2019) is another popular counterpart in tackling intransitive games; recent attempts improve its sample efficiency based on noisy

comparisons (Du et al. 2021; Rowland et al. 2019; Omidshafiei et al. 2019) and scalability by stochastic optimization (Yang et al. 2020). Despite various evaluation algorithms discussed, how to sample agent pairs at each round is of high value to realize these algorithms in large-scale evaluation tasks.

We consider dueling bandits for online match scheduling in the evaluation of players. The concept of dueling bandits was firstly proposed in (Yue and Joachims 2009). Compared to traditional bandits algorithms which pull one arm at each round and receive the reward of this arm directly, dueling bandits pull arm-pair at each round and only get a binary comparison result. DBGD (Yue and Joachims 2009) models a convex optimization problem as the dueling bandits problem which aims to find the best point in a convex space, and DBGD uses a random gradient as the direction of exploration for selecting the next arm-pair. Yue et al. (2012) formulates the best player identification as a dueling bandits problem with noisy comparison results and an underlying winning probability matrix, and proposes two algorithms as well as their corresponding regret bounds. These algorithms identify best arms based on the observed binary feedback however do not learn the player’s skills (ratings), which is helpful in predicting future competition outcomes. Szörényi et al. (2015) regards the ranking of  $M$  alternatives (e.g. human players or agents) as a dueling bandits problem. They introduce the confidence interval of rating or winning probability into the dueling bandits problem, and design algorithms to identify the close-to-optimal item or to obtain the close-to-optimal whole ranking respectively. Saha, Koren, and Mansour (2021) studies the adversarial setting, in which the winning probability is non-stationary because players’ skill may change over time. And they measure arms’ abilities by estimating Borda score, however Borda score does not possess predictive power of future competition results and the algorithm of (Saha, Koren, and Mansour 2021) estimates Borda score only by simply calculating the frequency of wins. Heckel et al. (2019) gives an active ranking algorithm that can solve the top- $k$  player identification problem and find the entire sequential ranking among all players.

While existing algorithms could rely on Borda score to update the rankings or design specific sort algorithms to obtain the ranking of items, they are not suitable for the Elo rating systems that adopt stochastic gradient descent to update ratings. Earlier attempt (Ding, Hsieh, and Sharpnack 2021) proposes SGD-TS for contextual bandit problem, which learns parameters in the generalized linear model through online SGD instead and employs Thompson Sampling (TS) (Thompson 1933; Agrawal and Goyal 2012, 2013) to encourage exploration in arm-pulling. Compared to UCB-GLM (Li, Lu, and Zhou 2017) that adopt maximum likelihood estimators, SGD-TS achieves a similar theoretical cumulative regret bound, but lower time and memory complexity. However, no prior work has studied the SGD update in dueling bandits setting.

In this work, we will tame dueling bandits for the online match scheduling in the Elo rating system that adopts SGD. Saha and Gopalan (2020) proposes the algorithm that selects Maximum-Informative-Pair (MaxInP) for  $K$ -armed context-

<sup>1</sup> $\tilde{O}$  ignores poly-logarithmic factors

tual dueling bandits. This algorithm utilizes the MLE to estimate parameter  $\hat{\theta}$  and uses UCB (Auer, Cesa-Bianchi, and Fischer 2002) estimator to narrow down the set of candidate pairs from which the pair of arms with the maximum uncertainty is selected at each round. Our algorithms adopt a similar design as MaxInP Saha and Gopalan (2020) in calculating the uncertainty of a pair. However, we use an online batch SGD instead of MLE to update Elo rating, which is more time and memory efficient.

## Methodology

### Background

Suppose there are  $n$  players, the Elo rating system (Elo 1978) assigns a rating  $r_x, x \in [n]$  to each player representing its skill. Let  $r^*$  denote the true ratings of  $n$  players. Our aim is to identify the best player among all  $n$  players:

$$x^* = \arg \max_{x \in [n]} r_x^* \quad (1)$$

Denote  $P$  as the true winning probability matrix,  $p_{xy}$  as the underlying groundtruth probability of  $x$  beating  $y$ . Based on the Bradley-Terry model (Hunter et al. 2004), the predicted probability of player  $x$  winning  $y$  is

$$\hat{p}_{xy} = \sigma(r_x - r_y). \quad (2)$$

$\sigma(x)$  is a sigmoid function with  $\sigma(x) = \frac{1}{1+e^{-x}}$ . Elo ratings are updated by maximizing the likelihood of win-loss predictions which corresponds to minimizing the loss:

$$\ell_{\text{Elo}}(p_{xy}, \hat{p}_{xy}) = -p_{xy} \log \hat{p}_{xy} - (1 - p_{xy}) \log(1 - \hat{p}_{xy}). \quad (3)$$

At time  $t$  player  $x$  compete with player  $y$  with outcome  $o_{xy}^t$ :  $o_{xy}^t = 1$  if  $x$  wins and  $o_{xy}^t = 0$  otherwise. We can use  $o_{xy}^t$  to compute the gradient of Eq. (3) and update Elo by gradient descent:

$$r_x^{t+1} \leftarrow r_x^t - \eta \cdot \nabla_{r_x} \ell_{\text{Elo}}(o_{xy}^t, \hat{p}_{xy}^t) = r_x^t + \eta \cdot (o_{xy}^t - \hat{p}_{xy}^t). \quad (4)$$

Let  $T$  denote the total number of rounds, at each round  $t \in [T]$ , we adopt a system that will pull a pair of players  $(x_t, y_t) \in [n] \times [n]$  and get comparison result  $o_t(x_t, y_t) \sim \text{Bern}(p_{xy})$ . The cumulative regret of  $T$  rounds is defined as

$$R(T) = \sum_{t=1}^T [r_{x^*}^* - \frac{1}{2}(r_{x_t}^* + r_{y_t}^*)]. \quad (5)$$

The definition is consistent with (Saha and Gopalan 2020). It measures the reward difference between the best arm and the two selected arms at each round.

### MaxIn-Elo Algorithm

We use the notations below in the followed presentations.

- $n$ : the number of players.
- $\tau$ : the batch size.
- $r$ : a vector of  $n$  players' ratings.
- $r^*$ : the true ratings of  $n$  players.
- $\hat{r}_t$ : estimator by MLE with  $t$  round comparisons.
- $\tilde{r}_j$ : the SGD estimator at batch  $j$ .

---

Algorithm 1: MaxIn-Elo: Dueling bandits with online SGD for top player identification.

---

**Input:** batch size  $\tau$ , maximum number of rounds  $T$ ,  $N$  players' strategies, parameters  $\alpha, \gamma$ .

**Output:** output  $r$

- 1: Randomly choose a pair to compare and record as  $x_t, y_t, o_t$  for  $t \in [\tau]$
  - 2:  $V_{\tau+1} = \sum_{t=1}^{\tau} (e_{x_t} - e_{y_t})(e_{x_t} - e_{y_t})^T$
  - 3: Calculate the maximum-likelihood estimator  $\hat{r}_\tau$  by solving  $\nabla_r \sum_{t=1}^{\tau} \ell_{\text{Elo}}(o_t, \hat{p}(x_t, y_t)) = 0$
  - 4: Maintain convex set  $\mathcal{C} = \{r : \|r - \hat{r}_\tau\| \leq 2\}$
  - 5: **for**  $t = \tau + 1, \tau + 2, \dots, T$  **do**
  - 6:   **if**  $t \% \tau = 1$  **then**
  - 7:      $j \leftarrow \lfloor (t-1)/\tau \rfloor$  and  $\eta_j = \frac{1}{\alpha_j}$
  - 8:     Calculate gradient  $\nabla_r l_{j,\tau}(\tilde{r}_{j-1})$  through Eq. (4)
  - 9:     Update ratings  $\tilde{r}_j$  through Eq. (8)
  - 10:     Compute  $\bar{r} = \frac{1}{j} \sum_{q=1}^j \tilde{r}_q$
  - 11:   **end if**
  - 12:   Define a candidate optimal set  $\mathcal{S} = \{x \mid \bar{r}_x - \bar{r}_y + \gamma \|e_x - e_y\|_{V_t^{-1}} > 0, \forall y \in [n]/\{x\}\}$
  - 13:   Select a pair as:  $(x_t, y_t) = \arg \max_{(x,y) \in \mathcal{S}} \|e_x - e_y\|_{V_t^{-1}}$ .
  - 14:   Let players  $(x_t, y_t)$  compete and observe  $o_t(x_t, y_t)$
  - 15:   Compute  $V_{t+1} = V_t + (e_{x_t} - e_{y_t})(e_{x_t} - e_{y_t})^T$
  - 16: **end for**
- 

- $\bar{r} = \sum_{q=1}^j \tilde{r}_q$ : the average of previous SGD iterations.
- $\|x\| = \sqrt{x^\top x}$ : the standard  $\ell_2$  norm.
- $e_i$ : the  $i$ -th unit base vector, i.e., the  $i$ -dimension equals 1 and all other components equal 0.
- $V_t$ : the history matrix recording previous  $t$  pulling information defined by  $V_t = \sum_{i=1}^{t-1} (e_{x_i} - e_{y_i})(e_{x_i} - e_{y_i})^\top$ .
- $\|x\|_V$ : a special  $\ell_2$ -norm associated with matrix  $V$  defined by  $\|x\|_V = \sqrt{x^\top V x}$ .
- $\mathcal{B}$  a neighborhood of  $r^*$  with  $\mathcal{B} = \{r \mid \|r - r^*\| \leq 3\}$ .
- $\mathcal{C}$ : a neighborhood of  $\hat{r}_\tau$  with  $\mathcal{C} = \{r \mid \|r - \hat{r}_\tau\| \leq 2\}$ .
- $\prod_{\mathcal{C}}(\cdot)$ : the projection operation defined by:

$$\prod_{\mathcal{C}}(r) = \hat{r}_\tau + \frac{2 * (r - \hat{r}_\tau)}{\min\{2, \|r - \hat{r}_\tau\|\}} \quad (6)$$

**Algorithm Overview** The main idea of our MaxIn-Elo algorithm is to maintain a candidate set of promising items via UCB and select the most informative pairs out of the set to evaluate at each round. Firstly, the ratings are initialized by maximizing likelihood of Eq. (3) on a batch of randomly sampled pairs with batch size  $\tau$ . The solution is denoted as  $\hat{r}_\tau$  and  $\tilde{r}_0 = \hat{r}$ . Then starting from round  $t = \tau + 1$ , we update  $\tilde{r}_j$  every  $\tau$  rounds by solving the following objective function

$$l_{j,\tau}(r) = \sum_{t=(j-1)\tau+1}^{j\tau} \ell_{\text{Elo}}(o_t, \hat{p}(x_t, y_t)). \quad (7)$$

The stochastic gradient update of  $\tilde{r}_j$  reads

$$\tilde{r}_j \leftarrow \prod_{\mathcal{C}} (\tilde{r}_{j-1} - \eta_j \nabla_r l_{j,\tau}(\tilde{r}_{j-1})). \quad (8)$$

First, the strong convexity of the objective function is required for fast convergence, and if we select a suitable  $\tau$  through Eq. (14), the aggregated objective function  $l_{j,\tau}(r)$  is a  $\alpha$ -strong convex function when  $r \in \mathcal{B}$ . Second, to ensure  $\tilde{r}_j \in \mathcal{B}$ ,  $\tilde{r}_j$  is projected into the convex set  $\mathcal{C}$  (also discussed in the proof of Lemma 2).

For each update, a batch of pairs are selected that lead to maximal information gain. The UCB score of a pair is defined by:

$$h(x_t, y_t) = \bar{r}_{x_t} - \bar{r}_{y_t} + \gamma \|e_{x_t} - e_{y_t}\|_{V_t^{-1}}, \quad (9)$$

with the balance parameter  $\gamma$ . The specific  $V_t^{-1}$  norm  $\gamma \|e_{x_t} - e_{y_t}\|_{V_t^{-1}}$  measures the uncertainty between two arms. The UCB estimator balances the exploitation and exploration through combining ratings estimation  $\bar{r}$  and the uncertainty term.

At each round  $t$ , we obtain a set of optimal player candidates  $\mathcal{S}$  with positive UCB scores:

$$\mathcal{S} = \{x | h(x, y) > 0, \forall y \in [n] \setminus \{x\}\}. \quad (10)$$

From the candidate set  $\mathcal{S}$ , we then pull a pair of arms with highest uncertainty by

$$(x_t, y_t) = \arg \max_{(x,y) \in \mathcal{S} \times \mathcal{S}} \|e_x - e_y\|_{V_t^{-1}} \quad (11)$$

to induce sufficient exploration. A detailed algorithm of our MaxIn-Elo is shown in Algorithm 1.

Compared to MaxInP which uses MLE at each iteration, MaxIn-Elo uses SGD to update the Elo rating  $r$  as traditional Elo does. Thus, our method is more efficient in both computation and time, and simple to implement. Compared to RG-UCB which randomly selects a pair to evaluate, our MaxIn-Elo selects the maximum informative pair and trade-off exploration and exploitation. The online SGD update rating according to a batch of comparisons to ensuring the  $\alpha$ -strong convexity of the objective function  $l_{j,\tau}$ , thus the selection of the batch size  $\tau$  is important for balancing the  $\alpha$ -strong convexity and the computation complexity of mini-batch update. To our best knowledge, this is the first algorithm that allows stochastic gradient descent update in dueling bandits settings. See Table 1 for a comparison on time and memory.

### MaxIn-mElo Algorithm

To enable the rating system to handle the intransitive skills, we extend the online sampling algorithm to multidimensional Elo ratings (mElo) (Balduzzi et al. 2018). Baking in the Hodge decomposition theory (Jiang et al. 2011), mElo proposed to decompose the antisymmetric logits matrix of win-loss probabilities into a transitive component, i.e. gradient flow of rating vector, and a cyclic component to capture the intransitive relations. By learning a  $2k$ -dimensional vector  $c_x$  and a rating  $r_x$  per player, the win-loss prediction for mElo $_{2k}$  is defined as:

$$\hat{p}_{xy} = \sigma(r_x - r_y + c_x^\top \cdot \Omega_{2k \times 2k} \cdot c_y). \quad (12)$$

Algorithms	Regret	Time Complexity	Memory
DBGD	$O(T^{2/3})$	$O(T)$	$O(n)$
RG-UCB	No	$O(T)$	$O(n)$
Random	No	$O(T)$	$O(n)$
MaxInP	$\tilde{O}(\sqrt{T})$	$O(nT^2 + n^2T)$	$O(nT)$
<b>MaxIn-Elo</b>	$\tilde{O}(\sqrt{T})$	$O(n^2T)$	$O(n^2)$

Table 1: Comparison of regret, time complexity and memory with other algorithms. Our MaxIn-Elo and the MaxInP achieve the lowest regret bound  $\tilde{O}(\sqrt{T})$ , but our MaxIn-Elo has lower time and memory complexity than the MaxInP.

where  $\Omega_{2k \times 2k} = \sum_{i=1}^k (e_{2i-1} e_{2i}^\top - e_{2i} e_{2i-1}^\top)$ .

The UCB estimate of a pair  $(x_t, y_t)$  for mElo then becomes:

$$h(x_t, y_t) = \bar{r}_{x_t} - \bar{r}_{y_t} + \bar{c}_x^\top \Omega \bar{c}_y + \gamma \|e_{x_t} - e_{y_t}\|_{V_t^{-1}}. \quad (13)$$

Notice that compared to Elo ratings with  $k = 0$  (Eq. (2)), mElo ratings assign a feature vector per player to approximate intransitive interactions. We present the details for the mElo ratings and Algorithm 2 for MaxIn-mElo in Appendix.

## Regret Analysis

We give the cumulative regret bound of MaxIn-Elo, as far as we know, this is the first work that combines the online gradient update with dueling bandits and gives the cumulative regret of dueling bandits while being updated with SGD. We make a mild assumption on the link function  $\sigma$ .

**Assumption 1.** Define  $c_\eta = \inf_{\{\|r-r^*\| \leq \eta\}} \sigma'(r_x - r_y)$ , where  $(x, y) \in [n] \times [n]$ , and we assume  $c_3 > 0$ .

This assumption is similar to that in (Ding, Hsieh, and Sharpnack 2021). Our main results rely on the following concentration events and the proofs of which are deferred to Appendix.

**Lemma 1.** Suppose we sample a sequence of arm pairs  $\{(x_1, y_1), (x_2, y_2), \dots, (x_t, y_t)\}$  through Algorithm 1 up to round  $t$ , and assume the selected batch size  $\tau$  satisfy that  $\lambda_{\min}(V_{\tau+1}) \geq 1$ , where  $\lambda_{\min}(V_{\tau+1})$  means the minimum eigenvalue of the matrix  $(V_{\tau+1})$ , Then  $\forall t > 0$ ,

$$\sum_{i=\tau+1}^{\tau+t} \|(e_{x_i} - e_{y_i})\|_{V_i^{-1}} < \sqrt{2nt \log \left( \frac{2\tau + t}{n} \right)}.$$

Lemma 1 gives the bound of the sum of selected pair's uncertainty from round  $\tau + 1$  to  $t$ . And this lemma will be adopted to derive the cumulative regret bound. In the following Lemma 2, we show that when the batch size  $\tau$  is chosen as Eq. (14), we have the concentration property of the averaged SGD estimator  $\bar{r}$ .

**Lemma 2.** Assume that there exists a positive constant  $\lambda_f$  such that  $\lambda_{\min}(\mathbb{E}[(e_{x_t} - e_{y_t})(e_{x_t} - e_{y_t})^\top]) \geq \lambda_f$  holds at each round  $t > \tau$ , where  $(x_t, y_t)$  is sampled through Algo-

Algorithm 1. Let the batch size  $\tau$  satisfies

$$\begin{aligned}\tau_1 &= 2 \left( \frac{C_1 \sqrt{n} + C_2 \sqrt{2 \log T}}{\lambda_{\min}(B)} \right)^2 + \frac{16(n + 2 \log T)}{c_1^2 \lambda_{\min}(B)}, \\ \tau_2 &= 2 \left( \frac{C_1 \sqrt{n} + C_2 \sqrt{2 \log T}}{\lambda_f} \right)^2 + \frac{4\alpha}{c_3 \lambda_f}, \\ \tau &= \lceil \max\{\tau_1, \tau_2\} \rceil,\end{aligned}\quad (14)$$

where  $B = \mathbb{E}_{(x,y) \stackrel{iid}{\sim} [n] \times [n]} [(e_x - e_y)(e_x - e_y)^T]$ . Define  $g_1(t)$  and  $g_2(j)$ ,

$$g_1(t) = \frac{1}{2c_1} \sqrt{\frac{n}{2} \log \left( 1 + \frac{2t}{n} \right) + 2 \log T}, \quad (15)$$

$$g_2(j) = \frac{\tau}{\alpha} \sqrt{1 + \log j}. \quad (16)$$

For a constant  $\alpha \geq c_3$ , there exists two positive constants  $C_1, C_2$  such that if the batch size  $\tau$  is chosen as Eq. (14), then we have that at each round  $t > \tau$  corresponding to batch  $j = \lfloor \frac{t-1}{\tau} \rfloor$ , event  $E_1(t)$  holds with probability at least  $1 - \frac{5}{T^2}$ , where  $E_1(t) = \{ \forall (x, y) : |(e_x - e_y)^T (\bar{r}_j - r^*)| \leq g_1(j\tau) \|e_x - e_y\|_{V_{j\tau+1}^{-1}} + g_2(j) \frac{\sqrt{2}}{\sqrt{j}} \}$ .

The following Lemma 3 shows how to select a suitable balanced parameter  $\gamma$  of UCB score that ensures the best player is always in the candidate set.

**Lemma 3.** Define the constant  $C = \sqrt{2nT \log \left( \frac{T+\tau}{n} \right)}$ . At each round  $t > \tau$ , let UCB balanced parameter  $\gamma = 2g_1(t)$  and assume  $\Delta > g_1(T)C$ , if  $\alpha$  satisfies that  $\alpha \geq \frac{\sqrt{2\tau} \sqrt{1+\log j}}{(\Delta - g_1(T)C) \sqrt{j}}$ , then we have  $x^* \in \mathcal{S}$  holds with probability at least  $1 - \frac{5}{T^2}$ , where  $j = \lfloor \frac{t-1}{\tau} \rfloor$ ,  $\Delta$  is the difference between ratings of optimal player  $x^*$  and sub-optimal player  $x'$ . Recall  $x^* = \arg \max_{x \in [n]} r_{x^*}^*$ , and define  $x' = \arg \max_{x \in [n]/x^*} r_{x'}^*$ ,  $\Delta = r_{x^*}^* - r_{x'}^*$ .

Lemma 3 shows that if we properly select UCB balanced parameter  $\gamma$  and parameter  $\alpha$  which describes objective function  $l_{j,\tau}$  as a  $\alpha$ -strongly convex, then it is promised that the best player  $x^*$  is in candidates set  $\mathcal{S}$  with high probability. This property is helpful for the top-1 identification because the candidate set  $\mathcal{S}$  will become tighter with the time, and  $x^*$  always in  $\mathcal{S}$ , thus candidate set  $\mathcal{S}$  only contains  $x^*$  eventually. Together we are ready to present our main results in Theorem 1.

**Theorem 1.** We run our Algorithm 1 to get a sequence of arm-pair, and let the learning rate parameter  $\alpha \geq \max\{c_3, \frac{\sqrt{2\tau} \sqrt{1+\log j}}{(\Delta - g_1(T)C) \sqrt{j}}\}$  with assumption that  $\Delta > g_1(T)C$ , the balanced parameter  $\gamma = 2g_1(t)$ , there exists two positive parameter  $C_1, C_2$  such that if the batch size  $\tau$  is chosen as Eq. (14), then we have the cumulative regret satisfies that:

$$R(T) \leq \tau * \Delta_{max} + (2 + \tau)g_1(T) \sqrt{2nT \log \left( \frac{2\tau + T}{n} \right)} + 4g_2(J) \sqrt{\tau T},$$

with probability at least  $1 - \frac{10}{T}$ , where  $J = \lfloor \frac{T}{\tau} \rfloor$ ,  $\Delta_{max} = \max_i r_i^* - \min_i r_i^*$ ,  $g_1(T), g_2(J)$  is defined in Eq. (15) and  $C$  is a constant defined as  $C = \sqrt{2nT \log \left( \frac{T+\tau}{n} \right)}$ .

Note that  $\tau \sim O(\max\{n, \log T\})$  (Eq. (14)),  $g_1(T) \sim O(\sqrt{n \log T})$ ,  $g_2(J) \sim O(\sqrt{\log T})$ . Combining the above analysis, we have  $R(T) \sim O(n \log T \sqrt{T})$  (or  $\tilde{O}(\sqrt{T})$ ). This regret upper bound is equivalent to that in (Saha and Gopalan 2020) which employs MLE estimators. However, our algorithm improves the efficiency in terms of memory and time. The memory cost is constant with respect to  $T$  while MaxInP's memory cost is linear in the time horizon  $T$ . The time complexity of our MaxIn-Elo is  $O(n^2 T)$ , while MaxInP's time complexity is  $O(nT^2 + n^2 T)$ . See Table 1 for a detailed comparison. Detailed proofs are referred to Appendix.

## Experiments

We consider the following two batteries of experiments to evaluate the performance of our algorithms in the scenarios of transitive and intransitive real world meta-games. Ablation studies of parameter  $\gamma$ , dimension of mElo and the batch size  $\tau$  can be found in Appendix.

### Baselines

**Random:** The pairwise matching scheme of the classical Round-robin (Rasmussen and Trick 2008) tournament is based on random sampling. We construct a simple baseline that randomly select a pair from all  $n * (n - 1) / 2$  pairs with replacement. After sampling a pair, we use the Elo/mElo model to update the ratings.

**RG-UCB** (Rowland et al. 2019): This algorithm adopts a pure exploration sampling scheme, which uniformly samples a pair from the set containing pairs that need to be estimated. And the stopping condition  $C(\delta)$  controls the total number of comparisons of each pair, where  $\delta$  is a hyper parameter deciding the confidence level of estimated competitive results.

**DBGD** (Yue and Joachims 2009): This dueling bandits algorithm is popular in ranking tasks when only pair-wise binary feedback is available. It maintains one winning arm at each round, and randomly synthesizes a gradient to obtain the opponent arm in the contextual bandit setting. In our feature free setting, this is equivalent to randomly selecting a player as the opponent.

**$\alpha$ -IG** (Rashid, Zhang, and Ciosek 2021): This is an active sampling algorithm used for estimating the  $\alpha$ -rank (Omidshafiei et al. 2019). This algorithm selects a pair with largest information gain at each round. In the transitive case, the top player has an  $\alpha$ -rank score equal to 1. Due to the high computation cost at each round ( computing  $\alpha$ -rank for 80000 times in a  $4 \times 4$  game), we only compare with it in a  $4 \times 4$  transitive game: the '2 Good, 2Bad' game given by  $\alpha$ -IG.

**MaxInP** (Saha and Gopalan 2020): This algorithm is for the generalized linear contextual dueling bandits problem, in which arms are represented as feature vectors. At each round  $t$ , it uses MLE to estimate model parameters  $\theta$  relying on all historical comparisons. This algorithm calculates a candidate set containing advanced arms and pulls an arm-pair with the largest uncertainty. In order to fit their model, each player is described as a one-hot vector, and the estimated parameters  $\theta$  correspond to players' ratings in our setting.

**MaxIn-Elo:** Our first algorithm adopts dueling bandits to adaptively sample pairs for Elo rating update in Eq. (2). The aim of our MaxIn-Elo is to identify the advanced players gradually, and to minimize the cumulative regret described in Eq. (5) simultaneously.

**MaxIn-mElo:** Our second algorithm tames the intransitive scenarios. Different to MaxIn-Elo, there is an extra vector  $c$  to capture intransitive relationship in competition outcome prediction. The dimension of  $c$  is set to 8 in experiments. For the MaxIn-mElo algorithm, we hope to identify players with superior mElo ratings and to minimize cumulative regret on mElo ratings.

## Experiments Setting

**Real world games** We do our experiments on twelve real-games released by Czarnecki et al. (2020), most of which are implemented on the OpenSpiel framework (Lanctot et al. 2019). The six games used for evaluating Elo are Triangular game, Transitive game, Elo game, and three noisy variants of Elo games. The first three are transitive games; the three variants of Elo game are Elo games with additive Gaussian noises. The six intransitive games used for evaluating mElo are Kuhn-poker, AlphaStar, tic\_tac\_toe, hex, Blotto and 5,3-Blotto game.

The intransitivity of games can be revealed by sink strongly connected components (SSCCs) (Omidshafiei et al. 2019), which is a set of strategies that cannot be defeated by external strategies and all internal strategies become a circle, such as Rock, Paper, Scissors. The statistics of these games is shown in Table 2 in Appendix.

**Metrics** Except the cumulative regret defined in Eq. (5), we introduce three other metrics for Reciprocal Rank (RR), Normalized Discounted Cumulative Gain (NDCG), and Hit Ratio (HR). RR is used for the results on generating top-1 players. NDCG and HR report discrete performance for top-1 performance and are thus used in top- $k$  results.

Reciprocal Rank (RR) (Donmez, Svore, and Burges 2009) give the reciprocal of predicted ranking of the best player  $x^*$ . Define  $RR = 1/R(x^*)$ , where  $R(x)$  returns the ranking of player  $x$  relying on currently predicted ratings  $\bar{r}$ . Larger RR corresponds to better performance on the top-1 player identification.

Hit Ratio@ $K$  (He et al. 2015) is defined as the ratio of the predicted top- $k$  that belong to the true top- $k$ . Since hit ratio does not consider the positions of correct predictions, we also adopt NDCG (Donmez, Svore, and Burges 2009) which assigns higher importance to top ranks. Normalized Discounted Cumulative Gain (NDCG) is widely used in the evaluation of rankings and NDCG@ $k$  measures the importance of predicted top- $k$  players. It is given by

$$NDCG@K = \frac{1}{N_K} \sum_{i=1}^K \frac{2^{l(d_i)} - 1}{\log(i + 1)},$$

where  $N_k$  is a normalizer to ensure that the perfect ranking would result in  $NDCG@K = 1$ .  $d_i$  denote the index of predicted  $i$ -th player, and  $l(x) \in \{0, 1\}$  is the relevance level about top- $k$  identification, we set  $l(x) = 1$  if player  $x$  in true top- $k$  otherwise 0.

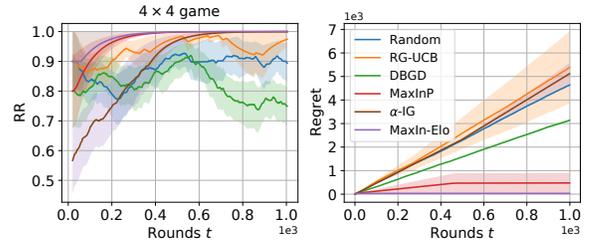


Figure 1: Results on  $4 \times 4$  game (2 Good 2 Bad).

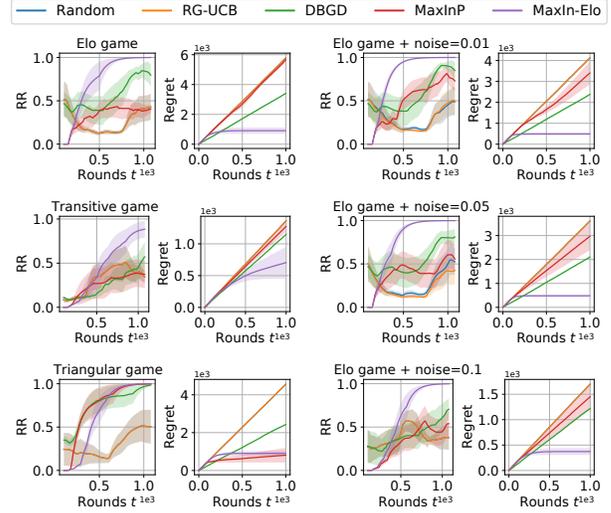


Figure 2: Results of Elo on transitive games.

**Parameters setting** For Random, DBGD, and RG-UCB baseline, we perform a grid search for the initial step size  $\eta$  in the range  $\{0.01, 0.05, 0.1, 0.5, 1, 5, 10\}$ . For RG-UCB, stopping confidence  $\delta = 0.2$ . For MaxInP, we tune the UCB balanced parameter  $\gamma \in \{0.2, 0.4, 0.6, \dots, 2.0\}$ . For MaxIn-Elo and MaxIn-mElo, we tune the initialized learning rate  $\eta \in \{0.01, 0.05, 0.1, 0.5, 1, 5, 10\}$ , and the learning rate at batch  $j$  is set as  $\frac{\eta}{j}$ . And the UCB balanced parameter  $\gamma \in \{0.2, 0.4, 0.6, \dots, 2.0\}$ . The batch size  $\tau$  of MaxInP, MaxIn-Elo and MaxIn-mElo is set to  $0.7 * n$ . When baselines use mElo model to calculate ratings, we set the dimension of the extra vector  $c$  as 8. We use the parameters that report the best performance for  $\alpha$ -IG. We repeat experiments 5 times with different random seeds and plot the averaged performance with standard deviations.

All experiments were run in a single x86\_64 GNU/Linux machine with 256 AMD EPYC 7742 64-Core Processor and 2 A100 PCIe 40GB GPU. We use sklearn(0.24.2) to solve the MLE.

## Results

Figure 1, 2, 3 show the results of top-1 identification on 13 games. To ensure a fair comparison between all baselines, we perform a grid search to select parameters with the best RR performance for each random seed. If the winning probability matrix can be fitted into the Elo model, then we cal-

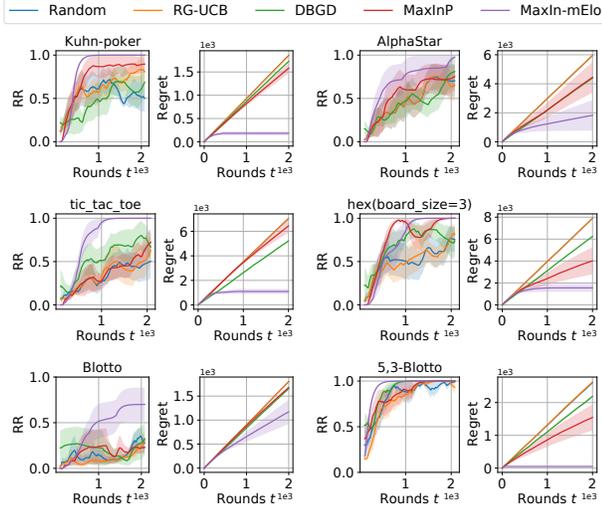


Figure 3: Results of mElo on intransitive games.

culate the true ratings through Eq. (2), otherwise we use the mElo ratings as the true ratings through Eq. (12).

**Evaluation of MaxIn-Elo** Figure 1 shows the results of a  $4 \times 4$  transitive game. MaxIn-Elo has the highest convergence rate on both RR and cumulative regret metrics, and MaxIn-Elo has the lowest cumulative regret close to 0. As shown in Figure 2, MaxIn-Elo significantly outperforms all other baselines on five games and achieves similar performance on Triangular game. Regarding the RR metric, MaxIn-Elo can converge to 1 on four games. Even on Transitive game and Elo game + noise=0.1, RR scores as up to 0.6 and 0.8 respectively, which indicates that the rank of the top player is no more than 2. Thus we think MaxIn-Elo has the ability to effectively identify the top player. On the Elo game, Elo game + noise=0.01, and Elo game + noise=0.05, the cumulative regret is closed to convergence at around 500 rounds. When the cumulative regret meets convergence, the candidate optimal set  $\mathcal{S}$  only contains the top player, and no regret increasing.

Different from the other 5 stochastic games, Triangular game is a deterministic game with all winning probabilities are equal to 1 or 0, thus it is easy to evaluate. For DBGD baseline, it maintains the current best player and randomly selects an opponent, so it could find the best player more quickly, but has a large cumulative regret because of randomly selected opponents.

**Evaluation of MaxIn-mElo** Figure 3 shows the results of baselines on six real-world intransitive games. MaxInP is based on the Elo model for it is a special generalized linear model only with rating parameter  $r$  without cyclic vector parameter  $c$ , but all other baselines are based on the mElo model. As the Figure 3 shows, MaxIn-mElo has the lowest cumulative regret and the highest RR on all six games. With regard to the RR, MaxIn-mElo can be up to 1 on all games except for Blotto. One possible reason why MaxIn-mElo cannot be up to 1 on Blotto may be that its size of top SSCC is very large. The other reason is that we use the low-

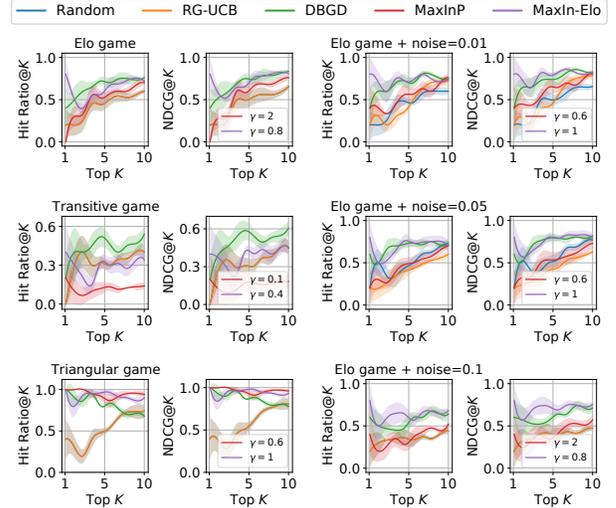


Figure 4: Results of Top- $k$  player identification on transitive games.  $\gamma$  in red and purple indicates that reports best performance for MaxInP and MaxIn-Elo respectively.

rank approximation of the probability matrix's rotation on the mElo model. Although we misidentified the top-1 player, we are still better than all other baselines.

**Results of Top- $k$  player identification** Figure 4 gives the results of top- $k$  predictions on transitive games. MaxIn-Elo and MaxInP both have a parameter  $\gamma$  used to balance exploration and exploitation, larger  $\gamma$  can lead to a larger candidate set then lead to better top- $k$  performance. We keep other parameters fixed and run experiments with different  $\gamma \in \{0.2, 0.4, 0.6, \dots, 2.0\}$ , and we report the performance of MaxInP and MaxIn-Elo under the best  $\gamma$ . Figure 4 shows that MaxIn-Elo has the best performance of the top-1 identification on all games, and it achieves the comparable performance of top- $k$  identification on most games. Results of different  $\gamma$  can be found in Appendix.

## Discussions

This work studied the problem of multi-agent evaluation with Elo ratings. We have adopted an online match scheduling framework to improve the sample efficiency of the Elo rating system and its extension mElo for the intransitive settings. Both empirical and theoretical results justify that our algorithms can achieve higher sample efficiency and lower regret on most of the tasks.

We consider two limitations of this work. Firstly, the match outcome prediction in our algorithm is based on only ratings without considering features that describe players. Future work may consider adding features into the match prediction. Secondly, our algorithm focuses more on identifying the best player without being tailored for identifying top- $k$  players. Future work can consider active sampling that achieves better results on both top-1 and top- $k$  cases.

## Ethics Statement

This work proposes algorithms for online match scheduling that improve the efficiency in identifying top players in competitive games such as chess. While empirical studies in this work, which are based on AI agents, have demonstrated the superior gain of using our proposed methods, there is a caveat that our algorithms assume that the all players' skill levels remain unchanged throughout the repeated competition rounds. This assumption likely does not hold for human players whose playing strengths will be affected by energy consumption due to frequent matches. Therefore, extra caution needs to be taken when deploying our methods to schedule real-world competitions involving human players and an interesting research extension would be to model such performance strength changes explicitly in designing the match scheduling algorithms.

## Acknowledgements

Co-author Haifeng Zhang is supported in part by the Strategic Priority Research Program of Chinese Academy of Sciences, Grant No. XDA27030401.

## References

- Agrawal, S.; and Goyal, N. 2012. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, 39–1.
- Agrawal, S.; and Goyal, N. 2013. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning (ICML)*, 127–135.
- Arneson, B.; Hayward, R. B.; and Henderson, P. 2010. Monte Carlo tree search in Hex. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4): 251–258.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2): 235–256.
- Balduzzi, D.; Tuyls, K.; Perolat, J.; and Graepel, T. 2018. Re-evaluating evaluation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3268–3279.
- Bubeck, S.; Munos, R.; and Stoltz, G. 2011. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19): 1832–1852.
- Czarnecki, W. M.; Gidel, G.; Tracey, B.; Tuyls, K.; Omidshafiei, S.; Balduzzi, D.; and Jaderberg, M. 2020. Real World Games Look Like Spinning Tops. *Advances in Neural Information Processing Systems (NeurIPS)*, 33.
- Ding, Q.; Hsieh, C.-J.; and Sharpnack, J. 2021. An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1585–1593.
- Donmez, P.; Svore, K. M.; and Burges, C. J. 2009. On the local optimality of lambdarank. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, 460–467.
- Du, Y.; Han, L.; Fang, M.; Dai, T.; Liu, J.; and Tao, D. 2019. LIIR: learning individual intrinsic reward in multi-agent reinforcement learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS)*, 4403–4414.
- Du, Y.; Yan, X.; Chen, X.; Wang, J.; and Zhang, H. 2021. Estimating  $\alpha$ -Rank from A Few Entries with Low Rank Matrix Completion. In *International Conference on Machine Learning (ICML)*, 2870–2879. PMLR.
- Elo, A. E. 1978. *The rating of chessplayers, past and present*. Arco Pub.
- Groh, C.; Moldovanu, B.; Sela, A.; and Sunde, U. 2012. Optimal seedings in elimination tournaments. *Economic Theory*, 49(1): 59–80.
- Gruslys, A.; Dabney, W.; Azar, M. G.; Piot, B.; Bellemare, M.; and Munos, R. 2018. The Reactor: A fast and sample-efficient Actor-Critic agent for Reinforcement Learning. In *International Conference on Learning Representations (ICLR)*.
- Han, L.; Sun, P.; Du, Y.; Xiong, J.; Wang, Q.; Sun, X.; Liu, H.; and Zhang, T. 2019. Grid-wise control for multi-agent reinforcement learning in video game ai. In *International Conference on Machine Learning (ICML)*, 2576–2585. PMLR.
- He, X.; Chen, T.; Kan, M.-Y.; and Chen, X. 2015. Tri-rank: Review-aware explainable recommendation by modeling aspects. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM)*, 1661–1670.
- Heckel, R.; Shah, N. B.; Ramchandran, K.; Wainwright, M. J.; et al. 2019. Active ranking from pairwise comparisons and when parametric assumptions do not help. *The Annals of Statistics*, 47(6): 3099–3126.
- Hunter, D. R.; et al. 2004. MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1): 384–406.
- Jiang, X.; Lim, L.-H.; Yao, Y.; and Ye, Y. 2011. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1): 203–244.
- Lai, M. 2015. Giraffe: Using deep reinforcement learning to play chess. *arXiv preprint arXiv:1509.01549*.
- Lanctot, M.; Lockhart, E.; Lepiauw, J.-B.; Zambaldi, V.; Upadhyay, S.; Pérolat, J.; Srinivasan, S.; Timbers, F.; Tuyls, K.; Omidshafiei, S.; et al. 2019. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*.
- Li, L.; Lu, Y.; and Zhou, D. 2017. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning (ICML)*, 2071–2080.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; and Ostrovski, G. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529.
- Muller, P.; Omidshafiei, S.; Rowland, M.; Tuyls, K.; Perolat, J.; Liu, S.; Hennes, D.; Marris, L.; Lanctot, M.; Hughes,

E.; et al. 2020. A Generalized Training Approach for Multiagent Learning. In *International Conference on Learning Representations (ICLR)*, 1–35.

Omidshafiei, S.; Papadimitriou, C.; Piliouras, G.; Tuyls, K.; Rowland, M.; Lespiau, J.-B.; Czarnecki, W. M.; Lanctot, M.; Perolat, J.; and Munos, R. 2019.  $\alpha$ -rank: Multi-agent evaluation by evolution. *Scientific reports*, 9(1): 1–29.

Rashid, T.; Zhang, C.; and Ciosek, K. 2021. Estimating  $\alpha$ -Rank by Maximizing Information Gain. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, 5673–5681.

Rasmussen, R. V.; and Trick, M. A. 2008. Round robin scheduling—a survey. *European Journal of Operational Research*, 188(3): 617–636.

Rowland, M.; Omidshafiei, S.; Tuyls, K.; Perolat, J.; Valko, M.; Piliouras, G.; and Munos, R. 2019. Multiagent evaluation under incomplete information. In *Advances in Neural Information Processing Systems (NeurIPS)*, 12291–12303.

Saha, A.; and Gopalan, A. 2020. Regret Minimization in Stochastic Contextual Dueling Bandits. *arXiv preprint arXiv:2002.08583*.

Saha, A.; Koren, T.; and Mansour, Y. 2021. Adversarial dueling bandits. In *International Conference on Machine Learning (ICML)*, 9235–9244.

Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; and Bolton, A. 2017. Mastering the game of Go without human knowledge. *Nature*, 550(7676): 354.

Szörényi, B.; Busa-Fekete, R.; Paul, A.; and Hüllermeier, E. 2015. Online rank elicitation for plackett-luce: A dueling bandits approach. In *Proceedings of the 19th international Conference on Neural Information Processing systems (NeurIPS)*.

Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4): 285–294.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Yang, Y.; Tutunov, R.; Sakulwongtana, P.; and Ammar, H. B. 2020.  $\alpha^\alpha$ -Rank: Practically Scaling  $\alpha$ -Rank through Stochastic Optimisation. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 1575–1583.

Yue, Y.; Broder, J.; Kleinberg, R.; and Joachims, T. 2012. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5): 1538–1556.

Yue, Y.; and Joachims, T. 2009. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, 1201–1208.