

Regret Lower Bound and Optimal Algorithm in Dueling Bandit Problem

Junpei Komiyama

The University of Tokyo

Junya Honda

The University of Tokyo

Hisashi Kashima

Kyoto University

Hiroshi Nakagawa

The University of Tokyo

JUNPEI@KOMIYAMA.INFO

HONDA@STAT.T.U-TOKYO.AC.JP

KASHIMA@I.KYOTO-U.AC.JP

NAKAGAWA@DL.ITC.U-TOKYO.AC.JP

Abstract

We study the K -armed dueling bandit problem, a variation of the standard stochastic bandit problem where the feedback is limited to relative comparisons of a pair of arms. We introduce a tight asymptotic regret lower bound that is based on the information divergence. An algorithm that is inspired by the Deterministic Minimum Empirical Divergence algorithm (Honda and Takemura, 2010) is proposed, and its regret is analyzed. The proposed algorithm is found to be the first one with a regret upper bound that matches the lower bound. Experimental comparisons of dueling bandit algorithms show that the proposed algorithm significantly outperforms existing ones.

Keywords: multi-armed bandit problem, dueling bandit problem, online learning

1. Introduction

A multi-armed bandit problem is a crystallized instance of a sequential decision-making problem in an uncertain environment, and it can model many real-world scenarios. This problem involves conceptual entities called arms, and a forecaster who tries to identify good arms from bad ones. At each round, the forecaster draws one of the K arms and receives a corresponding reward. The aim of the forecaster is to maximize the cumulative reward over rounds, which is achieved by running an algorithm that balances the exploration (acquisition of information) and the exploitation (utilization of information).

While it is desirable to obtain direct feedback from an arm, in some cases such direct feedback is not available. In this paper, we consider a version of the standard stochastic bandit problem called the K -armed dueling bandit problem (Yue et al., 2009), in which the forecaster receives relative feedback, which specifies which of two arms is preferred. Although the original motivation of the dueling bandit problem arose in the field of information retrieval, learning under relative feedback is universal to many fields, such as recommender systems (Gemmis et al., 2009), graphical design (Brochu et al., 2010), and natural language processing (Zaidan and Callison-Burch, 2011), which involve explicit or implicit feedback provided by humans.

Related work: Here, we briefly discuss the literature of the K -armed dueling bandit problem. The problem involves a preference matrix $M = \{\mu_{i,j}\} \in \mathbb{R}^{K \times K}$, whose ij entry $\mu_{i,j}$ corresponds to the probability that arm i is preferred to arm j .

Most algorithms assume that the preference matrix has certain properties. Interleaved Filter (IF) (Yue et al., 2012) and Beat the Mean Bandit (BTM) (Yue and Joachims, 2011), early algorithms proposed for solving the dueling bandit problem, require the arms to be totally ordered, that is, $i \succ j \Leftrightarrow \mu_{i,j} > 1/2$. Moreover, IF assumes *stochastic transitivity*: for any triple (i, j, k) with $i \succ j \succ k$, $\mu_{i,k} \geq \max\{\mu_{i,j}, \mu_{j,k}\}$. Unfortunately, stochastic transitivity does not hold in many real-world settings (Yue and Joachims, 2011). BTM relaxes this assumption by introducing *relaxed stochastic transitivity*: there exists $\gamma \geq 1$ such that for all pairs (j, k) with $1 \succ j \succ k$, $\gamma\mu_{1,k} \geq \max\{\mu_{1,j}, \mu_{j,k}\}$ holds. The drawback of BTM is that it requires the explicit value of γ on which the performance of the algorithm depends. Urvoy et al. (2013) considered a wide class of sequential learning problems with bandit feedback that includes the dueling bandit problem. They proposed the Sensitivity Analysis of VArIables for Generic Exploration (SAVAGE) algorithm, which empirically outperforms IF and BTM for moderate K . Among the several versions of SAVAGE, the one called Condorcet SAVAGE makes the *Condorcet assumption* and performed the best in their experiment. The Condorcet assumption is that there is a unique arm that is superior to the others. Unlike the two transitivity assumptions, the Condorcet assumption does not require the arms to be totally ordered and is less restrictive. IF, BTM, and SAVAGE either explicitly require the number of rounds T , or implicitly require T to determine the confidence level δ .

Recently, an algorithm called Relative Upper Confidence Bound (RUCB) (Zoghi et al., 2014b) was proven to have an $O(K^2 \log T)$ regret bound under the Condorcet assumption. RUCB is based on the upper confidence bound index (Lai and Robbins, 1985; Agrawal, 1995; Auer et al., 2002) that is widely used in the field of bandit problems. RUCB is *horizonless*: it does not require T beforehand and runs for any duration. Zoghi et al. (2015) extended RUCB into the mergeRUCB algorithm with an $O(K \log T)$ regret bound under the Condorcet assumption as well as the assumption that a portion of the preference matrix is informative (i.e., different from $1/2$). Ailon et al. (2014) proposed three algorithms named Doubler, MultiSBM, and Sparring. MultiSBM is endowed with an $O(K \log T)$ regret bound and Sparring was reported to outperform IF and BTM in their simulation. These algorithms assume that the pairwise feedback is generated from the non-observable utilities of the selected arms. The existence of the utility distributions associated with individual arms restricts the structure of the preference matrix.

In summary, only $O(K^2 \log T)$ regret is known to be achievable under the Condorcet assumption (SAVAGE and RUCB) and all known algorithms require additional assumptions to achieve $O(K \log T)$ regret (IF, BTM, MultiSBM, and mergeRUCB). These two regrets can be related to the difficulty of the dueling bandit problem in that there are $K - 1$ candidates of actions to test “how good” each arm i is. A naive use of the confidence bound requires every pair of arms to be compared $O(\log T)$ times and yields an $O(K^2 \log T)$ regret bound.

Contribution: In this paper, we propose an algorithm called Relative Minimum Empirical Divergence (RMED). This paper contributes to our understanding of the dueling bandit problem in the following three respects.

- **The regret lower bound:** Some studies (e.g., Yue et al., 2012) have shown that the K -armed dueling bandit problem has a $\Omega(K \log T)$ regret lower bound. In this paper, we further analyze this lower bound to obtain the optimal constant factor for models satisfying the Condorcet assumption. Furthermore, we show that the lower bound is the same under the total

order assumption. This means that optimal algorithms under the Condorcet assumption also achieve a lower bound of regret under the total order assumption even though such algorithms do not know that the arms are totally ordered.

- **An optimal algorithm:** The regret of RMED is not only $O(K \log T)$, but also optimal in the sense that its constant factor matches the asymptotic lower bound under the Condorcet assumption. RMED is the first optimal algorithm in the study of the dueling bandit problem.
- **Empirical performance assessment:** The performance of RMED is extensively evaluated by using five datasets: two synthetic datasets, one including preference data, and two including ranker evaluations in the information retrieval domain.

2. Problem Setup

The K -armed dueling bandit problem involves K arms that are indexed as $[K] = \{1, 2, \dots, K\}$. Let $M \in \mathbb{R}^{K \times K}$ be a preference matrix whose ij entry $\mu_{i,j}$ corresponds to the probability that arm i is preferred to arm j . At each round $t = 1, 2, \dots, T$, the forecaster selects a pair of arms $(l(t), m(t)) \in [K]^2$, then receives a relative feedback $\hat{X}_{l(t),m(t)}(t) \sim \text{Bernoulli}(\mu_{l(t),m(t)})$ that indicates which of $(l(t), m(t))$ is preferred. By definition, $\mu_{i,j} = 1 - \mu_{j,i}$ holds for any $i, j \in [K]$ and $\mu_{i,i} = 1/2$.

Let $N_{i,j}(t)$ be the number of comparisons of pair (i, j) and $\hat{\mu}_{i,j}(t)$ be the empirical estimate of $\mu_{i,j}$ at round t . In building statistics by using the feedback, we treat pairs without taking their order into consideration. Therefore, for $i \neq j$, $N_{i,j}(t) = \sum_{t'=1}^{t-1} (\mathbf{1}\{l(t') = i, m(t') = j\} + \mathbf{1}\{l(t') = j, m(t') = i\})$ and $\mu_{i,j} = (\sum_{t'=1}^{t-1} (\mathbf{1}\{l(t') = i, m(t') = j, \hat{X}_{l(t'),m(t')}(t') = 1\} + \mathbf{1}\{l(t') = j, m(t') = i, \hat{X}_{l(t'),m(t')}(t') = 0\})) / N_{i,j}(t)$, where $\mathbf{1}[\cdot]$ is the indicator function. For $j \neq i$, let $N_{i>j}(t)$ be the number of times i is preferred over j . Then, $\hat{\mu}_{i,j}(t) = N_{i>j}(t) / N_{i,j}(t)$, where we set $0/0 = 1/2$ here. Let $\hat{\mu}_{i,i}(t) = 1/2$.

Throughout this paper, we will assume that the preference matrix has a Condorcet winner (Urvoy et al., 2013). Here we call an arm i the Condorcet winner if $\mu_{i,j} > 1/2$ for any $j \in [K] \setminus \{i\}$. Without loss of generality, we will assume that arm 1 is the Condorcet winner. The set of preference matrices which have a Condorcet winner is denoted by \mathcal{M}_C . We also define the set of preference matrices satisfying the total order by $\mathcal{M}_o \subset \mathcal{M}_C$; that is, the relation $i \prec j \Leftrightarrow \mu_{i,j} < 1/2$ induces a total order iff $\{\mu_{i,j}\} \in \mathcal{M}_o$.

Let $\Delta_{i,j} = \mu_{i,j} - 1/2$. We define the regret per round as $r(t) = (\Delta_{1,i} + \Delta_{1,j})/2$ when the pair (i, j) is compared. The expectation of the cumulative regret, $\mathbb{E}[R(T)] = \mathbb{E} \left[\sum_{t=1}^T r(t) \right]$ is used to measure the performance of an algorithm. The regret increases at each round unless the selected pair is $(l(t), m(t)) = (1, 1)$.

2.1. Regret lower bound in the K -armed dueling bandits

In this section we provide an asymptotic regret lower bound when $T \rightarrow \infty$. Let the superiors of arm i be a set $\mathcal{O}_i = \{j | j \in [K], \mu_{i,j} < 1/2\}$, that is, the set of arms that is preferred to i on average. The essence of the K -armed dueling bandit problem is how to eliminate each arm $i \in [K] \setminus \{1\}$ by making sure that arm i is not the Condorcet winner. To do so, the algorithm uses some of the arms in \mathcal{O}_i and compares i with them.

A dueling bandit algorithm is strongly consistent for model $\mathcal{M} \subset \mathcal{M}_C$ iff it has $\mathbb{E}[R(T)] = o(T^a)$ regret for any $a > 0$ and any $M \in \mathcal{M}$. The following lemma is on the number of comparisons of suboptimal arm pairs.

Lemma 1 (The lower bound on the number of suboptimal arm draws) (i) *Let an arm $i \in [K] \setminus \{1\}$ and preference matrix $M \in \mathcal{M}_C$ be arbitrary. Given any strongly consistent algorithm for model \mathcal{M}_C , we have*

$$\mathbb{E} \left\{ \sum_{j \in \mathcal{O}_i} d(\mu_{i,j}, 1/2) N_{i,j}(T) \right\} \geq (1 - o(1)) \log T, \quad (1)$$

where $d(p, q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ is the KL divergence between two Bernoulli distributions with parameters p and q . (ii) *Furthermore, inequality (1) holds for any $M \in \mathcal{M}_o$ given any strongly consistent algorithm for \mathcal{M}_o .*

Lemma 1 states that, for arbitrary arm $j \in \mathcal{O}_i$, an algorithm needs to make $\log T / d(\mu_{i,j}, 1/2)$ comparisons between arms i and j to be convinced that arm i is inferior to arm j and thus i is not the Condorcet winner. Since the regret increase per round of comparing arm i with j is $(\Delta_{1,i} + \Delta_{1,j})/2$, eliminating arm i by comparing it with j incurs a regret of

$$\frac{(\Delta_{1,i} + \Delta_{1,j}) \log T}{2d(\mu_{i,j}, 1/2)}. \quad (2)$$

Therefore, the total regret is bounded from below by comparing each arm i with an arm j that minimizes (2) and the regret lower bound is formalized in the following theorem.

Theorem 2 (The regret lower bound) (i) *Let the preference matrix $M \in \mathcal{M}_C$ be arbitrary. For any strongly consistent algorithm for model \mathcal{M}_C ,*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log T} \geq \sum_{i \in [K] \setminus \{1\}} \min_{j \in \mathcal{O}_i} \frac{\Delta_{1,i} + \Delta_{1,j}}{2d(\mu_{i,j}, 1/2)} \quad (3)$$

holds. (ii) Furthermore, inequality (3) holds for any $M \in \mathcal{M}_o$ given any strongly consistent algorithm for \mathcal{M}_o .

The proof of Lemma 1 and Theorem 2 can be found in Appendix B. The proof of Lemma 1 is similar to that of Lai and Robbins (1985, Theorem 1) for the standard multi-armed bandit problem but differs in the following point that is characteristic to the dueling bandit. To achieve a small regret in the dueling bandit, it is necessary to compare the arm i with itself if i is the Condorcet winner. However, we trivially know that $\mu_{i,i} = 1/2$ without sampling and such a comparison yields no information to distinguish possible preference matrices. We can avoid this difficulty by evaluating $N_{i,j}$ and $N_{i,i}$ in different ways.

3. RMED1 Algorithm

In this section, we first introduce the notion of empirical divergence. Then, on the basis of the empirical divergence, we formulate the RMED1 algorithm.

Algorithm 1 Relative Minimum Empirical Divergence (RMED) Algorithm

1: **Input:** K arms, $f(K) \geq 0$. $\alpha > 0$ (RMED2FH, RMED2). T (RMED2FH).
 2: $L \leftarrow \begin{cases} 1 & \text{(RMED1, RMED2)} \\ \lceil \alpha \log \log T \rceil & \text{(RMED2FH)} \end{cases}$.
 3: **Initial phase:** draw each pair of arms L times. At the end of this phase, $t = L(K - 1)K/2$.
 4: **if** RMED2FH **then**
 5: For each arm $i \in [K]$, fix $\hat{b}^*(i)$ by (6).
 6: **end if**
 7: $L_C, L_R \leftarrow [K], L_N \leftarrow \emptyset$.
 8: **while** $t \leq T$ **do**
 9: **if** RMED2 **then**
 10: Draw all pairs (i, j) until it reaches $N_{i,j}(t) \geq \alpha \log \log t$. $t \leftarrow t + 1$ for each draw.
 11: **end if**
 12: **for** $l(t) \in L_C$ in an arbitrarily fixed order **do**
 13: Select $m(t)$ by using $\begin{cases} \text{Algorithm 2} & \text{(RMED1)} \\ \text{Algorithm 3} & \text{(RMED2, RMED2FH)} \end{cases}$.
 14: Draw arm pair $(l(t), m(t))$.
 15: $L_R \leftarrow L_R \setminus \{l(t)\}$.
 16: $L_N \leftarrow L_N \cup \{j\}$ (without a duplicate) for any $j \notin L_R$ such that $\mathcal{J}_j(t)$ holds.
 17: $t \leftarrow t + 1$.
 18: **end for**
 19: $L_C, L_R \leftarrow L_N, L_N \leftarrow \emptyset$.
 20: **end while**

Algorithm 2 RMED1 subroutine for selecting $m(t)$

1: $\hat{\mathcal{O}}_{l(t)}(t) \leftarrow \{j \in [K] \setminus \{l(t)\} \mid \hat{\mu}_{l(t),j}(t) \leq 1/2\}$
 2: **if** $i^*(t) \in \hat{\mathcal{O}}_{l(t)}(t)$ or $\hat{\mathcal{O}}_{l(t)}(t) = \emptyset$ **then**
 3: $m(t) \leftarrow i^*(t)$.
 4: **else**
 5: $m(t) \leftarrow \arg \min_{j \neq l(t)} \hat{\mu}_{l(t),j}(t)$.
 6: **end if**

3.1. Empirical divergence and likelihood function

In inequality (1) of Section 2.1, we have seen that $\sum_{j \in \mathcal{O}_i} d(\mu_{i,j}, 1/2) N_{i,j}(T)$, the sum of the divergence between $\mu_{i,j}$ and $1/2$ multiplied by the number of comparisons between i and j , is the characteristic value that defines the minimum number of comparisons. The empirical estimate of this value is fundamentally useful for evaluating how unlikely arm i is to be the Condorcet winner. Let the opponents of arm i at round t be the set $\hat{\mathcal{O}}_i(t) = \{j \mid j \in [K] \setminus \{i\}, \hat{\mu}_{i,j}(t) \leq 1/2\}$. Note that, unlike the superiors \mathcal{O}_i , the opponents $\hat{\mathcal{O}}_i(t)$ for each arm i are defined in terms of the empirical averages, and thus the algorithms know who the opponents are. Let the empirical divergence be

$$I_i(t) = \sum_{j \in \hat{\mathcal{O}}_i(t)} N_{i,j}(t) d(\hat{\mu}_{i,j}(t), 1/2).$$

The value $\exp(-I_i(t))$ can be considered as the ‘‘likelihood’’ that arm i is the Condorcet winner. Let $i^*(t) = \arg \min_{i \in [K]} I_i(t)$ (ties are broken arbitrarily) and $I^*(t) = I_{i^*(t)}(t)$. By definition, $I^*(t) \geq 0$. RMED is inspired by the Deterministic Minimum Empirical Divergence (DMED) algorithm (Honda and Takemura, 2010). DMED, which is designed for solving the standard K -armed bandit problem, draws arms that may be the best one with probability $\Omega(1/t)$, whereas RMED in the dueling bandit problem draws arms that are likely to be the Condorcet winner with probability $\Omega(1/t)$. Namely, any arm i that satisfies

$$\mathcal{J}_i(t) = \{I_i(t) - I^*(t) \leq \log t + f(K)\} \quad (4)$$

is the candidate of the Condorcet winner and will be drawn soon. Here, $f(K)$ can be any non-negative function of K that is independent of t . Algorithm 1 lists the main routine of RMED. There are several versions of RMED. First, we introduce RMED1. RMED1 initially compares all pairs once (initial phase). Let $T_{\text{init}} = (K - 1)K/2$ be the last round of the initial phase. From $t = T_{\text{init}} + 1$, it selects the arm by using a loop. $L_C = L_C(t)$ is the set of arms in the current loop, and $L_R = L_R(t) \subset L_C(t)$ is the remaining arms of L_C that have not been drawn yet in the current loop. $L_N = L_N(t)$ is the set of arms that are going to be drawn in the next loop. An arm i is put into L_N when it satisfies $\{\mathcal{J}_i(t) \cap \{i \notin L_R(t)\}\}$. By definition, at least one arm (i.e. $i^*(t)$ at the end of the current loop) is put into L_N in each loop. For arm $l(t)$ in the current loop, RMED1 selects $m(t)$ (i.e. the comparison target of $l(t)$) determined by Algorithm 2.

The following theorem, which is proven in Section 5, describes a regret bound of RMED1.

Theorem 3 *For any sufficiently small $\delta > 0$, the regret of RMED1 is bounded as:*

$$\mathbb{E}[R(T)] \leq \sum_{i \in [K] \setminus \{1\}} \frac{((1 + \delta) \log T + f(K)) \Delta_{1,i}}{2d(\mu_{i,1}, 1/2)} + O(K^2) + O\left(\frac{K}{\delta^2}\right) + O(Ke^{AK-f(K)}),$$

where $A = A(\{\mu_{i,j}\}_{i,j \in [K]})$ is a constant as a function of T . Therefore, by letting $\delta = \log^{-1/3} T$ and choosing an $f(K) = cK^{1+\epsilon}$ for arbitrary $c, \epsilon > 0$, we obtain

$$\mathbb{E}[R(T)] \leq \sum_{i \in [K] \setminus \{1\}} \frac{\Delta_{1,i} \log T}{2d(\mu_{i,1}, 1/2)} + O(K^{2+\epsilon}) + O(K \log^{2/3} T).$$

3.2. Gap between the constant factor of RMED1 and the lower bound

From the lower bound of Theorem 2, the $O(K \log T)$ regret bound of RMED1 is optimal up to a constant factor. Moreover, the constant factor matches the regret lower bound of Theorem 2 if $b^*(i) = 1$ for all $i \in [K] \setminus \{1\}$ where

$$b^*(i) = \arg \min_{j \in \mathcal{O}_i} \frac{\Delta_{1,i} + \Delta_{1,j}}{d(\mu_{i,j}, 1/2)}. \quad (5)$$

Here we define $d^+(p, q) = d(p, q)$ if $p < q$ and 0 otherwise, and $x/0 = +\infty$. Note that, there can be ties that minimize the RHS of (5). In that case, we may choose any of the ties as $b^*(i)$ to eliminate arm i . For ease of explanation, we henceforth will assume that $b^*(i)$ is unique, but our results can be easily extended to the case of ties.

We claim that $b^*(i) = 1$ holds in many cases for the following mathematical and practical reasons. (i) The regret of drawing a pair (i, j) , $j \neq 1$, is $(\Delta_{1,i} + \Delta_{1,j})/2$, whereas it is simply

Algorithm 3 Subroutine for selecting $m(t)$ in RMED2 and RMED2FH

```

1: if RMED2 then
2:   Update  $\hat{b}^*(l(t))$  by (6).
3: end if
4:  $\hat{\mathcal{O}}_{l(t)}(t) \leftarrow \{j \in [K] \setminus \{l(t)\} | \hat{\mu}_{l(t),j}(t) \leq 1/2\}$ .
5: if  $\hat{b}^*(l(t)) \in \hat{\mathcal{O}}_{l(t)}(t)$  and  $\begin{cases} N_{l(t),i^*(t)}(t) \geq N_{l(t),\hat{b}^*(l(t))}(t) / \log \log t & \text{(RMED2)} \\ N_{l(t),i^*(t)}(t) \geq N_{l(t),\hat{b}^*(l(t))}(t) / \log \log T & \text{(RMED2FH)} \end{cases}$  then
6:    $m(t) \leftarrow \hat{b}^*(l(t))$ .
7: else
8:   Select  $m(t)$  by using Algorithm 2.
9: end if
    
```

$\Delta_{1,i}/2$ for the pair $(i, 1)$. Thus, $d^+(\mu_{i,j}, 1/2)$ has to be much larger than $d^+(\mu_{i,1}, 1/2)$ in order to satisfy $b^*(i) = j$. (ii) The Condorcet winner usually wins over the other arms by a large margin, and therefore, $d^+(\mu_{i,1}, 1/2) \geq d^+(\mu_{i,j}, 1/2)$. For example, in the preference matrix of Example 1 (Table 1(a)), $b^*(3) = 1$ as long as $q < 0.79$. Example 2 (Table 1(b)) is a preference matrix based on six retrieval functions in the full-text search engine of ArXiv.org (Yue and Joachims, 2011)¹. In Example 2, $b^*(i) = 1$ holds for all i , even though $\mu_{1,4} < \mu_{2,4}$. In the case of a 16-ranker evaluation based on the Microsoft Learning to Rank dataset (details are given in Section 4), occasionally $b^*(i) \neq 1$ occurs, but the difference between the regrets of drawing arm 1 and $b^*(i)$ is fairly small (smaller than 1.2% on average). Nevertheless, there are some cases in which comparing arm i with 1 is not such a clever idea. Example 3 (Table 1(c)) is a toy example in which comparing arm i with $b^*(i) \neq 1$ makes a large difference. In Example 3, it is clearly better to draw pairs (2, 4), (3, 2) and (4, 3) to eliminate arms 2, 3, and 4, respectively. Accordingly, it is still interesting to consider an algorithm that reduces regret by comparing arm i with $b^*(i)$.

 Table 1: Three preference matrices. In each example, the value at row i , column j is $\mu_{i,j}$.

(a) Example 1				(b) Example 2							(c) Example 3				
	1	2	3		1	2	3	4	5	6		1	2	3	4
1	0.5	0.7	0.7	1	0.50	0.55	0.55	0.54	0.61	0.61	1	0.5	0.6	0.6	0.6
2	0.3	0.5	q	2	0.45	0.50	0.55	0.55	0.58	0.60	2	0.4	0.5	0.9	0.1
3	0.3	$1-q$	0.5	3	0.45	0.45	0.50	0.54	0.51	0.56	3	0.4	0.1	0.5	0.9
				4	0.46	0.45	0.46	0.50	0.54	0.50	4	0.4	0.9	0.1	0.5
				5	0.39	0.42	0.49	0.46	0.50	0.51					
				6	0.39	0.40	0.44	0.50	0.49	0.50					

3.3. RMED2 Algorithm

We here propose RMED2, which gracefully estimates $b^*(i)$ during a bandit game and compares arm i with $b^*(i)$. RMED2 and RMED1 share the main routine (Algorithm 1). The subroutine of RMED2 for selecting $m(t)$ is shown in Algorithm 3. Unlike RMED1, RMED2 keeps drawing pairs of arms (i, j) at least $\alpha \log \log t$ times (Line 10 in Algorithm 1). The regret of this exploration is insignificant since $O(\log \log T) = o(\log T)$. Once all pairs have been explored more than $\alpha \log \log t$

1. In the original preference matrix of Yue and Joachims (2011), $\mu_{2,4} \neq 1 - \mu_{4,2}$. To satisfy $\mu_{2,4} = 1 - \mu_{4,2}$, we replaced $\mu_{2,4}$ and $\mu_{4,2}$ of the original with $(\mu_{2,4} - \mu_{4,2} + 1)/2$ and $(\mu_{4,2} - \mu_{2,4} + 1)/2$, respectively.

times, RMED2 goes to the main loop. RMED2 determines $m(t)$ by using Algorithm 2 based on the estimate of $b^*(i)$ given by

$$\hat{b}^*(i) = \arg \min_{j \in [K] \setminus \{i\}} \frac{\hat{\Delta}_{i^*(t),i} + \hat{\Delta}_{i^*(t),j}}{d^+(\hat{\mu}_{i,j}(t), 1/2)}, \quad (6)$$

where ties are broken arbitrarily, $\hat{\Delta}_{i,j} = 1/2 - \hat{\mu}_{i,j}$ and we set $x/0 = +\infty$. Intuitively, RMED2 tries to select $m(t) = \hat{b}^*(i)$ for most rounds, and occasionally explores $i^*(t)$ in order to reduce the regret increase when RMED2 fails to estimate the true $b^*(i)$ correctly.

3.4. RMED2FH algorithm

Although we believe that the regret of RMED2 is optimal, the analysis of RMED2 is a little bit complicated since it sometimes breaks the main loop and explores from time to time. For ease of analysis, we here propose RMED2 Fixed Horizon (RMED2FH, Algorithm 1 and 3), which is a “static” version of RMED2. Essentially, RMED2 and RMED2FH have the same mechanism. The differences are that (i) RMED2FH conducts an $\alpha \log \log T$ exploration in the initial phase. After the initial phase (ii) $\hat{b}^*(i)$ for each i is fixed throughout the game. Note that, unlike RMED1 and RMED2, RMED2FH requires the number of rounds T beforehand to conduct the initial $\alpha \log \log T$ draws of each pair. The following Theorem shows the regret of RMED2FH that matches the lower bound of Theorem 2.

Theorem 4 *For any sufficiently small $\delta > 0$, the regret of RMED2FH is bounded as:*

$$\begin{aligned} \mathbb{E}[R(T)] \leq \sum_{i \in [K] \setminus \{1\}} \frac{(\Delta_{1,i} + \Delta_{1,b^*(i)})((1 + \delta) \log T)}{2d(\mu_{i,b^*(i)}, 1/2)} + O(\alpha K^2 \log \log T) + O(K e^{AK - f(K)}) \\ + O\left(\frac{K \log T}{\log \log T}\right) + O\left(\frac{K}{\delta^2}\right) + O(K f(K)), \quad (7) \end{aligned}$$

where $A = A(\{\mu_{i,j}\}) > 0$ is a constant as a function of T . By setting $\delta = O((\log T)^{-1/3})$ and choosing an $f(K) = cK^{1+\epsilon}$ ($c, \epsilon > 0$) we obtain

$$\mathbb{E}[R(T)] \leq \sum_{i \in [K] \setminus \{1\}} \frac{(\Delta_{1,i} + \Delta_{1,b^*(i)}) \log T}{2d(\mu_{i,b^*(i)}, 1/2)} + O(\alpha K^2 \log \log T) + O\left(\frac{K \log T}{\log \log T}\right) + O(K^{2+\epsilon}). \quad (8)$$

Note that all terms except the first one in (8) are $o(\log T)$. From Theorems 2 and 4 we see that (i) RMED2FH is asymptotically optimal under the Condorcet assumption and (ii) the logarithmic term on the regret bound of RMED2FH cannot be improved even if the arms are totally ordered and the forecaster knows of the existence of the total order. The proof sketch of Theorem 4 is in Section 5.

4. Experimental Evaluation

To evaluate the empirical performance of RMED, we conducted simulations² with five bandit datasets (preference matrices). The datasets are as follows:

2. The source code of the simulations is available at <https://github.com/jkomiyama/duelingbanditlib>.

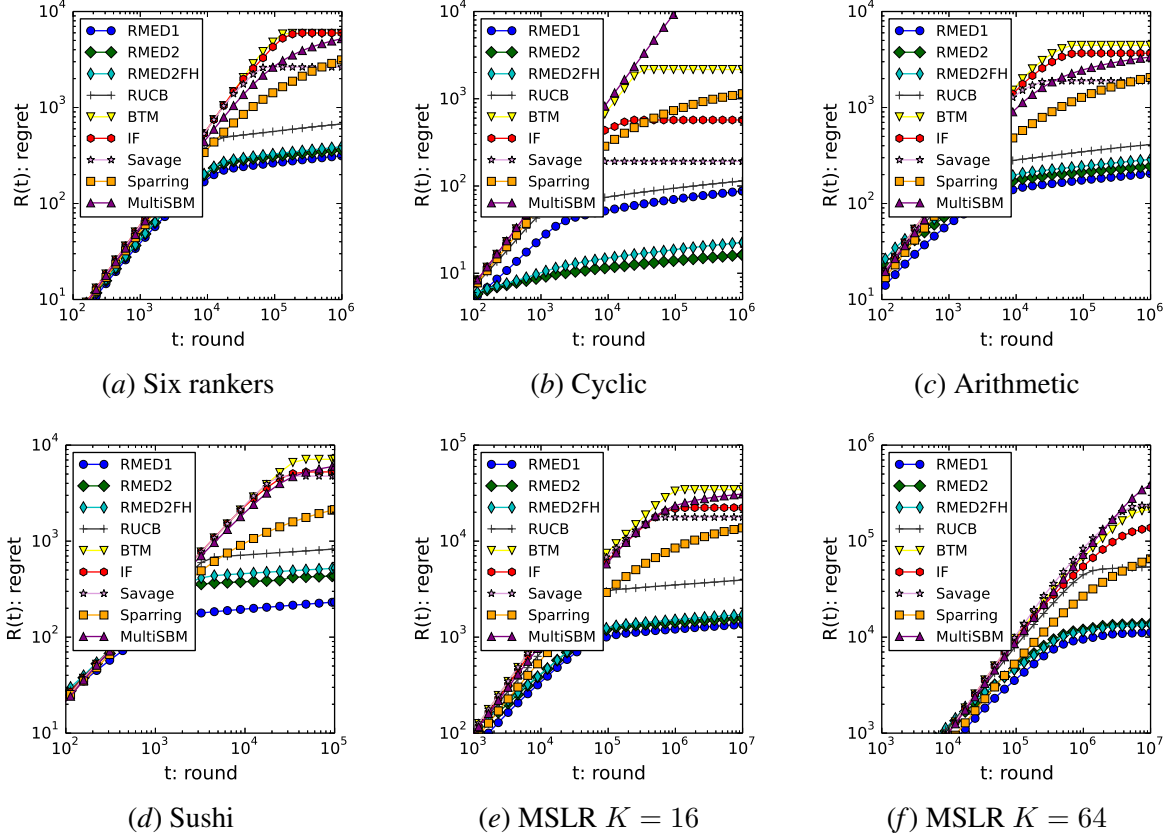


Figure 1: Regret-round log-log plots of algorithms.

Six rankers is the preference matrix based on the six retrieval functions in the full-text search engine of ArXiv.org (Table 1(b)).

Cyclic is the artificial preference matrix shown in Table 1(c). This matrix is designed so that the comparison of i with 1 is not optimal.

Arithmetic dataset involves eight arms with $\mu_{i,j} = 0.5 + 0.05(j - i)$ and has a total order.

Sushi dataset is based on the Sushi preference dataset (Kamishima, 2003) that contains the preferences of 5,000 Japanese users as regards 100 types of sushi. We extracted the 16 most popular types of sushi and converted them into arms with $\mu_{i,j}$ corresponding to the ratio of users who prefer sushi i over j . The Condorcet winner is the mildly-fatty tuna (chu-toro).

MSLR: We tested submatrices of a 136×136 preference matrix from Zoghi et al. (2015), which is derived from the Microsoft Learning to Rank (MSLR) dataset (Microsoft Research, 2010; Qin et al., 2010) that consists of relevance information between queries and documents with more than 30K queries. Zoghi et al. (2015) created a finite set of rankers, each of which corresponds to a ranking feature in the base dataset. The value $\mu_{i,j}$ is the probability that the ranker i beats ranker j based on the navigational click model (Hofmann et al., 2013). We randomly extracted $K = 16, 64$ rankers in our experiments and made sub preference matrices. The probability that the Condorcet winner exists in the subset of the rankers is high (more than 90%, c.f. Figure 1 in Zoghi et al. (2014a)), and we excluded the relatively small case where the Condorcet winner does not exist.

A Condorcet winner exists in all datasets. In the experiments, the regrets of the algorithms were averaged over 1,000 runs (Six rankers, Cyclic, Arithmetic, and Sushi), or 100 runs (MSLR).




Figure 2: Regret-round semilog plots of RMED compared with theoretical bounds. We set $f(K) = 0.3K^{1.01}$ for all algorithms, and $\alpha = 3$ for RMED2.

4.1. Comparison among algorithms

We compared the IF, BTM with $\gamma = 1.2$, RUCB with $\alpha = 0.51$, Condorcet SAVAGE with $\delta = 1/T$, MultiSBM and Sparring with $\alpha = 3$, and RMED algorithms. We set $f(K) = 0.3K^{1.01}$ for all RMED algorithms and set $\alpha = 3$ for RMED2 and RMED2FH. The effect of $f(K)$ is studied in Appendix A. Note that IF and BTM assume a total order among arms, which is not the case with the Cyclic, Sushi, and MSLR datasets. MultiSBM and Sparring assume the existence of the utility of each arm, which does not allow a cyclic preference that appears in the Cyclic dataset.

Figure 1 plots the regrets of the algorithms. In all datasets RMED significantly outperforms RUCB, the next best excluding the different versions of RMED. Notice that the plots are on a base 10 log-log scale. In particular, regret of RMED1 is more than twice smaller than RUCB on all datasets other than Cyclic, in which RMED2 performs much better. Among the RMED algorithms, RMED1 outperforms RMED2 and RMED2FH on all datasets except for Cyclic, in which comparing arm $i \neq 1$ with arm 1 is inefficient. RMED2 outperforms RMED2FH in the five of six datasets: this could be due to the fact that RMED2FH does not update $\hat{b}^*(i)$ for ease of analysis.

4.2. RMED and asymptotic bound

Figure 2 compares the regret of RMED with two asymptotic bounds. LB1 denotes the regret bound of RMED1. TrueLB is the asymptotic regret lower bound given by Theorem 2.

RMED1 and RMED2: When $T \rightarrow \infty$, the slope of RMED1 should converge to LB1, and the ones of RMED2 and RMED2FH should converge to TrueLB. On Six rankers, LB1 is exactly the same as TrueLB, and the slope of RMED1 converges to this TrueLB. In Cyclic, the slope of RMED2 converges to TrueLB, whereas that of RMED1 converges to LB1, from which we see that RMED2 is actually able to estimate $b^*(i) \neq 1$ correctly. In MSLR $K = 16$, LB1 and TrueLB are very close (the difference is less than 1.2%), and RMED1 and RMED2 converge to these lower bounds.

RMED2FH with different values of α : We also tested RMED2FH with several values of α . On the one hand, with $\alpha = 1$, the initial phase of RMED2FH is too short to identify $b^*(i)$; as a result it performs poorly on the Cyclic dataset. On the other hand, with $\alpha = 10$, the initial phase was too long, which incurs a practically non-negligible regret on the MSLR $K = 16$ dataset. We also tested several values of parameter α in RMED2FH. We omit plots of RMED2 with $\alpha = 1, 10$ for the sake of readability, but we note that in our datasets the performance of RMED2 is always better than or comparable with the one of RMED2FH under the same choice of α , although the optimality of RMED2 is not proved unlike RMED2FH.

5. Regret Analysis

This section provides two lemmas essential for the regret analysis of RMED algorithms and proves the asymptotic optimality of RMED1 based on these lemmas. A proof sketch on the optimal regret of RMED2FH is also given.

The crucial property of RMED is that, by constantly comparing arms with the opponents, the true Condorcet winner (arm 1) actually beats all the other arms with high probability. Let

$$\mathcal{U}(t) = \bigcap_{i \in [K] \setminus \{1\}} \{\hat{\mu}_{1,i}(t) > 1/2\}.$$

Under $\mathcal{U}(t)$, $\hat{\mu}_{i,1}(t) = 1 - \hat{\mu}_{1,i}(t) < 1/2$ for all $i \in [K] \setminus \{1\}$, and thus, $I_i(t) > 0$. Therefore, $\mathcal{U}(t)$ implies that $i^*(t) = \arg \min_{i \in [K]} I_i(t)$ is unique with $i^*(t) = 1$ and $I^*(t) = I_1(t) = 0$. Lemma 5 below shows that the average number of rounds that $\mathcal{U}^c(t)$ occurs is constant in T , where the superscript c denotes the complement.

Lemma 5 *When RMED1 or RMED2FH is run, the following inequality holds:*

$$\mathbb{E} \left[\sum_{t=T_{\text{init}}+1}^T \mathbf{1}\{\mathcal{U}^c(t)\} \right] = O(e^{AK-f(K)}), \quad (9)$$

where $A = A(\{\mu_{i,j}\}) > 0$ is a constant as a function of T .

Note that, since RMED2FH draws each pair $\lceil \alpha \log \log T \rceil$ times in the initial phase, we define $T_{\text{init}} = \lceil \alpha \log \log T \rceil (K-1)K/2$ for RMED2FH. We give a proof of this lemma in Appendix C. Intuitively, this lemma can be proved from the facts that arm 1 is drawn within roughly $e^{I_1(t)-f(K)}$ rounds and $I_1(t)$ is not very large with high probability.

Next, for $i \in [K] \setminus \{1\}$ and $j \in \mathcal{O}_i$, let

$$N_{i,j}^{\text{Suf}}(\delta) = \frac{(1+\delta) \log T + f(K)}{d(\mu_{i,j}, 1/2)} + 1,$$

which is a sufficient number of comparisons of i with j to be convinced that the arm i is not the Condorcet winner. The following lemma states that if pair (i, j) is drawn $N_{i,j}^{\text{Suf}}(\delta)$ times then i is rarely selected as $l(t)$ again.

Lemma 6 *When RMED1 or RMED2FH is run, for $i \in [K] \setminus \{1\}$, $j \in \mathcal{O}_i$,*

$$\mathbb{E} \left[\sum_{t=T_{\text{init}}+1}^T \mathbf{1}\{l(t) = i, N_{i,j}(t) \geq N_{i,j}^{\text{Suf}}(\delta)\} \right] = O\left(\frac{1}{\delta^2}\right) + O(e^{AK-f(K)}) + K.$$

We prove this lemma in Appendix D based on the Chernoff bound.

Now we can derive the regret bound of RMED1 based on these lemmas.

Proof of Theorem 3: Since $\mathcal{U}(t)$ implies $m(t) = 1$ in RMED1, the regret increase per round can be decomposed as:

$$r(t) = \mathbf{1}\{\mathcal{U}^c(t)\} + \sum_{i \in [K] \setminus \{1\}} \frac{\Delta_{1,i}}{2} \mathbf{1}\{l(t) = i, m(t) = 1, \mathcal{U}(t)\}. \quad (10)$$

Using Lemmas 5 and 6, we obtain

$$\begin{aligned}
\mathbb{E}[R(T)] &\leq T_{\text{init}} + \sum_{t=T_{\text{init}}+1}^T [r(t)] \\
&\leq \frac{K(K-1)}{2} + \mathbb{E} \left[\sum_{t=T_{\text{init}}+1}^T \mathbf{1}\{\mathcal{U}^c(t)\} \right] + \sum_{i \in [K] \setminus \{1\}} \frac{\Delta_{1,i}}{2} \left(N_{i,1}^{\text{Suf}}(\delta) + \sum_{t=1}^T \mathbf{1}[l(t) = i, m(t) = 1, N_{i,1}(t) \geq N_{i,1}^{\text{Suf}}(\delta)] \right) \\
&\leq \frac{K(K-1)}{2} + O(e^{AK-f(K)}) + \sum_{i \in [K] \setminus \{1\}} \frac{\Delta_{1,i}}{2} \left(N_{i,1}^{\text{Suf}}(\delta) + O\left(\frac{1}{\delta^2}\right) + O(e^{AK-f(K)}) + K \right),
\end{aligned}$$

which immediately completes the proof of Theorem 3. \blacksquare

We also prove Theorem 4 on the optimality of RMED2FH based on Lemmas 5 and 6. Because the full proof in Appendix E is a little bit lengthy, here we give its brief sketch.

Proof sketch of Theorem 4 (RMED2FH): Similar to Theorem 3, we use the fact that the $\mathcal{U}^c(t)$ does not occur very often (i.e., Lemma 5). Under $\mathcal{U}(t)$, we decompose the regret into the contributions of each arm $i \in [K] \setminus \{1\}$. There exists $C_2 > 0$ such that, for each $l(t) = i$, (i) with probability $1 - O((\log T)^{-C_2})$ RMED2FH successfully estimates $\hat{b}^*(i) = b^*(i)$ and selects $m(t) = b^*(i)$ for most rounds. The optimal $O(\log T)$ term comes from the comparison of i and $b^*(i)$. Arm 1 is also drawn for $O(\log T / \log \log T) = o(\log T)$ times. On the other hand, (ii) with probability $O((\log T)^{-C_2})$, RMED2FH fails to estimate $b^*(i)$ correctly. By occasionally comparing arm i with arm 1, we can bound the regret increase by $O(\log T \log \log T)$. Since $O((\log T)^{-C_2} \times \log T \log \log T) = o(\log T)$, this regret does not affect the $O(\log T)$ factor.

6. Discussion

We proved the regret lower bound in the dueling bandit problem. The RMED algorithm is based on the likelihood that the arm is the Condorcet winner. RMED is proven to have the matching regret upper bound. The empirical evaluation revealed that RMED significantly outperforms the state-of-the-art algorithms. To conclude this paper, we mention three directions of future work.

First, when a Condorcet winner does not necessarily exist, the Copeland bandits (Urvoy et al., 2013) are a natural extension of our problem. Thus, seeking an effective algorithm for solving this problem will be interesting. As is well known in the field of voting theory, there are several other criteria of winners that are incompatible with the Condorcet / Copeland bandits, such as the Borda winner (Urvoy et al., 2013). Comparing several criteria or developing an algorithm that outputs more than one of these winners should be interesting directions of future work.

Second, another direction is sequential preference elicitation problems under relative feedback that goes beyond the binary preference over pairs, such as multiscale feedback and/or preferences among three or more items.

Third, in the standard bandit problem, it is reported that KL-UCB+ (Lai, 1987; Garivier and Cappé, 2011) performs better than DMED. A study of a UCB-based optimal algorithm for the dueling bandits can yield an algorithm that outperforms RMED.

Acknowledgements

We thank the anonymous reviewers for their useful comments. This work was supported in part by JSPS KAKENHI Grant Number 26106506.

References

- R. Agrawal. Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27:1054–1078, 1995.
- Nir Ailon, Zohar Shay Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *ICML*, pages 856–864, 2014.
- Peter Auer, Nicoló Cesa-bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47:235–256, 2002.
- Eric Brochu, Tyson Brochu, and Nando de Freitas. A bayesian interactive optimization approach to procedural animation design. In *Proceedings of the 2010 Eurographics/ACM SIGGRAPH Symposium on Computer Animation, SCA 2010, Madrid, Spain, 2010*, pages 103–112, 2010.
- Sébastien Bubeck. *Bandits Games and Clustering Foundations*. Theses, Université des Sciences et Technologie de Lille - Lille I, June 2010.
- Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *COLT*, pages 359–376, 2011.
- Marco De Gemmis, Leo Iaquinta, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Giovanni Semeraro. Preference learning in recommender systems. In *In Preference Learning (PL-09) ECML/PKDD-09 Workshop*, 2009.
- Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. Fidelity, soundness, and efficiency of interleaved comparison methods. *Transactions on Information Systems*, 31(4):17:1–43, 2013.
- Junya Honda and Akimichi Takemura. An Asymptotically Optimal Bandit Algorithm for Bounded Support Models. In *COLT*, pages 67–79, 2010.
- Toshihiro Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In *KDD*, pages 583–588, 2003.
- T. L. Lai. Adaptive treatment allocation and the multi-armed bandit problem. *Ann. Statist.*, 15(3): 1091–1114, 09 1987.
- T. L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Microsoft Research. Microsoft Learning to Rank Datasets, 2010. URL <http://research.microsoft.com/en-us/projects/mslr/>.
- Tao Qin, Tie-Yan Liu, Jun Xu, and Hang Li. LETOR: A benchmark collection for research on learning to rank for information retrieval. *Inf. Retr.*, 13(4):346–374, 2010.
- Tanguy Urvoy, Fabrice Clérot, Raphaël Feraud, and Sami Naamane. Generic exploration and k-armed voting bandits. In *ICML*, pages 91–99, 2013.
- Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *ICML*, pages 241–248, 2011.

- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. In *COLT*, 2009.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *J. Comput. Syst. Sci.*, 78(5):1538–1556, 2012.
- Omar Zaidan and Chris Callison-Burch. Crowdsourcing translation: Professional quality from non-professionals. In *The 49th Annual Meeting of the Association for Computational Linguistics (ACL): Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*, pages 1220–1229, 2011.
- Masrour Zoghi, Shimon Whiteson, Maarten de Rijke, and Rémi Munos. Relative confidence sampling for efficient on-line ranker evaluation. In *WSDM*, pages 73–82, 2014a.
- Masrour Zoghi, Shimon Whiteson, Rémi Munos, and Maarten de Rijke. Relative upper confidence bound for the k-armed dueling bandit problem. In *ICML*, pages 10–18, 2014b.
- Masrour Zoghi, Shimon Whiteson, and Maarten de Rijke. MergeRUCB: A method for large-scale online ranker evaluation. In *WSDM*, 2015.