



HAL
open science

Unimodal Mono-Partite Matching in a Bandit Setting

Matthieu Rodet, Romaric Gaudel

► **To cite this version:**

Matthieu Rodet, Romaric Gaudel. Unimodal Mono-Partite Matching in a Bandit Setting. Complex Feedback in Online Learning Workshop at the 39th International Conference on Machine Learning, Jul 2022, Baltimore, United States. hal-03740961

HAL Id: hal-03740961

<https://inria.hal.science/hal-03740961>

Submitted on 2 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unimodal Mono-Partite Matching in a Bandit Setting

Matthieu Rodet¹ Romaric Gaudel²

Abstract

We tackle a new emerging problem, which is finding an optimal monopartite matching in a weighted graph. The semi-bandit version, where a full matching is sampled at each iteration, has been addressed by (Sentenac et al., 2021), creating an algorithm with an expected regret matching $O(\frac{L \log(L)}{\Delta} \log(T))$ with $2L$ players, T iterations and a minimum reward gap Δ . We reduce this bound in two steps. First, as in (Gauthier et al., 2021) and (Gauthier et al., 2022) we use the unimodality property of the expected reward on the appropriate graph to design an algorithm with a regret in $O(L \frac{1}{\Delta} \log(T))$. Secondly, we show that by moving the focus towards the main question ‘*Is user i better than user j ?*’ this regret becomes $O(L \frac{\Delta}{\Delta^2} \log(T))$, where $\Delta > \Delta$ derives from a better way of comparing users. Some experimental results finally show these theoretical results are corroborated in practice.

1. Introduction

In some online games, when all servers are running games, players wait in a queue to find a freed server to start a game. As finding a game could take some time, a player may not be available when her game is found, leading to its cancellation and causing extra wait to her opponent. Thus, game developers need to find a solution to cancel as less games as possible. This problem can be formulate as a monopartite matching problem: given a set of players which are each defined by their unknown inner probability of accepting the match, developers have to propose a matching among players maximizing the number of effectively played games.

This question boils down to the longstanding problem of finding matchings in a weighted graph, i.e. a subset of

^{*}Equal contribution ¹ENS Rennes, F-35000 Rennes, France
²Univ Rennes, Ensai, CNRS, CREST - UMR 9194, F-35000 Rennes, France. Correspondence to: Romaric Gaudel <romaric.gaudel@ensai.fr>.

edges without common vertices (Lovász & Plummer, 2009). Indeed, the graph could be constituted of the players as vertices and an edge could exists between two players only if they can compete against each other, weighted by the probability of a match occurring between them. In addition to this concrete application, this problem knows various applications in operations research (Wheaton, 1990), economics (Roth et al., 2004) and Machine Learning (Mehta, 2012). We consider the online version, where the algorithm is in charge of submitting a sequence of matchings (one per iteration) and gets in return the evaluation of these submissions. This problem is part of the combinatorial semi-bandits ones (Cesa-Bianchi & Lugosi, 2012).

As in (Komiyama et al., 2017; Gauthier et al., 2021; Sentenac et al., 2021), in this paper we focus on the setting where the matrix carrying the probability of match between two players is a rank-1 matrix. We use this setting as a toy example to first showcase the power of unimodality when facing a combinatorial bandit setting related to a ranking problem, and to show that we can go beyond current state of the art by fully embracing the main question ‘*Is user i better than user j ?*’.

Concretely, first we handle that setting by applying the strategy we introduced in (Gauthier et al., 2021) and (Gauthier et al., 2022): (i) we define a graph on the (ordered) matchings such that the expected reward is unimodal on this graph, and (ii) we apply a variant of OUSB (Combes & Proutière, 2014) on this graph. This again demonstrate the interest of the unimodality point of view to pull out the intrinsic complexity of a combinatorial semi-bandit setting. Indeed, by denoting $2L$ the number of players, T the number of iterations, and Δ the minimum reward gap, the naive application of the combinatorial bandit literature leads to an algorithm with a $O(\frac{L^2 \log^2(L)}{\Delta} \log(T))$ regret¹, while (Sentenac et al., 2021) gets an almost linear regret by accounting for the strong link between rank-1 matching and ranking. In this paper, we show that the unimodality enables an exact linear regret by focusing on the minimal set of comparisons (to the price of an additional $\log \log T$ term).

Secondly, the direct application of GRAB (Gauthier et al., 2021) compares two players by comparing the expected regret of two matchings. While this strategy seems obvious,

¹Using by example ESCB (Combes et al., 2015).

we show in this paper that it is suboptimal. Indeed, we exhibit another criterion with a gap $\tilde{\Delta}$ greater than the gap Δ between expected rewards, which drastically reduces the exploration budget of our algorithm and therefore of its expected regret.

This paper is divided in six sections. We start with a formal description of the addressed problem in Section 2. Then, we introduce in Section 3 the related work attached to this problem, its extensions and its problem realm. Section 4 is about our methodology, concerning the assumptions made and their consequences. After defining the precise context of the problem, we propose two versions of GRAB in Section 5 and we prove that they enjoy a regret linear in L in Section 6. Section 7 finally presents numerical results against a state-of-the-art algorithm.

2. Setting

We consider the problem of delivering monopartite matchings in a weighted graph, and give our settings. For any integer n , let $[n]$ be the set $\{1, \dots, n\}$. An instance of our problem is a couple $(L, (\rho_{i,j})_{i,j \in [2L]^2})$ where $2L$ is the number of elements to be matched, and for all elements $i, j \in [2L]^2$, $\rho_{i,j}$ is the probability of success associated to the elements i, j . A success could be, if i and j are players, that a game occurs, meaning that both player i and player j chose to play the game. Moreover, we assume that there exists a vector $\theta \in \mathbb{R}^{2L}$ such that $\rho_{i,j} = \theta_i \theta_j$, meaning that ρ is of rank 1. For any element i , θ_i represent the success rate associated to player i alone. Finally, without loss of generality, we add the hypothesis that $\theta_1 \geq \dots \geq \theta_{2L}$ to simplify the notations in the rest of the paper.

A matching algorithm is only aware of L , and has to deliver T consecutive matchings. A matching is a partition \mathbf{a} of $[2L]$ made of couples, and we note $\mathbf{a}(t)$ the matching returned by the algorithm at each iteration t . We note that couples are not ordered, which mean that $\forall i, j \in [2L], i \neq j, \{i, j\} = \{j, i\}$. We denote \mathcal{A} the set of all partitions of $[2L]$ made of couples, which correspond to the set of arms of the bandit setting. Formally, \mathcal{A} contains every \mathbf{a} such that:

$$\begin{cases} \bigcup_{\{i,j\} \in \mathbf{a}} \{i, j\} = [2L] \\ \forall c \in \mathbf{a}, |c| = 2 \\ |\mathbf{a}| = L \end{cases}$$

We name *optimum matching* the matching $\mathbf{a}^* \stackrel{\text{def}}{=} \{\{2k-1, 2k\} : k \in [L]\}$.

At iteration t , after recommending the matching $\mathbf{a}(t)$, the algorithm receives a semi-bandit feedback $\mathbf{c}(t)$ which contains the result of a Bernoulli associated to each couple $\{i, j\} \in \mathbf{a}(t)$, noted $\mathbf{c}_{i,j}(t)$, of parameter $\rho_{i,j}$. We address

the stochastic case, meaning that the way of generating $\mathbf{c}(t)$ is fixed over iterations and will not change. We focus on the *expected reward*

$$\mathbb{E} \left[\sum_{\{i,j\} \in \mathbf{a}(t)} \mathbf{c}_{i,j}(t) \right] = \sum_{\{i,j\} \in \mathbf{a}(t)} \rho_{i,j}.$$

Due to the rank-1 property, \mathbf{a}^* maximizes the expected reward, as expressed by Lemma 1

Lemma 1 (Maximum Expected Reward).

$$\sum_{\{i,j\} \in \mathbf{a}^*} \rho_{i,j} = \max_{\mathbf{a} \in \mathcal{M}} \sum_{\{i,j\} \in \mathbf{a}} \rho_{i,j}$$

We address the standard bandit objective, we minimize the cumulative expected regret which value is the sum over iterations of the difference between the maximum expected reward and the expected reward of the proposed matching:

$$R(T) = \sum_{t=1}^T \sum_{\{i,j\} \in \mathbf{a}^*} \rho_{i,j} - \sum_{\{i,j\} \in \mathbf{a}(t)} \rho_{i,j}.$$

3. Related Work

The problem addressed in this paper is a combinatorial bandit problem with semi-bandit feedback (Cesa-Bianchi & Lugosi, 2012). Combinatorial bandits have recently known several improvements (Combes et al., 2015; Cuvelier et al., 2021; Degenne & Perchet, 2016; Perrault et al., 2020; Wang & Chen, 2018), but these improvements do not account for a strong property of our setting. When applied to our, the state of the art combinatorial bandit algorithm, ESCB (Combes et al., 2015), suffers a regret of $\mathcal{O}(\frac{L^2 \log^2(L) \log(T)}{\Delta})$, Δ being the expected reward gap between an optimal matching and the best sub-optimal matching. A recent contribution to both monopartite and bipartite cases (Sentenac et al., 2021) has, among other things, lowered the monopartite regret bound to a scale of $\mathcal{O}(\frac{L \log(L)}{\Delta} \log(T))$, through the algorithm simple adaptative matching (SAM). The principle of this algorithm is, knowing L and T , to initially group all players together and, at each iteration, divide the current groups in two if the matching algorithm is certain with high probability that every players in one group are better than the other group players. The criterion used to split a group in two parts requires the knowledge of both the total number of players $2L$ and the horizon T .

The bipartite case has also been recently addressed (Gauthier et al., 2021), proposing the algorithm *parametric Graph for unimodal RAnking Bandit* (GRAB) and lowering the expected regret bound to $\mathcal{O}(\frac{L}{\Delta} \log(T))$. The principle of GRAB is based on a graph w.r.t. which the expected reward is unimodal. At each iteration, GRAB elects a node of this

graph as leader and choose from this leader the proposed recommendation. By acquiring knowledge from semi-bandit feedback, GRAB is able to search through nodes and to converge to the optimal recommendation. We show in this paper that we may adapt GRAB to the monopartite case and therefore get a lower bound smaller than the one proposed by SAM. Moreover, we show that by using a new criteria to select the played matching, we increase the value of Δ from $(\theta_i - \theta_{i'}) (\theta_j - \theta_{j'})$ for SAM and the original GRAB to $\theta_i (\theta_{i'} - \theta_{j'})$ and $\theta_j (\theta_{i'} - \theta_{j'})$, where i, i', j, j' are players.

3.1. Unimodality

Our approach builds upon the unimodality of the expected reward, so let us briefly remind corresponding main results.

The *unimodality* has been defined in (Cope, 2009; Yu & Mannor, 2011) and refined in (Combes & Proutière, 2014). Let $G = (V, E)$ be an undirected graph whose vertices are the arms of the bandit, noted $V = \{1, \dots, K\}$, and edges carries a partial order on expected rewards. We assume that there exists a unique arm $k^* \in V$ called optimum such that its expected reward μ_{k^*} is maximal. Finally, $\forall k \in V, k \neq k^*$ there exists a path $p = (k_1 = k, k_2, \dots, k_{m-1}, k_m = k^*)$ of length $m \in [m]$ depending of k , with $(k_i)_{i \in [m]} \in V^{[m]}$, such that $\forall i \in [m-1], (k_i, k_{i+1}) \in E$ and $\mu_i < \mu_{i+1}$. The bandit algorithm is aware of G but ignores the partial order induced by the edges. By relying on G , the algorithm is able of browsing efficiently among the arms and to converge to the optimum.

Both GRAB (Gauthier et al., 2021) and OSUB (Combes & Proutière, 2014) are designed to benefit from the unimodality. At each iteration, they recommend either the best arm sofar (a.k.a. *leader*) or an arm in its neighborhood. The restriction of the exploration to this neighborhood induces a regret which scales as $\mathcal{O}(\frac{\gamma}{\Delta} \log T)$, with γ the maximum degree of G , instead of $\mathcal{O}(\frac{|V|}{\Delta} \log T)$ with independent arms.

4. Methodology

In this section we add an assumption to the setting and clarify its consequences. While this assumption is not required by GRAB to get a logarithmic regret, our theoretical analysis needs it to get the linear-in- L regret. Note that (Sentenac et al., 2021) assumes the same property in its analysis. Whether this assumption may be discarded or not remains an opened question.

Our matching algorithm is based on an undirected graph provided with useful properties. We propose in the following a definition of this graph and its properties.

Each vertex is a matching provided with an order over its couples. We call them ordered matchings, note them $\tilde{\mathbf{a}}$, and we are now able to refer to the k -th couple as $\tilde{\mathbf{a}}_k$.

An edge exists between two nodes $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{a}}'$ if, and only if, the first can be obtained by switching one element of two successive couples. We note this property $\tilde{\mathbf{a}}' = \tilde{\mathbf{a}}_{[i \leftrightarrow j]}$, i, j being the two elements switched, and we formally define it as

$\exists k \in [L-1]$, and let $\{i, i'\} = \tilde{\mathbf{a}}_k, \{j, j'\} = \tilde{\mathbf{a}}_{k+1}$,

$$\forall n \in [L], \begin{cases} \tilde{\mathbf{a}}'_n = \{j, i'\} & , \text{ if } n = k \\ \tilde{\mathbf{a}}'_n = \{i, j'\} & , \text{ if } n = k + 1 \\ \tilde{\mathbf{a}}'_n = \tilde{\mathbf{a}}_n & , \text{ otherwise} \end{cases}$$

We also add a method to switch from ordered matching to matching, that we call *Set*, defined for any given ordered matching $\tilde{\mathbf{a}}$ as

$$\text{Set}(\tilde{\mathbf{a}}) \stackrel{\text{def}}{=} \{\tilde{\mathbf{a}}_k : k \in [L]\},$$

the matching constituted of the elements of $\tilde{\mathbf{a}}$.

We introduce to the problem setting a new assumption of *inter-pair strict order*, used later in the proof of Lemma 2 and stating that

Assumption 1 (inter-pair strict order). $\forall k, k' \in [L]^2, k < k' \implies \min(\theta_{2k-1}, \theta_{2k}) > \max(\theta_{2k'-1}, \theta_{2k'})$

We define the property π over an ordered matching $\tilde{\mathbf{a}}$ as

$$\forall k \in [L-1], \{i, j\} = \tilde{\mathbf{a}}_k, \{i', j'\} = \tilde{\mathbf{a}}_{k+1}, \rho_{i,j} \geq \rho_{i',j'}.$$

This property, noted $\pi(\tilde{\mathbf{a}})$, ensures that an order over ρ is respected within $\tilde{\mathbf{a}}$.

From this property, we construct the optimum ordered matching, called *optimum leader* and noted $\tilde{\mathbf{a}}^*$, as

$$\text{Set}(\tilde{\mathbf{a}}^*) = \mathbf{a}^* \quad \pi(\tilde{\mathbf{a}}^*) \quad (1)$$

Lemma 2 (*optimum leader uniqueness*). *Under Assumption 1, $\tilde{\mathbf{a}}^*$ is unique.*

From all the previous definitions and lemmas, we can deduce a relaxed unimodality property over the graph.

Lemma 3 (relaxed unimodality). *Under Assumption 1, for any ordered matching $\tilde{\mathbf{a}}$ satisfying $\pi(\tilde{\mathbf{a}})$, if $\tilde{\mathbf{a}} \neq \tilde{\mathbf{a}}^*$, we have*

$$\exists k \in [L-1], \exists i \in \tilde{\mathbf{a}}_k, \exists j \in \tilde{\mathbf{a}}_{k+1}, \mu_{\tilde{\mathbf{a}}} < \mu_{\tilde{\mathbf{a}}_{[i \leftrightarrow j]}}$$

Proof. Proving this lemma is equivalent to proving that for any ordered matching $\tilde{\mathbf{a}}$, one of the following properties is satisfied:

$$\neg \pi(\tilde{\mathbf{a}}) \quad (2)$$

$$\tilde{\mathbf{a}} = \tilde{\mathbf{a}}^* \quad (3)$$

$$\exists k \in [L-1], \exists i \in \tilde{\mathbf{a}}_k, \exists j \in \tilde{\mathbf{a}}_{k+1}, \mu_{\tilde{\mathbf{a}}} < \mu_{\tilde{\mathbf{a}}_{[i \leftrightarrow j]}} \quad (4)$$

To start with, $\tilde{\mathbf{a}}$ satisfy one of the following properties, because the last is the conjunction of the negation of the previous ones.

$$\exists k \in [L-1], \{i, i'\} = \mathbf{a}_k, \{j, j'\} = \mathbf{a}_{k+1}, \rho_{i, i'} < \rho_{j, j'} \quad (5)$$

$$\exists k \in [L-1], \exists i \in \mathbf{a}_k, \exists j \in \mathbf{a}_{k+1}, \theta_i < \theta_j \quad (6)$$

$$\forall k \in [L-1], \begin{cases} \{i, i'\} = \mathbf{a}_k, \{j, j'\} = \mathbf{a}_{k+1}, \rho_{i, i'} \geq \rho_{j, j'} \\ \exists i \in \mathbf{a}_k, \exists j \in \mathbf{a}_{k+1}, \theta_i \geq \theta_j \end{cases} \quad (7)$$

If $\tilde{\mathbf{a}}$ satisfies Property (5), π is not respected. Thus, Property (2) is verified.

If $\tilde{\mathbf{a}}$ satisfies Property (6), we can note $\{i, i'\} = \tilde{\mathbf{a}}_k$ and $\{j, j'\} = \tilde{\mathbf{a}}_{k+1}$. Without loss of generality, we assume that $\theta_i \geq \theta_{i'}$ and $\theta_j \geq \theta_{j'}$, Property (6) then can be precised to $\theta_{i'} < \theta_j$. Therefor,

$$\begin{aligned} \mu_{\tilde{\mathbf{a}}} - \mu_{\tilde{\mathbf{a}}_{[i' \leftrightarrow j]}} &= \theta_i \theta_{i'} + \theta_j \theta_{j'} - \theta_i \theta_j - \theta_{i'} \theta_{j'} \\ &= (\theta_i - \theta_{j'}) (\theta_{i'} - \theta_j) \end{aligned}$$

We deduce $\theta_{i'} - \theta_j < 0$ from Property (6). We then can assume that $\pi(\tilde{\mathbf{a}})$, otherwise Property (2) would have been satisfied and the lemma would still be verified. Thus we have $\theta_i \theta_{i'} \geq \theta_j \theta_{j'}$. However, by Property (6), $\theta_i \theta_{i'} < \theta_i \theta_j$, meaning that $\theta_i \theta_j > \theta_j \theta_{j'}$ and finally $\theta_i > \theta_{j'}$. To conclude, $\mu_{\tilde{\mathbf{a}}} - \mu_{\tilde{\mathbf{a}}_{[i' \leftrightarrow j]}} < 0$, fulfilling Property (4)

If $\tilde{\mathbf{a}}$ satisfies Property (7) and assuming $\pi(\tilde{\mathbf{a}})$, we have $\forall k \in [L-1]$, let $\{i, i'\} = \tilde{\mathbf{a}}_k$, and $\{j, j'\} = \tilde{\mathbf{a}}_{k+1}$, $i \geq i' \geq j \geq j'$. Thus, $\tilde{\mathbf{a}}$ is, by construction, $\tilde{\mathbf{a}}^*$, fulfilling Property (3).

In every cases, one of Properties (2), (3) and (4) is fulfilled, proving that for any ordered matching $\tilde{\mathbf{a}}$ satisfying π , if $\tilde{\mathbf{a}} \neq \tilde{\mathbf{a}}^*$,

$$\exists k \in [L-1], \exists i \in \tilde{\mathbf{a}}_k, \exists j \in \tilde{\mathbf{a}}_{k+1}, \mu_{\tilde{\mathbf{a}}} < \mu_{\tilde{\mathbf{a}}_{[i \leftrightarrow j]}}$$

□

This property is widely inspired from the unimodality described in Section 3.1. Indeed, when $\tilde{\mathbf{a}}$ satisfies π , either $\tilde{\mathbf{a}}$ is optimal or, if $\tilde{\mathbf{a}}$ is sub-optimal, there is at least one of its neighbor that has a better regret. We will use this property in the proof of Theorem 1.

5. GRAB Algorithm

Our algorithm, GRAB, is inspired by the unimodal ranking bandit algorithm GRAB (Gauthier et al., 2021). At each iteration t , we select an arm $\mathbf{a}(t)$ in the neighborhood of the leader, the current best arm according to the bandit w.r.t. $\hat{\rho}$ defined latter, which is learnt online. The main difference

is that we don't use the true argmax, but instead a greedy approximation, g-argmax, when electing the leader. The second difference is that the leader is an ordered matching, but not a matching. This is materialized by the notation *Set*. GRAB is described in Algorithm 1, it has two versions, GRAB and GRAB⁺, and uses the following notations.

When applied to two elements i and j and at each iteration t , we denote

$$\hat{\rho}_{i,j}(t) \stackrel{def}{=} \frac{1}{T_{i,j}(t)} \sum_{s=1}^{t-1} \mathbb{1}\{\{i, j\} \in \mathbf{a}(s)\} \mathbf{c}_{i,j}(s)$$

the average number of games played before iteration t while pairing i with j , where $\mathbf{c}_{i,j}(s)$ equals 1 if $\{i, j\}$ did lead to a game at iteration s and 0 otherwise, and

$$T_{i,j}(t) \stackrel{def}{=} \sum_{s=1}^{t-1} \mathbb{1}\{\{i, j\} \in \mathbf{a}(s)\}$$

is the number of time i has been paired with j until iteration $t-1$. We set $\hat{\rho}_{i,j}(t)$ to 0 when $T_{i,j}(t)$ equals 0.

At each iteration t we denote $\tilde{\ell}(t)$ the *leader*, the ordered matching returned by g-argmax. For any ordered matching $\tilde{\mathbf{a}}$, we denote

$$\tilde{T}_{\tilde{\mathbf{a}}}(t) \stackrel{def}{=} \sum_{s=1}^{t-1} \mathbb{1}\{\tilde{\ell}(s) = \tilde{\mathbf{a}}\}$$

the number of time $\tilde{\mathbf{a}}$ has been leader until iteration $t-1$.

Finally, we denote the optimistic probability of a couple $\tilde{\mathbf{a}}_k = \{i, j\}$ by

$$q_{i,j}(t) \stackrel{def}{=} f(\hat{\rho}_{i,j}(t), T_{i,j}(t), \tilde{T}_{\tilde{\ell}(t)}(t) + 1)$$

where $f(\hat{\rho}, s, t)$ stands for

$$\sup\{p \in [\hat{\rho}, 1] : s \times \text{kl}(\hat{\rho}, p) \leq \log(t) + 3 \log(\log(t))\},$$

with

$$\text{kl}(p, q) \stackrel{def}{=} p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right)$$

the *Kullback Leibler divergence* from a Bernoulli distribution of mean p to a Bernoulli distribution of mean q . By definition, we set $f(\hat{\rho}, s, t)$ to ∞ when $t = 0$, prioritizing exploration.

Let now explain GRAB. First, GRAB uses g-argmax to elicit the leader. This algorithm is described in Algorithm 2 and returns a greedy approximation of

$$\arg\max_{\tilde{\mathbf{a}}} \sum_{k=1}^L \hat{\rho}_{\tilde{\mathbf{a}}_k}(t). \quad (8)$$

Algorithm 1 GRAB for monopartite matching

Require: number of items $2L$, optimistic criterion V1 (original) or V2 (new)

- 1: **for** $t = 1, 2, \dots$ **do**
- 2: $\tilde{\ell}(t) \leftarrow \text{g-argmax}([2L])$
- 3: recommend

$$\mathbf{a}(t) = \begin{cases} \text{Set}(\tilde{\ell}(t)) & , \text{ if } \frac{\tilde{T}_{\tilde{\ell}(t)}(t)}{2L-1} \in \mathbb{N}, \\ \underset{\mathbf{a} \in \{\text{Set}(\tilde{\ell}(t))\} \cup \mathcal{N}_{\text{Set}}(\tilde{\ell}(t))}{\text{argmax}} \sum_{\{i, i'\} \in \mathbf{a}} q_{i, i'}(t) & , \text{ if } \frac{\tilde{T}_{\tilde{\ell}(t)}(t)}{2L-1} \notin \mathbb{N} \text{ and V1} \\ \underset{\mathbf{a} = \text{Set}(\tilde{\ell}(t))}{\text{argmax}} \max(0, q_{i, j'}(t) - q_{i, i'}(t), q_{j, i'}(t) - q_{i, i'}(t)) & , \text{ if } \frac{\tilde{T}_{\tilde{\ell}(t)}(t)}{2L-1} \notin \mathbb{N} \text{ and V2} \\ \text{or } \mathbf{a} \in \mathcal{N}_{\text{Set}}(\tilde{\ell}(t)) \text{ with } \mathbf{a} = \text{Set}(\tilde{\ell}(t)_{[i \leftrightarrow j]}), & \\ \{i, i'\} \in \text{Set}(\tilde{\ell}(t)) \text{ and } \{j, j'\} \in \text{Set}(\tilde{\ell}(t)) & \end{cases}$$

$$\text{where } \mathcal{N}_{\text{Set}}(\tilde{\mathbf{a}}) = \bigcup_{k=1}^{L-1} \{\text{Set}(\tilde{\mathbf{a}}_{[e_1 \leftrightarrow e_2]}) : e_1 \in \tilde{\mathbf{a}}_k, e_2 \in \tilde{\mathbf{a}}_{k+1}\}$$

- 4: observe the games vector $\mathbf{c}(t)$
- 5: **end for**

Algorithm 2 g-argmax: a greedy approximation of argmax

Require: Set \mathcal{E} of items to chose

Ensure: An approximation of the solution of Equation (8)

- 1: $\tilde{\mathbf{a}} \leftarrow$ empty list
- 2: **for** $k \in \left[\frac{|\mathcal{E}|}{2} \right]$ **do**
- 3: $e_1, e_2 \leftarrow \underset{e_1 \in \mathcal{E}, e_2 \in \mathcal{E} \setminus \{e_1\}}{\text{argmax}} \hat{p}_{e_1, e_2}(t)$
- 4: $\tilde{\mathbf{a}}.\text{append}(\{e_1, e_2\})$
- 5: $\mathcal{E} \leftarrow \mathcal{E} \setminus \{e_1, e_2\}$
- 6: **end for**

The couples of the leader are elected one by one. Algorithm g-argmax choses among the remaining couples the current best possible one w.r.t. $\hat{p}_{i, j}(t)$. Note that by Lemma 5, the returned leader $\tilde{\ell}(t)$ satisfies most of the time $\pi(\tilde{\ell}(t))$.

Finally, GRAB, given in Algorithm 1, identifies the leader $\tilde{\ell}(t)$ and recommends either the matching associated to $\tilde{\ell}(t)$ each $(2L - 1)$ -th iterations, or the best matching in the leader's neighborhood in regards to an optimistic criterion. With the original version, the criterion is an optimistic estimate of the number of games to be played, while with the new version it is an optimistic estimate of the increase of the game probability when replacing one of both players in the couple $\{i, i'\}$.

To conclude the presentation of our algorithm, here we discuss its initialization, its time complexity, and the utility of some choices.

Remark 1 (Utility of unordered matching). *The neighborhood $\mathcal{N}_{\text{Set}}(\tilde{\mathbf{a}})$ is defined after unordered matchings to divide its size by a factor 2. In our recommendations, couples do not need to be ordered, thus some recommen-*

dations in our leader's neighborhood are strictly equivalents w.r.t. the expected reward. For example, for any leader $\tilde{\ell}$ any $k \in [L - 1]$, by noting $\{i, i'\} = \tilde{\ell}_k$ and $\{j, j'\} = \tilde{\ell}_{k+1}$, $\text{Set}(\tilde{\ell}(t)_{[i \leftrightarrow j]}) = \text{Set}(\tilde{\ell}(t)_{[i' \leftrightarrow j']})$ because couples are not ordered anymore and the permutation of i with j produces the same couples as the permutation of i' with j' , leaving other couples unchanged. Therefore, $\{\text{Set}(\tilde{\ell}(t)_{[i \leftrightarrow j]}) : i \in \tilde{\mathbf{a}}_k, j \in \tilde{\mathbf{a}}_{k+1}\}$ is composed of both matchings $\text{Set}(\tilde{\ell}(t)_{[i \leftrightarrow j]})$ and $\text{Set}(\tilde{\ell}(t)_{[i' \leftrightarrow j']})$ while we may expect 4 matchings given the definition.

Remark 2 (Initialisation). *For all $i, j \in [L]^2, i \neq j$, we initialize $q_{i, j}$ to ∞ to ensure that every neighbor of an arm which is often the leader is played at least once.*

Remark 3 (Algorithmic Complexity). *The computation time of GRAB is polynomial in L .*

First, to elect the leader, g-argmax carry out an iteration over the number of couples L . At each iteration, it performs an argmax over at most $L \times L$ values. Thus, the elicitation of the leader is done in $\mathcal{O}(L^3)$ operations.

Then, the maximisation when recommending $\mathbf{a}(t)$ is over a set of $2L - 2$ recommendations. Indeed, for all $k \in [L - 1]$, $\{\text{Set}(\tilde{\ell}(t)_{[i \leftrightarrow j]}) : i \in \tilde{\mathbf{a}}_k, j \in \tilde{\mathbf{a}}_{k+1}\}$ is of size 2, according to Remark 1. The union of those disjointed sets of size 2 finally gives us a set of $2L - 2$ matchings.

Finally for each matching \mathbf{a} , the computation of the optimistic criterion cost $\mathcal{O}(1)$ operations. This is obvious with the new criterion, and it becomes clear with the original criterion as soon as we remark that the maximization of

$\sum_{\{i,i'\} \in \mathbf{a}} q_{i,i'}(t)$ is equivalent to the maximization of

$$B_{\mathbf{a}}(t) = \sum_{\{i,j\} \in \mathbf{a}} q_{i,j}(t) - \sum_{\{i,j\} \in \text{Set}(\tilde{\ell}(t))} q_{i,j}(t)$$

which, by elimination of common terms, reduces to the sum of at most four $q_{i,j}(t)$ values.

Overall, each recommendation is done in $\mathcal{O}(L^3)$.

6. Theoretical Analysis

Before looking at the detailed theoretical analysis, let us give Theorem 1 and Theorem 2 which respectively express the regret of GRAB with the original criterion and the regret GRAB with the new criterion (denoted GRAB⁺).

Theorem 1 (Upper-bound on the regret of GRAB assuming inter-pair strict order). *Let $(L, (\rho_{i,j})_{i,j \in [2L]^2})$ be an instance of online matching satisfying Assumption 1. Then, the expected regret of GRAB satisfies*

$$\begin{aligned} R(T) &\leq \sum_{\mathbf{a} \in \mathcal{N}_{\text{Set}}(\tilde{\mathbf{a}}^*)} \frac{8}{\Delta_{\mathbf{a}}} \log T + \mathcal{O}(\log \log T) \\ &= \mathcal{O}\left(\frac{L}{\Delta} \log T\right) \end{aligned}$$

where $\Delta_{\mathbf{a}} \stackrel{\text{def}}{=} \mu_{\mathbf{a}^*} - \mu_{\mathbf{a}}$ and

$$\Delta \stackrel{\text{def}}{=} \min_{k \in [L-1], \{i,i'\} = a_k, \{j,j'\} = a_{k+1}} (\theta_i - \theta_j)(\theta_{i'} - \theta_{j'})$$

Theorem 2 (Upper-bound on the regret of GRAB⁺ assuming inter-pair strict order). *Let $(L, (\rho_{i,j})_{i,j \in [2L]^2})$ be an instance of online matching satisfying Assumption 1. Then, the expected regret of GRAB⁺ satisfies*

$$\begin{aligned} R(T) &\leq \sum_{k \in [L-1], \{i,i'\} = a_k, \{j,j'\} = a_{k+1}} \frac{8\Delta_{\tilde{\mathbf{a}}^*_{[i' \leftrightarrow j]}}}{\tilde{\Delta}_{i,i',j'}^2} \log T + \mathcal{O}(\log \log T) \\ &= \mathcal{O}\left(\frac{L\tilde{\Delta}}{\tilde{\Delta}^2} \log T\right) \end{aligned}$$

where $\Delta_{\mathbf{a}} \stackrel{\text{def}}{=} \mu_{\mathbf{a}^*} - \mu_{\mathbf{a}}$, $\Delta_{i,i',j} \stackrel{\text{def}}{=} \theta_i(\theta_{i'} - \theta_j)$

$$\Delta \stackrel{\text{def}}{=} \min_{k \in [L-1], \{i,i'\} = a_k, \{j,j'\} = a_{k+1}} (\theta_i - \theta_j)(\theta_{i'} - \theta_{j'})$$

$$\text{and } \tilde{\Delta} \stackrel{\text{def}}{=} \min_{k \in [L-1], \{i,i'\} = a_k, \{j,j'\} = a_{k+1}} \theta_i(\theta_{i'} - \theta_{j'})$$

Both Theorems are proven following the same path as the one in (Gauthier et al., 2021). The idea is to (1) apply standard bandit analysis to limit the expected regret when the leader $\tilde{\ell}(t)$ is the optimum leader $\tilde{\mathbf{a}}^*$, and to (2) upper-bound the number of iterations t in which $\tilde{\ell}(t) \neq \tilde{\mathbf{a}}^*$ by a $\mathcal{O}(\log(\log(T)))$. To assume $\pi(\tilde{\ell}(t))$ in both previous cases,

we also show that (3) we can upper-bound the number of iterations such that $\neg\pi(\tilde{\ell}(t))$ by a $\mathcal{O}(1)$.

Despite the change of the setting, the Theorem and the lemmas used to handle these 3 points are similar (in their expression) to the one in (Gauthier et al., 2021). We give them here for the original criterion and we defer their proof to a complete version of the paper.

Step (1) consists in applying Theorem 2 from (Gauthier et al., 2021) which we remind hereafter.

Theorem 3 (Upper-Bound on the Regret of KL-CombUCB (Theorem 2 of from (Gauthier et al., 2021))). *We consider a combinatorial semi-bandit setting. Let E be a set of elements and $\mathcal{A} \subseteq 0, 1^E$ be a set of arms, where each arm \mathbf{a} is a subset of E . Let's assume that the reward when drawing the arm $\mathbf{a} \in \mathcal{A}$ is $\sum_{e \in \mathbf{a}} c_e$, where for each element $e \in E$, c_e is an independant draw of a Bernoulli distribution of mean $\rho_e \in [0, 1]$. Therefore, the expected reward when drawing the arm $\mathbf{a} \in \mathcal{A}$ is $\mu_{\mathbf{a}} = \sum_{e \in \mathbf{a}} \rho_e$.*

When facing this bandit setting, KL-CombUCB (CombUCB1 equipped with Kullback-Leibler indices) fullfils

$$\forall \mathbf{a} \in \mathcal{A} \text{ s.t. } \mu_{\mathbf{a}} \neq \mu^*,$$

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\{\mathbf{a}(t) = \mathbf{a}\}\right] \leq \frac{2K_{\mathbf{a}}^2}{\Delta_{\mathbf{a}}^2} \log T + \mathcal{O}(\log \log T),$$

and hence

$$\begin{aligned} R(T) &\leq \sum_{\mathbf{a} \in \mathcal{A}: \mu_{\mathbf{a}} \neq \mu^*} \frac{2K_{\mathbf{a}}^2}{\Delta_{\mathbf{a}}} \log T + \mathcal{O}(\log \log T) \\ &= \mathcal{O}\left(\frac{|\mathcal{A}|K_{\max}^2}{\Delta_{\min}} \log T\right) \end{aligned}$$

Moreover, the number of iterations where GRAB elects a sub-optimal leader, is upper-bounded by Lemma 4.

Lemma 4 (Upper-Bound on the number of iterations where GRAB elect a sub-optimal leader satisfying π). *Under the hypothesis of Theorem 1 and using its notation,*

$$\forall \tilde{\ell} \in \mathcal{A}, \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\{\tilde{\ell}(t) = \tilde{\ell}\}\right] = \mathcal{O}(\log \log T).$$

Finally, the number of iterations where GRAB elects a leader that does not satisfy π is controlled by Lemma 5.

Lemma 5 (Upper-Bound on the number of iterations where GRAB elect a leader that does not satisfy π). *Under the hypothesis of Theorem 1 and using its notation,*

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\{\neg\pi(\tilde{\ell}(t))\}\right] = \mathcal{O}(1).$$

From these results, we are able to prove Theorem 1.

Proof of Theorem 1. The proof is based on a decomposition of the set of the iterations $[T]$. The goal is to apply the previous results to make the proof easier.

$$[T] = \bigcup_{\mathbf{a} \in \{\mathbf{a}^*\} \cup \mathcal{N}_{Set}(\tilde{\mathbf{a}}^*)} \{t \in [T] : \tilde{\ell}(t) = \tilde{\mathbf{a}}^*, \mathbf{a}(t) = \mathbf{a}\} \\ \cup \{t \in [T] : \tilde{\ell}(t) \neq \tilde{\mathbf{a}}^*, \pi(\tilde{\ell}(t))\} \cup \{t \in [T] : \neg \pi(\tilde{\ell}(t))\}$$

As for any recommendation \mathbf{a} , $\Delta_{\mathbf{a}} \leq K_{\mathbf{a}}$, this decomposition gives us the following inequality:

$$R(T) \leq \sum_{\mathbf{a} \in \mathcal{N}_{Set}(\tilde{\mathbf{a}}^*)} \Delta_{\mathbf{a}} A_{\mathbf{a}} + KB + KC$$

with

$$A_{\mathbf{a}} = \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \{ \tilde{\ell}(t) = \tilde{\mathbf{a}}^*, \mathbf{a}(t) = \mathbf{a} \} \right], \\ B = \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \{ \tilde{\ell}(t) \neq \tilde{\mathbf{a}}^*, \pi(\tilde{\ell}(t)) \} \right], \\ C = \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \{ \neg \pi(\tilde{\ell}(t)) \} \right].$$

The terme $A_{\mathbf{a}}$ is smaller than the number of time the arm \mathbf{a} is chosen by KL-CombUCB playing on the set of arms $\{\tilde{\mathbf{a}}^*\} \cup \mathcal{N}_{Set}(\tilde{\mathbf{a}}^*)$. As any arms differs with $\tilde{\mathbf{a}}^*$ at at most two positions, Theorem 3 upper-bounds $A_{\mathbf{a}}$ by

$$\frac{8}{\Delta_{\mathbf{a}}^2} \log T + \mathcal{O}(\log \log T)$$

and thus, as $|\mathcal{N}_{Set}(\tilde{\mathbf{a}}^*)| = 2L - 2$, $\sum_{\mathbf{a} \in \mathcal{N}_{Set}(\tilde{\mathbf{a}}^*)} \Delta_{\mathbf{a}} A_{\mathbf{a}} = \mathcal{O}(\frac{L}{\Delta} \log T)$.

By Lemma 4, we upper-bound the terme B by

$$B = \mathcal{O}(\log \log T),$$

and by Lemma 5, we upper-bound the terme C by

$$C = \mathcal{O}(1),$$

The expected regret of GRAB is finally the sum of those three terms, which end the proof. \square

7. Experimental Results

In this section, we compare empirically GRAB and GRAB⁺ to SAM (Sentenac et al., 2021). To evaluate the performances, we use simulated data over two experiments with settings similar to the ones used in (Sentenac et al., 2021). Each figure compares the three algorithms when varying the parameters of these settings:

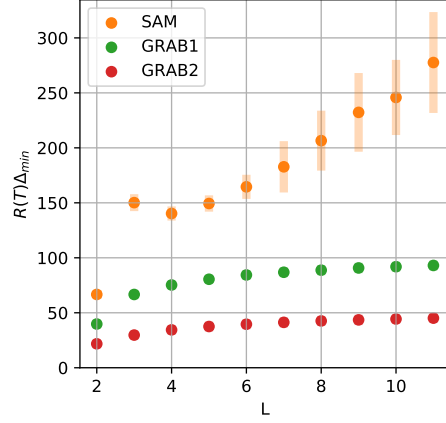


Figure 1. Normalized regret comparison of SAM to GRAB and GRAB⁺ versus L in the settings of the experience 1.

- L , the number of couples to recommend (thus, there are $2L$ players);
- T , the horizon of the simulations;
- δ , the gap between each player of consecutive couples in the optimal matching, it can be interpreted as $\delta = \sqrt{\Delta}$;
- μ , the mean of all θ_i .

Remark 4 (GRAB simplification). *For the sake of fairness, our implementation of GRAB use the same optimistic formula than SAM:*

$$q_{i,j}(t) \stackrel{def}{=} \hat{\rho}_{i,j}(t) + \sqrt{\frac{2 \log(t)}{T_{i,j}(t)}},$$

instead of $f(\hat{\rho}, s, t)$.

In the first experiment, the problems are instantiated with a tuple (L, δ) such that $0 < (L - 1)\delta \leq 1$. We fluctuate L between 2 and 11, while δ and T are respectively fixed to 0.1 and 10^6 . Finally, we construct a set of $2L$ players whose inner probability is, for each $i \in [L]$, $\theta_{2i-1} = \theta_{2i} = (L - i)\delta$, thus $\mu = \frac{L-1}{2}\delta$. As a result, shown in Figure 1, not only we scale, as established, as $\frac{L}{\Delta}$, but we also have a normalized regret divided by a factor 2 for GRAB and more than 5 for GRAB⁺ w.r.t. the normalized regret of SAM. In Figure 2, we focus on the expected regret evolution along iterations. The difference of strategy between SAM and our algorithms become clear: SAM explores as much as it can at the beginning, and when the algorithm is able to set an order between players with high probability, it divides them in groups, setting a partial order and defining its optimal matching. Meanwhile, our algorithms takes advantage of every information they have to immediately choose better matchings and to limit the comparison between players to

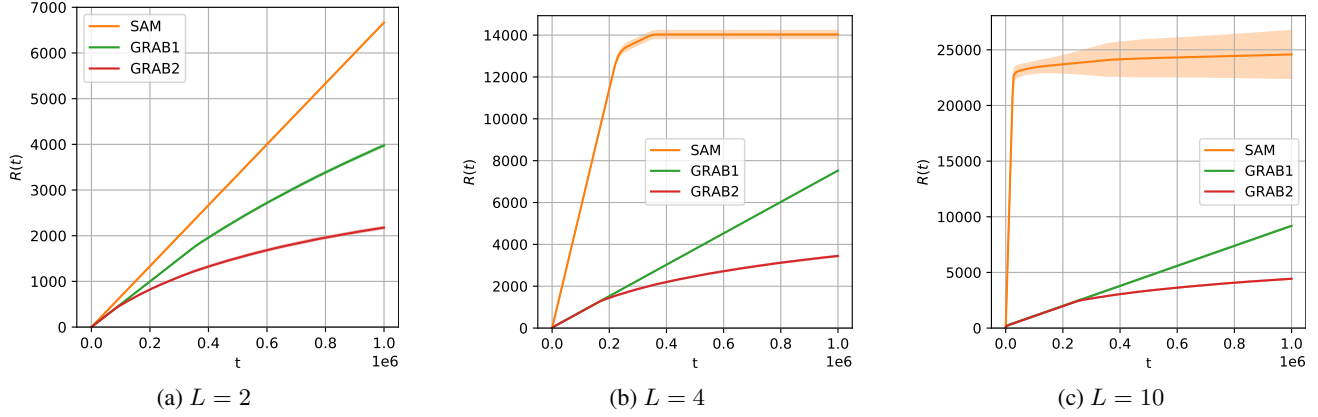


Figure 2. Regret comparison of SAM to GRAB and GRAB⁺ versus iteration t




Figure 3. Regret comparison of SAM to GRAB and GRAB⁺ versus μ in the settings of the experience 2.

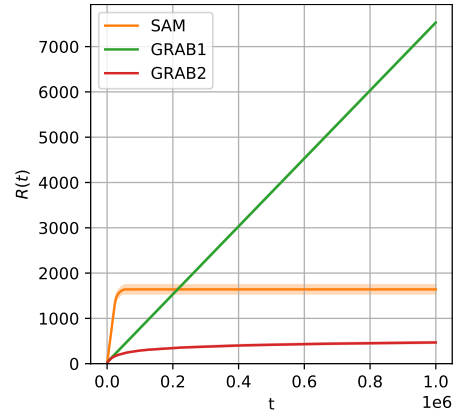


Figure 4. Regret evolution comparison of SAM to GRAB and GRAB⁺ among iterations in the settings of the experience 2 and with $\mu = 0.5$.

as few couples as possible. The early decision leads to a lower regret at first iterations, and the smaller set of couples to evaluate keep the regret low thereafter.

In the second experiment, we add μ to the tuple, therefore problems are instantiated by a tuple (L, μ, δ) such that $0 \leq \mu - (L - 1)\delta$ and $\mu + (L + 1)\delta \leq 1$. We then fluctuate μ between 0.5 and 0.8, while L , δ and T are respectively fixed to 4, 0.1 and 10^6 . Although it is disadvantageous for GRAB (as we see in Figure 1, others value of θ would have make it better than SAM), we decided to take this setting for the sake of comparison, since SAM did use this setting in its experimental evaluation. As we can see in Figure 3, results show that GRAB⁺ has still a lower regret than SAM. Figure 4 focuses on the case $\mu = 0.5$ and shows the regret evolution over time, highlighting, as Figure 2, the difference of strategy between algorithms. This figure also explains the lack of efficiency of GRAB by showing a linear regret, which means that it has not the time to learn properly the partial order between arms.

8. Conclusion

We tackle the problem of finding an optimal matching in a monopartite weighted graph. In order to solve this problem, we define and use a graph made of all the possible matchings, and use a property of unimodality over the expected reward of each matching. Our algorithms learn online a partial order between each player, by taking advantage of a semi-bandit feedback at each iteration, and give regret upper-bound in $\mathcal{O}(\frac{L}{\Delta} \log T)$ and $\mathcal{O}(\frac{L\Delta}{\Delta^2} \log T)$, reducing at least by a factor $\log L$ the bound w.r.t. the state-of-the-art algorithms. Moreover, unlike SAM, GRAB does not need T as an initial information.

References

- Cesa-Bianchi, N. and Lugosi, G. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.

- Combes, R. and Proutière, A. Unimodal bandits: Regret lower bounds and optimal algorithms. In *proc. of the 31st Int. Conf. on Machine Learning, ICML'14*, 2014.
- Combes, R., Talebi Mazraeh Shahi, M. S., Proutiere, A., et al. Combinatorial bandits revisited. *Advances in neural information processing systems*, 28, 2015.
- Cope, E. W. Regret and convergence bounds for a class of continuum-armed bandit problems. *IEEE Transactions on Automatic Control*, 54(6):1243–1253, 2009. doi: 10.1109/TAC.2009.2019797.
- Cuvelier, T., Combes, R., and Gourdin, E. Statistically efficient, polynomial-time algorithms for combinatorial semi-bandits. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 5(1):1–31, 2021.
- Degenne, R. and Perchet, V. Combinatorial semi-bandit with known covariance. *Advances in Neural Information Processing Systems*, 29, 2016.
- Gauthier, C.-S., Gaudel, R., Fromont, E., and Lompo, B. A. Parametric graph for unimodal ranking bandit. In *Proc. of the 38th Int. Conf. on Machine Learning, ICML'21*, pp. 3630–3639, 2021.
- Gauthier, C.-S., Gaudel, R., and Fromont, E. Unirank: Unimodal bandit algorithms for online ranking. In *Proc. of the 39th Int. Conf. on Machine Learning, ICML'22*, 2022.
- Komiyama, J., Honda, J., and Takeda, A. Position-based multiple-play bandit problem with unknown position bias. In *Advances in Neural Information Processing Systems 30, NIPS'17*, 2017.
- Lovász, L. and Plummer, M. D. *Matching theory*, volume 367. American Mathematical Soc., 2009.
- Mehta, A. Online matching and ad allocation. *Theoretical Computer Science*, 8(4):265–368, 2012.
- Perrault, P., Boursier, E., Valko, M., and Perchet, V. Statistical efficiency of thompson sampling for combinatorial semi-bandits. *Advances in Neural Information Processing Systems*, 33:5429–5440, 2020.
- Roth, A. E., Sönmez, T., and Ünver, M. U. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.
- Sentenac, F., Yi, J., Calauzènes, C., Perchet, V., and Vojnovic, M. Pure exploration and regret minimization in matching bandits. In *Proc. of the 38th Int. Conf. on Machine Learning, ICML'21*, 2021.
- Wang, S. and Chen, W. Thompson sampling for combinatorial semi-bandits. In *International Conference on Machine Learning*, pp. 5114–5122. PMLR, 2018.
- Wheaton, W. C. Vacancy, search, and prices in a housing market matching model. *Journal of political Economy*, 98(6):1270–1292, 1990.
- Yu, J. Y. and Mannor, S. Unimodal bandits. In *proc. of the 28th Int. Conf. on Machine Learning, ICML'11*, 2011.