
Efficient Online Learning for Dynamic k -Clustering

Dimitris Fotakis¹ Georgios Piliouras² Stratis Skoulakis²

Abstract

We study dynamic clustering problems from the perspective of online learning. We consider an online learning problem, called *Dynamic k -Clustering*, in which k centers are maintained in a metric space over time (centers may change positions) such as a dynamically changing set of r clients is served in the best possible way. The connection cost at round t is given by the p -norm of the vector consisting of the distance of each client to its closest center at round t , for some $p \geq 1$ or $p = \infty$. We present a $\Theta(\min(k, r))$ -regret polynomial-time online learning algorithm and show that, under some well-established computational complexity conjectures, *constant-regret* cannot be achieved in polynomial-time. In addition to the efficient solution of Dynamic k -Clustering, our work contributes to the long line of research on combinatorial online learning.

1. Introduction

Clustering problems are widely studied in Combinatorial Optimization literature due to their vast applications in Operational Research, Machine Learning, Data Science and Engineering (Williamson & Shmoys, 2011; Lin & Vitter, 1992; Charikar et al., 1999; Arya et al., 2001; Charikar & Guha, 1999; Jain & Vazirani, 2001; Kumar, 2012; Young, 2000; Li & Svensson, 2016; Charikar & Li, 2012; Kumar et al., 2010; Alamdari & Shmoys, 2018). Typically a fixed number of centers must be placed in a metric space such that a set of clients is served the best possible way. The quality of a clustering solution is captured through the p -norm of the vector consisting of the distance of each client to its closest center, for some $p \geq 1$ or $p = \infty$. For example k -median and k -means assume $p = 1$ and 2 respectively,

^{*}Equal contribution ¹Departement of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece ²Pillar of Engineering Systems and Design, Singapore University of Technology and Design, Singapore. Correspondence to: Dimitris Fotakis <fotakis@cs.ntua.gr>, Georgios Piliouras <georgios@sutd.edu.sg>, Stratis Skoulakis <efstratios@sutd.edu.sg>.

while k -center assumes $p = \infty$ (Lin & Vitter, 1992; Kumar et al., 2010; Alamdari & Shmoys, 2018).

Today's access on vast data (that may be frequently updated over time) has motivated the study of clustering problems in case of *time-evolving clients*, which dynamically change positions over time (de Keijzer & Wojtczak; Fotakis et al., 2019; Eisenstat et al.; An et al., 2017). In time-evolving clustering problems, centers may also change position over time so as to better capture the clients' trajectories. For example, a city may want to reallocate the units performing rapid tests for Covid-19 so as to better serve neighborhoods with more cases, the distribution of which may substantially change from day to day. Other interesting applications of dynamic clustering include viral marketing, epidemiology, facility location (e.g. schools, hospitals), conference planning etc. (Stehl e et al., 2011; Eisenstat et al.; Newman, 2003; Pastor-Satorras & Vespignani, 2001; Tanipathananandh et al., 2007).

Our work is motivated by the fact that in most settings of interest, clients can move in fairly complicated and unpredictable ways, and thus, an *a-priori knowledge* on such trajectories is heavily under question (most of the previous work assumes perfect knowledge on clients' positions over time (Eisenstat et al.; An et al., 2017; de Keijzer & Wojtczak; Fotakis et al., 2019)). To capture this lack of information we cast clustering problems under the perspective of *online learning* (Hazan, 2016). We study an online learning problem called *Dynamic k -Clustering* in which a *learner* selects at each round t , the positions of k centers trying to minimize the connection cost of some clients, the positions of which are unknown to the learner prior to the selection of the centers.

Online Learning Problem 1 (Dynamic k -Clustering).
Given a metric space $d : V \times V \mapsto \mathbb{R}_{\geq 0}$. At each round t ,

1. The learner selects a set $F_t \subseteq V$, with $|F_t| = k$, at which centers are placed.
2. The adversary selects the positions of the clients, denoted as R_t (after the selection of the positions of the centers by the learner).


3. The learner suffers the connection cost of the clients,

$$C_{R_t}(F_t) = \left(\sum_{j \in R_t} d(j, F_t)^p \right)^{1/p}$$

where $d(j, F_t)$ is the distance of client j to the closest center, $d(j, F_t) = \min_{i \in F_t} d_{ij}$.

Based on the past positions of the clients R_1, R_2, \dots, R_{t-1} an *online learning algorithm* must select at each round t , a set of k centers $F_t \subseteq V$ such that the connection cost of the clients over time is close to the connection cost of the *optimal (static) solution* F^* . If the cost of the online learning algorithm is at most α times the cost of F^* , the algorithm is called α -regret, whereas in case $\alpha = 1$, the algorithm is called *no-regret* (Hazan, 2016). Intuitively, a low-regret online learning algorithm converges to the optimal positions of the centers (with respect to the overall trajectories of the clients) by just observing the clients' dynamics.

Example 1. The clients are randomly generated according to a time-varying uniform distribution with radius 0.3 and center following the periodic trajectory $(\sin(\frac{2\pi \cdot t}{T}), \cos(\frac{2\pi \cdot t}{T}))$ for $t = 1, \dots, T$. The centers placed



by a (sufficiently) low-regret algorithm would converge to positions similar in structure to the ones illustrated in Figure 1 (for $k = 1, 2, 4$ and $k = 8$) which are clearly close to the optimal (static) solution for the different values of k .

Efficient Online Learning for Dynamic k -Clustering.

The existence of no-regret online learning algorithms for Dynamic k -Clustering immediately follows by standard results in online learning literature (Hazan, 2016). Dynamic k -Clustering is a special case of *Learning from Expert Advice* problem for which the famous *Multiplicative Weights Update Algorithm* achieves no-regret (Hazan, 2016). Unfortunately using the MWU for Dynamic k -Clustering is not really an option due to the huge time and space complexity that MWU requires. In particular MWU keeps a different weight (probability) for each of the possible $\binom{|V|}{k}$ possible placements of the k centers, rendering it inapplicable even for small values of $|V|$ and k .

Our work aims to shed light on the following question.

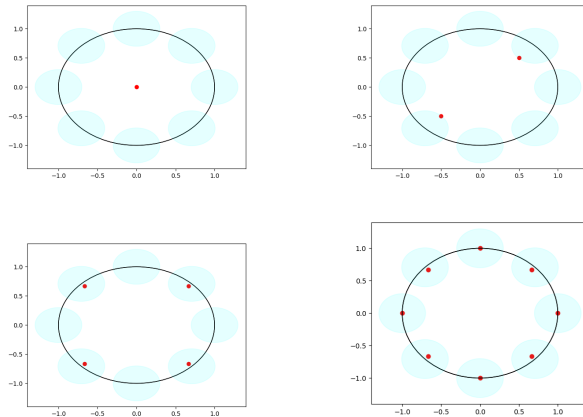


Figure 1: The figure depicts the actual centers at which a low-regret algorithm, that we subsequently propose, converges. For further details see Section 6.

Question 1. Is there an online learning algorithm for Dynamic k -Clustering that runs in polynomial time and achieves α -regret?

Our Contribution and Techniques. We first show that constant regret cannot be achieved in polynomial time for Dynamic k -Clustering. In particular we prove that any $O(1)$ -regret polynomial-time online learning algorithm for Dynamic k -Clustering implies the existence of an $O(1)$ -approximation algorithm for the *Minimum- p -Union problem* (Chlamtác et al., 2016). Recent works on the theory of computational complexity establish that unless well-established cryptographic conjectures fail, there is no $O(1)$ -approximation algorithm for *Min- p -Union* (Chlamtác et al., 2016; Applebaum, 2012; Chlamtác et al., 2017). This result narrows the plausible regret bounds achievable in polynomial time, and reveals an interesting gap between Dynamic k -Clustering and its offline counterparts, which admit polynomial-time $O(1)$ -approximation algorithms.

Our main technical contribution consists of polynomial-time online learning algorithms for Dynamic k -Clustering with non trivial regret bounds. We present a $\Theta(k)$ -regret polynomial-time deterministic online learning algorithm and a $\Theta(r)$ -regret polynomial-time randomized online learning algorithm, where r is the maximum number of clients appearing in a single round ($r = \max_{1 \leq t \leq T} |R_t|$). Combining these algorithms, one can achieve $\Theta(\min(k, r))$ -regret for Dynamic k -Clustering, which (to the best of our knowledge) is the first guarantee on the regret achievable in polynomial time. The regret bounds above are independent of the selected p -norm, and hold for any $p \geq 1$ and for $p = \infty$.

At a technical level, our approach consists of two major steps. In the first step, we consider an online learning prob-

lem, that can be regarded as the *fractional relaxation* of the Dynamic k -Clustering (see Section 3), where the *fractional connection cost* is given by the optimal value of an appropriate convex program and the action space of the learner is the $|V|$ -dimensional simplex. For this intermediate problem, we design a *no-regret* polynomial-time online learning algorithm through the use of the subgradients of the fractional connection cost. Since the *fractional connection cost* comes as the solution of an appropriate convex program, the subgradient vectors do not admit a closed form and it is not clear whether they can be efficiently computed. A key idea of ours was to show that the subgradient vectors can be computed in polynomial time via the solution of the dual program of the fractional connection cost. We remark that the latter idea extends from the context of the *online k -clustering* and it may be used in the design of efficient low-regret algorithms in other combinatorial domains.

In the second step of our approach (see Section 4 and Section 5), we provide computationally efficient online (deterministic and randomized) rounding schemes converting a vector lying in the $|V|$ -dimensional simplex (the action space of Fractional Dynamic k -Clustering) into k locations for the centers on the metric space V (the action space of Dynamic k -Clustering).

In Section 4, we present a deterministic rounding scheme that, combined with the no-regret algorithm for Fractional Dynamic k -Clustering, leads to a $\Theta(k)$ -regret polynomial-time deterministic online learning algorithm for the original Dynamic k -Clustering. Interestingly, this regret bound is approximately optimal for all deterministic algorithms. In Section 5, we show that combining the no-regret algorithm for Fractional Dynamic k -Clustering with a randomized rounding scheme proposed in (Charikar & Li, 2012)¹ leads to a $\Theta(r)$ -regret randomized algorithm running in polynomial time. Combining these two online learning algorithms, we obtain a $\Theta(\min(k, r))$ -regret polynomial-time online learning algorithm for Dynamic k -Clustering, which is the main technical contribution of this work. Finally, in Section 6, we present the results of an experimental evaluation, indicating that for client locations generated in a variety of natural and practically relevant ways, the realized regret of the proposed algorithms is way smaller than $\Theta(\min(k, r))$.

Remark 1. *Our two-step approach provides a structured framework for designing polynomial-time low-regret algorithms in various combinatorial domains. The first step extends far beyond the context of Dynamic k -Clustering and provides a systematic approach to the design of polynomial-time no-regret online learning algorithms for the **fractional relaxation** of the combinatorial online learning problem of interest. Combining such no-regret algorithms with on-*

line rounding schemes, which convert fractional solutions into integral solutions of the original online learning problem, may lead to polynomial time low-regret algorithms for various combinatorial settings. Obviously, designing such rounding schemes is usually far from trivial, since the specific combinatorial structure of each specific problem must be taken into account.

Related Work. Our work relates with the research line of Combinatorial Online Learning. There exists a long line of research studying low-regret online learning algorithms for various combinatorial domains such that online routing (Helmbold et al., 1997; Awerbuch & Kleinberg, 2008), selection of permutations (Takimoto & Warmuth, 2000; Yasutake et al.; Fotakis et al., 2020; Ailon; Helmbold & Warmuth), selection of binary search trees (Takimoto & Warmuth, 2003), submodular optimization (Hazan & Kale, 2012; Jegelka & Bilmes; Streeter & Golovin), matrix completion (Hazan et al., b), contextual bandits (Agarwal et al.; Dudík et al., b) and many more. Finally, in combinatorial games agents need to learn to play optimally against each other over complex domains (Immorlica et al., 2011; Dehghani et al., 2016). As in the case of Dynamic k -Clustering in all the above online learning problems, MWU is not an option, due to the exponential number of possible actions.

Another research direction of Combinatorial Online Learning studies *black-box reductions* converting polynomial time offline algorithm (full information on the data) into polynomial time online learning algorithms. (Kalai & Vempala, 2003) showed that any (offline) algorithm solving optimally and in polynomial time the objective function, that the *Follow the Leader framework* suggests, can be converted into a no-regret online learning algorithm. (Kakade et al.) extended the previous result for specific class of online learning problems called *linear optimization problems* for which they showed that any α -approximation (offline) can be converted into an α -regret online learning algorithm. They also provide a surprising counterexample showing that such black-box reductions do not hold for general combinatorial online learning problems. Both the time efficiency and the regret bounds of the reductions of (Kalai & Vempala, 2003) and (Kakade et al.) were subsequently improved by (Rahmanian & Warmuth; Suehiro et al.; Koolen et al., 2010; Balcan & Blum, 2006; Dudík et al., a; Hazan & Koren; Fujita et al.; Garber; Hazan et al., a). We remark that the above results do not apply in our setting since Dynamic k -Clustering can neither be optimally solved in polynomial-time nor is a linear optimization problem.

Our works also relates with the more recent line of research studying clustering problems with *time-evolving clients*. (Eisenstat et al.) and (An et al., 2017) respectively provide $\Theta(\log(nT))$ and $O(1)$ -approximation algorithm for a generalization of the facility location problem in which

¹This randomized rounding scheme was part of a 4-approximation algorithm for k -median (Charikar & Li, 2012)

clients change their positions over time. The first difference of Dynamic k -Clustering with this setting is that in the former case there is no constraint on the number of centers that can open and furthermore, crucially perfect knowledge of the positions of the clients is presumed. More closely related to our work are (de Keijzer & Wojtczak; Fotakis et al., 2019), where the special case of Dynamic k -Clustering on a line is studied (the clients move on a line over time). Despite the fact that both works study online algorithms, which do not require knowledge on the clients' future positions, they only provided positive results for $k = 1$ and 2. Finally our work relates with (Cohen-Addad et al., 2019) which studies efficient low-regret online learning algorithms for *Online k -means Clustering*. *Online k -means Clustering* is the special case of *Online k -Clustering* at which the 2-norm ($p = 2$) is considered and the clients lie on the Euclidean space. Moreover in *Online k -means Clustering* only one client arrives at each round ($r = 1$ in our notation). Our results are much more general, since they hold for any p -norm (even for $p = \infty$) and general metric spaces while our regret-bound $O(\min(k, r))$ shows how regret scales as $r = \#\text{clients per round}$ increases. A dependence on r is inevitable since (as we show) no-constant regret is possible in poly-time for general r .

2. Preliminaries and Our Results

In this section we introduce notation and several key notions as long as present the formal statements of our results.

We denote by D the diameter of the metric space, $D = \max_{i \in V, j \in V} d_{ij}$. We denote with n the cardinality of the metric space ($|V| = n$) and with r the maximum number of clients appearing in a single round, $r = \max_{1 \leq t \leq T} |R_t|$. Finally we denote with Δ_n^k the n -dimensional simplex, $\Delta_n^k = \{y \in \mathbb{R}^n : \sum_{i \in V} y_i = k \text{ and } y_i \geq 0\}$.

Following the standard notion of regret in online learning (Hazan, 2016), we provide the formal definition of an α -regret online learning algorithm for *Dynamic k -Clustering*.

Definition 1. *An online learning algorithm for the Dynamic k -Clustering is α -regret if and only if for any sequence of clients' positions $R_1, \dots, R_T \subseteq V$,*

$$\sum_{t=1}^T C_{R_t}(F_t) \leq \alpha \cdot \min_{|F^*| \leq k} \sum_{t=1}^T C_{R_t}(F^*) + \Theta(\text{poly}(n, D) \cdot T^\beta)$$

where F_1, \dots, F_T are the positions of the centers produced by the algorithm for the sequence R_1, \dots, R_T and $\beta < 1$.

Next, we introduce the *Minimum- p -Union* problem, the inapproximability results of which allow us to establish that constant regret cannot be achieved in polynomial time for Dynamic k -Clustering.

Problem 1 (Min- p -Union). *Given a universe of elements \mathbb{E} and a collection of sets $\mathbb{U} = \{S_1, \dots, S_m\}$ where $S_i \subseteq \mathbb{E}$. Select $\mathbb{U}' \subseteq \mathbb{U}$ such that $|\mathbb{U}'| = p$ and $|\cup_{S_i \in \mathbb{U}'} S_i|$ is minimized.*

As already mentioned, the existence of an $O(1)$ -approximation algorithm for Min- p -Union violates several widely believed conjectures in computational complexity theory (Chlamtác et al., 2016; Applebaum, 2012; Chlamtác et al., 2017). In Theorem 1 we establish the fact that the exact same conjectures are violated in case there exists an online learning algorithm for *Dynamic k -Clustering* that runs in polynomial-time and achieves $O(1)$ -regret.

Theorem 1. *Any c -regret polynomial-time online learning algorithm for the Dynamic k -Clustering implies a $(c + 1)$ -approximation polynomial-time algorithm for Min- p -Union.*

In Section 4, we present a polynomial-time deterministic online learning algorithm achieving $\Theta(k)$ -regret.

Theorem 2. *There exists a $6k$ -regret deterministic online learning algorithm for Dynamic k -Clustering that runs in polynomial time (Algorithm 4). More precisely,*

$$\sum_{t=1}^T C_{R_t}(F_t) \leq 6k \cdot \min_{|F^*|=k} \sum_{t=1}^T C_{R_t}(F^*) + \Theta(kDn\sqrt{\log nT})$$

where F_1, \dots, F_T are the positions in which Algorithm 4 places the centers for the sequence of clients' positions R_1, \dots, R_T .

In Theorem 3 we prove that the $\Theta(k)$ bound on the regret of Algorithm 4 cannot be significantly ameliorated with deterministic online learning algorithm even if the algorithm uses exponential time and space.

Theorem 3. *For any deterministic online learning algorithm for Dynamic k -Clustering problem, there exists a sequence of clients R_1, \dots, R_T such as the regret is at least $k + 1$.*

Proof. Let the metric space be composed by $k + 1$ points with the distance between any pair of (different) points being 1. At each round t , there exists a position at which the learner has not placed a facility (there are $k + 1$ positions and k facilities). If the adversary places one client at the empty position of the metric space, then the deterministic online learning algorithm admits overall connection cost equal to T . However the optimal static solution that leaves empty the position with the least requests pays at most $T/(k + 1)$. \square

In Section 5 we present a randomized online learning algorithm the regret of which depends on the parameter r .

Theorem 4. *There exists a $\Theta(r)$ -regret randomized algorithm that runs in polynomial time (Algorithm 5). For any sequence of clients' positions R_1, \dots, R_T with $|R_t| \leq r$,*

$$\sum_{t=1}^T \mathbb{E}[C_{R_t}(F_t)] = 4r \cdot \min_{|F^*|=k} \sum_{t=1}^T C_{R_t}(F^*) + \Theta\left(kDn\sqrt{\log nT}\right)$$

where F_t is the random variable denoting the k positions at which Algorithm 5 places the centers at round t .

By combining Algorithm 4 and Algorithm 5 we can achieve $\Theta(\min(k, r))$ -regret in polynomial time.

Theorem 5. *There exists an online learning algorithm for Dynamic k -Clustering that runs in polynomial-time and achieves $\min(6k, 4r)$ -regret.*

Remark 2. *In case the value $r = \min_{1 \leq t \leq T} |R_t|$ is initially known to the learner, then Theorem 5 follows directly by Theorem 2 and 4. However even if r is not initially known, the learner can run a Multiplicative Weight Update Algorithm that at each round follows either Algorithm 4 or Algorithm 5 with some probability distribution depending on the cost of each algorithm so far. By standard results for MWU (Hazan, 2016), this meta-algorithm admits time-average cost less than the best of Algorithm 4 and 5.*

3. Fractional Dynamic k -Clustering

In this section we present the *Fractional Dynamic k -Clustering* problem for which we provide a polynomial-time no-regret online learning algorithm. This online learning algorithm serves as a primitive for both Algorithm 4 and Algorithm 5 of the subsequent sections concerning the original Dynamic k -Clustering.

The basic difference between Dynamic k -Clustering and Fractional Dynamic k -Clustering is that in the second case the learner can *fractionally* place a center at some point of the metric space V . Such a fractional opening is described by a vector $y \in \Delta_n^k$.

Online Learning Problem 2. [Fractional Dynamic k -Clustering] At each round $t \geq 1$,

1. The learner selects a vector $y_t \in \Delta_n^k$. The value y_i^t stands for the fractional amount of center that the learner opens in position $i \in V$.
2. The adversary selects the positions of the clients denoted by $R_t \subseteq V$ (after the selection of the vector y_t).
3. The learner incurs fractional connection cost $\text{FC}_{R_t}(y_t)$ described in Definition 2.

Definition 2 (Fractional Connection Cost). *Given the positions of the clients $R \subseteq V$, we define the fractional connection cost $\text{FC}_R(\cdot)$ of a vector $y \in \Delta_n^k$ as the optimal value of the following convex program.*

$$\begin{aligned} & \text{minimize} \left(\sum_{j \in R} \beta_j^p \right)^{1/p} \\ & \text{s.t.} \quad \beta_j = \sum_{i \in V} d_{ij} \cdot x_{ij} \quad \forall j \in R \\ & \quad \sum_{i \in V} x_{ij} = 1 \quad \forall j \in R \\ & \quad x_{ij} \leq y_i \quad \forall j \in R, \forall i \in V \\ & \quad x_{ij} \geq 0 \quad \forall j \in R, \forall i \in V \end{aligned} \quad (1)$$

It is not hard to see that once the convex program of Definition 2 is formulated with respect to an *integral* vector $y \in \Delta_n^k$ (y_i is either 0 or 1) the fractional connection cost $\text{FC}_R(y)$ equals the original connection cost $C_R(y)$. As a result, the cost of the optimal solution $y^* \in \Delta_n^k$ of Fractional Dynamic k -Clustering is upper bounded by the cost of the optimal positioning of the centers F^* in the original Dynamic k -Clustering.

Lemma 1. *For any sequence of clients' positions R_1, \dots, R_T , the cost of the optimal fractional solution y^* for Fractional Dynamic k -Clustering is smaller than the cost of the optimal positioning F^* for Dynamic k -Clustering,*

$$\min_{y^* \in \Delta_n^k} \sum_{t=1}^T \text{FC}_{R_t}(y^*) \leq \min_{|F^*|=k} \sum_{t=1}^T C_{R_t}(F^*)$$

Lemma 1 will be used in the next sections where the online learning algorithms for the original Dynamic k -Clustering are presented. To this end, we dedicate the rest of this section to design a polynomial time no-regret algorithm for Fractional Dynamic k -Clustering. A key step towards this direction is the use of the subgradient vectors of $\text{FC}_{R_t}(\cdot)$.

Definition 3 (Subgradient). *Given a function $f : \mathbb{R}^n \mapsto \mathbb{R}$, a vector $g \in \mathbb{R}^n$ belongs in the subgradient of f at point $x \in \mathbb{R}^n$, $g \in \partial f(x)$, if and only if $f(y) \geq f(x) + g^\top(y - x)$, for all $y \in \mathbb{R}^n$.*

Computing the subgradient vectors of functions, as complicated as $\text{FC}_{R_t}(\cdot)$, is in general a computationally hard task. One of our main technical contributions consists in showing that the latter can be done through the solution of an adequate convex program corresponding to the dual of the convex program of Definition 2.

Lemma 2. *Consider the convex program of Definition 2 formulated with respect to a vector $y \in \Delta_n^k$ and the clients'*

positions R . Then the following convex program is its dual.

$$\begin{aligned} & \text{maximize} && \sum_{j \in R} A_j - \sum_{i \in V} \sum_{j \in R} k_{ij} \cdot y_i \\ & \text{s.t.} && \|\lambda\|_p^* \leq 1 \\ & && d_{ij} \cdot \lambda_j + k_{ij} \geq A_j \quad \forall i \in V, j \in R \\ & && k_{ij} \geq 0 \quad \forall i \in V, j \in R \end{aligned} \quad (2)$$

where $\|\cdot\|_p^*$ is the dual norm of $\|\cdot\|_p$

In the following lemma we establish the fact that a sub-gradient vector of $\partial \text{FC}_{R_t}(\cdot)$ can be computed through the optimal solution of the convex program in Lemma 2.

Lemma 3. *Let k_{ij}^* denote the value of the variables k_{ij} in the optimal solution of the convex program in Lemma 2 formulated with respect to vector $y \in \Delta_n^k$ and the clients' positions R . Then for any vector $y' \in \Delta_n^k$,*

$$\text{FC}_{R_t}(y') \geq \text{FC}_{R_t}(y) + \sum_{i \in V} \left(- \sum_{j \in R} k_{ij}^* \right) \cdot (y'_i - y_i)$$

Moreover there exists an $\Theta(r \cdot |V|)$ algorithm for solving the dual program (Algorithm 1) and additionally $|k_{ij}^*| \leq D$.

Algorithm 1 A time-efficient algorithm for solving the dual program of Lemma 2

- 1: **Input:** A vector $y \in \Delta_n^k$ and a set of clients $R \subseteq V$.
 - 2: **Output:** An optimal solution for the convex program of Lemma 2.
 - 3: **for** each client $j \in R$, **do**
 - 4: Sort the nodes $i \in V$ in increasing order according to d_{ij} .
 - 5: $\text{Rem} \leftarrow 1$
 - 6: **for** each $i \in V$ **do**
 - 7: $x_{ij} \leftarrow \min(y_i, \text{Rem})$.
 - 8: $\text{Rem} \leftarrow \text{Rem} - x_{ij}$.
 - 9: **end for**
 - 10: **end for**
 - 11: **for** each client $j \in R$ **do**
 - 12: $V_j^+ \leftarrow \{i \in V : x_{ij} > 0\}$ and $D_j \leftarrow \max_{i \in V_j^+} d_{ij}$.
 - 13: $\beta_j \leftarrow \sum_{i \in V} d_{ij} \cdot x_{ij}$
 - 14: $\lambda_j \leftarrow \left[\frac{\beta_j}{\|\beta\|_p} \right]^{p-1}$
 - 15: $A_j \leftarrow \lambda_j \cdot D_j$
 - 16: $k_{ij} \leftarrow \min \left(\lambda_j \cdot \frac{x_{ij}}{y_i} \cdot (D_j - d_{ij}), 0 \right)$
 - 17: **end for**
-

Remark 3. *Algorithm 1 is not only a computationally efficient way to solve the convex program of Lemma 2, but most importantly guarantees that the value k_{ij}^* are bounded by D (this is formally stated and proven in Lemma 2). The latter property is crucial for developing the no-regret algorithm for Fractional Dynamic k-Clustering.*

Up next we present the no-regret algorithm for Fractional Dynamic k-Clustering.

Algorithm 2 A no-regret algorithm for Fractional Dynamic k-Clustering

- 1: Initially, the learner selects $y_i^1 = k/n$ for all $i \in V$.
- 2: **for** rounds $t = 1 \dots T$ **do**
- 3: The learner selects $y_t \in \Delta_n^k$.
- 4: The adversary selects the positions of the clients $R_t \subseteq V$.
- 5: The learner receives cost, $\text{FC}_{R_t}(y_t)$.
- 6: The learner runs Algorithm 1 with input y_t and R_t and sets $g_i^t = - \sum_{j \in R_t} k_{ij}^t$
- 7: **for** each $i \in V$ **do**
- 8:

$$y_i^{t+1} = \frac{y_i^t \cdot e^{-\epsilon g_i^t}}{\sum_{i \in V} y_i^t \cdot e^{-\epsilon g_i^t}}$$

$$\text{where } \epsilon = \frac{\sqrt{\log n}}{Dr\sqrt{T}}$$

- 9: **end for**
 - 10: **end for**
-

We conclude the section with Theorem 6 that establishes the no-regret property of Algorithm 2 and the proof of which is deferred to the Appendix B.

Theorem 6. *Let y_1, \dots, y_T be the sequence of vectors in Δ_n^k produced by Algorithm 2 for the clients' positions R_1, \dots, R_T . Then,*

$$\sum_{t=1}^T \text{FC}_{R_t}(y_t) \leq \min_{y^* \in \Delta_n^k} \sum_{t=1}^T \text{FC}_{R_t}(y^*) + \Theta \left(kDn\sqrt{\log nT} \right)$$

4. A $\Theta(k)$ -Regret Deterministic Online Learning Algorithm

In this section we show how one can use Algorithm 2 described in Section 3 to derive $\Theta(k)$ -regret for the Dynamic k-Clustering in polynomial-time.

The basic idea is to use a rounding scheme that given a vector $y \in \Delta_n^k$ produces a placement of the k centers $F_y \subseteq V$ (with $|F_y| \leq k$) such that for any set of clients' positions R , the connection cost $C_R(F_y)$ is approximately bounded by the fractional connection cost $\text{FC}_R(y)$. This rounding scheme is described in Algorithm 3.

Lemma 4. *[Rounding Lemma] Let F_y denote the positions of the centers produced by Algorithm 3 for input $y \in \Delta_n^k$. Then the following properties hold,*

- For any set of clients R ,

$$C_R(F_y) \leq 6k \cdot \text{FC}_R(y)$$

Algorithm 3 Deterministic Rounding Scheme

1: **Input:** A vector $y \in \Delta_n^k$.
 2: **Output:** A set $F_y \subseteq V$ at which centers are opened.
 3: Run Algorithm 1 with input y and $R = V$.
 4: Sort the positions $i \in V$ according to the values β_i produced by Algorithm 1.
 5: $F_y \leftarrow \emptyset$
 6: **for** $i = 1$ **to** V **do**
 7: **if** $\min_{j \in F_y} d_{ij} > 6k \cdot \beta_i$ **then**
 8: $F_y \leftarrow F_y \cup \{i\}$
 9: **end if**
 10: **end for**

- The cardinality of F_y is at most k , $|F_y| \leq k$.

Up next we show how the deterministic rounding scheme described in Algorithm 3 can be combined with Algorithm 2 to produce an $\Theta(k)$ -regret deterministic online learning algorithm that runs in polynomial-time. The overall online learning algorithm is described in Algorithm 4 and its regret bound is formally stated and proven in Theorem 2.

Algorithm 4 A $\Theta(k)$ -regret deterministic online learning algorithm for Dynamic k -Clustering

1: **for** rounds $t = 1 \dots T$ **do**
 2: The learner computes the vector $y_t \in \Delta_n^k$ by running Algorithm 2 for the sequence of clients' positions (R_1, \dots, R_{t-1}) .
 3: The learner places centers to the positions F_{y_t} produced by Algorithm 3 given input y_t .
 4: The adversary selects the clients' positions $R_t \subseteq V$.
 5: The learner suffers connection cost $C_{R_t}(F_{y_t})$
 6: **end for**

We conclude the section with the proof of Theorem 2 in which the regret bounds of Algorithm 4 are established.

Proof of Theorem 2. The second case of Lemma 4 ensures that $|F_t| \leq k$ and thus Algorithm 4 opens at most k facilities at each round. Applying the first case of Lemma 4 for $R = R_t$ we get that $C_{R_t}(F_t) \leq 6k \cdot \text{FC}_{R_t}(y_t)$. As a result,

$$\begin{aligned}
 \sum_{t=1}^T C_{R_t}(F_t) &\leq \sum_{t=1}^T 6k \cdot \text{FC}_{R_t}(y_t) \\
 &\leq 6k \min_{y^* \in \Delta_k} \sum_{t=1}^T \text{FC}_{R_t}(y^*) + \Theta\left(kDn\sqrt{\log nT}\right)
 \end{aligned}$$

where the last inequality follows by Theorem 6. However Lemma 1 ensures that

$$\min_{y^* \in \Delta_k} \sum_{t=1}^T \text{FC}_{R_t}(y^*) \leq \min_{F^*: |F^*|=k} \sum_{t=1}^T C_{R_t}(F^*)$$

□

5. A $\Theta(r)$ -Regret Randomized Online Learning Algorithm

In this section we present a $\Theta(r)$ -regret randomized online learning algorithm. This algorithm is described in Algorithm 5 and is based on the randomized rounding developed by Charikar and Li for the k -median problem (Charikar & Li, 2012).

Lemma 5 ((Charikar & Li, 2012)). *There exists a polynomial-time randomized rounding scheme that given a vector $y \in \Delta_n^k$ produces a probability distribution, denoted as $\text{CL}(y)$, over the subsets of V such that,*

1. with probability 1 exactly k facilities are opened, $\mathbb{P}_{F \sim \text{CL}(y)}[|F| = k] = 1$.
2. for any position $j \in V$,

$$\mathbb{E}_{F \sim \text{CL}(y)} [C_{\{j\}}(F_y)] \leq 4 \cdot \text{FC}_{\{j\}}(y).$$

Similarly with the previous section, combining the randomized rounding of Charikar-Li with Algorithm 1 produces a $\Theta(r)$ -regret randomized online learning algorithm that runs in polynomial-time.

Algorithm 5 A $\Theta(r)$ -regret randomized online learning algorithm

1: **for** rounds $t = 1 \dots T$ **do**
 2: The learner computes the vector $y_t \in \Delta_n^k$ by running Algorithm 2 for the sequence of clients' positions (R_1, \dots, R_{t-1}) .
 3: The learner places centers to the positions $F_t \subseteq V$ produced by the Charikar-Li randomized rounding with input y_t , $F_t \sim \text{CL}(y_t)$.
 4: The adversary selects a request $R_t \subseteq V$.
 5: The learner suffers connection cost $C_{R_t}(F_t)$
 6: **end for**

The proof of Theorem 4 that establishes the regret bound of Algorithm 5 follows by Lemma 5 and Theorem 6 and is deferred to the Appendix D.

6. Experimental Evaluations

In this section we evaluate the performance of our online learning algorithm against adversaries that select the positions of the clients according to time-evolving probability distributions. We remark that the regret bounds established in Theorem 2 and Theorem 4 hold even if the adversary *maliciously* selects the positions of the clients at each round so as to maximize the connection cost. As a result, in case

clients arrive according to some (unknown and possibly time-varying) probability distribution that does not depend on the algorithm’s actions, we expect the regret of to be way smaller.

In this section we empirically evaluate the regret of Algorithm 4 for Dynamic k -Clustering in case $p = \infty$. We assume that at each round t , 20 clients arrive according to several static or time-varying two-dimensional probability distributions with support on the $[-1, 1] \times [-1, 1]$ square and the possible positions for the centers being the discretized grid with $\epsilon = 0.1$. In order to monitor the quality of the solutions produced by Algorithm 4, we compare the time-average connection cost of Algorithm 4 with the time-average *fractional connection cost* of Algorithm 2. Theorem 6 ensures that for $T = \Theta(k^2 D^2 / \epsilon^2)$ the time-average fractional connection cost of Algorithm 2 is at most ϵ greater than the time-average connection cost of the optimal static solution for Dynamic k -Clustering. In the following simulations we select $\epsilon = 0.1$ and track the ratio between the time-average cost of Algorithm 4 and of Algorithm 2 which acts as an upper bound on the regret.

Uniform Square In this case the 20 clients arrive *uniformly at random* in the $[-1, 1] \times [-1, 1]$ square. Figure 2 illustrates the solutions at which Algorithm 4 converges for $k = 2, 3$ and 8 as long as the achieved regret.

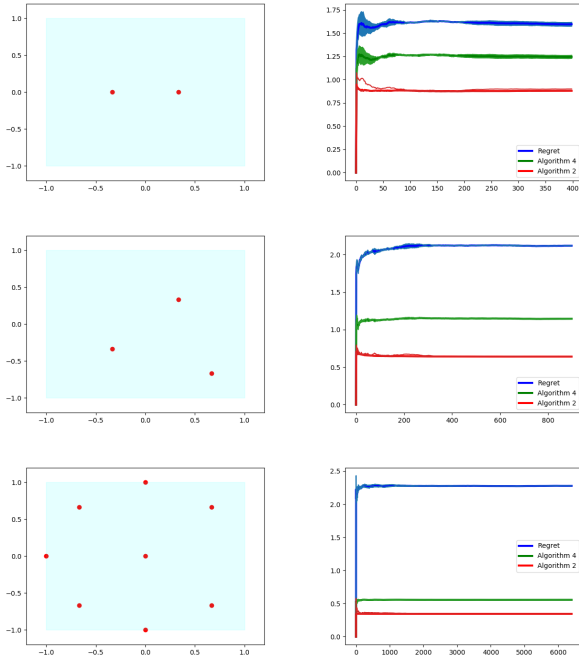


Figure 2: The **green curve** depicts the time-average connection cost Algorithm 4, the **red curve** depicts the time-average fractional connection cost of Algorithm 2 and the **blue curve** depicts their ratio that acts as an upper bound on the regret.

Uniform Distribution with Time-Evolving Centers In this case the 20 clients arrive according to the uniform distribution with radius 0.3 and a time-varying center that periodically follows the trajectory described in Example 1. Figure 1 depicts the centers at which Algorithm 4 converges after $100k^2$ rounds which are clearly close to the optimal ones.

Moving-Clients on the Ellipse In this case the 20 clients move in the ellipse $(\frac{x}{1.2})^2 + (\frac{y}{0.6})^2 = 1$ with different speeds and initial positions. The position of client i is given by $(x_i(t), y_i(t)) = (1.2 \cos(2\pi f_i t + \theta_i), 0.6 \sin(2\pi f_i t + \theta_i))$ where each f_i, θ_i was selected uniformly at random in $[0, 1]$. Figure 3 illustrates how Algorithm 4 converges to the underlying ellipse as the number of rounds increases.




Figure 3: The solution produced by Algorithm 4 for $k = 8$ after 100, 1000 and 10000 rounds.

Mixture of Multivariate Gaussians In this case 15 clients arrive according to the Gaussian with $\mu_1 = (-0.7, 0.7)$ and $\Sigma_1 = [[0.3, 0], [0, 0.3]]$ and 5 according to the Gaussian with $\mu_2 = (0.7, -0.7)$ and $\Sigma_2 = [[0.3, 0], [0, 0.3]]$. All the clients outside the $[-1, 1] \times [-1, 1]$ are projected back to the square. Figure 4 illustrates the solutions at which Algorithm 4 converges for $k = 2, 8$ and 16.

7. Conclusion

This work studies polynomial-time low-regret online learning algorithms for Dynamic k -Clustering, an online learning problem capturing clustering settings with time-evolving clients for which no information on their locations over time is available. We show that, under some well-established conjectures, $O(1)$ -regret cannot be achieved in polynomial time and we provide a $\Theta(\min(k, r))$ -regret polynomial time algorithm with r being the maximum number of clients appearing in a single round. At a technical level, we present a

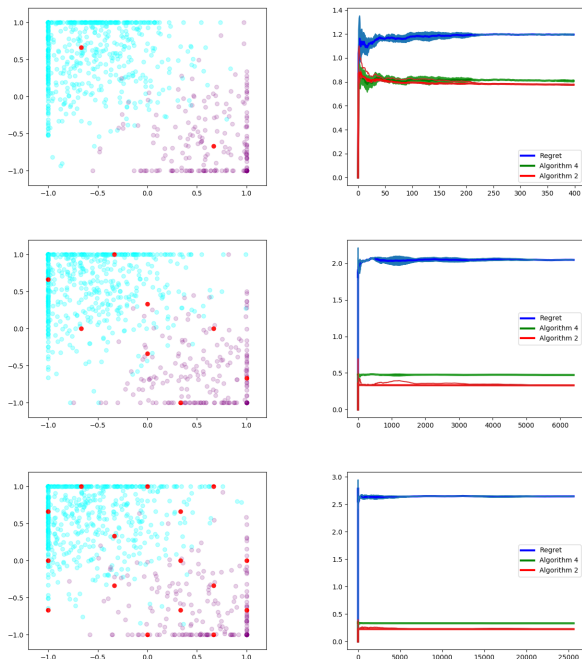


Figure 4: On the left, the solutions which Algorithm 4 converges for $k = 2, 8$ and $k = 16$. On the right, the time-average cost of Algorithm 4, Algorithm 2 and the regret bounds.

two-step approach where in the first step we provide a no-regret algorithm for the Fractional Dynamic k -Clustering while in the second step we provide online rounding scheme converting the sequence of fractional solutions, produced by the no-regret algorithm, into solutions of Dynamic k -Clustering. Applying the same approach to other combinatorial online learning problems is an interesting research direction.

8. Acknowledgments

This research/project is supported in part by NRF2019-NRF-ANR095 ALIAS grant, grant PIE-SGP-AI-2018-01, NRF 2018 Fellowship NRF-NRFF2018-07, AME Programmatic Fund (Grant No. A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) the National Research Foundation, Singapore, under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), and the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “First Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant”, project BALSAM, HFRI-FM17-1424.

References

- Agarwal, A., Hsu, D. J., Kale, S., Langford, J., Li, L., and Schapire, R. E. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of the 31th International Conference on Machine Learning*, ICML 2014.
- Ailon, N. Improved bounds for online learning over the permutahedron and other ranking polytopes. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, AISTATS 2014.
- Alamdari, S. and Shmoys, D. *A Bicriteria Approximation Algorithm for the k -Center and k -Median Problems*, pp. 66–75. 01 2018.
- An, H., Norouzi-Fard, A., and Svensson, O. Dynamic facility location via exponential clocks. *ACM Trans. Algorithms*, 13(2):21:1–21:20, 2017.
- Applebaum, B. Pseudorandom generators with long stretch and low locality from random local one-way functions. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pp. 805–816. Association for Computing Machinery, 2012. ISBN 9781450312455.
- Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., and Pandit, V. Local search heuristic for k -median and facility location problems. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC ’01, pp. 21–29. Association for Computing Machinery, 2001.
- Awerbuch, B. and Kleinberg, R. Online linear optimization and adaptive routing. *J. Comput. Syst. Sci.*, 2008.
- Balcan, M.-F. and Blum, A. Approximation algorithms and online mechanisms for item pricing. In *ACM Conference on Electronic Commerce*, 2006.
- Charikar, M. and Guha, S. Improved combinatorial algorithms for the facility location and k -median problems. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, FOCS ’99. IEEE Computer Society, 1999.
- Charikar, M. and Li, S. A dependent lp-rounding approach for the k -median problem. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012*, volume 7391 of *Lecture Notes in Computer Science*, pp. 194–205. Springer, 2012.
- Charikar, M., Guha, S., Tardos, E., and Shmoys, D. B. A constant-factor approximation algorithm for the k -median problem (extended abstract). In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*,

- STOC '99, pp. 1–10. Association for Computing Machinery, 1999. ISBN 1581130678.
- Chlamtác, E., Dinitz, M., Konrad, C., Kortsarz, G., and Rabanca, G. The densest k-subhypergraph problem. *Leibniz International Proceedings in Informatics (LIPIcs)*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2016.
- Chlamtác, E., Dinitz, M., and Makarychev, Y. Minimizing the union: Tight approximations for small set bipartite vertex expansion. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17*, pp. 881–899. Society for Industrial and Applied Mathematics, 2017.
- Cohen-Addad, V., Guedj, B., Kanade, V., and Rom, G. Online k-means clustering, 2019.
- de Keijzer, B. and Wojtczak, D. Facility reallocation on the line. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pp. 188–194.
- Dehghani, S., Hajiaghayi, M., Mahini, H., and Seddighin, S. Price of competition and dueling games. *arXiv preprint arXiv:1605.04004*, 2016.
- Dudík, M., Haghtalab, N., Luo, H., Schapire, R. E., Syrgkanis, V., and Vaughan, J. W. Oracle-efficient online learning and auction design. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, a.
- Dudík, M., Hsu, D. J., Kale, S., Karampatziakis, N., Langford, J., Reyzin, L., and Zhang, T. Efficient optimal learning for contextual bandits. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2011*, b.
- Eisenstat, D., Mathieu, C., and Schabanel, N. Facility location in evolving metrics. In *Automata, Languages, and Programming - 41st International ColloquiumICALP 2014*, volume 8573 of *Lecture Notes in Computer Science*, pp. 459–470. Springer.
- Fotakis, D., Kavouras, L., Kostopanagiotis, P., Lazos, P., Skoulakis, S., and Zarifis, N. Reallocating multiple facilities on the line. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pp. 273–279, 2019.
- Fotakis, D., Lianas, T., Piliouras, G., and Skoulakis, S. Efficient online learning of optimal rankings: Dimensionality reduction via gradient descent. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.
- Fujita, T., Hatano, K., and Takimoto, E. Combinatorial online prediction via metarounding. In *24th International Conference on Algorithmic Learning Theory, ALT 2013*.
- Garber, D. Efficient online linear optimization with approximation algorithms. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS 2017*.
- Hazan, E. Introduction to online convex optimization. *Found. Trends Optim.*, 2(3–4):157–325, 2016.
- Hazan, E. and Kale, S. Online submodular minimization. *J. Mach. Learn. Res.*, 2012.
- Hazan, E. and Koren, T. The computational power of optimization in online learning. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing, STOC 2016*.
- Hazan, E., Hu, W., Li, Y., and Li, Z. Online improper learning with an approximation oracle. In *Advances in Neural Information Processing Systems, NeurIPS 2018*. a.
- Hazan, E., Kale, S., and Shalev-Shwartz, S. Near-optimal algorithms for online matrix prediction. In *25th Annual Conference on Learning Theory, COLT 2012*, b.
- Helmbold, D. P. and Warmuth, M. K. Learning permutations with exponential weights. In *Proceedings of the 20th Annual Conference on Learning Theory, COLT 2007*.
- Helmbold, D. P., Schapire, R. E., and Long, M. Predicting nearly as well as the best pruning of a decision tree. In *Machine Learning*, 1997.
- Immorlica, N., Kalai, A. T., Lucier, B., Moitra, A., Postlewaite, A., and Tennenholtz, M. Dueling algorithms. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11*, pp. 215–224, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306911. doi: 10.1145/1993636.1993666. URL <https://doi.org/10.1145/1993636.1993666>.
- Jain, K. and Vazirani, V. V. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *J. ACM*, 48(2):274–296, 2001.
- Jegelka, S. and Bilmes, J. A. Online submodular minimization for combinatorial structures. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*.
- Kakade, S., Kalai, A. T., and Ligett, K. Playing games with approximation algorithms. In *Proceedings of the 39th*

- Annual ACM Symposium on Theory of Computing, STOC 2007.
- Kalai, A. and Vempala, S. Efficient algorithms for online decision problems. In *J. Comput. Syst. Sci.* Springer, 2003.
- Koolen, W. M., Warmuth, M. K., and Kivinen, J. Hedging structured concepts. In *the 23rd Conference on Learning Theory, COLT 2010*, 2010.
- Kumar, A. Constant factor approximation algorithm for the knapsack median problem. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pp. 824–832. Society for Industrial and Applied Mathematics, 2012.
- Kumar, A., Sabharwal, Y., and Sen, S. Linear-time approximation schemes for clustering problems in any dimensions. *J. ACM*, 57(2), 2010.
- Li, S. and Svensson, O. Approximating k-median via pseudo-approximation. *SIAM J. Comput.*, 45(2):530–547, 2016.
- Lin, J.-H. and Vitter, J. S. Approximation algorithms for geometric median problems. *Information Processing Letters*, 44(5):245 – 249, 1992. ISSN 0020-0190.
- Newman, M. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- Pastor-Satorras, R. and Vespignani, A. Epidemic Spreading in Scale-Free Networks. *Physical Review Letters*, 86(14): 3200–3203, 2001.
- Rahmanian, H. and Warmuth, M. K. K. Online dynamic programming. In *Advances in Neural Information Processing Systems*, NIPS 2017.
- Stehlé, J., Voirin, N., Barrat, A., Cattuto, C., Isella, L., Pinton, J., Quaghiotto, M., Van den Broeck, W., Régis, C., Lina, B., and Vanhems, P. High-resolution measurements of face-to-face contact patterns in a primary school. *PLOS ONE*, 6(8), 2011.
- Streeter, M. J. and Golovin, D. An online algorithm for maximizing submodular functions. In *22nd Annual Conference on Neural Information Processing Systems*, NIPS 2008.
- Suehiro, D., Hatano, K., Kijima, S., Takimoto, E., and Nagano, K. Online prediction under submodular constraints. In *Algorithmic Learning Theory*, ALT 2012.
- Takimoto, E. and Warmuth, M. K. Predicting nearly as well as the best pruning of a planar decision graph. In *Theoretical Computer Science*, 2000.
- Takimoto, E. and Warmuth, M. K. Path kernels and multiplicative updates. *J. Mach. Learn. Res.*, 2003.
- Tantipathananandh, C., Berger-Wolf, T., and Kempe, D. A framework for community identification in dynamic social networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pp. 717–726. Association for Computing Machinery, 2007.
- Williamson, D. P. and Shmoys, D. B. *The Design of Approximation Algorithms*. Cambridge University Press, USA, 1st edition, 2011.
- Yasutake, S., Hatano, K., Kijima, S., Takimoto, E., and Takeda, M. Online linear optimization over permutations. In *Proceedings of the 22nd International Conference on Algorithms and Computation*, ISAAC 2011.
- Young, N. E. K-medians, facility location, and the chernoff-wald bound. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '00, pp. 86–95. Society for Industrial and Applied Mathematics, 2000.