

---

# Fair Correlation Clustering

---

Sara Ahmadian

Alessandro Epasto

Ravi Kumar

Mohammad Mahdian

Google Research

{sahmadian,aepasto,mahdian}@google.com, ravi.k53@gmail.com

## Abstract

In this paper, we study correlation clustering under fairness constraints. Fair variants of  $k$ -median and  $k$ -center clustering have been studied recently, and approximation algorithms using a notion called fairlet decomposition have been proposed. We obtain approximation algorithms for fair correlation clustering under several important types of fairness constraints.

Our results hinge on obtaining a fairlet decomposition for correlation clustering by introducing a novel combinatorial optimization problem. We define a fairlet decomposition with cost similar to the  $k$ -median cost and this allows us to obtain approximation algorithms for a wide range of fairness constraints.

We complement our theoretical results with an in-depth analysis of our algorithms on real graphs where we show that fair solutions to correlation clustering can be obtained with limited increase in cost compared to the state-of-the-art (unfair) algorithms.

## 1 Introduction

There is a growing literature on fairness in various learning and optimization problems (Kamishima et al., 2011, 2012; Joseph et al., 2016; Celis et al., 2018b,a; Chierichetti et al., 2019; Yang and Stoyanovich, 2017). The goal of this literature is to develop criteria and algorithms to ensure that we can find solutions for optimization/learning problems that are fair with respect to a certain sensitive feature. In the case of clustering, a fundamental unsupervised learning and

optimization problem, the study of fairness was initiated by Chierichetti et al. (2017). They formulated the notion of proportional fairness, and developed approximation algorithms for fair  $k$ -median and fair  $k$ -center under this notion of fairness. Follow-up work generalized their results to other clustering problems, such as  $k$ -means and facility location, and to more relaxed notions of fairness (Ahmadian et al., 2019; Bercea et al., 2019; Bera et al., 2019; Backurs et al., 2019; Kleindessner et al., 2019a). Notably, the important graph problem of correlation clustering has been so far not addressed by this literature.

Correlation clustering uses information about both similarity and dissimilarity relationships among a set of objects in order to cluster them (Bansal et al., 2004). In contrast to other clustering problems such as  $k$ -median,  $k$ -means, and  $k$ -center, the number of clusters is not pre-specified but rather determined based on the outcome of an optimization. This, as well as the fact that correlation clustering uses both similarity and dissimilarity information, makes it a desirable clustering model in many applications (Kushagra et al., 2019; Pouget-Abadie et al., 2019; Ailon et al., 2008). Therefore, it is natural to study this problem with fairness.

The main tool introduced by Chierichetti et al. (2017) for solving fair  $k$ -center and  $k$ -median problems is the notion of *fairlets*. A fairlet is a small set of elements that satisfies the fairness property. Chierichetti et al. (2017) showed that fair  $k$ -median and  $k$ -center can be solved by first decomposing an instance into fairlets and then solving the clustering problem on the set of centers of these fairlets. To the best of our knowledge, this technique has been used only for *metric space* clustering problems such as  $k$ -center and  $k$ -median.

Our main result is developing a fairlet-based reduction for the *graph* clustering problem of correlation clustering. Whereas, in the case of  $k$ -center and  $k$ -median, the fairlet decomposition problem amounts to solving the same clustering problem on the same instance under the condition that each cluster is a fairlet, the situation for correlation clustering is complicated by the lack of

the properties of metric spaces. To tackle this problem, we introduce a novel cost function for the correlation clustering fairlet decomposition, and prove that this cost can be approximated by a median-type clustering cost function for a carefully defined metric space.

Given a solution to this fairlet decomposition problem, we show that the fair correlation clustering instance can be reduced to a regular correlation clustering instance through a graph transformation. Therefore, any approximation algorithm for fairlet decomposition with median cost yields an approximation algorithm for fair correlation clustering; the loss in the approximation ratio depends on the size of the fairlets. We show that in many natural cases, there is a fairlet decomposition with small fairlets, thereby bounding the approximation ratio of our algorithm for fair correlation clustering.

In addition to the theoretical bounds on our algorithm, we provide an empirical evaluation based on real data sets, showing that the algorithms often perform much better in practice than their worst-case guarantees and that they yield solutions of costs comparable to that of unfair clustering algorithms while substantially reducing the cluster imbalance. Our implementation is also publicly available (Ahmadian et al., 2020).

**Related work.** Clustering is a fundamental unsupervised machine learning task with a long history (Jain, 2010). Our paper spans the areas of correlation clustering, clustering in metric spaces, and fairness in clustering which are actively growing fields. For brevity, we will only focus of key works in these three areas.

*Correlation clustering.* Correlation clustering is a widely studied formulation of clustering with both similarity and dissimilarity information (Bansal et al., 2004), with many applications in machine learning (Kushagra et al., 2019; Bressan et al., 2019). Variants of the problem include complete signed graphs (Bansal et al., 2004; Ailon et al., 2008) and weighted graphs (Demaine et al., 2006). We focus on the complete graph case with  $\pm 1$  weights which is APX-hard (Charikar et al., 2005) but admits constant-factor algorithms (Ailon et al., 2008; Chawla et al., 2015). Distributed and streaming algorithms are also known (Pan et al., 2015; Ahn et al., 2015).

*Metric space clustering.* The most widely studied clustering setting is clustering in metric spaces consisting in minimizing the  $\ell_p$ -norm of the distances between points in a cluster and their center. For  $p \in \{1, 2, \infty\}$  this corresponds to  $k$ -median,  $k$ -means, and  $k$ -center, respectively, which are NP-hard problems but admit constant-factor approximations (Gonzalez, 1985; Hochbaum and Shmoys, 1985; Li and Svensson, 2016; Ahmadian et al., 2017; Kanungo et al., 2004).

*Fairness in clustering.* Fairness in machine learning is new area with a fast growing literature. Fundamental work in this area is devoted to defining notions of fairness (Calders and Verwer, 2010; Dwork et al., 2012; Feldman et al., 2015; Kamishima et al., 2012) and solving fairness-constrained problems (Celis et al., 2018a,b; Chierichetti et al., 2017; Joseph et al., 2016; Kamishima et al., 2012; Yang and Stoyanovich, 2017; Backurs et al., 2019; Kamishima et al., 2011; Fish et al., 2016; Ahmadian et al., 2019; Chierichetti et al., 2019; Har-Peled and Mahabadi, 2019).

Chierichetti et al. (2017) introduced a notion of disparate impact for clustering and provided fair  $k$ -center algorithms for the case of two colors (or groups); see Section 2. Following this work, the problem has been later generalized in many directions including allowing many colors (Rösner and Schmidt, 2018), allowing upper bounds on the fraction of points of a given color (Ahmadian et al., 2019) and both upper and lower bounds (Bera et al., 2019; Bercea et al., 2019). Backurs et al. (2019) designed near-linear algorithms for finding  $k$ -median fairlets, and Huang et al. (2019) designed core-sets for the problem. Other variants include clustering with diversity constraints (Li et al., 2010), proportionality constraints (Chen et al., 2019), and fair center selection (Kleindessner et al., 2019a). Fairness has been studied in spectral clustering as well (Kleindessner et al., 2019b; Ziko et al., 2019). From an application point, fair clustering can be seen through the lenses of fair allocation (Elzayn et al., 2019), Medicaid eligibility (Fang et al., 2019), and ensuring protected group representations (Dash et al., 2019).

To the best of our knowledge no prior work has addressed correlation clustering with fairness. Kalhan et al. (2019) recently studied a fairness notion in which the maximum error for a vertex is bounded.

## 2 Problem Statement

**Correlation clustering.** Let  $G = (V, E)$  be a complete undirected graph on  $|V| = n$  vertices and  $\sigma : E \mapsto \mathbb{R}$  be a function that assigns a label to each edge. The label  $\sigma(e)$  for each  $e$  is either positive (indicating that the two endpoints of  $e$  are *similar*) or non-positive (indicating that they are *dissimilar*). In the *unweighted* version of the problem (Bansal et al., 2004),  $\sigma(e) \in \{-1, +1\}$  for each  $e$ . Our focus in this paper is on the unweighted version, although we will use the weighted version in the proofs. Let  $E^+ = \{e \in E \mid \sigma(e) > 0\}$  be the set of positive edges and  $E^- = E \setminus E^+$  be the set of non-positive edges. For subsets  $S, T \subseteq V$ , let  $E(S) = E \cap S^2$  denote the edges inside  $S$  and  $E(S, T) = E \cap (S \times T)$  denote the edges between  $S$  and  $T$ . Let  $E^+(S, T) = E^+ \cap E(S, T)$  and

$$E^-(S, T) = E^- \cap E(S, T).$$

A *clustering* is a partitioning  $\mathcal{C} = \{C_1, C_2, \dots\}$  of  $V$  into disjoint subsets. The sets of intra-cluster and inter-cluster edges in a clustering  $\mathcal{C}$  are defined as  $\text{intra}(\mathcal{C}) = \bigcup_{C \in \mathcal{C}} E(C)$  and  $\text{inter}(\mathcal{C}) = E \setminus \text{intra}(\mathcal{C})$ . The *correlation clustering cost* of  $\mathcal{C}$  is defined as:

$$\text{COST}(G, \mathcal{C}) = \sum_{e \in \text{intra}(\mathcal{C}) \cap E^-} |\sigma(e)| + \sum_{e \in \text{inter}(\mathcal{C}) \cap E^+} |\sigma(e)|.$$

In the unweighted version of the problem, this is simply  $\text{COST}(G, \mathcal{C}) = |\text{intra}(\mathcal{C}) \cap E^-| + |\text{inter}(\mathcal{C}) \cap E^+|$ . The goal of correlation clustering<sup>1</sup> is to find  $\mathcal{C}$  to minimize  $\text{COST}(G, \mathcal{C})$ . For unweighted correlation clustering, there are constant-factor approximation algorithms (Bansal et al., 2004; Ailon et al., 2008; Chawla et al., 2015), with the best known constant being 2.06. For the weighted case, the best known algorithm obtains an  $O(\log n)$ -approximation (Demaine et al., 2006).

**Fairness constraints.** In the fair version of any clustering problem, each vertex  $v \in V$  has a color  $c(v)$ . Proportional fairness (Chierichetti et al., 2017) requires that in every cluster, the number of vertices of each color is proportional to the corresponding number in the whole graph. In particular, in the symmetric case where each color appears the same number of times in the graph, we require the same in each cluster. Ahmadian et al. (2019) relaxed this property by requiring that each color constitutes at most an  $\alpha$ -fraction of each cluster, for a given  $\alpha \in (0, 1)$ . Bera et al. (2019) further generalized this notion to include lower bounds on the number of vertices of each color in each cluster.

We give a general reduction from fair correlation clustering to a median fairlet decomposition that works for any of these definitions of fairness, and in fact for a more general class of constraints. As long as the fairlet decomposition problem can be solved with small fairlets (holds for the above fairness definitions; see Section 4), this will give us an approximation algorithm for the corresponding fair correlation clustering problems.

### 3 Overview of Results

In this section, we give a high-level overview of our algorithm and our main result. A key ingredient of our algorithm is a general reduction from the given constrained correlation clustering problem (as defined below) to a *fairlet decomposition* problem. We then show how the cost of a fairlet decomposition can be approximated by a median clustering cost function.

<sup>1</sup>This is the minimizing disagreements variant of correlation clustering. A maximizing agreements version can also be defined similarly. In this paper we focus on minimizing disagreements, since the maximization version admits a trivial randomized 2-approximation that can be made fair.

This allows us to use previous results on the fair median problem to solve fairlet decomposition for the standard notions of fairness defined in the previous section. Finally, given an approximately optimal fairlet decomposition, we use our reduction to reduce the constrained correlation clustering instance to a standard correlation clustering instance, and apply known algorithms (Bansal et al., 2004; Ailon et al., 2008; Chawla et al., 2015) to solve this problem.

**Constrained correlation clustering.** We start by defining a general class of *constrained correlation clustering* problems. Consider an unweighted correlation clustering instance  $G$  and let  $\mathcal{F}$  be a family of subsets of  $V$ . We treat  $\mathcal{F}$  as the family of feasible clusters, and assume it has the following *composability* property: for every  $F_1, F_2 \in \mathcal{F}$ , we have  $F_1 \cup F_2 \in \mathcal{F}$ . Note this property is satisfied when  $\mathcal{F}$  is the collection of all fair sets under any of the definitions of fairness given in Section 2. The constrained correlation clustering problem is to define a correlation clustering  $\mathcal{C}$  with minimum  $\text{COST}(G, \mathcal{C})$  such that for all  $C \in \mathcal{C}$ , we have  $C \in \mathcal{F}$ .

**Fairlet decomposition.** Next, we define the notion of *fairlet decomposition* used in our reduction. A fairlet decomposition for a constrained correlation clustering problem is simply a partition  $\mathcal{P} = \{P_1, P_2, \dots\}$  of  $V$  into subsets in  $\mathcal{F}$ , i.e.,  $P_i \in \mathcal{F}$  for all  $i$ . We call each  $P_i$  a *fairlet*. The key in our reduction is a cost function  $\text{FCOST}$  that evaluates  $\mathcal{P}$ 's usefulness in building a correlation clustering of  $G$ . Here we define this cost function, and in Section 4.1 we show how it can be approximated by the standard median clustering cost function in a carefully defined metric space.

**Fairlet decomposition cost.** Consider a fairlet decomposition  $\mathcal{P} = \{P_1, P_2, \dots\}$ . For each fairlet  $P_i$ , we let  $\text{FCOST}^{\text{in}}(P_i)$  be the number of negative edges inside  $P_i$ , i.e.,  $\text{FCOST}^{\text{in}}(P_i) = |E^- \cap \text{intra}(P_i)|$ . For fairlets  $P_i, P_j$ , we let  $\text{FCOST}^{\text{out}}(P_i, P_j)$  be the number of edges between them with the minority sign, i.e.,  $\text{FCOST}^{\text{out}}(P_i, P_j) = \min(|E^-(P_i, P_j)|, |E^+(P_i, P_j)|)$ .

Finally, we let  $\text{FCOST}^{\text{in}}(\mathcal{P}) = \sum_i \text{FCOST}^{\text{in}}(P_i)$ ,  $\text{FCOST}^{\text{out}}(\mathcal{P}) = \sum_{i < j} \text{FCOST}^{\text{out}}(P_i, P_j)$ , and  $\text{FCOST}(\mathcal{P}) = \text{FCOST}^{\text{in}}(\mathcal{P}) + \text{FCOST}^{\text{out}}(\mathcal{P})$ .

**Reduced instance.** Given a constrained correlation clustering instance  $G$  and a fairlet decomposition  $\mathcal{P}$  for  $G$ , we define a *reduced* correlation clustering instance as follows. Let  $G^{\mathcal{P}}$  be a complete graph on  $\{p_1, \dots, p_{|\mathcal{P}|}\}$ , where vertex  $p_i$  corresponds to fairlet  $P_i \in \mathcal{P}$ . The label  $\sigma(p_i, p_j)$  of the edge between  $p_i$  and  $p_j$  is the majority sign of the edges in  $E(P_i, P_j)$  (with ties broken arbitrarily) multiplied by a weight that is equal to the number of edges in  $E(P_i, P_j)$  with the majority sign.

Note that the instance  $G^{\mathcal{P}}$  defined above is an instance of *weighted* correlation clustering, although as we will observe, the edges have weights that are within a constant factor each other, and therefore the problem can be solved using unweighted correlation clustering algorithms. Given a solution to this problem, it can be expanded into a solution of the original constrained problem. The final algorithm is sketched below.

---

**Algorithm 1** Constrained Correlation Clustering
 

---

- 1:  $\mathcal{P} \leftarrow$  approx. fairlet decomp. (Lemmas 4.5, 4.7).
  - 2:  $G^{\mathcal{P}}$ :  $(p_i, p_j)$  gets majority sign in  $E(P_i, P_j)$  and weight  $\max(|E^+(P_i, P_j)|, |E^-(P_i, P_j)|)$ .
  - 3: Let  $\mathcal{C}$  be an approximate (non-constrained) correlation clustering solution of  $G^{\mathcal{P}}$ .
  - 4: Output the clustering  $\{\bigcup_{p_j \in C_i} P_j : C_i \in \mathcal{C}\}$ .
- 

To prove that the Algorithm 1 produces an approximately optimal solution to the constrained correlation clustering problem, we need the following lemmas. The first two lemmas prove that a solution of  $G$  can be transformed to a solution of  $G^{\mathcal{P}}$  and vice versa, and these transformations do not increase the cost by more than the cost of the fairlet decomposition. The third lemma bounds the cost of a fairlet decomposition in terms of the cost of the optimal solution to the constrained correlation clustering problem. The proofs of these lemmas are presented in Supplementary Material.

**Lemma 3.1.** *Given a correlation clustering instance  $G$ , a fairlet decomposition  $\mathcal{P}$  for  $G$ , and a clustering  $\mathcal{C}$  of  $G$ , there exists a clustering  $\mathcal{C}'$  of  $G^{\mathcal{P}}$  such that  $\text{COST}(G^{\mathcal{P}}, \mathcal{C}') \leq \text{COST}(G, \mathcal{C}) + \text{FCOST}^{\text{out}}(\mathcal{P})$ .*

**Lemma 3.2.** *Let  $\mathcal{C}$  be a clustering of  $G^{\mathcal{P}}$  and  $\mathcal{C}'$  be the clustering computed in line 4 of Algorithm 1. Then,  $\text{COST}(G, \mathcal{C}') \leq \text{COST}(G^{\mathcal{P}}, \mathcal{C}) + \text{FCOST}(\mathcal{P})$ .*

**Lemma 3.3.** *For any constrained correlation clustering instance  $G$ , and any constrained clustering  $\mathcal{C}$  of  $G$ , there is a fairlet decomposition  $\mathcal{P}$  of  $G$  satisfying  $\text{FCOST}(\mathcal{P}) \leq \text{COST}(G, \mathcal{C})$ .*

Putting these together, we have the following:

**Theorem 3.4.** *Assume there is an  $\eta$ -approximation algorithm  $A_1$  for finding the minimum cost fairlet decomposition  $\mathcal{P}$  and a  $\beta$ -approximation algorithm  $A_2$  for the unconstrained correlation clustering instance  $G^{\mathcal{P}}$ . Then Algorithm 1 produces a  $(\beta(1 + \eta) + \eta)$ -approximation for the constrained correlation clustering instance  $G$ .*

*Proof.* Let OPT be an optimal solution to the constrained correlation clustering instance  $G$ . By Lemma 3.3, the fairlet decomposition problem has a solution of cost at most  $\text{COST}(G, \text{OPT})$ , and therefore, algorithm  $A_1$  for this problem must find a decomposition  $\mathcal{P}$  with  $\text{FCOST}(\mathcal{P}) \leq \eta \cdot \text{COST}(G, \text{OPT})$ .

Also, by Lemma 3.1, the instance  $G^{\mathcal{P}}$  has a solution of cost at most  $(1 + \eta) \cdot \text{COST}(G, \text{OPT})$ . Therefore, algorithm  $A_2$  can find a clustering  $\mathcal{C}$  of cost at most  $\beta(1 + \eta) \cdot \text{COST}(G, \text{OPT})$ . Thus, by Lemma 3.2, the cost of the clustering produced by Algorithm 1 is at most  $(\beta(1 + \eta) + \eta) \cdot \text{COST}(G, \text{OPT})$ . Finally, by the composability property of the constraints, we know that this clustering satisfies the constraints, since each of its clusters is a union of fairlets in  $\mathcal{P}$ .  $\square$

In the following, we explain the approximation factor  $\beta$  we can get for solving unconstrained correlation clustering instance  $G^{\mathcal{P}}$  and we dedicate the next section to approximation ratios  $\eta$  that we can get for minimum cost fairlet decomposition problem depending on the fairness parameter  $\alpha$  and the number of colors in a given fair correlation clustering instance.

**Lemma 3.5.** *There exists an approximation algorithm for unconstrained correlation clustering of  $G^{\mathcal{P}}$  with approximation ratio of  $\beta = \min(\log n, 2\rho r^2)$  where  $r = \frac{\max_{P \in \mathcal{P}} |P|}{\min_{P \in \mathcal{P}} |P|}$  and  $\rho$  is the approximation factor of unweighted correlation clustering<sup>2</sup>.*

*Proof.* Since the reduced correlation clustering instance is a weighted instance, there exists an  $O(\log n)$ -approximation (Demaine et al., 2006). Now since the weight of the edge between  $p_i$  and  $p_j$  in  $G^{\mathcal{P}}$  is at least  $|P_i| \cdot |P_j|/2$  and at most  $|P_i| \cdot |P_j|$ , any two edges weights are within  $2r^2$  of each other. So if we remove the weights from  $G^{\mathcal{P}}$  and solve the resulting unweighted instance, we will get a  $(2\rho r^2)$ -approximation.  $\square$

## 4 Fairlet Decomposition

In this section we show how to solve the fairlet decomposition problem by reducing it to a fair clustering problem with the median cost function in an appropriate metric space. The reduction (Section 4.1), loses a factor that is proportional to the size of the largest fairlet, but as we show in Section 4.2, in cases that we know how to solve the fairlet decomposition problem, the size of the fairlets can be guaranteed to be small.

### 4.1 Reduction to median cost

Consider a correlation clustering instance  $G$  and let  $d$  be a distance function defined on a metric space  $M$  containing the set of vertices  $V$ . For a fairlet decomposition  $\mathcal{P} = \{P_1, P_2, \dots\}$ , we define the following median cost:  $\text{MCOST}(P_i) = \min_{u \in M} \sum_{v \in P_i} d(u, v)$  and  $\text{MCOST}(\mathcal{P}) = \sum_{P_i \in \mathcal{P}} \text{MCOST}(P_i)$ . Notice that the problem of finding the fairlet decomposition with minimum

<sup>2</sup>Currently best known approximation factor is 2.06 (Chawla et al., 2015)

$\text{MCOST}(\mathcal{P})$  is precisely the fairlet decomposition problem for fair  $k$ -median, as studied by Chierichetti et al. (2017); Bera et al. (2019).

We now define a metric space  $(M, d)$  such that the median cost  $\text{MCOST}$  can approximate the fairlet cost  $\text{FCOST}$ . We first define an embedding  $\phi : V \rightarrow [0, 1]^n$  as follows. For a vertex  $u \in V$ , let

$$\phi(u)_v = \begin{cases} 1 & \text{if } u = v \text{ or } (u, v) \in E^+ \\ 0 & \text{if } (u, v) \in E^- . \end{cases}$$

In other words,  $\phi(v)$  is the  $v$ th row of the adjacency matrix of  $G(V, E^+)$  after adding a positive self-loop at every vertex. Now we let  $M = [0, 1]^n$ , and place the vertices in  $V$  in this space using the mapping  $\phi$ . We let  $d(\cdot, \cdot)$  be the Hamming distance between the points in  $M$ . In other words, for vertices  $u, v \in V$ , we have  $d(u, v) = |\phi(u) - \phi(v)|$ . Intuitively,  $d(u, v)$  measures the ‘‘cost’’ of committing to put  $u$  and  $v$  in one cluster in the correlation clustering instance.

We now prove the following two lemmas, which show that the  $\text{FCOST}$  of a fairlet decomposition is close to its  $\text{MCOST}$  with respect to the metric  $d$ . The proofs of these lemmas are in the Supplementary Material.

**Lemma 4.1.** *For any fairlet decomposition  $\mathcal{P}$ , we have  $\text{MCOST}(\mathcal{P}) \leq 2 \cdot \text{FCOST}(\mathcal{P})$ .*

**Lemma 4.2.** *Let  $\mathcal{P}$  be any fairlet decomposition and let  $f = \max_{P \in \mathcal{P}} |P|$ . Then,  $\text{FCOST}(\mathcal{P}) \leq 2f \cdot \text{MCOST}(\mathcal{P})$ .*

Using the above lemmas, we have the following.

**Theorem 4.3.** *Assume there is a  $\gamma$ -approximation algorithm for fairlet decomposition with median costs. Furthermore, assume that this algorithm always produces fairlets of size at most  $f$ . Then the solution produced by this algorithm is a  $(4f\gamma)$ -approximation to the problem of finding a fairlet decomposition with minimum  $\text{FCOST}$ .*

## 4.2 Algorithms for fairlet decomposition

In this section, we give algorithms for fairlet decomposition with median cost for several notions of fairness. Using Theorem 4.3, these algorithms imply algorithms for fairlet decomposition problem and provide the algorithm  $A_1$  in Theorems 3.4. We focus on three fairness constraints: an upper bound of  $\alpha = \frac{1}{2}$  on the fraction of vertices of each color in each cluster; an upper bound of  $\alpha = 1/C$  where  $C$  is the number of distinct colors (this corresponds to the proportional fairness property studied by Chierichetti et al. (2017)); and an upper bound of  $\alpha = 1/t$  for an integer  $t$  on the fraction of vertices of each color in each cluster. We give approximation algorithms for the first two cases and a bicriteria approximation for the third case, with

upper bounds of 3,  $C$ , and  $2t - 1$ , respectively, on the size of fairlets.

Throughout this subsection, when we speak of the cost of a fairlet decomposition, we mean its median cost.

### 4.2.1 $\alpha = 1/2$

This is probably the most common case of fair decomposition where clusters are required to not have a dominant color. In this case, we can show that a fairlets have size at most 3 and find these fairlets by solving a minimum weight 2-factor problem in a graph. Recall that a 2-factor is a subgraph where each vertex has degree 2 and edges may be used multiple times. This problem can be solved polynomially (Schrijver, 2003, Chapter 21). Define a graph  $H$  on points in  $V$  as follows: two vertices  $u, v$  are connected by an edge if they have distinct colors; the weight of the edge is  $d(u, v)$ . We first bound the cost of the optimal 2-factor and then explain our approximation factor in the lemma.

**Lemma 4.4.** *The cost of an optimal 2-factor in  $H$  can be bounded by  $2 \cdot \text{MCOST}(\mathcal{P}^*)$ , where  $\mathcal{P}^*$  is the optimal fairlet decomposition.*

*Proof.* We construct a feasible 2-factor by constructing a 2-factor for each fairlet  $P \in \mathcal{P}^*$  with center  $\mu$ . Since there are at most  $|P|/2$  vertices of any color, depending on the parity of  $P$ , vertices of  $P$  can be covered by matching and a possible multi-color triangle in  $H$ . Doubling the matching edges, we can get a 2-factor for covering  $\mathcal{P}^*$ . It remains to bound the cost of this 2-factor. For a matching edge  $(u, v)$  for  $u, v \in P$ , by triangle inequality,  $d(u, v) \leq d(u, \mu_i) + d(v, \mu_i)$ , and for a triangle  $(u, v, w)$ , the sum of pairwise distances can be bounded by  $2(d(u, \mu) + d(v, \mu) + d(w, \mu))$ . Hence the cost of the proposed 2-factor for covering  $\mathcal{P}^*$  is at most  $2 \cdot \text{MCOST}(\mathcal{P}^*)$  and the overall cost of the optimal 2-factor is at most  $2 \cdot \text{MCOST}(\mathcal{P}^*)$ .  $\square$

**Lemma 4.5.** *For  $\alpha = 1/2$ , there is an algorithm for fairlet decomposition that returns a solution with*

- median cost at most  $2 \cdot \text{MCOST}(\mathcal{P}^*)$ .
- the size of largest fairlet is at most 3.
- the size of smallest fairlet is at least 2.

*Proof.* Consider an optimal 2-factor in  $H$ . Define a fairlet decomposition as follows. For each cycle of even length, consider a set of alternating edges and let each alternating edge be a fairlet with one of the endpoints chosen as the center. For a cycle of odd length, there must exist three consecutive vertices of pairwise distinct colors. In this case, let one fairlet be these three vertices with the middle vertex chosen as the center and for the (unique) alternating edges covering the remaining vertices, let each edge be a

fairlet with one of the endpoints chosen as the center. The median cost of these fairlets is at most the weight of the original 2-factor, which is at most  $2 \cdot \text{MCOST}(\mathcal{P}^*)$  by Lemma 4.4; the proof follows by construction.  $\square$

Lemma 4.5 and Lemma 3.5 yield the following.

**Theorem 4.6.** *For  $\alpha = 1/2$ , there is a 256-approximation algorithm for fair correlation clustering.*

#### 4.2.2 $\alpha = 1/C$ for $C$ colors

In the case of  $\alpha = 1/C$ , each fairlet has equal number of points of each color and so these points can be matched together. We use this observation and devise an algorithm based on solving repeated matching problems in graph  $H$  (construction explained in Section 3).

**Lemma 4.7.** *For  $\alpha = 1/C$ , there is an algorithm for fairlet decomposition that returns a solution with*

- median cost at most  $C \cdot \text{MCOST}(\mathcal{P}^*)$ ,
- the size of each fairlet is  $C$ .

*Proof.* Consider an arbitrary ordering of the colors and solve a mincost matching problem between points of color  $c$  and  $c + 1$  in the graph  $H$ . The union of these matchings yields a partition of  $V$  into paths of length  $C$ . Each such path is a fairlet; let  $\mathcal{P}$  denote this fairlet decomposition. It remains to bound the cost of  $\mathcal{P}$ .

Let us fix two colors  $c$  and  $c + 1$  and let  $M$  be an arbitrary matching between vertices of color  $c$  and  $c + 1$  such that point  $u$  is matched to a point  $v$  only if  $u$  and  $v$  belong to the same partition of  $\mathcal{P}^*$ . Since each part in  $\mathcal{P}^*$  has equal number of vertices of each color and there is an edge between any two vertices of different colors in  $H$ , matching  $M$  exists. Now since the cost of each matching edge  $(u, v)$  can be bounded by  $d(u, \mu) + d(v, \mu)$  where  $\mu$  is the center of the partition containing  $u$  and  $v$ , the cost of  $M$  can be bounded by median cost of serving clients of colors  $c$  and  $c + 1$ . Since each color is matched twice, the total cost of each path corresponding to a partition is at most  $2 \cdot \text{MCOST}(\mathcal{P}^*)$ . Now for each path we pick the middle vertex as center and the cost of assigning vertices of the path to the center is at most  $C/2$  cost of the path as each edge is charged at most  $C/2$  times. Hence  $\text{MCOST}(\mathcal{P}) \leq C \cdot \text{MCOST}(\mathcal{P}^*)$ .  $\square$

So in this case, we get a 2-approximation with fairlets of size at most  $C$ . Note that in the motivating application of fair clustering, the color of each vertex corresponds to one possible value of a sensitive feature like race or gender, and therefore the value  $C$  tends to be small in such applications.

Lemma 4.7 and Lemma 3.5 yield the following.

**Theorem 4.8.** *For  $\alpha = 1/C$ , there is a  $(16.48C^2)$ -approximation algorithm for fair correlation clustering.*

#### 4.2.3 $\alpha = 1/t$

The fair decomposition with median cost is not well studied in the literature and in this paper, we were able to devise algorithms for the case of  $\alpha = 1/2$  and  $\alpha = 1/C$ . Next we consider the case where  $1/\alpha \in \mathbb{Z}_+$  and we argue how to utilize any approximation algorithm for fairlet decomposition as a black-box to build an algorithm for fair correlation clustering. While we allow the black-box algorithm to produce fairlets of arbitrary size, the following lemma ensures that the size of the fairlets can be bounded.

**Lemma 4.9.** *For any set  $P$  that satisfies fairness constraint with  $\alpha = 1/t$ , there exists a partition of  $P$  into sets  $(P_1, P_2, \dots)$  where each  $P_i$  satisfies the fairness constraint and  $t \leq |P_i| < 2t$ .*

*Proof.* Let  $p = m \times t + r$  with  $0 \leq r < t$ . Then, the fairness constraints ensures that there are at most  $m$  elements of each color. Consider the partitioning obtained through the following process: consider an ordering of elements where points of the same color are in consecutive places, assign points to sets  $P_1, \dots, P_m$  in a round-robin fashion. So each set  $P_i$  gets at least  $t$  elements and at most  $t + r < 2t$  elements assigned to it. Since there are at most  $m$  elements of each color, each set gets at most one point of any color and hence all sets satisfy the fairness constraint as  $1 \leq \frac{1}{t} \cdot |P_i|$ .  $\square$

**Theorem 4.10.** *For  $\alpha = 1/t$ , given an  $\gamma$ -approximation algorithm for fairlet decomposition with median cost, there is an  $O(t\gamma)$ -approximation algorithm for fair correlation clustering.*

## 5 Experiments

In this section we present our experiments demonstrating that our algorithm solves the correlation clustering problem with fairness constraints with only a limited loss in the cost when compared to the vanilla (unfair) solution. We describe the datasets used, the algorithms evaluated, the quality measures, and our results. Our code is available open source (Ahmadian et al., 2020).

**Datasets.** We use publicly-available datasets from the UCI Repository<sup>3</sup> and from the SNAP Datasets<sup>4</sup>.

The datasets represent complete signed graphs from different domains, including both co-purchasing relationships among products, and semantic similarities among texts learned with embedding methods. The graphs are represented by complete signed matrices up to 1600x1600 in size, up to 0.9 million positive edges,

<sup>3</sup>[archive.ics.uci.edu/ml](http://archive.ics.uci.edu/ml)

<sup>4</sup>[snap.stanford.edu/data/](http://snap.stanford.edu/data/)

Dataset	Unfair Alg. ERROR		Unfair Alg. IMBALANCE		Fair Alg. ERROR		
	LOCAL	PIVOT	LOCAL	PIVOT	MATCH + LOCAL	SINGLE	RAND
amazon	0.010	0.011	0.40	0.39	0.064	0.786	0.215
reuters, $\theta = 0.25$	0.096	0.161	0.64	0.59	0.230	0.754	0.255
reuters, $\theta = 0.50$	0.181	0.231	0.50	0.40	0.350	0.504	0.502
reuters, $\theta = 0.75$	0.188	0.241	0.15	0.25	0.199	0.252	0.746
victorian, $\theta = 0.25$	0.109	0.158	0.53	0.46	0.212	0.753	0.251
victorian, $\theta = 0.50$	0.183	0.268	0.31	0.23	0.348	0.502	0.499
victorian, $\theta = 0.75$	0.203	0.280	0.12	0.12	0.237	0.251	0.747
mean over datasets	0.139	0.193	0.38	0.35	0.234	0.459	0.543

Table 1: ERROR and IMBALANCE in  $C = 2$  color case for various datasets and different threshold  $\theta$  for the quantile used for positive edges. Notice how our algorithm MATCH + LOCAL has cost comparable to PIVOT and not much higher than LOCAL while reducing the imbalance from the up to 65% of the unfair algorithms to 0.

and up to  $C = 16$  colors. Here we only briefly describe the datasets; for details see Supplementary Material.

**amazon:** Vertices represent products on Amazon (Leskovec et al., 2007), the color is the item category, and two co-reviewed items have a +1 weight edge (all non-co-reviewed items have  $-1$  weights). We use 1000 vertices equally distributed in two popular categories.

**reuters** and **victorian:** These are extracted from text data used in previous fair clustering work (Ahmadian et al., 2019). The datasets include between 50 and 100 English language texts from each of up to 16 authors.<sup>5</sup> Each vertex represents a text and the color represents the author. For each text we obtain a semantic embedding vector with standard methods. We use a threshold on the dot product of the embedding vectors for the edges. Through this operation, we set the top  $\theta \in \{0.25, 0.50, 0.75\}$  fraction of edges via dot products as +1, and the remaining edges as  $-1$ .

**Algorithms.** We evaluate Algorithm 1 in two fairness scenarios: with an upper bound of  $\alpha = 1/2$  of the vertices for each color, and with  $\alpha = 1/C$  for equal color representation (for the two-color case, the two are equivalent). For the  $\alpha = 1/2$  case, in our experiments we simplify the algorithm of Section 4.2.2 to compute a minimum-cost perfect matching (1-factor) instead of a 2-factor decomposition. This can be formally shown to be sufficient for the  $C = 2$  case, or when all optimal clusters are even sized, and we observe it works well in practice in our experiments. For  $\alpha = 1/C$ , we implement the repeated matching algorithm in 4.2.2 to obtain the fairlets in a similar fashion. After finding the fairlets, we use an in-house correlation clustering solver based on local search (LOCAL). We refer to our algorithms as MATCH + LOCAL for the  $\alpha = 1/2$  case

<sup>5</sup>The datasets are available at [archive.ics.uci.edu/ml/datasets/Reuter\\_50\\_50](http://archive.ics.uci.edu/ml/datasets/Reuter_50_50) and [archive.ics.uci.edu/ml/datasets/Victorian+Era+Authorship+Attribution](http://archive.ics.uci.edu/ml/datasets/Victorian+Era+Authorship+Attribution)

Algorithm	ERROR	IMBALANCE for 1/2	IMBALANCE for equality
LOCAL	0.249	0.011	0.218
PIVOT	0.345	0.008	0.191
MATCH + LOCAL	0.255	0	0.180
REP. MATCH + LOCAL	0.321	0	0
SINGLE	0.5	0	0
RAND	0.5	0	0

Table 2: Results for victorian,  $\theta = 0.50$ ,  $C = 8$  colors.

and as REP. MATCH + LOCAL for the  $\alpha = 1/C$  case.

We also consider the following (unfair) baseline algorithms: the standard PIVOT algorithm of Ailon et al. (2008) for correlation clustering (for PIVOT, we repeat the randomized algorithm 10 times and use the best result), and the (unfair) local search heuristic LOCAL used as part of our algorithm. In addition, as no prior work has addressed the fair correlation clustering problem, we compare our algorithm with two simple fair baselines: the whole graph as one cluster SINGLE, and a random fairlet decomposition RAND.

**Quality measures.** For each algorithm, we report the following measures. The ERROR of the correlation cost of the clustering obtained by the algorithm, presented as the ratio of the edges of the graph that are in disagreement with the clustering (i.e., inter-cluster positive edges and intra-cluster negative edges) over the number of edges (i.e., 0 corresponds to a perfect solution, whereas 1 corresponds to a completely incorrect solution). For fairness, we report the IMBALANCE as the total fraction of vertices that violate the  $\alpha$  color representation constraint, i.e., vertices for each cluster and color that are above the  $\alpha$  fraction for the size of the cluster (Ahmadian et al., 2019). More precisely, let  $P$  be a cluster in the solution of an  $\alpha$ -color constraint instance. The maximum allowed number of

points of a certain color in the cluster  $P$  is  $\lfloor |P|\alpha \rfloor$ . Let  $V_c$  be the vertices of color  $c$  on the graph and  $\Delta_P = \sum_{P,c} \max(|P \cap V_c| - \lfloor |P|\alpha \rfloor, 0)$  be the vertices violating the constraint in  $P$ . We report  $\sum_P \Delta_P / |V|$  as the IMBALANCE for  $\alpha \in \{1/2, 1/C\}$  (i.e., 0 corresponds to no imbalance, and 1 corresponds to complete imbalance). We repeat all algorithms 10 times and report the mean results for all measures.

### 5.1 Results for $C = 2$

Table 1 shows the results for the  $C = 2$  color case with  $\alpha = 1/2$ , i.e., equal representation over the various datasets. We use the MATCH + LOCAL as a fair algorithm, which has an IMBALANCE of 0 by construction.

The table shows clearly that our algorithm MATCH + LOCAL obtains clusters that are fair and has costs comparable to the unfair PIVOT baseline (and sometimes better) and slightly worse than the LOCAL baseline. On average, over all datasets, our algorithm has an average cost of 0.234 vs 0.193 for PIVOT and 0.139 for LOCAL. On the other hand, our algorithm is significantly better than both SINGLE and RAND, which, while satisfying fairness, have exorbitantly high costs.

Notice how the LOCAL and the PIVOT baselines have very high IMBALANCE values of up to 65% of the vertices (as they are oblivious to colors), showing the importance of developing novel algorithms for the problem. This result is not surprising. If pairs of vertices of the same color are more likely to be similar, it is expected that many clusters will contain vast majorities of points with a single color.

### 5.2 Results for $C > 2$

Here, we study the behavior of our algorithm in the case when more than two colors are present in the dataset. We use MATCH + LOCAL to obtain a  $\alpha = 1/2$  fair solution and REP. MATCH + LOCAL to obtain a  $\alpha = 1/C$  (equal representation) solution.

We report an overview of our experimental results in Table 2 for the dataset victorian with threshold  $\theta = 0.50$  and  $C = 8$  colors (note that this is a different graph than that produced with the previous  $C = 2$  dataset). More experimental results are available in Supplementary Material.

It is easy to see that all the earlier trends continue to hold. Notice how the algorithm for the  $1/2$  case MATCH + LOCAL is only marginally worse than the best unfair solution LOCAL and much better than all other baselines. Our algorithm for the more difficult equal representation case is again better than the PIVOT baseline. Notice how all unfair algorithms are significantly far from being equally balanced. However, the presence




Figure 1: ERROR of our algorithms over that of the unfair LOCAL algorithms for  $\alpha = 1/2$  and  $\alpha = 1/C$ , on a series of graphs from victorian,  $\theta = 0.50$ , and using  $C = 2$  to  $C = 16$  colors.

of many colors makes it easier to get closer to the  $1/2$  threshold.

We confirm this observation in Figure 1. The figure shows a comparison of the ERROR for our algorithms for  $\alpha = 1/2$  and  $\alpha = 1/C$  vs the ERROR for the vanilla LOCAL algorithms as the number of colors goes from  $C = 2$  to  $C = 16$ . We report the ratio of MATCH + LOCAL over LOCAL as a solid line, and the ratio for REP. MATCH + LOCAL as a dashed line. Notice how the error for our algorithm for the  $\alpha = 1/2$  case gets closer and closer to the LOCAL output error for more colors. Again, this result is obtained because the presence of many colors makes the problem easier for the  $\alpha = 1/2$  case. The performance of the algorithm for the  $\alpha = 1/C$  case has a less stable pattern, but it confirms that the algorithm is quite competitive with the unfair solution (between 30–80% higher error) even for quite a few colors. These results are significantly better than what is expected from a worst-case analysis.

Finally, we report having observed that using just MATCH or REP. MATCH fairlets without the re-clustering part of the algorithm is not sufficient for obtaining good results, as the re-clustering step is needed for obtaining clusters that do not have large errors.

## 6 Conclusions

In this paper we initiated the study of correlation clustering with fairness constraints. We showed a reduction to the fairlet decomposition problem with a standard median cost function, for a carefully chosen distance function. Using this, and old and new results on the fairlet decomposition problem with a median cost function, we obtained provable constant-factor approximation algorithms for fair correlation clustering for various notions of fairness. Our experimental evaluation shows that these algorithms perform well not only in theory but also in practice.



## References

- S. Ahmadian, A. Norouzi-Fard, O. Svensson, and J. Ward. Better guarantees for  $k$ -means and Euclidean  $k$ -median by primal-dual algorithms. In *FOCS*, pages 61–72, 2017.
- S. Ahmadian, A. Epasto, R. Kumar, and M. Mahdian. Clustering without over-representation. In *KDD*, pages 267–275, 2019.
- S. Ahmadian, A. Epasto, R. Kumar, and M. Mahdian. Fair correlation clustering: Open-source code, 2020. URL [https://github.com/google-research/google-research/tree/master/correlation\\_clustering](https://github.com/google-research/google-research/tree/master/correlation_clustering).
- K. Ahn, G. Cormode, S. Guha, A. McGregor, and A. Wirth. Correlation clustering in data streams. In *ICML*, pages 2237–2246, 2015.
- N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5):23:1–23:27, 2008.
- A. Backurs, P. Indyk, K. Onak, B. Schieber, A. Vakilian, and T. Wagner. Scalable fair clustering. In *ICML*, pages 405–413, 2019.
- N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Mach. Learn.*, 56(1-3):89–113, 2004.
- S. K. Bera, D. Chakrabarty, N. Flores, and M. Negahbani. Fair algorithms for clustering. In *NeurIPS*, pages 4955–4966, 2019.
- I. O. Bercea, M. Gross, S. Khuller, A. Kumar, C. Rösner, D. R. Schmidt, and M. Schmidt. On the cost of essentially fair clusterings. In *APPROX-RANDOM*, pages 18:1–18:22, 2019.
- M. Bressan, N. Cesa-Bianchi, A. Paudice, and F. Vitale. Correlation clustering with adaptive similarity queries. In *NeurIPS*, pages 12510–12519, 2019.
- T. Calders and S. Verwer. Three naive bayes approaches for discrimination-free classification. *DMKD*, 21(2):277–292, 2010.
- L. E. Celis, L. Huang, and N. K. Vishnoi. Multiwinner voting with fairness constraints. In *IJCAI*, pages 144–151, 2018a.
- L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. In *ICALP*, pages 28:1–28:15, 2018b.
- M. Charikar, V. Guruswami, and A. Wirth. Clustering with qualitative information. *JCSS*, 71(3):360–383, 2005.
- S. Chawla, K. Makarychev, T. Schramm, and G. Yaroslavtsev. Near optimal LP rounding algorithm for correlation clustering on complete and complete  $k$ -partite graphs. In *STOC*, pages 291–228, 2015.
- X. Chen, B. Fain, C. Lyu, and K. Munagala. Proportionally fair clustering. In *ICML*, pages 1032–1041, 2019.
- F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *NIPS*, pages 5029–5037, 2017.
- F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Matroids, matchings, and fairness. In *AISTATS*, pages 2212–2220, 2019.
- A. Dash, A. Shandilya, A. Biswas, K. Ghosh, S. Ghosh, and A. Chakraborty. Summarizing user-generated textual content: Motivation and methods for fairness in algorithmic summaries. In *CSCW*, pages 1–28, 2019.
- E. D. Demaine, D. Emanuel, A. Fiat, and N. Immerlica. Correlation clustering in general weighted graphs. *TCS*, 361(2-3):172–187, 2006.
- C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *ITCS*, pages 214–226, 2012.
- H. Elzayn, S. Jabbari, C. Jung, M. Kearns, S. Neel, A. Roth, and Z. Schutzman. Fair algorithms for learning in allocation problems. In *FAT*, pages 170–179, 2019.
- B. Fang, M. Jiang, and J. Shen. Achieving fairness in determining medicaid eligibility through fairgroup construction. *arXiv:1906.00128*, 2019.
- M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *KDD*, pages 259–268, 2015.
- B. Fish, J. Kun, and Á. D. Lelkes. A confidence-based approach for balancing fairness and accuracy. In *SDM*, pages 144–152, 2016.
- T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *TCS*, 38:293–306, 1985.
- A. Gungor. Fifty victorian era novelists authorship attribution data. 2018.
- S. Har-Peled and S. Mahabadi. Near neighbor: Who is the fairest of them all? In *NeurIPS*, pages 13176–13187, 2019.
- D. S. Hochbaum and D. B. Shmoys. A best possible heuristic for the  $k$ -center problem. *MOR*, 10(2):180–184, 1985.
- L. Huang, S. H. C. Jiang, and N. K. Vishnoi. Coresets for clustering with fairness constraints. In *NeurIPS*, pages 7587–7598, 2019.
- A. K. Jain. Data clustering: 50 years beyond  $k$ -means. *Pattern Recognition Letters*, 31(8):651–666, 2010.

- M. Joseph, M. Kearns, J. H. Morgenstern, and A. Roth. Fairness in learning: Classic and contextual bandits. In *NIPS*, pages 325–333, 2016.
- S. Kalhan, K. Makarychev, and T. Zhou. Improved algorithms for correlation clustering with local objectives. In *NeurIPS*, pages 9341–9350, 2019.
- T. Kamishima, S. Akaho, and J. Sakuma. Fairness-aware learning through regularization approach. In *ICDM Workshops*, pages 643–650, 2011.
- T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *ECML/PKDD*, pages 35–50, 2012.
- T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. A local search approximation algorithm for  $k$ -means clustering. *Computational Geometry*, 28(2-3):89–112, 2004.
- M. Kleindessner, P. Awasthi, and J. Morgenstern. Fair  $k$ -center clustering for data summarization. In *ICML*, pages 3448–3457, 2019a.
- M. Kleindessner, S. Samadi, P. Awasthi, and J. Morgenstern. Guarantees for spectral clustering with fairness constraints. In *ICML*, pages 3448–3457, 2019b.
- S. Kushagra, S. Ben-David, and I. Ilyas. Semi-supervised clustering for de-duplication. In *AIS-TATS*, pages 1659–1667, 2019.
- J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *TWEB*, 1(1):5, 2007.
- J. Li, K. Yi, and Q. Zhang. Clustering with diversity. In *ICALP*, pages 188–200, 2010.
- S. Li and O. Svensson. Approximating  $k$ -median via pseudo-approximation. *SICOMP*, 45(2):530–547, 2016.
- X. Pan, D. Papailiopoulos, S. Oymak, B. Recht, K. Ramachandran, and M. I. Jordan. Parallel correlation clustering on big graphs. In *NIPS*, pages 82–90, 2015.
- J. Pouget-Abadie, K. Aydin, W. Schudy, K. Brodersen, and V. Mirrokni. Variance reduction in bipartite experiments through correlation clustering. In *NeurIPS*, pages 13288–13298, 2019.
- C. Rösner and M. Schmidt. Privacy preserving clustering with constraints. In *ICALP*, pages 96:1–96:14, 2018.
- A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24. Springer Science & Business Media, 2003.
- K. Yang and J. Stoyanovich. Measuring fairness in ranked outputs. In *SSDM*, pages 1–6, 2017.
- I. M. Ziko, E. Granger, J. Yuan, and I. B. Ayed. Clustering with fairness constraints: A flexible and scalable approach. *arXiv:1906.08207*, 2019.