

# Generalizing Fair Clustering to Multiple Groups: Algorithms and Applications

Diptarka Chakraborty<sup>1</sup>, Kushagra Chatterjee<sup>1</sup>, Debarati Das<sup>2</sup>, Tien-Long Nguyen<sup>2</sup>

<sup>1</sup> National University of Singapore

<sup>2</sup> Pennsylvania State University

diptarka@comp.nus.edu.sg, kushagra.chatterjee@u.nus.edu, {dxd5606,tfn5179}@psu.edu

## Abstract

Clustering is a fundamental task in machine learning and data analysis, but it frequently fails to provide fair representation for various marginalized communities defined by multiple protected attributes – a shortcoming often caused by biases in the training data. As a result, there is a growing need to enhance the fairness of clustering outcomes, ideally by making minimal modifications, possibly as a post-processing step after conventional clustering. A recent work initiated the study of *closest fair clustering*, though in a restricted scenario where data points belong to only two groups. In practice, however, data points are typically characterized by many groups, reflecting diverse protected attributes such as age, ethnicity, gender, etc.

In this work, we generalize the study of the *closest fair clustering* problem to settings with an arbitrary number (more than two) of groups. We begin by showing that the problem is NP-hard even when all groups are of equal size – a stark contrast with the two-group case, for which an exact algorithm exists. Next, we propose near-linear time approximation algorithms that efficiently handle arbitrary-sized multiple groups. Leveraging our closest fair clustering algorithms, we further achieve improved approximation guarantees for the *fair correlation clustering* problem, advancing the state-of-the-art results. Additionally, we are the first to provide approximation algorithms for the *fair consensus clustering* problem involving multiple (more than two) groups.

**Extended version** — <https://arxiv.org/pdf/2511.11539>

## Introduction

Clustering, the task of partitioning a set of data points into groups based on their mutual similarity or dissimilarity, stands as a fundamental problem in unsupervised learning and is ubiquitous in applications of machine learning and data analysis. Often, each data point carries certain protected attributes, which can be encoded by assigning a specific color to each point. While traditional clustering algorithms typically succeed in optimizing their target objectives, they frequently fail to ensure *fairness* in their results. This shortfall can introduce or perpetuate biases against marginalized groups defined by sensitive attributes such as gender

or race (Kay, Matuszek, and Munson 2015; Bolukbasi et al. 2016). These potential biases arise not necessarily from the algorithms themselves, but from historical marginalization inherent in the data used for training. Addressing and mitigating such biases to achieve fair outcomes has emerged as a central topic in the field, with considerable attention given in recent years to the development of algorithms that guarantee *demographic parity* (Dwork et al. 2012) and/or *equal opportunity* (Hardt, Price, and Srebro 2016).

In the context of clustering, (Chierichetti et al. 2017) initiated the study of *fair clustering* to address the issue of disparate impact and promote fair representation. Their work initially focused on datasets with two groups, each point colored either red or blue, and aimed to partition the data such that the blue-to-red ratio in every cluster matched that of the overall dataset. However, restricting attention to only two colors is limiting, as real-world data often involves multiple (and sometimes non-binary) protected attributes, such as age, race, or gender, resulting in several disjoint colored groups. Subsequent research generalized the fair clustering framework to accommodate more than two colors (Rösner and Schmidt 2018), requiring that the proportion of each colored group within clusters reflects the global proportions of colors. Further studies have considered scenarios where each color group is of equal size (Böhm et al. 2020). We refer to the related works for different variants of the clustering problems that have been studied under fairness constraints.

As previously highlighted, a range of effective clustering algorithms is available for various clustering paradigms; however, these methods may yield unfair or biased results when the training data itself is biased. Such skewed clustering outcomes may lead to inequitable treatment, particularly if the clusters serve as the basis for decision-making or analysis. To counteract this, post-processing existing clustering solutions to mitigate bias and achieve fairness is often necessary, ideally with only minimal adjustments to the cluster assignments. Despite the fundamental nature of this problem, it has received limited attention within the broader context of clustering in prior research. However, it has been studied for specific metric spaces, such as ranking (Celis, Straszak, and Vishnoi 2018; Chakraborty et al. 2022; Kliachkin et al. 2024). Only recently (Chakraborty et al. 2025a) did introduce the problem of obtaining a *closest fair clustering*, where given an existing clustering, finding a fair clustering

by altering as few assignments as possible. Their study, however, was confined to the special case of two colored groups, which, as previously noted, is restrictive in many practical scenarios. In their work, (Chakraborty et al. 2025a) presented  $O(1)$ -approximation algorithms for the case of blue and red groups with arbitrary ratios, and demonstrated that the problem can be solved exactly in near-linear time when the groups are of equal size. In the end, they posed the general case of more than two colored groups as an intriguing open direction. In this paper, we address this open problem by devising approximation algorithms and establishing computational hardness that holds even for equi-proportioned multiple-colored groups, thereby showing a stark distinction from the two-group case.

Building upon our findings for the closest fair clustering problem, we further investigate their implications for other prominent clustering variants, specifically *correlation clustering* and *consensus clustering*. In correlation clustering, one is presented with a labeled (typically complete) graph where each edge is marked as either  $+$  or  $-$ . The cost of a clustering is calculated as the total number of  $+$  edges that span across clusters and  $-$  edges that fall within clusters. This problem has widespread applications in fields such as data mining, social network analysis, computational biology, and marketing analysis (Bonchi, Garcia-Soriano, and Liberty 2014; Hou et al. 2016; Veldt, Gleich, and Wirth 2018; Bressan et al. 2019; Kushagra, Ben-David, and Ilyas 2019). The *fair correlation clustering* problem, introduced in (Ahmadian et al. 2020; Ahmadi et al. 2020), seeks a fair clustering that minimizes this cost. In their work, approximation algorithms were proposed for cases involving multiple colored groups, with the approximation factor depending on the ratios of the group sizes. In this paper, we improved upon their approximation bound and are the first to achieve an approximation guarantee independent of the group size ratio.

In consensus clustering, the objective is to derive a single, representative (consensus) clustering from a collection of clusterings over the same set of data points, to minimize a chosen objective function. The objective often depends on the application, with the most common being the *median* that minimizes the total distance to all input clusterings and the *center* that minimizes the maximum distance. Here, the distance between two clusterings is typically defined by the number of point pairs that are co-clustered (together) in one clustering but not in the other. Consensus clustering is widely applicable across various fields, including bioinformatics (Filkov and Skiena 2004b,a), data mining (Topchy, Jain, and Punch 2003), and community detection (Lancichinetti and Fortunato 2012). Recently, (Chakraborty et al. 2025a) extended the study of consensus clustering to incorporate fairness constraints, providing constant-factor approximations but only in the special case of two colored groups. In this work, we expand upon these results by establishing approximation guarantees for consensus clustering problems involving more than two colored groups.

**Our Contribution** In this work, we present both new algorithms and hardness results for the *Closest Fair Clustering* problem in settings where the input data points come

from multiple groups, and also study two well-known applications, namely, the *Fair Correlation Clustering* and *Fair Consensus Clustering* problems, and provide new algorithm results. Below, we summarize our main contributions.

**Closest Fair Clustering with Multiple Groups.** In this problem, we are given a clustering  $\mathcal{D} = \{D_1, D_2, \dots, D_m\}$  defined on a dataset  $V$  where  $V$  is classified into a set  $\chi$  of disjoint groups, each represented by a unique color. The objective is to compute a fair clustering of  $V$ , where the proportion of data points from each group (or color) within any output cluster should reflect their overall proportions in the entire dataset, while also minimizing the distance to the input clustering  $\mathcal{D}$ .

- Considering  $|\chi|$  to be the total number of colors, our first algorithm handles the case where each output cluster must contain the different colored points according to global ratio  $p_1 : p_2 : \dots : p_{|\chi|}$ , and achieves an  $O(|\chi|^{3.81})$ -approximation.

This study significantly generalizes the work of (Chakraborty et al. 2025a), which focused solely on the binary (two-group) setting.

- Next, we consider the special case where each output cluster must contain an equal number of data points from each color group, and we provide an  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$ -approximation for the Closest Fair Clustering problem. Furthermore, when  $|\chi|$  is a power of two, we present an improved algorithm with an  $O(|\chi|^{1.6})$ -approximation.

- Finally, we show that the problem is NP-hard for any setting involving more than two colors, even when all color groups are equally represented in an output cluster.

This shows a clear hardness gap between the two-color setting, where an exact algorithm exists, and the multi-color case, where the problem becomes computationally intractable; thus underscoring the necessity of our approximation algorithms.

**Application to Fair Correlation Clustering.** Building on the above result, we study their implications for another key clustering variant: the *Fair Correlation Clustering* problem. In correlation clustering, the input is a graph with edges labeled as  $+$  or  $-$ , and the goal is to produce a clustering minimizing disagreements;  $+$  edges between clusters and  $-$  edges within clusters. In the fair version, the clustering must also satisfy group fairness constraints. This problem has been shown to be NP-hard (Ahmadi et al. 2020).

- We begin by designing an algorithm for the setting where each output cluster must contain points from different color groups as per the global ratio  $p_1 : p_2 : \dots : p_{|\chi|}$ , achieving an  $O(|\chi|^{3.81})$ -approximation.

This eliminates the dependence on the max-min color ratio  $q = \frac{\max(p_j)}{\min(p_j)}$  that appeared in the previous  $O(q^2 |\chi|^2)$  bound of (Ahmadian et al. 2020; Ahmadi et al. 2020), where  $q$  can be as large as a polynomial in  $|V|$ .

- For the special case where each output cluster must contain an equal number of data points from each color group, we improve the approximation to

$O(|\chi|^{1.6} \log^{2.81} |\chi|)$ , and further to  $O(|\chi|^{1.6})$  when  $|\chi|$  is a power of two.

This improves upon the previous  $O(|\chi|^2)$  bound given by (Ahmadian et al. 2020; Ahmadi et al. 2020).

**Application to Fair Consensus Clustering.** Next, we turn to another application: the *consensus clustering* problem. The goal here is to compute a single representative (consensus) clustering from a collection of input clusterings over the same dataset, minimizing a specified objective function (e.g., median or center) while also satisfying group fairness constraints. By combining a triangle inequality argument with our results for Closest Fair Clustering, we obtain new approximation guarantees for this problem. These results generalize the work of (Chakraborty et al. 2025a), which was limited to the binary (two-group) setting, to the more general multi-group case.

- Analogous to our results for fair correlation clustering, we obtain an  $O(|\chi|^{3.81})$ -approximation algorithm for the general case with arbitrary group proportions.
- For the equi-proportion case, we again improve the approximation to  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$ , and further to  $O(|\chi|^{1.6})$  when  $|\chi|$  is a power of two.

The details are provided in the extended version.

**Related Works** Since the introduction of the fair clustering in (Chierichetti et al. 2017), recent years have witnessed a significant increase in research focused on different aspects of fair clustering problems. The literature so far encompasses numerous variants of the fair clustering problem with multiple colored groups, such as  $k$ -center/median/means clustering (Chierichetti et al. 2017; Huang, Jiang, and Vishnoi 2019), scalable clustering (Backurs et al. 2019), proportional clustering (Chen et al. 2019), fair representational clustering (Bera et al. 2019; Bercea et al. 2019), pairwise fair clustering (Bandyapadhyay, Fomin, and Simonov 2024; Bandyapadhyay et al. 2024, 2025), correlation clustering (Ahmadian et al. 2020; Ahmadi et al. 2020; Ahmadian and Negahbani 2023), 1-clustering over rankings (Wei et al. 2022; Chakraborty et al. 2022, 2025b), and consensus clustering (Chakraborty et al. 2025a), among others.

A comprehensive study of the correlation clustering problem was first undertaken in (Bansal, Blum, and Chawla 2004). Since then, correlation clustering has been studied across various graph settings, including the extensively examined complete graphs (Ailon, Charikar, and Newman 2008; Chawla et al. 2015) and weighted graphs (Demaine et al. 2006). The problem, even when restricted to complete graphs, is known to be APX-hard (Charikar, Guruswami, and Wirth 2005), with the best-known approximation algorithm currently achieving a 1.438-approximation factor (Cao et al. 2024). Its fair variant remains NP-hard even in the case of two color groups of equal sizes (Ahmadi et al. 2020), and several approximation algorithms have been developed for both the exact fairness notion (Ahmadian et al. 2020; Ahmadi et al. 2020) and the relaxed fairness notion (Ahmadian and Negahbani 2023).

The consensus clustering problem, under both the median and center objectives, is known to be NP-hard (Křivánek

and Morávek 1986; Swamy 2004) and in fact, APX-hard (that is it is unlikely to have an  $(1 + \epsilon)$ -factor algorithm for any  $\epsilon > 0$ ) even with as few as three input clusterings (Bonizzoni et al. 2008). Currently, the best-known algorithms include an 11/7-approximation for the median objective (Ailon, Charikar, and Newman 2008) and an approximation slightly better than 2 for the center objective (Das and Kumar 2025). Apart from that, various heuristics have been proposed to produce reasonable solutions (e.g., (Goder and Filkov 2008; Monti et al. 2003; Wu et al. 2014)). More recently, (Chakraborty et al. 2025a) began examining fair consensus clustering, focusing on only two colored groups.

## Preliminaries

In this section, we define key terms and concepts that are essential for understanding our proofs and algorithms.

**Definition 1 (Fair Clustering).** *Given a set of points  $V$  and a set of colors  $\chi = \{c_1, c_2, \dots, c_k\}$ , suppose  $c_j(V) \subseteq V$  be the set of points of color  $c_j$  in  $V$ . We call a clustering  $\mathcal{F}$  of  $V$  a Fair Clustering if for all clusters  $F \in \mathcal{F}$  we have*

$$|c_1(F)| : \dots : |c_k(F)| = |c_1(V)| : \dots : |c_k(V)|.$$

For two clustering  $\mathcal{C}$  and  $\mathcal{C}'$  of  $V$  we define  $\text{dist}(\mathcal{C}, \mathcal{C}')$  as the distance between two clustering  $\mathcal{C}$  and  $\mathcal{C}'$ . The distance is measured by the number of pairs  $(u, v)$  that are together in  $\mathcal{C}$  but separated by  $\mathcal{C}'$  and the number of pairs  $(u, v)$  that are separated by  $\mathcal{C}$  but together in  $\mathcal{C}'$ . More specifically,

$$\text{dist}(\mathcal{C}, \mathcal{C}') = \left| \left\{ \{u, v\} \mid u, v \in V, [u \sim_{\mathcal{C}} v \wedge u \not\sim_{\mathcal{C}'} v] \vee [u \not\sim_{\mathcal{C}} v \wedge u \sim_{\mathcal{C}'} v] \right\} \right|$$

where  $u \sim_{\mathcal{C}} v$  denotes whether both  $u$  and  $v$  belong to the same cluster in  $\mathcal{C}$  or not.

**Definition 2 (Closest Fair Clustering).** *Given an arbitrary clustering  $\mathcal{D}$ , a clustering  $\mathcal{F}_{\mathcal{D}}^*$  is called a closest Fair Clustering to  $\mathcal{D}$  if for all Fair Clustering  $\mathcal{F}$  we have  $\text{dist}(\mathcal{D}, \mathcal{F}) \geq \text{dist}(\mathcal{D}, \mathcal{F}_{\mathcal{D}}^*)$ .*

We use  $\mathcal{F}_{\mathcal{D}}^*$  to denote a closest Fair Clustering to  $\mathcal{D}$ .

**$\gamma$ -close Fair Clustering** We call a Fair Clustering  $\mathcal{F}$  a  $\gamma$ -close Fair Clustering to a clustering  $\mathcal{D}$  if

$$\text{dist}(\mathcal{D}, \mathcal{F}) \leq \gamma \text{dist}(\mathcal{D}, \mathcal{F}_{\mathcal{D}}^*).$$

## Approximate Closest Fair Clustering for Equi-Proportion Groups

In this section, we provide an approximation algorithm to find a closest Fair Clustering when all the groups are of equal size.

**Theorem 1.** *There exists an algorithm that, given a clustering  $\mathcal{D}$  where each color group contains an equal number of points, computes an  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$ -close Fair Clustering in  $O(|V| \log |V|)$  time, where  $\chi$  denotes the set of colors. Moreover, when  $|\chi|$  is a power of two, the algorithm computes a  $O(|\chi|^{1.6})$ -close Fair Clustering.*

First, let us handle the case when  $|\chi|$  is a power of 2. To do that, we provide an algorithm `fairpower-of-two`, which produces an  $O(|\chi|^{1.6})$ -close fair clustering  $\mathcal{F}_{\text{fpt}}$  to a clustering  $\mathcal{D}$  when  $|\chi|$  is a power of 2.

### Algorithm fairpower-of-two

Let the input be a clustering  $\mathcal{D} = \{D_1, D_2, \dots, D_m\}$ , where each point is colored from a color set  $\chi = \{c_1, \dots, c_k\}$ , and assume  $|\chi|$  is a power of two. The goal is to output a clustering  $\mathcal{F}_{\text{fpt}}$  in which every cluster contains an equal number of points of each color, i.e.,  $c_p(F_a) = c_q(F_a)$  for all  $p \neq q$ ,  $F_a \in \mathcal{F}_{\text{fpt}}$ .

The algorithm proceeds in  $\log |\chi|$  iterations. At iteration  $i$ , the color set  $\chi$  is partitioned into  $|\chi|/2^i$  disjoint blocks of size  $2^i$ , defined as:

$$B_j^i = \{c_{(j-1) \cdot 2^{i+1}}, \dots, c_{j \cdot 2^i}\}, \quad \text{for } j = 1, \dots, |\chi|/2^i.$$

Let  $\mathcal{N}^i$  be the clustering at iteration  $i$ , with  $\mathcal{N}^0 := \mathcal{D}$ . The algorithm maintains the invariant that, in every cluster  $N_a^i \in \mathcal{N}^i$ , the colors within each block  $B_j^i$  appear equally.

For adjacent blocks  $B_j^i$  and  $B_{j+1}^i$  in a cluster  $N_a^i$ , the surplus  $T_a^j$  is the excess of the larger bucket over the smaller. The surplus is chosen so that all colors in the surplus are equally represented. Let us now describe the algorithm fairpower-of-two.

1. Initialization: Set  $\mathcal{N}^0 \leftarrow \mathcal{D}$ .
2. Iterative Refinement: For each iteration  $i = 1$  to  $\log |\chi|$ :
  - Initialize  $\mathcal{N}^i \leftarrow \mathcal{N}^{i-1}$ .
  - For each pair of disjoint adjacent blocks  $(B_j^i, B_{j+1}^i)$ :
    - For each cluster  $N_a^i \in \mathcal{N}^i$  and odd  $j$ , compute the surplus  $T_a^j$  between adjacent blocks  $B_j^i(N_a^i)$  and  $B_{j+1}^i(N_a^i)$ , and remove it. Here for a color block  $B_j^i$ ,  $B_j^i(N_a^i)$  is the set of points in  $N_a^i$  that has a color from the block  $B_j^i$ .
    - Store removed surpluses into sets  $S_j$  or  $S_{j+1}$ , depending on which block had the surplus.
    - Call `multi-GM`( $S_j, S_{j+1}$ ) to form new fair clusters and add them to  $\mathcal{N}^i$ .
3. Output: Return  $\mathcal{N}^{\log |\chi|} \rightarrow \mathcal{F}_{\text{fpt}}$ .

Let us now describe the subroutine `multi-gm`. Given two collections of point groups from blocks  $B_j^i$  and  $B_{j+1}^i$ , the `multi-GM` procedure greedily merges pairs of subsets into fair subsets in which each color from  $B_j^i \cup B_{j+1}^i$  is equally represented.

The subroutine:

- Iteratively selects one subset from each collection.
- Trims the larger subset to match the smaller, preserving equal color counts.
- Merges the trimmed subsets into a fair set and adds it to the output.

Continue this until no further fair subsets can be formed.

We provide the pseudocode of the algorithms fairpower-of-two and multi-gm in the extended version.

**Proof of theorem 1** We now analyze the algorithm fairpower-of-two by first establishing the following.

**Lemma 1.** *Given a clustering  $\mathcal{D}$  as input, the algorithm fairpower-of-two computes a  $O(|\chi|^{1.6})$ -close Fair Clustering, where  $|\chi|$  is a power of 2.*

We show the above result by generating a sequence of  $\log |\chi|$  intermediate clusterings, with each intermediate step involving an approximation factor of 2, and thus finally achieving a factor of  $3^{\log |\chi|} = |\chi|^{1.6}$ . We provide the proof in the extended version.

The above lemma 1 proves theorem 1 when  $|\chi|$  is a power of 2. Now, we prove theorem 1 for any values of  $|\chi|$ . To do that, we need to describe the algorithm `make-pdc-fair`.

### Algorithm make-pdc-fair

Given a clustering  $\mathcal{I} = \{I_1, I_2, \dots, I_s\}$  over a point set  $V$ , where each point is colored from  $\zeta = \{z_1, \dots, z_r\}$  and satisfies the global proportion:

$$z_1(V) : z_2(V) : \dots : z_r(V) = p_1 : p_2 : \dots : p_r,$$

the goal is to construct a fair clustering  $\mathcal{F}_{\text{mpf}}$  such that every cluster  $F \in \mathcal{F}_{\text{mpf}}$  satisfies this ratio. We assume w.l.o.g. that  $p_1 > p_2 > \dots > p_r$ . Also assume that each input cluster is  $p$ -divisible, i.e.,  $z_j(I_i)$  is a multiple of  $p_j$ . We call such clustering  $p$ -divisible clustering (pdc for short). Let us now describe the algorithm `make-pdc-fair`.

1. The algorithm proceeds in  $T = \lceil \log_2 r \rceil$  iterations. Let  $\mathcal{F}^0 := \mathcal{I}$ , and define:

$$\mathcal{I} = \mathcal{F}^0 \rightarrow \mathcal{F}^1 \rightarrow \dots \rightarrow \mathcal{F}^T = \mathcal{F}_{\text{mpf}},$$

where each  $\mathcal{F}^t$  enforces proportionality within blocks of colors.

2. Hierarchical Block Structure: At iteration  $t$ , the color set is partitioned into blocks  $\{B_1^t, B_2^t, \dots, B_{m_t}^t\}$ , constructed hierarchically:

$$B_i^t = B_{2i-1}^{t-1} \cup B_{2i}^{t-1}, \text{ with singleton blocks } B_j^0 = \{z_j\}.$$

If  $m_{t-1}$  is odd, the last block is carried forward unchanged.

3. Balancing Rule: To merge two sub-blocks  $A = B_{2i-1}^{t-1}$  and  $B = B_{2i}^{t-1}$ , consider a cluster  $F^{t-1} \in \mathcal{F}^{t-1}$ , where

$$\begin{aligned} z_c(F^{t-1}) &= p_c \cdot x \text{ for } z_c \in A, \\ z_d(F^{t-1}) &= p_d \cdot y \text{ for } z_d \in B. \end{aligned}$$

To equalize the scaling factors  $x$  and  $y$ , we do:

- If  $x > y$ : merge  $p_d \cdot (x - y)$  points of color  $z_d \in B$  into  $F^{t-1}$ .
- If  $x < y$ : cut  $p_d \cdot (y - x)$  points of color  $z_d \in B$  from  $F^{t-1}$ .

This ensures that the merged block  $B_i^t = A \cup B$  in each cluster  $F^t \in \mathcal{F}^t$  satisfies the combined proportionality.

4. Output: After  $T = \lceil \log_2 r \rceil$  iterations, the final clustering  $\mathcal{F}_{\text{mpf}}$  satisfies: For all  $F \in \mathcal{F}_{\text{mpf}}$ ,

$$z_1(F) : z_2(F) : \dots : z_r(F) = p_1 : p_2 : \dots : p_r.$$

We provide the pseudocode of `make-pdc-fair` in the extended version.

To analyze the algorithm `make-pdc-fair`, we need to prove the following lemma.

**Lemma 2.** *The algorithm `make-pdc-fair` outputs a clustering  $\mathcal{F}$  that is  $O(r^{2.81})$ -close Fair Clustering to the input clustering  $\mathcal{I}$ , where  $r$  is the number of colors.*

We show the above result by generating a sequence of  $\log r$  intermediate clusterings, with each intermediate step involving an approximation factor of 6, and thus finally achieving a factor of  $7^{\log r} = r^{2.81}$ . We provide a detailed proof in the extended version.

Using lemma 1 and lemma 2, we can prove theorem 1. We provide the full proof of theorem 1 in the extended version.

*Proof Sketch of theorem 1.* We design an algorithm `fair-equi` to convert an arbitrary clustering  $\mathcal{D}$ , where the color set  $\chi$  is not necessarily a power of two, into a fair clustering  $\mathcal{F}$  in which every color appears equally in each cluster. The algorithm proceeds in two main stages:

1. Color Grouping and Intermediate Fairness:
  - Partition the color set  $\chi$  into  $\log |\chi|$  disjoint groups  $G_1, \dots, G_{\log |\chi|}$ , where each group's size is a power of two. This is done greedily by assigning group sizes according to the binary representation of  $|\chi|$ .
  - Apply the algorithm `fairpower-of-two` to obtain an intermediate clustering  $\mathcal{I}$ , in which each group  $G_\ell$  satisfies intra-group fairness: all colors in  $G_\ell$  appear equally in each cluster.
2. Global Fairness via Multi-Group Merging:
  - Treat each group  $G_\ell$  as a single meta-color and apply the algorithm `make-pdc-fair` on  $\mathcal{I}$  to obtain the final clustering  $\mathcal{F}$ .
  - `make-pdc-fair` ensures that across all clusters, the meta-colors (i.e., groups) are in proportion to their sizes, and also restores uniformity within each group, thus achieving full color-wise fairness.

Approximation Bound:

- By lemma 1, the intermediate clustering  $\mathcal{I}$  is  $O(|\chi|^{1.6})$ -close to  $\mathcal{D}$ .
- By lemma 2, the final clustering  $\mathcal{F}$  is  $O(\log^{2.81} |\chi|)$ -close to  $\mathcal{I}$ .
- Combining via triangle inequality yields:

$$\text{dist}(\mathcal{D}, \mathcal{F}) \leq O(|\chi|^{1.6} \log^{2.81} |\chi|) \cdot \text{dist}(\mathcal{D}, \mathcal{F}_{\mathcal{D}}^*)$$

where  $\mathcal{F}_{\mathcal{D}}^*$  is the closest fair clustering to  $\mathcal{D}$ .

This completes the proof sketch. □

### Approximate Closest Fair Clustering for Arbitrary-Proportion Groups

In this section, we prove the following theorem.

**Theorem 2.** *There is an algorithm that, given an arbitrary clustering  $\mathcal{D}$  over a vertex set  $V$  where each vertex  $v \in V$  has a color in  $\chi = \{c_1, \dots, c_k\}$ , finds a  $O(|\chi|^{3.81})$ -close Fair Clustering  $\mathcal{F}$  in time  $O(|V| \log |V|)$ .*

To prove the above theorem, we take a 2-step approach similar to (Chakraborty et al. 2025a), which was constrained to only two colors.

- (i) First, we convert the clustering  $\mathcal{D}$  to a clustering  $\mathcal{M}$  such that for each cluster  $M_i \in \mathcal{M}$ ,  $c_j(M_i)$  is divisible by  $p_j$ . Recall we call such a clustering a  $p$ -divisible clustering.
- (ii) In the second step, we will provide the clustering  $\mathcal{M}$  as input to the algorithm `make-pdc-fair` and get a fair clustering  $\mathcal{F}$  as output.

Now, we provide an algorithm `create-pdc` to convert a clustering  $\mathcal{D}$  on a vertex set  $V$  to a  $p$ -divisible clustering  $\mathcal{M}$ .

#### Algorithm `create-pdc`

Given a clustering  $\mathcal{D} = \{D_1, \dots, D_m\}$  over vertex set  $V$ , color set  $\chi = \{c_1, \dots, c_k\}$ , and a proportion vector  $\mathbf{p} = (p_1, \dots, p_k)$  satisfying:

$$c_1(V) : c_2(V) : \dots : c_k(V) = p_1 : p_2 : \dots : p_k.$$

The goal of `create-pdc` is to convert  $\mathcal{D}$  into a  $p$ -divisible clustering  $\mathcal{M}$  where each cluster contains a multiple of  $p_j$  points of color  $c_j$ .

Key Definitions:

- Surplus:  $\sigma(D_i, c_j) \subseteq D_i$  of size  $c_j(D_i) \bmod p_j$  if  $p_j \nmid c_j(D_i)$ , else it has size  $p_j$ , denoting excess  $c_j$ -colored vertices in  $D_i$ .
- Total surplus:  $\sigma_j = \sum_{D_i \in \mathcal{D}} |\sigma(D_i, c_j)|$  (always a multiple of  $p_j$ ).
- Deficit:  $\delta(D_i, c_j) \subseteq V \setminus D_i$ , of size  $p_j - |\sigma(D_i, c_j)|$ . Represents the number of  $c_j$  colored points required to make it a multiple of  $p_j$ .
- Cut and merge costs:

$$\kappa^j(D_i) = |\sigma(D_i, c_j)| \cdot (|D_i| - |\sigma(D_i, c_j)|)$$

$$\mu^j(D_i) = |\delta(D_i, c_j)| \cdot |D_i|$$

Let us now describe the algorithm `create-pdc`

1. For each color  $c_j$ , initialize  $\sigma_j/p_j$  empty auxiliary clusters  $\{P_1, \dots, P_{\sigma_j/p_j}\}$ .
2. Classify each cluster  $D_i \in \mathcal{D}$  into:
  - CUT, if  $|\sigma(D_i, c_j)| \leq p_j/2$
  - MERGE, otherwise.
3. Cut and redistribute: While CUT is non-empty:
  - For  $D_i \in \text{CUT}$ , remove  $\sigma(D_i, c_j)$  from  $D_i$ .
  - Try to donate surplus to deficits in  $D_\ell \in \text{MERGE}$ .
  - If no deficit remains, assign surplus to an available extra cluster  $P_m$ , ensuring each  $P_m$  reaches size  $p_j$ .
4. Handle remaining merges: While deficits remain:
  - Pick the cluster with minimal  $\kappa^j(D_k) - \mu^j(D_k)$ , redistribute its surplus as above. Note that in this step, we can pick a cluster multiple times.
  - Remove a cluster  $D_\ell \in \text{MERGE}$  if its deficit is satisfied.
  - Repeat until all deficits are filled.
5. Output: The final clustering  $\mathcal{M}$  where each cluster is  $p$ -divisible for all colors.

We provide the pseudocode of the algorithm `create-pdc` in the extended version.

Now, we analyze the algorithm `create-pdc`. To analyze and state the lemma, let us define some terms

- **Optimal- $p$ -divisible clustering:** Given a clustering  $\mathcal{D}$ , we call a clustering  $\mathcal{M}^*$  an optimal- $p$ -divisible clustering if  $\text{dist}(\mathcal{D}, \mathcal{M}^*) \leq \text{dist}(\mathcal{D}, \mathcal{M})$ .
- **$\alpha$ -close- $p$ -divisible clustering:** Given a clustering  $\mathcal{D}$ , we call a clustering  $\mathcal{M}$  an  $\alpha$ -close- $p$ -divisible clustering if  $\text{dist}(\mathcal{D}, \mathcal{M}) \leq \alpha \text{dist}(\mathcal{D}, \mathcal{M}^*)$ .

Now, we state the lemma for analyzing the algorithm `create-pdc`.

**Lemma 3.** *Given a clustering  $\mathcal{D}$ , the algorithm `create-pdc` outputs a clustering  $\mathcal{M}$  s.t.  $\mathcal{M}$  is  $O(|\chi|)$ -close- $p$ -divisible clustering to  $\mathcal{D}$ .*

We provide the proof of lemma 3 in the extended version. Next, we prove theorem 2 assuming lemma 3.

*Proof of theorem 2.* To prove the theorem, we describe an algorithm `fair-general` that proceeds in two stages. First, we apply the algorithm `create-pdc` to the input clustering  $\mathcal{D}$  to obtain a  $p$ -divisible clustering  $\mathcal{M}$  that is  $O(|\chi|)$ -close to  $\mathcal{D}$ . Then, we apply the algorithm `make-pdc-fair` on  $\mathcal{M}$  to obtain the final fair clustering  $\mathcal{F}$ , which is  $O(|\chi|^{2.81})$ -close to  $\mathcal{M}$  (see lemma 2). By combining the guarantees from both steps, we conclude that  $\mathcal{F}$  is  $O(|\chi|^{3.81})$ -close *Fair Clustering* to the input  $\mathcal{D}$ .

$$\begin{aligned} \text{dist}(\mathcal{D}, \mathcal{F}) &\leq \text{dist}(\mathcal{D}, \mathcal{M}) + \text{dist}(\mathcal{M}, \mathcal{F}) \\ &\text{(by triangle inequality)} \\ &\leq O(|\chi|) \text{dist}(\mathcal{D}, \mathcal{F}^*) \\ &\quad + O(|\chi|^{2.81}) \text{dist}(\mathcal{M}, \mathcal{F}^*) \\ &\leq O(|\chi|) \text{dist}(\mathcal{D}, \mathcal{F}^*) \\ &\quad + O(|\chi|^{2.81})(\text{dist}(\mathcal{D}, \mathcal{M}) + \text{dist}(\mathcal{D}, \mathcal{F}^*)) \\ &\leq O(|\chi| + |\chi|^{3.81} + |\chi|^{2.81}) \text{dist}(\mathcal{D}, \mathcal{F}^*) \\ &\leq O(|\chi|^{3.81}) \text{dist}(\mathcal{D}, \mathcal{F}^*) \end{aligned}$$

This completes the proof of theorem 2.  $\square$

## Hardness for Three Equi-Proportion Groups

In this section, we show that finding a closest fair clustering to a given clustering where each point is assigned one color from a set of  $k \geq 3$  colors is hard even when the number of points in each color class is equal. Our reduction also extends to arbitrary color ratios.

We begin by defining the decision version of closest fair clustering with multiple colors and equal representation.

**Definition 3** ( $k$ -CLOSEST EQUIFAIR). *Given a clustering  $\mathcal{H}$  over a set of points  $V$  where each point is assigned with one of the colors from a set of  $k \geq 3$  colors, and the number of points of each color is equal, together with a non-negative integer  $\tau$ , decide between the following:*

- **YES:** *There exists a Fair Clustering (on input point set)  $\mathcal{F}$  such that  $\text{dist}(\mathcal{H}, \mathcal{F}) \leq \tau$ ;*

- **NO:** *For every Fair Clustering (on input point set)  $\mathcal{F}$ ,  $\text{dist}(\mathcal{H}, \mathcal{F}) > \tau$ .*

We show the following theorem.

**Theorem 3.** *For any integer  $k \geq 3$ ,  $k$ -CLOSEST EQUIFAIR is NP-hard.*

We present a polynomial-time reduction from the 3-PARTITION problem (defined below) to  $k$ -CLOSEST EQUIFAIR.

**Definition 4** (3-PARTITION). *Given a (multi)set of positive integers  $S = \{x_1, \dots, x_d\}$ , decide whether (YES:) there exists a partition of  $S$  into  $m$  disjoint subsets  $S_1, S_2, \dots, S_f \subseteq S$  where  $f = d/3$ , such that*

- *For all  $i$ ,  $|S_i| = 3$ ; and*
- *For all  $i$ ,  $\sum_{x_j \in S_i} x_j = T$ , where  $T = \frac{\sum_{x_j \in S} x_j}{n/3}$ ,*

*or (NO:) no such partition exists.*

We reduce from a restricted version of 3-PARTITION, in which each  $x_i \in S$  satisfies  $x_i \in (T/4, T/2)$ , where  $T = \frac{3}{n} \sum_{x_i \in S} x_i$ , and additionally  $x_i \leq d^b$ , for some non-negative constant  $b$ . Note that, this variant remains NP-complete, as established by (Garey and Johnson 1975), which shows that 3-PARTITION as defined in definition 4 is *strongly NP-complete*. Hence, for the rest of this section, we refer to this restricted version simply as 3-PARTITION.

**Description of Reduction.** Given a 3-PARTITION instance  $S = \{x_1, x_2, \dots, x_d\}$  we create a  $k$ -CLOSEST EQUIFAIR instance  $(\mathcal{H}, \tau)$  as follows:

- $\mathcal{H} = \{\text{GB}_1, \text{GB}_2, \dots, \text{GB}_{d/3}, \text{R}_1, \text{R}_2, \dots, \text{R}_d\}$ , where for each  $i \in \{1, \dots, d/3\}$ ,  $\text{GB}_i$  is a cluster of size  $(k-1)T$  with  $T$  points of color  $c_t$  ( $2 \leq t \leq k$ ), and for each  $j \in \{1, \dots, n\}$ ,  $\text{R}_j$  is a monochromatic  $c_1$  cluster (i.e., containing only points with color  $c_1$ ) of size  $x_j$  (i.e.,  $|\text{R}_j| = x_j$ );
- Set  $\tau = \frac{n}{3}(k-1)T^2 + \frac{1}{2} \sum_{i=1}^n x_i(T-x_i)$ , if  $k > 3$ , and,  $\tau = 2 \sum_{i=1}^n x_i^2 + 2 \sum_{i=1}^n x_i(T-x_i)$ , if  $k = 3$ .

Note that for each  $x_i \leq d^b$ , for some non-negative constant  $b$ , the size of the instance  $(\mathcal{H}, \tau)$  is polynomial in  $d$ . Moreover, it is straightforward to see that the reduction runs in polynomial time.

The following lemma argues that the above reduction maps a YES instance of the 3-PARTITION to a YES instance of the  $k$ -CLOSEST EQUIFAIR, the proof of which is deferred to the extended version.

**Lemma 4.** *For any integer  $k \geq 3$ , if  $S$  is a YES instance of the 3-PARTITION, then  $(\mathcal{H}, \tau)$  is also a YES instance of the  $k$ -CLOSEST EQUIFAIR.*

It remains to demonstrate that our reduction maps a NO instance of the 3-PARTITION to a NO instance of the  $k$ -CLOSEST EQUIFAIR.

**Lemma 5.** *For  $k \geq 3$ , if  $S$  is a NO instance of the 3-PARTITION, then  $(\mathcal{H}, \tau)$  is also a NO instance of the  $k$ -CLOSEST EQUIFAIR.*

*Proof.* Assume to the contrary that  $(\mathcal{H}, \tau)$  is a YES instance. Then there exists a *Fair Clustering*  $\mathcal{F}$  such that  $\text{dist}(\mathcal{H}, \mathcal{F}) \leq \tau$ .

Without loss of generality, we refer to  $c_1$  as the color red, and refer to  $c_2$  as the color blue.

Let  $V'$  be the set of red-blue points obtained from  $V$  by recoloring every point with color  $c_j$  ( $j \geq 3$ ) to blue. Denote  $\mathcal{H}^c$  and  $\mathcal{F}^c$  the clusterings obtained from  $\mathcal{H}$  and  $\mathcal{F}$ , respectively, under this recoloring. Then,  $\text{dist}(\mathcal{H}^c, \mathcal{F}^c) = \text{dist}(\mathcal{H}, \mathcal{F})$ . We analyze the structure of  $\mathcal{H}^c$  and  $\mathcal{F}^c$ .

Observe that  $\mathcal{H}^c$  is a clustering over  $V'$ , in which  $\mathcal{H}^c = \{B_1, B_2, \dots, B_{n/3}, R_1, R_2, \dots, R_n\}$ , where each  $B_i$  is a monochromatic blue cluster of size  $(k-1)T$ , obtained from the original cluster  $GB_i$  in  $\mathcal{H}$  by recoloring every point to blue. Each  $R_i$  is a monochromatic red cluster, which remains unchanged from  $\mathcal{H}$ .

Now we claim that  $\mathcal{F}^c$  is a *Fair Clustering* clustering over  $V'$ . Indeed, recall that  $\mathcal{F}$  is a *Fair Clustering* over  $V$ . Hence, in every cluster  $F \in \mathcal{F}$ , the numbers of points of each color are equal, that is,  $|c_i(F)| = |c_j(F)|$ , for all  $1 \leq i, j \leq k$ . Note that each cluster  $F^c \in \mathcal{F}^c$  is obtained from a cluster in  $\mathcal{F}$  by recoloring every point with color  $c_i$  ( $i \geq 3$ ) to blue. Hence, the ratio between the number of blue points and the number of red points in  $F^c$  is  $(k-1)$ . This implies that  $\mathcal{F}^c$  is a *Fair Clustering* over  $V'$ .

Applying a result from (Chakraborty et al. 2025a, Lemma 45, Lemma 46) for the set of points  $V'$ , it follows that  $\text{dist}(\mathcal{H}^c, \mathcal{F}^c) > \tau$ , which is a contradiction since we have established that  $\text{dist}(\mathcal{H}^c, \mathcal{F}^c) = \text{dist}(\mathcal{H}, \mathcal{F}) \leq \tau$ . This concludes that  $(\mathcal{H}, \tau)$  is a NO instance.  $\square$

As our reduction runs in polynomial time, from lemma 4 and lemma 5, we conclude that  $k$ -CLOSEST EQUIFAIR is NP-hard. This completes the proof of theorem 3. In the extended version, we remark on how to generalize this reduction for arbitrary ratios.

## Implication to Fair Correlation Clustering

Given a complete undirected graph  $G(V, E)$  where each edge  $(u, v) \in E$  is labeled either “+” (similar) or “-” (dissimilar), let  $E^+$  and  $E^-$  denote the sets of “+” and “-” edges, respectively. A clustering  $\mathcal{C} = \{C_1, C_2, \dots, C_t\}$  partitions  $V$  into disjoint subsets.

Let  $\text{EXT}(\mathcal{C})$  denote the set of inter-cluster edges and  $\text{INT}(\mathcal{C})$  the set of intra-cluster edges. The cost of clustering  $\mathcal{C}$  is defined as:

$$\text{cost}(\mathcal{C}) = |\text{EXT}(\mathcal{C}) \cap E^+| + |\text{INT}(\mathcal{C}) \cap E^-|.$$

The goal is to find a clustering that minimizes  $\text{cost}(\mathcal{C})$ . In addition, when we want  $\mathcal{C}$  to be also a *Fair Clustering*, the problem is referred to as *fair correlation clustering*.

**Theorem 4.** *There exists an algorithm that, given a correlation clustering instance  $G$ , computes a  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$  approximate fair correlation clustering when there are equal number of data points from each color group, and a  $O(|\chi|^{3.81})$  approximate fair correlation clustering when the ratio between the number of points from different color groups is arbitrary.*

To prove theorem 4, we use the following.

**Lemma 6.** *Let  $\mathcal{C}$  be a clustering, and suppose there exists an algorithm  $\mathcal{A}$  that computes a  $\gamma$ -close fair clustering with respect to  $\mathcal{C}$ . Additionally, suppose there exists a  $\beta$ -approximation algorithm  $\mathcal{B}$  for the standard correlation clustering problem on a graph  $G$ . Then, there exists an algorithm that computes a fair correlation clustering of  $G$  with approximation factor  $(\gamma + \beta + \gamma\beta)$ .*

*Proof.* Let us first describe the algorithm `fairifyCC`, which computes a fair correlation clustering for a given instance  $G$ .

- Input: Correlation clustering instance  $G$ .
- Output: A fair clustering  $\mathcal{F}$ .
- Procedure:
  1. Compute a correlation clustering  $\mathcal{D}$  of  $G$  using a  $\beta$ -approximation algorithm  $\mathcal{B}$ .
  2. Apply the closest fair clustering algorithm  $\mathcal{A}$  to  $\mathcal{D}$  to obtain a fair clustering  $\mathcal{F}$  that is  $\gamma$ -close to  $\mathcal{D}$ .
  3. Return  $\mathcal{F}$ .

We argue that  $\mathcal{F}$  is  $(\gamma + \beta + \gamma\beta)$  approximate correlation clustering of  $G$  using triangle inequality. We defer the proof to the extended version.  $\square$

*Proof of theorem 4.* By (Cao et al. 2025), we get that there exists an  $O(1)$ -approximation algorithm to find a correlation clustering for a correlation clustering instance  $G$ . When each color group has the same size, the algorithm `fair-equi` produces an  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$ -close clustering to any input clustering  $\mathcal{D}$ . Hence, by lemma 6 we get that the algorithm `fairifyCC` produces an  $O(|\chi|^{1.6} \log^{2.81} |\chi|)$  approximate fair correlation clustering.

In the general case with arbitrary group size ratio  $p_1 : p_2 : \dots : p_{|\chi|}$ , the algorithm `fair-general` gives  $O(|\chi|^{3.81})$ -close clustering to any input clustering  $\mathcal{D}$ . Hence, again by lemma 6 we show algorithm `fairifyCC` produces an  $O(|\chi|^{3.81})$  approximate fair correlation clustering.  $\square$

## Conclusion

In this paper, we generalize the closest fair clustering problem originally proposed by (Chakraborty et al. 2025a) to scenarios involving any number of groups, thereby addressing settings with non-binary, multiple protected attributes. We demonstrate that the problem becomes NP-hard even when there are just three equal-sized groups, showing a strong separation with the two equi-proportion group case, where an exact solution exists. We further propose near-linear time approximation algorithms for clustering with multiple (potentially unequal-sized) groups, answering an open problem posed by (Chakraborty et al. 2025a). Leveraging these results, we achieve improved approximation guarantees for fair correlation clustering and, for the first time, provide approximation guarantees for fair consensus clustering involving more than two groups. Promising directions for future research include improving the approximation factors further and investigating alternative fairness criteria with similar approximation guarantees.

## Acknowledgments

Diptarka Chakraborty was supported in part by an MoE AcRF Tier 1 grant (T1 251RES2303), and a Google South & South-East Asia Research Award. Kushagra Chatterjee was supported by an MoE AcRF Tier 1 grant (T1 251RES2303). Debarati Das was supported in part by NSF grant 2337832.

## References

- Ahmadi, S.; Galhotra, S.; Saha, B.; and Schwartz, R. 2020. Fair correlation clustering. *arXiv preprint arXiv:2002.03508*.
- Ahmadian, S.; Epasto, A.; Kumar, R.; and Mahdian, M. 2020. Fair Correlation Clustering. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108, 4195–4205.
- Ahmadian, S.; and Negahbani, M. 2023. Improved approximation for fair correlation clustering. In *International Conference on Artificial Intelligence and Statistics*, 9499–9516. PMLR.
- Ailon, N.; Charikar, M.; and Newman, A. 2008. Aggregating inconsistent information: ranking and clustering. *Journal of the ACM (JACM)*, 55(5): 1–27.
- Backurs, A.; Indyk, P.; Onak, K.; Schieber, B.; Vakilian, A.; and Wagner, T. 2019. Scalable Fair Clustering. In *International Conference on Machine Learning (ICML)*, volume 97, 405–413.
- Bandyapadhyay, S.; Chen, T.; Friggstad, Z.; and Jamshidian, M. 2025. A Constant-Factor Approximation for Pairwise Fair  $k$ -Center Clustering. In *International Conference on Integer Programming and Combinatorial Optimization*, 43–57. Springer.
- Bandyapadhyay, S.; Chlamtáč, E.; Friggstad, Z.; Jamshidian, M.; Makarychev, Y.; and Vakilian, A. 2024. A Polynomial-Time Approximation for Pairwise Fair  $k$ -Median Clustering. *arXiv preprint arXiv:2405.10378*.
- Bandyapadhyay, S.; Fomin, F. V.; and Simonov, K. 2024. On coresets for fair clustering in metric and euclidean spaces and their applications. *Journal of Computer and System Sciences*, 142: 103506.
- Bansal, N.; Blum, A.; and Chawla, S. 2004. Correlation clustering. *Machine learning*, 56(1): 89–113.
- Bera, S.; Chakraborty, D.; Flores, N.; and Negahbani, M. 2019. Fair algorithms for clustering. *Advances in Neural Information Processing Systems*, 32.
- Bercea, I. O.; Groß, M.; Khuller, S.; Kumar, A.; Rösner, C.; Schmidt, D. R.; and Schmidt, M. 2019. On the Cost of Essentially Fair Clusterings. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*, 18–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Böhm, M.; Fazzone, A.; Leonardi, S.; and Schwiegelshohn, C. 2020. Fair clustering with multiple colors. *arXiv preprint arXiv:2002.07892*.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems (NeurIPS)*, 29.
- Bonchi, F.; Garcia-Soriano, D.; and Liberty, E. 2014. Correlation clustering: from theory to practice. In *KDD*, 1972.
- Bonizzoni, P.; Vedova, G. D.; Dondi, R.; and Jiang, T. 2008. On the Approximation of Correlation Clustering and Consensus Clustering. *J. Comput. Syst. Sci.*, 74(5): 671–696.
- Bressan, M.; Cesa-Bianchi, N.; Paudice, A.; and Vitale, F. 2019. Correlation clustering with adaptive similarity queries. *Advances in neural information processing systems*, 32.
- Cao, N.; Cohen-Addad, V.; Lee, E.; Li, S.; Lolck, D. R.; Newman, A.; Thorup, M.; Vogl, L.; Yan, S.; and Zhang, H. 2025. Solving the Correlation Cluster LP in Sublinear Time. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, 1154–1165.
- Cao, N.; Cohen-Addad, V.; Lee, E.; Li, S.; Newman, A.; and Vogl, L. 2024. Understanding the Cluster Linear Program for Correlation Clustering. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 1605–1616.
- Celis, L. E.; Straszak, D.; and Vishnoi, N. K. 2018. Ranking with Fairness Constraints. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107, 28:1–28:15.
- Chakraborty, D.; Chatterjee, K.; Das, D.; Nguyen, T. L.; and Nobahari, R. 2025a. Towards Fair Representation: Clustering and Consensus. *38th Annual Conference on Learning Theory (COLT 2025); full version: arXiv preprint arXiv:2506.08673*.
- Chakraborty, D.; Das, H.; Dey, S.; and Yao, A. Y. H. 2025b. Improved Rank Aggregation Under Fairness Constraint. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2025, Montreal, Canada, August 16-22, 2025*, 330–338. ijcai.org.
- Chakraborty, D.; Das, S.; Khan, A.; and Subramanian, A. 2022. Fair rank aggregation. *Advances in Neural Information Processing Systems*, 35: 23965–23978. Full version: arXiv preprint arXiv:2308.10499.
- Charikar, M.; Guruswami, V.; and Wirth, A. 2005. Clustering with qualitative information. *Journal of Computer and System Sciences*, 71(3): 360–383.
- Chawla, S.; Makarychev, K.; Schramm, T.; and Yaroslavtsev, G. 2015. Near optimal lp rounding algorithm for correlation clustering on complete and complete  $k$ -partite graphs. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 219–228.
- Chen, X.; Fain, B.; Lyu, L.; and Munagala, K. 2019. Proportionally Fair Clustering. In *International Conference on Machine Learning (ICML)*, 1032–1041.
- Chierichetti, F.; Kumar, R.; Lattanzi, S.; and Vassilvitskii, S. 2017. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 5029–5037.

- Das, D.; and Kumar, A. 2025. Breaking the Two Approximation Barrier for Various Consensus Clustering Problems. In Azar, Y.; and Panigrahi, D., eds., *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, 323–372. SIAM.
- Demaine, E. D.; Emanuel, D.; Fiat, A.; and Immorlica, N. 2006. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2-3): 172–187.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science*, 214–226.
- Filkov, V.; and Skiena, S. 2004a. Heterogeneous data integration with the consensus clustering formalism. In *International Workshop on Data Integration in the Life Sciences*, 110–123. Springer.
- Filkov, V.; and Skiena, S. 2004b. Integrating microarray data by consensus clustering. *International Journal on Artificial Intelligence Tools*, 13(04): 863–880.
- Garey, M. R.; and Johnson, D. S. 1975. Complexity results for multiprocessor scheduling under resource constraints. *SIAM journal on Computing*, 4(4): 397–411.
- Goder, A.; and Filkov, V. 2008. Consensus clustering algorithms: Comparison and refinement. In *2008 Proceedings of the Tenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, 109–117. SIAM.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3315–3323.
- Hou, J. P.; Emad, A.; Puleo, G. J.; Ma, J.; and Milenkovic, O. 2016. A new correlation clustering method for cancer mutation analysis. *Bioinformatics*, 32(24): 3717–3728.
- Huang, L.; Jiang, S. H.; and Vishnoi, N. K. 2019. Coresets for Clustering with Fairness Constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, 7587–7598.
- Kay, M.; Matuszek, C.; and Munson, S. A. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *ACM conference on human factors in computing systems*, 3819–3828.
- Kliachkin, A.; Psaroudaki, E.; Mareček, J.; and Fotakis, D. 2024. Fairness in Ranking: Robustness through Randomization without the Protected Attribute. In *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*, 201–208. IEEE.
- Křivánek, M.; and Morávek, J. 1986. NP-hard problems in hierarchical-tree clustering. *Acta informatica*, 23: 311–323.
- Kushagra, S.; Ben-David, S.; and Ilyas, I. 2019. Semi-supervised clustering for de-duplication. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 1659–1667. PMLR.
- Lancichinetti, A.; and Fortunato, S. 2012. Consensus clustering in complex networks. *Scientific reports*, 2(1): 336.
- Monti, S.; Tamayo, P.; Mesirov, J.; and Golub, T. 2003. Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. *Machine learning*, 52: 91–118.
- Rösner, C.; and Schmidt, M. 2018. Privacy Preserving Clustering with Constraints. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, 96–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Swamy, C. 2004. Correlation Clustering: maximizing agreements via semidefinite programming. In *SODA*, volume 4, 526–527. Citeseer.
- Topchy, A.; Jain, A. K.; and Punch, W. 2003. Combining multiple weak clusterings. In *Third IEEE international conference on data mining*, 331–338. IEEE.
- Veldt, N.; Gleich, D. F.; and Wirth, A. 2018. A correlation clustering framework for community detection. In *Proceedings of the 2018 World Wide Web Conference*, 439–448.
- Wei, D.; Islam, M. M.; Schieber, B.; and Roy, S. B. 2022. Rank Aggregation with Proportionate Fairness. In *SIGMOD International Conference on Management of Data*, 262–275.
- Wu, J.; Liu, H.; Xiong, H.; Cao, J.; and Chen, J. 2014. K-means-based consensus clustering: A unified view. *IEEE transactions on knowledge and data engineering*, 27(1): 155–169.