
Fair Hierarchical Clustering

Sara Ahmadian
Google
sahmadian@google.com

Alessandro Epasto
Google
aepasto@google.com

Marina Knittel
University of Maryland
mknittel@cs.umd.edu

Ravi Kumar
Google
ravi.k53@gmail.com

Mohammad Mahdian
Google
mahdian@google.com

Benjamin Moseley
Carnegie Mellon University
moseleyb@andrew.cmu.edu

Philip Pham
Google
phillypham@google.com

Sergei Vassilvitskii
Google
sergeiv@google.com

Yuyan Wang
Carnegie Mellon University
yuyanw@andrew.cmu.edu

Abstract

As machine learning has become more prevalent, researchers have begun to recognize the necessity of ensuring machine learning systems are fair. Recently, there has been an interest in defining a notion of fairness that mitigates over-representation in traditional clustering.

In this paper we extend this notion to hierarchical clustering, where the goal is to recursively partition the data to optimize a specific objective. For various natural objectives, we obtain simple, efficient algorithms to find a provably good fair hierarchical clustering. Empirically, we show that our algorithms can find a fair hierarchical clustering, with only a negligible loss in the objective.

1 Introduction

Algorithms and machine learned models are increasingly used to assist in decision making on a wide range of issues, from mortgage approval to court sentencing recommendations [28]. It is clearly undesirable, and in many cases illegal, for models to be biased to groups, for instance to discriminate on the basis of race or religion. Ensuring that there is no bias is not as easy as removing these protected categories from the data. Even without them being explicitly listed, the correlation between sensitive features and the rest of the training data may still cause the algorithm to be biased. This has led to an emergent literature on computing provably fair outcomes (see the book [7]).

The prominence of clustering in data analysis, combined with its use for data segmentation, feature engineering, and visualization makes it critical that efficient fair clustering methods are developed. There has been a flurry of recent results in the ML research community, proposing algorithms for fair *flat* clustering, i.e., partitioning a dataset into a set of disjoint clusters, as captured by K-CENTER, K-MEDIAN, K-MEANS, correlation clustering objectives [3, 4, 6, 8, 9, 14, 18, 24, 25, 30, 31]. However, the same issues affect *hierarchical clustering*, which is the problem we study.

The input to the hierarchical clustering problem is a set of data points, with pairwise similarity or dissimilarity scores. A hierarchical clustering is a tree, whose leaves correspond to the individual datapoints. Each internal node represents a cluster containing all the points in the leaves of its subtree. Naturally, the cluster at an internal node is the union of the clusters given by its children. Hierarchical clustering is widely used in data analysis [21], social networks [32, 34], and image/text organization [26].

Hierarchical clustering is frequently used for flat clustering when the number of clusters is a priori unknown. A hierarchical clustering yields a set of clusterings at different granularities that are consistent with each other. Therefore, in all clustering problems where fairness is desired but the number of clusters is unknown, fair hierarchical clustering is useful. As concrete examples, consider a set of news articles organized by a topic hierarchy, where we wish to ensure that no single source or view point is over-represented in a cluster; or a hierarchical division of a geographic area, where the sensitive attribute is gender or race, and we wish to ensure balance in every level of the hierarchy. There are many such problems that benefit from fair hierarchical clustering, motivating its study.

Our contributions. We initiate an algorithmic study of fair hierarchical clustering. We build on Dasgupta’s seminal formal treatment of hierarchical clustering [20] and prove our results for the revenue [33], value [19], and cost [20] objectives in his framework.

To achieve fairness, we show how to extend the fairlets machinery, introduced by [16] and extended by [3], to this problem. We then investigate the complexity of finding a good fairlet decomposition, giving both strong computational lower bounds and polynomial time approximation algorithms.

Finally, we conclude with an empirical evaluation of our approach. We show that ignoring protected attributes when performing hierarchical clustering can lead to unfair clusters. On the other hand, adopting the fairlet framework in conjunction with the approximation algorithms we propose yields *fair* clusters with a *negligible* objective degradation.

Related work. Hierarchical clustering has received increased attention over the past few years. Dasgupta [20] developed a cost function objective for data sets with similarity scores, where similar points are encouraged to be clustered together lower in the tree. Cohen-Addad et al. [19] generalized these results into a class of optimization functions that possess other desirable properties and introduced their own value objective in the dissimilarity score context. In addition to validating their objective on inputs with known ground truth, they gave a theoretical justification for the average-linkage algorithm, one of the most popular algorithms used in practice, as a constant-factor approximation for value. Contemporaneously, Moseley and Wang [33] designed a revenue objective function based on the work of Dasgupta for point sets with similarity scores and showed the average-linkage algorithm is a constant approximation for this objective as well. This work was further improved by Charikar [13] who gave a tighter analysis of average-linkage for Euclidean data for this objective and [2, 5] who improved the approximation ratio in the general case.

In parallel to the new developments in algorithms for hierarchical clustering, there has been tremendous development in the area of fair machine learning. We refer the reader to a recent textbook [7] for a rich overview, and focus here on progress for fair clustering. Chierichetti et al. [16] first defined fairness for k -median and k -center clustering, and introduced the notion of *fairlets* to design efficient algorithms. Extensive research has focused on two topics: adapting the definition of fairness to broader contexts, and designing efficient algorithms for finding good fairlet decompositions. For the first topic, the fairness definition was extended to multiple values for the protected feature [3, 9, 35]. For the second topic, Backurs et al. [6] proposed a near-linear constant approximation algorithm for finding fairlets for k -median, Schmidt et al. [36] introduced a streaming algorithm for scalable computation of coresets for fair clustering, Kleindessner et al. [30] designed a linear time constant approximation algorithm for k -center where cluster centers are selected proportionally from a set of colors, Bercea et al. [9] developed methods for fair k -means, while Ahmadian et al. [4] and Ahmadi et al. [1] defined approximation algorithms for fair correlation clustering. Concurrently with our work, Chhabra et al. [15] introduced a possible approach to ensuring fairness in hierarchical clustering. However, their fairness definition differs from ours (in particular, they do not ensure that all levels of the tree are fair), and the methods they introduce are heuristic, without formal fairness or quality guarantees.

Beyond clustering, the same balance notion that we use has been utilized to capture fairness in other contexts, for instance: fair voting [10], fair optimization [17], as well as other problems [11].

2 Formulation

2.1 Objectives for hierarchical clustering

Let $G = (V, s)$ be an input instance, where V is a set of n data points, and $s : V^2 \rightarrow \mathbb{R}^{\geq 0}$ is a similarity function over vertex pairs. For two sets, $A, B \subseteq V$, we let $s(A, B) = \sum_{a \in A, b \in B} s(a, b)$

and $S(A) = \sum_{\{i,j\} \in A^2} s(i,j)$. For problems where the input is $G = (V, d)$, with d a distance function, we define $d(A, B)$ and $d(A)$ similarly. We also consider the *vertex-weighted* versions of the problem, i.e. $G = (V, s, m)$ (or $G = (V, d, m)$), where $m : V \rightarrow \mathbb{Z}^+$ is a weight function on the vertices. The vertex-unweighted version can be interpreted as setting $m(i) = 1, \forall i \in V$. For $U \subseteq V$, we use the notation $m(U) = \sum_{i \in U} m(i)$.

A *hierarchical clustering* of G is a tree whose leaves correspond to V and whose internal vertices represent the merging of vertices (or clusters) into larger clusters until all data merges at the root. The goal of hierarchical clustering is to build a tree to optimize some objective.

To define these objectives formally, we need some notation. Let T be a hierarchical clustering tree of G . For two leaves i and j , we say $i \vee j$ is their least common ancestor. For an internal vertex u in T , let $T[u]$ be the subtree in T rooted at u . Let $\text{leaves}(T[u])$ be the leaves of $T[u]$.

We consider three different objectives—*revenue*, *value*, and *cost*—based on the seminal framework of [20], and generalize them to the vertex-weighted case.

Revenue. Moseley and Wang [33] introduced the revenue objective for hierarchical clustering. Here the input instance is of the form $G = (V, s, m)$, where $s : V^2 \rightarrow \mathbb{R}^{\geq 0}$ is a *similarity* function.

Definition 1 (Revenue). *The revenue (rev) of a tree T for an instance $G = (V, s, m)$, where $s(\cdot, \cdot)$ denotes similarity between data points, is: $\text{rev}_G(T) = \sum_{i,j \in V} s(i,j) \cdot (m(V) - m(\text{leaves}(T[i \vee j])))$.*

Note that in this definition, each weight is scaled by (the vertex-weight of) the non-leaves. The goal is to find a tree of maximum revenue. It is known that average-linkage is a $1/3$ -approximation for vertex-unweighted revenue [33]; the state-of-the-art is a 0.585-approximation [5].

As part of the analysis, there is an upper bound for the revenue objective [19, 33], which is easily extended to the vertex-weighted setting:

$$\text{rev}_G(T) \leq \left(m(V) - \min_{u,v \in V, u \neq v} m(\{u, v\}) \right) \cdot s(V). \quad (1)$$

Note that in the vertex-unweighted case, the upper bound is just $(|V| - 2)s(V)$.

Value. A different objective was proposed by Cohen-Addad et al. [19], using distances instead of similarities. Let $G = (V, d, m)$, where $d : V^2 \rightarrow \mathbb{R}^{\geq 0}$ is a distance (or dissimilarity) function.

Definition 2 (Value). *The value (val) of a tree T for an instance $G = (V, d, m)$ where $d(\cdot, \cdot)$ denotes distance is: $\text{val}_G(T) = \sum_{i,j \in V} d(i,j) \cdot m(\text{leaves}(T[i \vee j]))$.*

As in revenue, we aim to find a hierarchical clustering to maximize value. Cohen-Addad et al. [19] showed that both average-linkage and a locally ϵ -densest cut algorithm achieve a $2/3$ -approximation for vertex-unweighted value. They also provided an upper bound for value, much like that in (1), which in the vertex-weighted context, is:

$$\text{val}_G(T) \leq m(V) \cdot d(V). \quad (2)$$

Cost. The original objective introduced by Dasgupta [20] for analyzing hierarchical clustering algorithms introduces the notion of cost.

Definition 3 (Cost). *The cost of a tree T for an instance $G = (V, s)$ where $s(\cdot, \cdot)$ denotes similarity is: $\text{cost}_G(T) = \sum_{i,j \in V} s(i,j) \cdot |\text{leaves}(T[i \vee j])|$.*

The objective is to find a tree of minimum cost. From a complexity point of view, cost is a harder objective to optimize. Charikar and Chatziafratis [12] showed that cost is not constant-factor approximable under the Small Set Expansion hypothesis, and the current best approximations are $O(\sqrt{\log n})$ and require solving SDPs.

Convention. Throughout the paper we adopt the following convention: $s(\cdot, \cdot)$ will always denote similarities and $d(\cdot, \cdot)$ will always denote distances. Thus, the inputs for the cost and revenue objectives will be instances of the form (V, s, m) and inputs for the value objective will be instances of the form (V, d, m) . All the missing proofs can be found in the Supplementary Material.

2.2 Notions of fairness

Many definitions have been proposed for fairness in clustering. We consider the setting in which each data point in V has a *color*; the color corresponds to the protected attribute.

Disparate impact. This notion is used to capture the fact that decisions (i.e., clusterings) should not be overly favorable to one group versus another. This notion was formalized by Chierichetti et al. [16] for clustering when the protected attribute can take on one of two values, i.e., points have one of two colors. In their setup, the *balance* of a cluster is the ratio of the minimum to the maximum number of points of any color in the cluster. Given a balance requirement t , a clustering is fair if and only if each cluster has a balance of at least t .

Bounded representation. A generalization of disparate impact, bounded representation focuses on mitigating the imbalance of the representation of protected classes (i.e., colors) in clusters and was defined by Ahmadian et al. [3]. Given an over-representation parameter α , a cluster is fair if the *fractional representation* of each color in the cluster is at most α , and a clustering is fair if each cluster has this property. This was further generalized by Bera et al. [8] and Bercea et al. [9]. They introduce vectors $\vec{\alpha}, \vec{\beta}$ such that for a cluster to be fair, for each color c_i , the fractional representation of c_i in the cluster must be between β_i and α_i . We discuss our results in terms of the over-representation constraint by [3], however many of these results extend to this more general setting given an appropriate fairlet decomposition. An interesting special case of this notion is when there are c total colors and $\alpha = 1/c$. In this case, we require that every color is equally represented in every cluster. We will refer to this as *equal representation*. These notions enjoy the following useful property:

Definition 4 (Union-closed). *A fairness constraint is union-closed if for any pair of fair clusters A and B , $A \cup B$ is also fair.*

This property is useful in hierarchical clustering: given a tree T and internal node u , if each child cluster of u is fair, then u must also be a fair cluster.

Definition 5 (Fair hierarchical clustering). *For any fairness constraint, a hierarchical clustering is fair if all of its clusters (besides the leaves) are fair.*¹

Thus, under any union-closed fairness constraint, this definition is equivalent to restricting the bottom-most clustering (besides the leaves) to be fair. Then given an objective (e.g., revenue), the goal is to find a fair hierarchical clustering that optimizes the objective. We focus on the bounded representation fairness notion with c colors and an over-representation cap α . However, the main ideas for the revenue and value objectives work under any notion of fairness that is union-closed.

3 Fairlet decomposition

Definition 6 (Fairlet [16]). *A fairlet Y is a fair set of points such that there is no partition of Y into Y_1 and Y_2 with both Y_1 and Y_2 being fair.*

In the bounded representation fairness setting, a set of points is fair if at most an α fraction of the points have the same color. We call this an α -capped fairlet. For $\alpha = 1/t$ with t an integer, the fairlet size will always be at most $2t - 1$. We will refer to the maximum size of a fairlet by m_f .

Recall that given a union-closed fairness constraint, if the bottom clustering in the tree is a layer of fairlets (which we call a *fairlet decomposition* of the original dataset) the hierarchical clustering tree is also fair. This observation gives an immediate algorithm for finding fair hierarchical clustering trees in a two-phase manner. (i) Find a fairlet decomposition, i.e., partition the input set V into clusters Y_1, Y_2, \dots that are all fairlets. (ii) Build a tree on top of all the fairlets. Our goal is to complete both phases in such a way that we optimize the given objective (i.e., revenue or value).

In Section 4, we will see that to optimize for the revenue objective, all we need is a fairlet decomposition with bounded fairlet size. However, the fairlet decomposition required for the value objective is more nuanced. We describe this next.

Fairlet decomposition for the value objective For the value objective, we need the total distance between pairs of points inside each fairlet to be small. Formally, suppose V is partitioned into fairlets

¹According to the definition, a hierarchical clustering tree might be fair even if every layer (apart from the root) is an unfair clustering. For example, consider a tree that splits off one singleton at its root. Every layer in the tree apart from the root will contain this singleton and thus is an unfair clustering. An alternative way of defining a fair tree is to enforce that the tree has to contain a layer of fairlets of some small size. The results of this paper extend to either definition.

$\mathcal{Y} = \{Y_1, Y_2, \dots\}$ such that Y_i is an α -capped fairlet. The cost of this decomposition is defined as:

$$\phi(\mathcal{Y}) = \sum_{Y \in \mathcal{Y}} \sum_{\{u,v\} \subseteq Y} d(u,v). \quad (3)$$

Unfortunately, the problem of finding a fairlet decomposition to minimize $\phi(\cdot)$ does not admit any constant-factor approximation unless $P = NP$.

Theorem 7. *Let $z \geq 3$ be an integer. Then there is no bounded approximation algorithm for finding $(\frac{z}{z+1})$ -capped fairlets optimizing $\phi(\mathcal{Y})$, which runs in polynomial time, unless $P = NP$.*

The proof proceeds by a reduction from the Triangle Partition problem, which asks if a graph $G = (V, E)$ on $3n$ vertices can be partitioned into three element sets, with each set forming a triangle in G . Fortunately, for the purpose of optimizing the value objective, it is not necessary to find an approximate decomposition.

4 Optimizing revenue with fairness

This section considers the revenue objective. We will obtain an approximation algorithm for this objective in three steps: (i) obtain a fairlet decomposition such that the maximum fairlet size in the decomposition is small, (ii) show that any β -approximation algorithm to (1) (i.e., any algorithm that achieves a β -factor approximation of (1) for some given β) plus this fairlet decomposition can be used to obtain a (roughly) β -approximation for fair hierarchical clustering under the revenue objective, and (iii) use average-linkage, which is known to be a $1/3$ -approximation to (1). (We note that the recent work [2, 5] on improved approximation algorithms compare to a bound on the optimal solution that differs from (1) and therefore do not fit into our framework.)

First, we address step (ii). Due to space, this proof can be found in Appendix B. Note that Theorem 8 extends to the fairness constraint defined by [8, 9]’s provided a fairlet decomposition in this setting.

Theorem 8. *Given an algorithm that obtains a β -approximation to (1) where $\beta \leq 1$, and a fairlet decomposition with maximum fairlet size m_f , there is a $\beta \left(1 - \frac{2m_f}{n}\right)$ -approximation for fair hierarchical clustering under the revenue objective.*

Prior work showed that average-linkage is a $1/3$ -approximation to (1) in the vertex-unweighted case; this proof can be easily modified to show that it is still $1/3$ -approximation even with vertex weights. This accounts for step (iii) in our process.

Combined with the fairlet decomposition methods for the two-color case [16], which has $m_f = b + r$ for b blue vertices and r red vertices, and for multi-color case (Supplementary Material), which has $m_f \leq 2t - 1$, to address step (i), we have the following.

Corollary 9. *There is polynomial time algorithm that constructs a fair tree that is a $\frac{1}{3} \left(1 - \frac{2m_f}{n}\right)$ -approximation for revenue objective, where m_f is the maximum size of fairlets.*

5 Optimizing value with fairness

In this section we consider the value objective. As in the revenue objective, we prove that we can reduce fair hierarchical clustering to the problem of finding a good fairlet decomposition for the proposed fairlet objective (3), and then use any approximation algorithm for weighted hierarchical clustering with the decomposition as the input. Again, our result applies to [8, 9]’s fairness constraint if we are given an appropriate fairness decomposition.

Theorem 10. *Given an algorithm that gives a β -approximation to (2) where $\beta \leq 1$, and a fairlet decomposition \mathcal{Y} such that $\phi(\mathcal{Y}) \leq \epsilon \cdot d(V)$, there is a $\beta(1 - \epsilon)$ approximation for (2).*

We complement this result with an algorithm that finds a good fairlet decomposition in polynomial time under the bounded representation fairness constraint with cap α .

Let R_1, \dots, R_c be the c colors and let $\mathcal{Y} = \{Y_1, Y_2, \dots\}$ be the fairlet decomposition. Let n_i be the number of points colored R_i in V . Let $r_{i,k}$ denote the number of points colored R_i in the k th fairlet.

Theorem 11. *There exists a local search algorithm that finds a fairlet decomposition \mathcal{Y} with $\phi(\mathcal{Y}) \leq (1 + \epsilon) \max_{i,k} \frac{r_{i,k}}{n_i} d(V)$ in time $\tilde{O}(n^3/\epsilon)$.*

We can now use the fact that both average-linkage and the $\frac{\epsilon}{n}$ -locally-densest cut algorithm give a $\frac{2}{3}$ - and $(\frac{2}{3} - \epsilon)$ -approximation respectively for vertex-weighted hierarchical clustering under the value objective. Finally, recall that fairlets are intended to be minimal, and their size depends only on the parameter α , and not on the size of the original input. Therefore, as long as the number of points of each color increases as input size, n , grows, the ratio $r_{i,k}/n_i$ goes to 0. These results, combined with Theorem 10 and Theorem 11, yield Corollary 12.

Corollary 12. *Given bounded size fairlets, the fairlet decomposition computed by local search combined with average-linkage constructs a fair hierarchical clustering that is a $\frac{2}{3}(1 - o(1))$ -approximation for the value objective. For the $\frac{\epsilon}{n}$ -locally-densest cut algorithm in [19], we get a polynomial time algorithm for fair hierarchical clustering that is a $(\frac{2}{3} - \epsilon)(1 - o(1))$ -approximation under the value objective for any $\epsilon > 0$.*

Given at most a small fraction of every color is in any cluster, Corollary 12 states that we can extend the state-of-the-art results for value to the α -capped, multi-colored constraint. Note that the preconditions will always be satisfied and the extension will hold in the two-color fairness setting or in the multi-colored equal representation fairness setting.

Fairlet decompositions via local search In this section, we give a local search algorithm to construct a fairlet decomposition, which proves Theorem 11. This is inspired by the ϵ -densest cut algorithm of [19]. To start, recall that for a pair of sets A and B we denote by $d(A, B)$ the sum of interpoint distances, $d(A, B) = \sum_{u \in A, v \in B} d(u, v)$. A fairlet decomposition is a partition of the input $\{Y_1, Y_2, \dots\}$ such that each color composes at most an α fraction of each Y_i .

We start by finding an arbitrary α -capped fairlet decomposition. For two colors with $\alpha = r/(b+r)$, we use the fairlet decomposition introduced by Chierichetti et al. [16]. For multiple colors with $\alpha = 1/t$, we defer to Lemma 24 in Appendix C.2. Our algorithm will then recursively subdivide the cluster of all data to construct a hierarchy by finding cuts. To search for a cut, we will use a *swap* method.

Definition 13 (Local optimality). *Consider any fairlet decomposition $\mathcal{Y} = \{Y_1, Y_2, \dots\}$ and $\epsilon > 0$. Define a swap of $u \in Y_i$ and $v \in Y_j$ for $j \neq i$ as updating Y_i to be $(Y_i \setminus \{u\}) \cup \{v\}$ and Y_j to be $(Y_j \setminus \{v\}) \cup \{u\}$. We say \mathcal{Y} is ϵ -locally-optimal if any swap with u, v of the same color reduces the objective value by less than a $(1 + \epsilon)$ factor.*

The algorithm constructs a (ϵ/n) -locally optimal algorithm for fairlet decomposition, which runs in $\tilde{O}(n^3/\epsilon)$ time. Consider any given instance (V, d) . Let d_{\max} denote the maximum distance, m_f denote the maximum fairlet size, and $\Delta = d_{\max} \cdot \frac{m_f}{n}$. The algorithm begins with an arbitrary decomposition. Then it swaps pairs of monochromatic points until it terminates with a locally optimal solution. By construction we have the following.

Claim 14. *Algorithm 1 finds a valid fairlet decomposition.*

We prove two things: Algorithm 1 optimizes the objective (3), and has a small running time. The following lemma gives an upper bound on \mathcal{Y} 's performance for (3) found by Algorithm 1.

Lemma 15. *The fairlet decomposition \mathcal{Y} computed by Algorithm 1 has an objective value for (3) of at most $(1 + \epsilon) \max_{i,k} \frac{r_{i,k}}{n_i} d(V)$.*

Finally we bound the running time. The algorithm has much better performance in practice than its worst-case analysis would indicate. We will show this later in Section 7.

Lemma 16. *The running time for Algorithm 1 is $\tilde{O}(n^3/\epsilon)$.*

Together, Lemma 15, Lemma 16, and Claim 14 prove Theorem 11. This establishes that there is a local search algorithm that can construct a good fairlet decomposition.

Algorithm 1 Algorithm for (ϵ/n) -locally-optimal fairlet decomposition.

Input: A set V with distance function $d \geq 0$, parameter α , small constant $\epsilon \in [0, 1]$.**Output:** An α -capped fairlet decomposition \mathcal{Y} .

- 1: Find $d_{\max}, \Delta \leftarrow \frac{m_f}{n} d_{\max}$.
 - 2: Arbitrarily find an α -capped fairlet decomposition $\{Y_1, Y_2, \dots\}$ such that each partition has at most an α fraction of any color. {See [16] or Appendix C.2}
 - 3: **while** $\exists u \in Y_i, v \in Y_j, i \neq j$ of the same color, such that for the decomposition \mathcal{Y}' after swapping $u, v, \frac{\sum_{Y_k \in \mathcal{Y}} d(Y_k)}{\sum_{Y_k \in \mathcal{Y}'} d(Y_k)} \geq (1 + \epsilon/n)$ **and** $\sum_{Y_k \in \mathcal{Y}} d(Y_k) > \Delta$ **do**
 - 4: Swap u and v by setting $Y_i \leftarrow (Y_i \setminus \{u\}) \cup \{v\}$ and $Y_j \leftarrow (Y_j \setminus \{v\}) \cup \{u\}$.
 - 5: **end while**
-

Table 1: Dataset description. Here (b, r) denotes the balance of the dataset.

Name	Sample size	# features	Protected feature	Color (blue, red)	(b, r)
CENSUSGENDER	30162	6	gender	(female, male)	(1, 3)
CENSUSRACE	30162	6	race	(non-white, white)	(1, 7)
BANKMARRIAGE	45211	7	marital status	(not married, married)	(1, 2)
BANKAGE	45211	7	age	(< 40, \geq 40)	(2, 3)

Table 2: Impact of Algorithm 1 on $\text{ratio}_{\text{value}}$ in percentage (mean \pm std. dev).

Samples	400	800	1600	3200	6400	12800
CENSUSGENDER, initial	88.17 \pm 0.76	88.39 \pm 0.21	88.27 \pm 0.40	88.12 \pm 0.26	88.00 \pm 0.10	88.04 \pm 0.13
final	99.01 \pm 0.60	99.09 \pm 0.58	99.55 \pm 0.26	99.64 \pm 0.13	99.20 \pm 0.38	99.44 \pm 0.23
CENSUSRACE, initial	84.49 \pm 0.66	85.01 \pm 0.31	85.00 \pm 0.42	84.88 \pm 0.43	84.84 \pm 0.16	84.89 \pm 0.20
final	99.50 \pm 0.20	99.89 \pm 0.32	100.0 \pm 0.21	99.98 \pm 0.21	99.98 \pm 0.11	99.93 \pm 0.31
BANKMARRIAGE, initial	92.47 \pm 0.54	92.58 \pm 0.30	92.42 \pm 0.30	92.53 \pm 0.14	92.59 \pm 0.14	92.75 \pm 0.04
final	99.18 \pm 0.22	99.28 \pm 0.33	99.59 \pm 0.14	99.51 \pm 0.17	99.46 \pm 0.10	99.50 \pm 0.05
BANKAGE, initial	93.70 \pm 0.56	93.35 \pm 0.41	92.95 \pm 0.25	93.28 \pm 0.13	93.36 \pm 0.12	93.33 \pm 0.12
final	99.40 \pm 0.28	99.40 \pm 0.51	99.61 \pm 0.13	99.64 \pm 0.07	99.65 \pm 0.08	99.59 \pm 0.06

6 Optimizing cost with fairness

This section considers the cost objective of [20]. Even without our fairness constraint, the difficulty of approximating cost is clear in its approximation hardness and the fact that all known solutions require an LP or SDP solver. We obtain the result in Theorem 17; extending this result to other fairness constraints, improving its bound, or even making the algorithm practical, are open questions.

Theorem 17. *Consider the two-color case. Given a β -approximation for cost and a γ_t -approximation for minimum weighted bisection² on input of size t , then for parameters t and ℓ such that $n \geq t\ell$ and $n > \ell + 108t^2/\ell^2$, there is a fair $O\left(\frac{n}{t} + t\ell + \frac{n\ell\gamma_t}{t} + \frac{nt\gamma_t}{\ell^2}\right)\beta$ -approximation for $\text{cost}(T_{\text{unfair}}^*)$.*

With proper parameterization, we achieve an $O\left(n^{5/6} \log^{5/4} n\right)$ -approximation. We defer our algorithm description, pseudocode, and proofs to the Supplementary Material. While our algorithm is not simple, it is an important (and non-obvious) step to show the existence of an approximation, which we hope will spur future work in this area.

7 Experiments

This section validates our algorithms from Sections 4 and 5 empirically. We adopt the disparate impact fairness constraint [16]; thus each point is either blue or red. In particular, we would like to:

- Show that running the standard average-linkage algorithm results in highly unfair solutions.
- Demonstrate that demanding fairness in hierarchical clustering incurs only a small loss in the hierarchical clustering objective.
- Show that our algorithms, including fairlet decomposition, are practical on real data.

In Appendix G we consider multiple colors and the same trends as the two color case occur.

²The minimum weighted bisection problem is to find a partition of nodes into two equal-sized subsets so that the sum of the weights of the edges crossing the partition is minimized.

Table 3: Impact of Algorithm 1 on $\text{ratio}_{\text{fairlets}}$.

Samples	100	200	400	800	1600	3200	6400	12800
CENSUSGENDER, initial	2.5e-2	1.2e-2	6.2e-3	3.0e-3	1.5e-3	7.5e-4	3.8e-4	1.9e-4
final	4.9e-3	1.4e-3	6.9e-4	2.5e-4	8.5e-5	3.6e-5	1.8e-5	8.0e-6
CENSUSRACE, initial	6.6e-2	3.4e-2	1.7e-2	8.4e-3	4.2e-3	2.1e-3	1.1e-3	5.3e-4
final	2.5e-2	1.2e-2	6.2e-3	3.0e-3	1.5e-3	7.5e-4	3.8e-4	1.9e-5
BANKMARRIAGE, initial	1.7e-2	8.2e-3	4.0e-3	2.0e-3	1.0e-3	5.0e-4	2.5e-4	1.3e-4
final	5.9e-3	2.1e-3	9.3e-4	4.1e-4	1.3e-4	7.1e-5	3.3e-5	1.4e-5
BANKAGE, initial	1.3e-2	7.4e-3	3.5e-3	1.9e-3	9.3e-4	4.7e-4	2.3e-4	1.2e-4
final	5.0e-3	2.2e-3	7.0e-4	3.7e-4	1.3e-4	5.7e-5	3.0e-5	1.4e-5




Figure 1: (i) $\text{ratio}_{\text{fairlets}}$, every 100 swaps. (ii) $\text{ratio}_{\text{value}}$, every 100 swaps. (iii) CENSUSGENDER: running time vs sample size on a log-log scale.

Datasets. We use two datasets from the UCI data repository.³ In each dataset, we use features with numerical values and leave out samples with empty entries. For value, we use the Euclidean distance as the dissimilarity measure. For revenue, we set the similarity to be $s(i, j) = \frac{1}{1+d(i, j)}$ where $d(i, j)$ is the Euclidean distance. We pick two different protected features for both datasets, resulting in four datasets in total (See Table 1 for details).

- *Census* dataset: We choose *gender* and *race* to be the protected feature and call the resulting datasets CENSUSGENDER and CENSUSRACE.
- *Bank* dataset: We choose *marital status* and *age* to be the protected features and call the resulting datasets BANKMARRIAGE and BANKAGE.

In this section, unless otherwise specified, we report results only for the value objective. Results for the revenue objective are qualitatively similar and are omitted here. We do not evaluate our algorithm for the cost objective since it is currently only of theoretical interest.

We sub-sample points of two colors from the original data set proportionally, while approximately retaining the original color balance. The sample sizes used are 100×2^i , $i = 0, \dots, 8$. On each, we do 5 experiments and report the average results. We set ϵ in Algorithm 1 to 0.1 in all of the experiments.

Implementation. The code is available in the Supplementary Material. In the experiments, we use Algorithm 1 for the fairlet decomposition phase, where the fairlet decomposition is initialized by randomly assigning red and blue points to each fairlet. We apply the average-linkage algorithm to create a tree on the fairlets. We further use average-linkage to create subtrees inside of each fairlet.

The algorithm selects a *random* pair of blue or red points in different fairlets to swap, and checks if the swap sufficiently improves the objective. We do not run the algorithm until all the pairs are checked, rather the algorithm stops if it has made a $2n$ failed attempts to swap a random pair. As we observe empirically, this does not have material effect on the quality of the overall solution.

Metrics. We present results for value here, the results for revenue are qualitatively similar. In our experiments, we track the following quantities. Let G be the given input instance and let T be the output of our fair hierarchical clustering algorithm. We consider the following ratio $\text{ratio}_{\text{value}} = \frac{\text{value}_G(T)}{\text{value}_G(T')}$, where T' is the tree obtained by the standard average-linkage algorithm. We consider the fairlet objective function where \mathcal{Y} is a fairlet decomposition. Let $\text{ratio}_{\text{fairlets}} = \frac{\phi(\mathcal{Y})}{d(V)}$.

³archive.ics.uci.edu/ml/index.php, Census: archive.ics.uci.edu/ml/datasets/census+income, Bank: archive.ics.uci.edu/ml/datasets/Bank+Marketing

Table 4: Clustering on fairlets found by local search vs. upper bound, at size 1600 (mean \pm std. dev).

Dataset	CENSUSGENDER	CENSUSRACE	BANKMARRIAGE	BANKAGE
Revenue vs. upper bound	81.89 \pm 0.40	81.75 \pm 0.83	61.53 \pm 0.37	61.66 \pm 0.66
Value vs. upper bound	84.31 \pm 0.15	84.52 \pm 0.22	89.17 \pm 0.29	88.81 \pm 0.18

Results. Average-linkage algorithm always constructs unfair trees. For each of the datasets, the algorithm results in monochromatic clusters at some level, strengthening the case for fair algorithms.

In Table 2, we show for each dataset the $\text{ratio}_{\text{value}}$ both at the time of initialization (Initial) and after using the local search algorithm (Final). We see the change in the ratio as the local search algorithm performs swaps. Fairness leads to almost no degradation in the objective value as the swaps increase. Table 3 shows the $\text{ratio}_{\text{fairlets}}$ between the initial initialization and the final output fairlets. As we see, Algorithm 1 significantly improves the fairness of the initial random fairlet decomposition. The more the locally-optimal algorithm improves the objective value of (3), the better the tree’s performance based on the fairlets. Figures 1(i) and 1(ii) show $\text{ratio}_{\text{value}}$ and $\text{ratio}_{\text{fairlets}}$ for every 100 swaps in the execution of Algorithm 1 on a subsample of size 3200 from Census data set. The plots show that as the fairlet objective value decreases, the value objective of the resulting fair tree increases. Such correlation are found on subsamples of all sizes.

Now we compare the objective value of the algorithm with the upper bound on the optimum. We report the results for both the revenue and value objectives, using fairlets obtained by local search, in Table 4. On all datasets, we obtain ratios significantly better than the theoretical worst case guarantee. In Figure 1(iii), we show the average running time on Census data for both the original average-linkage and the fair average-linkage algorithms. As the sample size grows, the running time scales almost as well as current implementations of average-linkage algorithm. Thus with a modest increase in time, we can obtain a fair hierarchical clustering under the value objective.

8 Conclusions

In this paper we extended the notion of fairness to the classical problem of hierarchical clustering under three different objectives (revenue, value, and cost). Our results show that revenue and value are easy to optimize with fairness; while optimizing cost appears to be more challenging.

Our work raises several questions and research directions. Can the approximations be improved? Can we find better upper and lower bounds for fair cost? Are there other important fairness criteria?

Broader Impact

Our work builds upon a long line of work of fairness in machine learning. See the excellent books by Kearns and Roth [27], and Barocas et al. [7] for a rich introduction to the field.

Our aim in this work is algorithmic in nature, finding near-optimal hierarchical clustering algorithms that attain certain fairness guarantees. Since these methods are common unsupervised learning primitives, it is important to develop tools for practitioners to use. At the same time we remark that just because an algorithm is proven to be “fair” under some definition, does not mean it can be applied blindly.

As is now well known, [29], different fairness notions can be incompatible with each other. Moreover, fairness in machine learning is necessarily problem specific, and depends on the goals and the values of the person invoking the algorithm. While these facts are well established in the research community, they are far from common knowledge outside of it. Thus work on algorithmic notions of fairness runs the risk of someone treating the results as a silver bullet, and eschewing the deeper analysis that is necessary in any real world application.

Funding Sources

B. Moseley and Y. Wang were supported in part by a Google Research Award, an Infor Research Award, a Carnegie Bosch Junior Faculty Chair and NSF grants CCF-1824303, CCF-1845146, CCF-1733873 and CMMI-1938909. B. Moseley additionally is a part time employee of Relational-AI. M. Knittel was supported in part by NSF BIGDATA grant IIS-1546108, and NSF SPX grant CCF-1822738 and some of the work was conducted while she was visiting Google.

References

- [1] S. Ahmadi, S. Galhotra, B. Saha, and R. Schwartz. Fair correlation clustering. *arXiv:2002.03508*, 2020.
- [2] S. Ahmadian, V. Chatziafratis, A. Epasto, E. Lee, M. Mahdian, K. Makarychev, and G. Yaroslavtsev. Bisect and conquer: Hierarchical clustering via max-uncut bisection. In *AISTATS*, 2020.
- [3] S. Ahmadian, A. Epasto, R. Kumar, and M. Mahdian. Clustering without over-representation. In *KDD*, pages 267–275, 2019.
- [4] S. Ahmadian, A. Epasto, R. Kumar, and M. Mahdian. Fair correlation clustering. In *AISTATS*, 2020.
- [5] N. Alon, Y. Azar, and D. Vainstein. Hierarchical clustering: a 0.585 revenue approximation. In *COLT*, 2020.
- [6] A. Backurs, P. Indyk, K. Onak, B. Schieber, A. Vakilian, and T. Wagner. Scalable fair clustering. In *ICML*, pages 405–413, 2019.
- [7] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning*. www.fairmlbook.org, 2019.
- [8] S. Bera, D. Chakrabarty, N. Flores, and M. Negahbani. Fair algorithms for clustering. In *NeurIPS*, pages 4955–4966, 2019.
- [9] 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.
- [10] L. E. Celis, L. Huang, and N. K. Vishnoi. Multiwinner voting with fairness constraints. In *IJCAI*, pages 144–151, 2018.
- [11] L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. In *ICALP*, pages 28:1–28:15, 2018.
- [12] M. Charikar and V. Chatziafratis. Approximate hierarchical clustering via sparsest cut and spreading metrics. In *SODA*, pages 841–854, 2017.

- [13] M. Charikar, V. Chatziafratis, and R. Niazadeh. Hierarchical clustering better than average-linkage. In *SODA*, pages 2291–2304, 2019.
- [14] X. Chen, B. Fain, C. Lyu, and K. Munagala. Proportionally fair clustering. In *ICML*, pages 1032–1041, 2019.
- [15] A. Chhabra and P. Mohapatra. Fair algorithms for hierarchical agglomerative clustering. *arXiv:2005.03197*, 2020.
- [16] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *NIPS*, pages 5029–5037, 2017.
- [17] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Matroids, matchings, and fairness. In *AISTATS*, pages 2212–2220, 2019.
- [18] A. Chiplunkar, S. Kale, and S. N. Ramamoorthy. How to solve fair k -center in massive data models. In *ICML*, 2020.
- [19] V. Cohen-Addad, V. Kanade, F. Mallmann-Trenn, and C. Mathieu. Hierarchical clustering: Objective functions and algorithms. In *SODA*, pages 378–397, 2018.
- [20] S. Dasgupta. A cost function for similarity-based hierarchical clustering. In *STOC*, pages 118–127, 2016.
- [21] R. Dubes and A. K. Jain. Clustering methodologies in exploratory data analysis. In *Advances in Computers*, volume 19, pages 113–228. Academic Press, 1980.
- [22] U. Feige and R. Krauthgamer. A polylogarithmic approximation of the minimum bisection. In *FOCS*, pages 105–115, 2000.
- [23] M. R. Garey and D. S. Johnson. *Computers and Intractability*. WH Freeman New York, 2002.
- [24] L. Huang, S. H. C. Jiang, and N. K. Vishnoi. Coresets for clustering with fairness constraints. In *NeurIPS*, pages 7587–7598, 2019.
- [25] M. Jones, T. Nguyen, and H. Nguyen. Fair k -centers via maximum matching. In *ICML*, 2020.
- [26] M. S. G. Karypis, V. Kumar, and M. Steinbach. A comparison of document clustering techniques. In *TextMining Workshop at KDD*, 2000.
- [27] M. Kearns and A. Roth. *The Ethical Algorithm*. Oxford University Press, 2020.
- [28] J. Kleinberg, H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan. Human decisions and machine predictions. *The Quarterly Journal of Economics*, 133(1):237–293, 2017.
- [29] J. M. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent trade-offs in the fair determination of risk scores. In *ITCS*, pages 43:1–43:23, 2017.
- [30] M. Kleindessner, P. Awasthi, and J. Morgenstern. Fair k -center clustering for data summarization. In *ICML*, pages 3448–3457, 2019.
- [31] M. Kleindessner, S. Samadi, P. Awasthi, and J. Morgenstern. Guarantees for spectral clustering with fairness constraints. In *ICML*, pages 3448–3457, 2019.
- [32] C. F. Mann, D. W. Matula, and E. V. Olinick. The use of sparsest cuts to reveal the hierarchical community structure of social networks. *Social Networks*, 30(3):223–234, 2008.
- [33] B. Moseley and J. Wang. Approximation bounds for hierarchical clustering: Average linkage, bisecting k -means, and local search. In *NIPS*, pages 3094–3103, 2017.
- [34] A. Rajaraman and J. D. Ullman. *Mining of Massive Datasets*. Cambridge University Press, 2011.
- [35] C. Rösner and M. Schmidt. Privacy preserving clustering with constraints. In *ICALP*, pages 96:1–96:14, 2018.
- [36] M. Schmidt, C. Schwiegelshohn, and C. Sohler. Fair coresets and streaming algorithms for fair k -means. In E. Bampis and N. Megow, editors, *WAOA*, pages 232–251, 2019.

Appendix

A Approximation algorithms for weighted hierarchical clustering

In this section we first prove that running constant-approximation algorithms on fairlets gives good solutions for value objective, and then give constant approximation algorithms for both revenue and value in weighted hierarchical clustering problem, as is mentioned in Corollary 9 and 12. That is, a weighted version of average-linkage, for both weighted revenue and value objective, and weighted (ϵ/n) -locally densest cut algorithm, which works for weighted value objective. Both proofs are easily adapted from previous proofs in [19] and [33].

A.1 Running constant-approximation algorithms on fairlets

In this section, we prove Theorem 10, which says if we run any β -approximation algorithm for the upper bound on weighted value on the fairlet decomposition, we get a fair tree with minimal loss in approximation ratio. For the remainder of this section, fix any hierarchical clustering algorithm A that is guaranteed on any *weighted* input (V, d, m) to construct a hierarchical clustering with objective value at least $\beta m(V)d(V)$ for the value objective on a weighted input. Recall that we extended the value objective to a weighted variant in the Preliminaries Section and $m(V) = \sum_{u \in V} m_u$. Our aim is to show that we can combine A with the fairlet decomposition \mathcal{Y} introduced in the prior section to get a fair hierarchical clustering that is a $\beta(1 - \epsilon)$ -approximation for the value objective, if $\phi(\mathcal{Y}) \leq \epsilon d(V)$.

In the following definition, we transform the point set to a new set of points that are weighted. We will analyze A on this new set of points. We then show how we can relate this to the objective value of the optimal tree on the original set of points.

Definition 18. Let $\mathcal{Y} = \{Y_1, Y_2, \dots\}$ be the fairlet decomposition for V that is produced by the local search algorithm. Define $V(\mathcal{Y})$ as follows:

- Each set Y_i has a corresponding point a_i in $V(\mathcal{Y})$.
- The weight m_i of a_i is set to be $|Y_i|$.
- For each partitions Y_i, Y_j , where $i \neq j$ and $Y_i, Y_j \in \mathcal{Y}$, $d(a_i, a_j) = d(Y_i, Y_j)$.

We begin by observing the objective value that A receives on the instance $V(\mathcal{Y})$ is large compared to the weights in the original instance.

Theorem 19. On the instance $V(\mathcal{Y})$ the algorithm A has a total weighted objective of $\beta(1 - \epsilon) \cdot nd(V)$.

Proof. Notice that $m(V(\mathcal{Y})) = |V| = n$. Consider the total sum of all the distances in $V(\mathcal{Y})$. This is $\sum_{a_i, a_j \in V(\mathcal{Y})} d(a_i, a_j) = \sum_{Y_i, Y_j \in \mathcal{Y}} d(Y_i, Y_j) = d(V) - \phi(\mathcal{Y})$. The upper bound on the optimal solution is $(\sum_{Y_i \in \mathcal{Y}} m_i)(d(V) - \phi(\mathcal{Y})) = n(d(V) - \phi(\mathcal{Y}))$. Since $\phi(\mathcal{Y}) \leq \epsilon d(V)$, this upper bound is at least $(1 - \epsilon)nd(V)$. Theorem 10 follows from the fact that the algorithm A archives a weighted revenue at least a β factor of the total weighted distances. \square

A.2 Weighted hierarchical clustering: Constant-factor approximation

For weighted hierarchical clustering with positive integral weights, we define the weighted average-linkage algorithm for input (V, d, m) and (V, s, m) . Define the *average distance* to be $Avg(A, B) = \frac{d(A, B)}{m(A)m(B)}$ for dissimilarity-based input, and $Avg(A, B) = \frac{s(A, B)}{m(A)m(B)}$ for similarity-based input. In each iteration, weighted average-linkage seeks to merge the clusters which minimizes this value, if dissimilarity-based, and maximizes this value, if similarity-based.

Lemma 20. Weighted average-linkage is a $\frac{2}{3}$ (resp., $\frac{1}{3}$) approximation for the upper bound on weighted value (resp., revenue) objective with positive, integral weights.

Proof. We prove it for weighted value first. This is directly implied by the fact that average-linkage is $\frac{2}{3}$ approximation for unweighted value objective, as is proved in [19]. We have already seen in the

last subsection that a unweighted input V can be converted into weighted input $V(\mathcal{Y})$. Vice versa, we can construct a weighted input (V, d, m) into unweighted input with same upper bound for value objective.

In weighted hierarchical clustering we treat each point p with integral weights as $m(p)$ duplicates of points with distance 0 among themselves, let's call this set $S(p)$. For two weighted points $(p, m(p))$ and $(q, m(q))$, if $i \in S(p), j \in S(q)$, let $d(i, j) = \frac{d(p, q)}{m(p)m(q)}$. This *unweighted* instance, composed of many duplicates, has the same upper bound as the weighted instance. Notice that running average-linkage on the unweighted instance will always choose to put all the duplicates $S(p)$ together first for each p , and then do hierarchical clustering on top of the duplicates. Thus running average-linkage on the unweighted input gives a valid hierarchical clustering tree for weighted input. Since unweighted value upper bound equals weighted value upper bound, the approximation ratio is the same.

Now we prove it for weighted revenue. In [33], average-linkage being $\frac{1}{3}$ approximation for unweighted revenue is proved by the following. Given any clustering \mathcal{C} , if average-linkage chooses to merge A and B in \mathcal{C} , we define a local revenue for this merge:

$$\text{merge-rev}(A, B) = \sum_{C \in \mathcal{C} \setminus \{A, B\}} |C||A||B| \text{Avg}(A, B).$$

And correspondingly, a local cost:

$$\text{merge-cost}(A, B) = \sum_{C \in \mathcal{C} \setminus \{A, B\}} (|B||A||C| \text{Avg}(A, C) + |A||B||C| \text{Avg}(B, C)).$$

Summing up the local revenue and cost over all merges gives the upper bound. [33] used the property of average-linkage to prove that at every merge, $\text{merge-cost}(A, B) \leq 2\text{merge-rev}(A, B)$, which guarantees the total revenue, which is the summation of $\text{merge-rev}(A, B)$ over all merges, is at least $\frac{1}{3}$ of the upper bound. For the weighted case, we define

$$\text{merge-rev}(A, B) = \sum_{C \in \mathcal{C} \setminus \{A, B\}} m(C)m(A)m(B) \text{Avg}(A, B).$$

And

$$\text{merge-cost}(A, B) = \sum_{C \in \mathcal{C} \setminus \{A, B\}} (m(B)m(A)m(C) \text{Avg}(A, C) + m(A)m(B)m(C) \text{Avg}(B, C)).$$

And the rest of the proof works in the same way as in [33], proving weighted average-linkage to be $\frac{1}{3}$ for weighted revenue. \square

Next we define the weighted (ϵ/n) -locally-densest cut algorithm. The original algorithm, introduced in [19], defines a cut to be $\frac{d(A, B)}{|A||B|}$. It starts with the original set as one cluster, at every step, it seeks the partition of the current set that locally maximizes this value, and thus constructing a tree from top to bottom. For the weighted input (V, d, m) , we define the cut to be $\frac{d(A, B)}{m(A)m(B)}$, and let $n = m(V)$. For more description of the algorithm, see Algorithm 4 in Section 6.2 in [19].

Lemma 21. *Weighted (ϵ/n) -locally-densest cut algorithm is a $\frac{2}{3} - \epsilon$ approximation for weighted value objective.*

Proof. Just as in the average-linkage proof, we convert each weighted point p into a set S of $m(p)$ duplicates of p . Notice that the converted unweighted hierarchical clustering input has the same upper bound as the weighted hierarchical clustering input, and the ϵ/n -locally-densest cut algorithm moves all the duplicate sets S around in the unweighted input, instead of single points as in the original algorithm in [19].

Focus on a split of cluster $A \cup B$ into (A, B) . Let S be a duplicate set. $\forall S \subseteq A$, where S is a set of duplicates, we must have

$$\left(1 + \frac{\epsilon}{n}\right) \frac{d(A, B)}{|A||B|} \geq \frac{d(A \setminus S, B \cup S)}{(|A| - |S|)(|B| + |S|)}.$$

Pick up a point $q \in S$,

$$\begin{aligned}
& \left(1 + \frac{\epsilon}{n}\right)d(A, B)|S|(|A| - 1)(|B| + 1) \\
&= \left(1 + \frac{\epsilon}{n}\right)d(A, B)(|A||B| + |A| - |B| - 1)|S| \\
&= \left(1 + \frac{\epsilon}{n}\right)d(A, B)(|A||B| + |A||S| - |B||S| - |S|) + \left(1 + \frac{\epsilon}{n}\right)d(A, B)(|A||B|)(|S| - 1) \\
&\geq \left(1 + \frac{\epsilon}{n}\right)d(A, B)(|A| - |S|)(|B| + |S|) + d(A, B)|A||B|(|S| - 1) \\
&\geq |A||B|d(A \setminus S, B \cup S) + d(A, B)|A||B|(|S| - 1) \\
&= |A||B|(d(A, B) + |S|d(q, A) - |S|d(q, B)) + |A||B|(|S| - 1)d(A, B) \\
&= |A||B||S|(d(A, B) + d(q, A) - d(q, B)).
\end{aligned}$$

Rearrange the terms and we get the following inequality holds for any point $q \in A$:

$$\left(1 + \frac{\epsilon}{n}\right) \frac{d(A, B)}{|A||B|} \geq \frac{d(A, B) + d(q, A) - d(q, B)}{(|A| - 1)(|B| + 1)}.$$

The rest of the proof goes exactly the same as the proof in [19, Theorem 6.5]. \square

B Proof of Theorem 8

Proof. Let \mathcal{A} be the β -approximation algorithm to (1). For a given instance $G = (V, s)$, let $\mathcal{Y} = \{Y_1, Y_2, \dots\}$ be a fairlet decomposition of V ; let $m_f = \max_{Y \in \mathcal{Y}} |Y|$. Recall that $n = |V|$.

We use \mathcal{Y} to create a weighted instance $G_{\mathcal{Y}} = (\mathcal{Y}, s_{\mathcal{Y}}, m_{\mathcal{Y}})$. For $Y, Y' \in \mathcal{Y}$, we define $s(Y, Y') = \sum_{i \in Y, j \in Y'} s(i, j)$ and we define $m_{\mathcal{Y}}(Y) = |Y|$.

We run \mathcal{A} on $G_{\mathcal{Y}}$ and let $T_{\mathcal{Y}}$ be the hierarchical clustering obtained by \mathcal{A} . To extend this to a tree T on V , we simply place all the points in each fairlet as leaves under the corresponding vertex in $T_{\mathcal{Y}}$.

We argue that $\text{rev}_G(T) \geq \beta \left(1 - \frac{2m_f}{n}\right) (n - 2)s(V)$.

Since \mathcal{A} obtains a β -approximation to hierarchical clustering on $G_{\mathcal{Y}}$, we have $\text{rev}_{G_{\mathcal{Y}}}(T_{\mathcal{Y}}) \geq \beta \cdot \sum_{Y, Y' \in \mathcal{Y}} s(Y, Y')(n - m(Y) - m(Y'))$.

Notice the fact that, for any pair of points u, v in the same fairlet $Y \in \mathcal{Y}$, the revenue they get in the tree T is $(n - m(Y))s(u, v)$. Then using $\text{rev}_G(T) = \sum_{Y \in \mathcal{Y}} (n - m(Y))s(Y) + \text{rev}(T_{\mathcal{Y}})$,

$$\begin{aligned}
\text{rev}_G(T) &\geq \sum_{Y \in \mathcal{Y}} \beta(n - m(Y))s(Y) + \beta \sum_{Y, Y' \in \mathcal{Y}} s(Y, Y')(n - m(Y) - m(Y')) \\
&\geq \beta(n - 2m_f) \left(\sum_{Y \in \mathcal{Y}} s(Y) + \sum_{Y, Y' \in \mathcal{Y}} s(Y, Y') \right) \geq \beta \left(1 - \frac{2m_f}{n}\right) (n - 2)s(V).
\end{aligned}$$

Thus the resulting tree T is a $\beta \left(1 - \frac{2m_f}{n}\right)$ -approximation of the upper bound. \square

C Proofs for (ϵ/n) -locally-optimal local search algorithm

In this section, we prove that Algorithm 1 gives a good fairlet decomposition for the fairlet decomposition objective 3, and that it has polynomial run time.

C.1 Proof for a simplified version of Lemma 15

In Subsection C.2, we will prove Lemma 15. For now, we will consider a simpler version of Lemma 15 in the context of [16]'s disparate impact problem, where we have red and blue points and strive to preserve their ratios in all clusters. Chierichetti et al. [16] provided a valid fairlet decomposition in this context, where each fairlet has at most b blue points and r red points. Before going deeper into the analysis, we state the following useful proposition.

Proposition 22. Let $r_t = |\text{red}(V)|$ be the total number of red points and $b_t = |\text{blue}(V)|$ the number of blue points. We have that, $\max\{\frac{r}{r_t}, \frac{b}{b_t}\} \leq \frac{2(b+r)}{n}$.

Proof. Recall that $\text{balance}(V) = \frac{b_t}{r_t} \geq \frac{b}{r}$, and wlog $b_t \leq r_t$. Since the fractions are positive and $\frac{b_t}{r_t} \geq \frac{b}{r}$ we know that $\frac{b_t}{b_t+r_t} \geq \frac{b}{b+r}$. Since $b_t + r_t = n$ we conclude that $b_t \geq \frac{b}{b+r}n$. Similarly, we conclude that $\frac{r_t}{b_t+r_t} \leq \frac{r}{b+r}$. Therefore $r_t \leq \frac{r}{b+r}n$.

Thus, $\frac{r}{r_t} \geq \frac{b+r}{n} \geq \frac{b}{b_t}$. However, since $b_t \leq r_t$ and $b_t + r_t = n$, $r_t \geq \frac{1}{2}n$, $\frac{r}{r_t} \leq \frac{2r}{n} \leq \frac{2(b+r)}{n}$. \square

Using this, we can define and prove the following lemma, which is a simplified version of Lemma 15.

Lemma 23. The fairlet decomposition \mathcal{Y} computed by Algorithm 1 has an objective value for (3) of at most $(1 + \epsilon)\frac{2(b+r)}{n}d(V)$.

Proof. Let $Y : V \mapsto \mathcal{Y}$ denote a mapping from a point in V to the fairlet it belongs to. Let $d_R(X) = \sum_{u \in \text{red}(X)} d(u, X)$, and $d_B(X) = \sum_{v \in \text{blue}(X)} d(v, X)$. Naturally, $d_R(X) + d_B(X) = 2d(X)$ for any set X . For a fairlet $Y_i \in \mathcal{Y}$, let r_i and b_i denote the number of red and blue points in Y_i .

We first bound the total number of intra-fairlet pairs. Let $x_i = |Y_i|$, we know that $0 \leq x_i \leq b + r$ and $\sum_i x_i = n$. The number of intra-fairlet pairs is at most $\sum_i x_i^2 \leq \sum_i (b+r)x_i = (b+r)n$.

The **While** loop can end in two cases: 1) if \mathcal{Y} is (ϵ/n) -locally-optimal; 2) if $\sum_{Y_k \in \mathcal{Y}} d(Y_k) \leq \Delta$. Case 2 immediately implies the lemma, thus we focus on case 1. By definition of the algorithm, we know that for any pair $u \in Y(u)$ and $v \in Y(v)$ where u, v have the same color and $Y(u) \neq Y(v)$ the swap does not increase objective value by a large amount. (The same trivially holds if the pair are in the same cluster.)

$$\begin{aligned} \sum_{Y_k} d(Y_k) &\leq (1 + \frac{\epsilon}{n}) \left(\sum_{Y_k} d(Y_k) - d(u, Y(u)) - d(v, Y(v)) + d(u, Y(v)) + d(v, Y(u)) - 2d(u, v) \right) \\ &\leq (1 + \frac{\epsilon}{n}) \left(\sum_{Y_k} d(Y_k) - d(u, Y(u)) - d(v, Y(v)) + d(u, Y(v)) + d(v, Y(u)) \right). \end{aligned}$$

After moving terms and some simplification, we get the following inequality:

$$\begin{aligned} &d(u, Y(u)) + d(v, Y(v)) \\ &\leq d(u, Y(v)) + d(v, Y(u)) + \frac{\epsilon/n}{1 + \epsilon/n} \sum_{Y_k \in \mathcal{Y}} d(Y_k) \\ &\leq d(u, Y(v)) + d(v, Y(u)) + \frac{\epsilon}{n} \sum_{Y_k \in \mathcal{Y}} d(Y_k). \end{aligned} \tag{4}$$

Then we sum up (4), $d(u, Y(u)) + d(v, Y(v)) \leq d(u, Y(v)) + d(v, Y(u)) + \frac{\epsilon}{n} \sum_{Y_k \in \mathcal{Y}} d(Y_k)$, over every pair of points in $\text{red}(V)$ (even if they are in the same partition).

$$r_t \sum_{Y_i} d_R(Y_i) \leq \left(\sum_{Y_i} r_i d_R(Y_i) \right) + \left(\sum_{u \in \text{red}(V)} \sum_{Y_i \neq Y(u)} r_i d(u, Y_i) \right) + r_t^2 \frac{\epsilon}{n} \sum_{Y_i} d(Y_i).$$

Divide both sides by r_t and use the fact that $r_i \leq r$ for all Y_i :

$$\sum_{Y_i} d_R(Y_i) \leq \left(\sum_{Y_i} \frac{r}{r_t} d_R(Y_i) \right) + \left(\sum_{u \in \text{red}(V)} \sum_{Y_i \neq Y(u)} \frac{r}{r_t} d(u, Y_i) \right) + \frac{r_t \epsilon}{n} \sum_{Y_i} d(Y_i). \tag{5}$$

For pairs of points in $\text{blue}(V)$ we sum (4) to similarly obtain:

$$\sum_{Y_i} d_B(Y_i) \leq \left(\sum_{Y_i} \frac{b}{b_t} d_B(Y_i) \right) + \left(\sum_{v \in \text{blue}(V)} \sum_{Y_i \neq Y(v)} \frac{b}{b_t} d(v, Y_i) \right) + \frac{b_t \epsilon}{n} \sum_{Y_i} d(Y_i). \tag{6}$$

Now we sum up (5) and (6). The LHS becomes:

$$\sum_{Y_i} (d_R(Y_i) + d_B(Y_i)) = \sum_{Y_i} \sum_{u \in Y_i} d(u, Y_i) = 2 \sum_{Y_i} d(Y_i).$$

For the RHS, the last term in (5) and (6) is $\frac{\epsilon(b_t+r_t)}{n} \sum_{Y_i} d(Y_i) = \epsilon \sum_{Y_i} d(Y_i)$.

The other terms give:

$$\begin{aligned} & \frac{r}{r_t} \sum_{Y_i} d_R(Y_i) + \frac{r}{r_t} \sum_{u \in \text{red}(V)} \sum_{Y_i \neq Y(u)} d(u, Y_i) + \frac{b}{b_t} \sum_{Y_i} d_B(Y_i) + \frac{b}{b_t} \sum_{v \in \text{blue}(V)} \sum_{Y_i \neq Y(v)} d(v, Y_i) \\ & \leq \max\left\{\frac{r}{r_t}, \frac{b}{b_t}\right\} \left\{ \sum_{Y_i} (d_R(Y_i) + d_B(Y_i)) + \sum_{u \in V} \sum_{Y_i \neq Y(u)} d(u, Y_i) \right\} \\ & = \max\left\{\frac{r}{r_t}, \frac{b}{b_t}\right\} \left\{ \sum_{Y_i} \sum_{u \in Y_i} d(u, Y_i) + \sum_{Y_i} \sum_{Y_j \neq Y_i} d(Y_i, Y_j) \right\} \\ & = 2 \max\left\{\frac{r}{r_t}, \frac{b}{b_t}\right\} d(V) \\ & \leq \frac{4(b+r)}{n} d(V). \end{aligned}$$

The last inequality follows from Proposition 22. All together, this proves that

$$2 \sum_{Y_k} d(Y_k) \leq \frac{4(b+r)}{n} d(V) + \epsilon \sum_{Y_k} d(Y_k).$$

Then, $\frac{\sum_{Y_k} d(Y_k)}{d(V)} \leq \frac{2(b+r)}{n} \cdot \frac{1}{1-\epsilon/2} \leq (1+\epsilon) \frac{2(b+r)}{n}$. The final step follows from the fact that $(1+\epsilon)(1-\epsilon/2) = 1 + \frac{\epsilon}{2}(1-\epsilon) \geq 1$. This proves the lemma. \square

C.2 Proof for the generalized Lemma 15

Next, we prove Lemma 15 for the more generalized definition of fairness, which is α -capped fairness.

Proof of [Lemma 15] The proof follows the same logic as in the two-color case: we first use the (ϵ/n) -local optimality of the solution, and sum up the inequality over all pairs of points with the same color.

Let $Y : V \mapsto \mathcal{Y}$ denote a mapping from a point in V to the fairlet it belongs to. Let $R_i(X)$ be the set of R_i colored points in a set X . Let $d_{R_i}(X) = \sum_{u \in R_i(X)} d(u, X)$. Naturally, $\sum_i d_{R_i}(X) = 2d(X)$ for any set X since the weight for every pair of points is repeated twice.

The **While** loop can end in two cases: 1) if \mathcal{Y} is (ϵ/n) -locally-optimal; 2) if $\sum_{Y_k \in \mathcal{Y}} d(Y_k) \leq \Delta$. Case 2 immediately implies the lemma, thus we focus on case 1.

By definition of the algorithm, we know that for any pair $u \in Y(u)$ and $v \in Y(v)$ where u, v have the same color and $Y(u) \neq Y(v)$ the swap does not increase objective value by a large amount. (The same trivially holds if the pair are in the same cluster.) We get the following inequality as in the two color case:

$$d(u, Y(u)) + d(v, Y(v)) \leq d(u, Y(v)) + d(v, Y(u)) + \frac{\epsilon}{n} \sum_{Y_k \in \mathcal{Y}} d(Y_k). \quad (7)$$

For any color R_i , we sum it over every pair of points in $R_i(V)$ (even if they are in the same partition).

$$n_i \sum_{Y_k} d_{R_i}(Y_k) \leq \left(\sum_{Y_k} r_{ik} d_{R_i}(Y_k) \right) + \left(\sum_{u \in R_i(V)} \sum_{Y_k \neq Y(u)} r_{ik} d(u, Y_k) \right) + n_i^2 \frac{\epsilon}{n} \sum_{Y_k} d(Y_k).$$

Divide both sides by n_i and we get:

$$\sum_{Y_k} d_{R_i}(Y_k) \leq \left(\sum_{Y_k} \frac{r_{ik}}{n_i} d_{R_i}(Y_k) \right) + \left(\sum_{u \in R_i(V)} \sum_{Y_k \neq Y(u)} \frac{r_{ik}}{n_i} d(u, Y_k) \right) + \frac{n_i \epsilon}{n} \sum_{Y_k} d(Y_k). \quad (8)$$

Now we sum up this inequality over all colors R_i . The LHS becomes:

$$\sum_{Y_k} \sum_i d_{R_i}(Y_k) = \sum_{Y_k} \sum_{u \in Y_k} d(u, Y_k) = 2 \sum_{Y_k} d(Y_k).$$

For the RHS, the last term sums up to $\frac{\epsilon(\sum_i n_i)}{n} \sum_{Y_k} d(Y_k) = \epsilon \sum_{Y_k} d(Y_k)$. Using the fact that $\frac{r_{ik}}{n_i} \leq \max_{i,k} \frac{r_{ik}}{n_i}$, the other terms sum up to :

$$\begin{aligned} & \sum_i \sum_{Y_k} \frac{r_{ik}}{n_i} d_{R_i}(Y_k) + \sum_i \sum_{u \in R_i(V)} \sum_{Y_k \neq Y(u)} \frac{r_{ik}}{n_i} d(u, Y_k) \\ & \leq \max_{i,k} \frac{r_{ik}}{n_i} \left\{ \sum_{Y_k} \sum_i d_{R_i}(Y_i) + \sum_{u \in V} \sum_{Y_k \neq Y(u)} d(u, Y_k) \right\} \\ & = \max_{i,k} \frac{r_{ik}}{n_i} \left\{ \sum_{Y_k} \sum_{u \in Y_k} d(u, Y_k) + \sum_{Y_k} \sum_{Y_j \neq Y_k} d(Y_j, Y_k) \right\} \\ & = 2 \max_{i,k} \frac{r_{ik}}{n_i} \cdot d(V). \end{aligned}$$

Therefore, putting LHS and RHS together, we get

$$2 \sum_{Y_k} d(Y_k) \leq 2 \max_{i,k} \frac{r_{ik}}{n_i} d(V) + \epsilon \sum_{Y_k} d(Y_k).$$

Then, $\frac{\sum_{Y_k} d(Y_k)}{d(V)} \leq \max_{i,k} \frac{r_{ik}}{n_i} \cdot \frac{1}{1-\epsilon/2} \leq (1+\epsilon) \cdot \max_{i,k} \frac{r_{ik}}{n_i}$. The final step follows from the fact that $(1+\epsilon)(1-\epsilon/2) = 1 + \frac{\epsilon}{2}(1-\epsilon) \geq 1$. \square

In the two-color case, the ratio $\max_{i,k} \frac{r_{ik}}{n_i}$ becomes $\max\{\frac{r}{r_i}, \frac{b}{b_i}\}$, which can be further bounded by $\frac{2(b+r)}{n}$ (see Proposition 22). If there exists a caplet decomposition such that $\max_{i,k} \frac{r_{ik}}{n_i} = o(1)$, Lemma 15 implies we can build a fair hierarchical clustering tree with $o(1)$ loss in approximation ratio for value objective.

Assuming for all color class R_i , $n_i \rightarrow +\infty$ as $n \rightarrow +\infty$, here we give a possible caplet decomposition for $\alpha = \frac{1}{t}$ ($t \leq c$) with size $O(t)$ for positive integer t , thus guaranteeing $\max_{i,k} \frac{r_{ik}}{n_i} = o(1)$ for any i .

Lemma 24. *For any set P of size 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 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, P_2, \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

C.3 Proof for the running time of (ϵ/n) -locally-optimal fairlet decomposition algorithm

Proof of [Lemma 16] Notice that finding the maximum pairwise distance takes $O(n^2)$ time. Thus, we focus on analyzing the time spent on the **While** loop.

Let t be the total number of swaps. We argue that $t = \tilde{O}(n/\epsilon)$. If $t = 0$ the conclusion trivially holds. Otherwise, consider the decomposition \mathcal{Y}_{t-1} before the last swap. Since the **While** loop does not terminate here, $\sum_{Y_k \in \mathcal{Y}_{t-1}} d(Y_k) \geq \Delta = \frac{b+r}{n} d_{max}$. However, at the beginning, we have $\sum_{Y_k \in \mathcal{Y}} d(Y_k) \leq (b+r)n \cdot d_{max} = n^2 \Delta \leq n^2 \sum_{Y_k \in \mathcal{Y}_{t-1}} d(Y_k)$. Therefore, it takes at most $\log_{1+\epsilon/n}(n^2) = \tilde{O}(n/\epsilon)$ iterations to finish the **While** loop.

It remains to discuss the running time of each iteration. We argue that there is a way to finish each iteration in $O(n^2)$ time. Before the **While** loop, keep a record of $d(u, Y_i)$ for each point u and each fairlet Y_i . This takes $O(n^2)$ time. If we know $d(u, Y_i)$ and the objective value from the last iteration, in the current iteration, it takes $O(1)$ time to calculate the new objective value after each swap (u, v) , and there are at most n^2 such calculations, before the algorithm either finds a pair to swap, or determines that no such pair is left. After the swap, the update for all the $d(u, Y_i)$ data takes $O(n)$ time. In total, every iteration takes $O(n^2)$ time.

Therefore, Algorithm 1 takes $\tilde{O}(n^3/\epsilon)$ time. \square

D Hardness of optimal fairlet decomposition

Before proving Theorem 7, we state that the PARTITION INTO TRIANGLES (PIT) problem is known to belong to the NP-complete class [23], defined as follows. In the definition, we call a clique k -clique if it has k nodes. A triangle is a 3-clique.

Definition 25. PARTITION INTO TRIANGLES

(PIT). Given graph $G = (V, E)$, where $V = 3n$, determine if V can be partitioned into 3-element sets S_1, S_2, \dots, S_n , such that each S_i forms a triangle in G .

The NP-hardness of PIT problem gives us a more general statement.

Definition 26. PARTITION INTO k -CLIQUES

(PIKC). For a fixed number k treated as constant, given graph $G = (V, E)$, where $V = kn$, determine if V can be partitioned into k -element sets S_1, S_2, \dots, S_n , such that each S_i forms a k -clique in G .

Lemma 27. For a fixed constant $k \geq 3$, the PIKC problem is NP-hard.

Proof. We reduce the PIKC problem from the PIT problem. For any graph $G = (V, E)$ given to the PIT problem where $|V| = 3n$, construct another graph $G' = (V', E')$. Let $V' = V \cup C_1 \cup C_2 \cup \dots \cup C_n$, where all the C_i 's are $(k-3)$ -cliques, and there is no edge between any two cliques C_i and C_j where $i \neq j$. For any C_i , let all points in C_i to be connected to all nodes in V .

Now let G' be the input to PIKC problem. We prove that G can be partitioned into triangles if and only if G' can be partitioned into k -cliques. If V has a triangle partition $V = \{S_1, \dots, S_n\}$, then $V' = \{S_1 \cup C_1, \dots, S_n \cup C_n\}$ is a k -clique partition. On the other hand, if V' has a k -clique partition $V' = \{S'_1, \dots, S'_n\}$ then C_1, \dots, C_n must each belong to different k -cliques since they are not connected to each other. Without loss of generality we assume $C_i \subseteq S'_i$, then $V = \{S'_1 \setminus C_1, \dots, S'_n \setminus C_n\}$ is a triangle partition. \square

We are ready to prove the theorem.

Proof of [Theorem 7] We prove Theorem 7 by proving that for given $z \geq 4$, if there exists a c -approximation polynomial algorithm \mathcal{A} for (3), it can be used to solve the PIKC problem where $k = z - 1$ for any instance as well. This holds for any finite c .

Given any graph $G = (V, E)$ that is input to the PIKC problem, where $|V| = kn = (z-1)n$, let a set V' with distances be constructed in the following way:

1. $V' = V \cup \{C_1, \dots, C_n\}$, where each C_i is a singleton.
2. Color the points in V red, and color all the C_i 's blue.
3. For a $e = (u, v)$, let $d(u, v) = 0$, if it satisfies one of the three conditions: 1) $e \in E$. 2) $u, v \in C_i$ for some C_i . 3) one of u, v is in V , while the other belong to some C_i .
4. All other edges have distance 1.

Obviously the blue points make up a $1/z$ fraction of the input so each fairlet should have exactly 1 blue point and $z-1$ red points.

We claim that G has a k -clique partition if and only if algorithm \mathcal{A} gives a solution of 0 for (3). The same argument as in the proof of Lemma 27 will show that G has a k -clique partition if and only if the optimal solution to (3) is 0. This is equal to algorithm \mathcal{A} giving a solution of 0 since otherwise the approximate is not bounded. \square

E Optimizing cost with fairness

In this section, we present our fair hierarchical clustering algorithm that approximates Dasgupta’s cost function and satisfies Theorem 17. Most of the proofs can be found in Section E.1. We consider the problem of equal representation, where vertices are red or blue and $\alpha = 1/2$. From now on, whenever we use the word “fair”, we are referring to this fairness constraint. Our algorithm also uses parameters t and ℓ such that $n \geq t\ell$ and $t > \ell + 108t^2/\ell^2$ for $n = |V|$, and leverages a β -approximation for cost and γ_t -approximation for minimum weighted bisection. We will assume these are fixed and use them throughout the section.

We will ultimately show that we can find a fair solution that is a sublinear approximation for the unfair optimum T_{unfair}^* , which is a lower bound of the fair optimum. Our main result is Theorem 17, which is stated in the body of the paper.

The current best approximations described in Theorem 17 are $\gamma_t = O(\log^{3/2} n)$ by [22] and $\beta = \sqrt{\log n}$ by both [20] and [12]. If we set $t = \sqrt{n}(\log^{3/4} n)$ and $\ell = n^{1/3}\sqrt{\log n}$, then we get Corollary 28.

Corollary 28. *Consider the equal representation problem with two colors. There is an $O\left(n^{5/6} \log^{5/4} n\right)$ -approximate fair clustering under the cost objective.*

The algorithm will be centered around a single clustering, which we call \mathcal{C} , that is extracted from an unfair hierarchy. We will then adapt this to become a similar, fair clustering \mathcal{C}' . To formalize what \mathcal{C}' must satisfy to be sufficiently “similar” to \mathcal{C} , we introduce the notion of a \mathcal{C} -good clustering. Note that this is not an intuitive set of properties, it is simply what \mathcal{C}' must satisfy in order

Definition 29 (Good clustering). *Fix a clustering \mathcal{C} whose cluster sizes are at most t . A fair clustering \mathcal{C}' is \mathcal{C} -good if it satisfies the following two properties:*

1. *For any cluster $C \in \mathcal{C}$, there is a cluster $C' \in \mathcal{C}'$ such that all but (at most) an $O(\ell\gamma_t/t + t\gamma_t/\ell^2)$ -fraction of the weight of edges in C is also in C' .*
2. *Any $C' \in \mathcal{C}'$ is not too much bigger, so $|C'| \leq 6t\ell$.*

The hierarchy will consist of a \mathcal{C} -good (for a specifically chosen \mathcal{C}) clustering \mathcal{C}' as its only nontrivial layer.

Lemma 30. *Let T be a β -approximation for cost and \mathcal{C} be a maximal clustering in T under the condition that all cluster sizes are at most t . Then, a fair two-tiered hierarchy T' whose first level consists of a \mathcal{C} -good clustering achieves an $O\left(\frac{n}{t} + t\ell + \frac{n\ell\gamma_t}{t} + \frac{nt\gamma_t}{\ell^2}\right)$ β -approximation for cost.*

Proof. Since T is a β -approximation, we know that:

$$\text{cost}(T) \leq \beta \text{cost}(T_{\text{unfair}}^*)$$

We will then utilize a scheme to account for the cost contributed by each edge relative to their cost in T in the hopes of extending it to T_{unfair}^* . There are three different types of edges:

1. An edge e that is merged into a cluster of size t or greater in T , thus contributing $t \cdot s(e)$ to the cost. At worst, this edge is merged in the top cluster in T' to contribute $n \cdot s(e)$. Thus, the factor increase in the cost contributed by e is $O(n/t)$. Then since the total contribution of all such edges in T is at most $\text{cost}(T)$, the total contribution of all such edges in T' is at most $O(n/t) \cdot \text{cost}(T)$.
2. An edge e that started in some cluster $C \in \mathcal{C}$ that does not remain in the corresponding cluster C' . We are given that the total weight removed from any such C is an $O(\ell\gamma_t/t + t\gamma_t/\ell^2)$ -fraction of the weight contained in C . If we sum across the weight in all clusters in \mathcal{C} , that is at most $\text{cost}(T)$. So the total amount of weight moved is at most $O(\ell\gamma_t/t + t\gamma_t/\ell^2) \cdot \text{cost}(T)$. These edges contributed at least $2s(e)$ in T as the smallest possible cluster size is two. In T' , these may have been merged at the top of the cluster, for a maximum cost contribution of $n \cdot s(e)$. Therefore, the total cost across all such edges is increased by at most a factor of $n/2$, which gives a total cost of at most $O(n\ell\gamma_t/t + nt\gamma_t/\ell^2) \cdot \text{cost}(T)$.
3. An edge e that starts in some cluster $C \in \mathcal{C}$ and remains in the corresponding $C' \in \mathcal{C}'$. Similarly, this must have contributed at least $2s(e)$ in T , but now we know that this edge is

merged within C' in T' , and that the size of C' is $|C'| \leq 6t\ell$. Thus its contribution increases at most by a factor of $3t\ell$. By the same reasoning from the first edge type we discussed, all these edges total contribute at most a factor of $O(t\ell) \cdot \text{cost}(T)$.

We can then put a conservative bound by putting this all together.

$$\text{cost}(T') \leq O\left(\frac{n}{t} + t\ell + \frac{n\ell\gamma_t}{t} + \frac{nt\gamma_t}{\ell^2}\right) \text{cost}(T).$$

Finally, we know T is a β -approximation for T_{unfair}^* .

$$\begin{aligned} \text{cost}(T') &\leq O\left(\frac{n}{t} + t\ell + \frac{n\ell\gamma_t}{t} + \frac{nt\gamma_t}{\ell^2}\right) \\ &\quad \cdot \beta \cdot \text{cost}(T_{\text{unfair}}^*). \square \end{aligned}$$

With this proof, the only thing left to do is find a \mathcal{C} -good clustering C' (Definition 29). Specifically, using the clustering \mathcal{C} mentioned in Lemma 30, we would like to find a \mathcal{C} -good clustering C' using the following.

Lemma 31. *There is an algorithm that, given a clustering \mathcal{C} with maximum cluster size t , creates a \mathcal{C} -good clustering.*

The proof is deferred to the Section E.1. With these two Lemmas, we can prove Theorem 17.

Proof. Consider our graph G . We first obtain a β -approximation for unfair cost, which yields a hierarchy tree T . Let \mathcal{C} be the maximal clustering in T under the constraint that the cluster sizes must not exceed t . We then apply the algorithm from Lemma 31 to get a \mathcal{C} -good clustering C' . Construct T' such that it has one layer that is C' . Then we can apply the results from Lemma 30 to get the desired approximation. \square

From here, we will only provide a high-level description of the algorithm for Lemma 31. For precise details and proofs, see Section E.1. To start, we need to propose some terminology.

Definition 32 (Red-blue matching). *A **red-blue matching** on a graph G is a matching M such that $M(u) = v$ implies u and v are different colors.*

Red-blue matchings are interesting because they help us ensure fairness. For instance, suppose M is a red-blue matching that is also perfect (i.e., touches all nodes). If the lowest level of a hierarchy consists of a clustering such that v and $M(v)$ are in the same cluster for all v , then that level of the hierarchy is fair since there is a bijection between red and blue vertices within each cluster. When these clusters are merged up in the hierarchy, fairness is preserved.

Our algorithm will modify an unfair clustering to be fair by combining clusters and moving a small number of vertices. To do this, we will use the following notion.

Definition 33 (Red-blue clustering graph). *Given a graph G and a clustering $\mathcal{C} = \{C_1, \dots, C_k\}$, we can construct a **red-blue clustering graph** $H_M = (V_M, E_M)$ that is associated with some red-blue matching M . Then H_M is a graph where $V_M = \mathcal{C}$ and $(C_i, C_j) \in E_M$ if and only if there is a $v_i \in C_i$ and $M(v_i) = v_j \in C_j$.*

Basically, we create a graph of clusters, and there is an edge between two clusters if and only if there is at least one vertex in one cluster that is matched to some vertex in the other cluster. We now show that the red-blue clustering graph can be used to construct a fair clustering based on an unfair clustering.

Proposition 34. *Let H_M be a red-blue clustering graph on a clustering \mathcal{C} with a perfect red-blue matching M . Let C' be constructed by merging all the clusters in each component of H_M . Then C' is fair.*

Proof. Consider some $C \in \mathcal{C}'$. By construction, this must correspond to a connected component in H_M . By definition of H_M , for any vertex $v \in C$, $M(v) \in C$. That means M , restricted to C , defines a bijection between the red and blue nodes in C . Therefore, C has an equal number of red and blue vertices and hence is fair. \square

We will start by extracting a clustering \mathcal{C} from an unfair hierarchy T that approximates cost. Then, we will construct a red-blue clustering graph H_M with a perfect red-blue matching M . Then we can use the components of H_M to define our first version of the clustering \mathcal{C}' . However, this requires a non-trivial way of moving vertices between clusters in \mathcal{C} .

We now give an overview of our algorithm in Steps (A)–(G). For a full description, see our pseudocode in Section H.

(A) **Get an unfair approximation T .** We start by running a β -approximation for cost in the unfair setting. This gives us a tree T such that $\text{cost}(T) \leq \beta \cdot \text{cost}(T_{\text{unfair}}^*)$.

(B) **Extract a t -maximal clustering.** Given T , we find the maximal clustering \mathcal{C} such that (i) every cluster in the clustering is of size at most t , and (ii) any cluster above these clusters in T is of size more than t .

(C) **Combine clusters to be size t to $3t$.** We will now slowly change \mathcal{C} into \mathcal{C}' during a number of steps. In the first step, we simply define \mathcal{C}_0 by merging small clusters $|C| \leq t$ until the merged size is between t and $3t$. Thus clusters in \mathcal{C} are contained within clusters in \mathcal{C}_0 , and all clusters are between size t and $3t$.

(D) **Find cluster excesses.** Next, we strive to make our clustering more fair. We do this by trying to find an underlying matching between red and blue vertices that agrees with \mathcal{C}_0 (matches are in the same cluster). If the matching were perfect, then the clusters in \mathcal{C}_0 would have equal red and blue representation. However, this is not guaranteed initially. We start by conceptually matching as many red and blue vertices within clusters as we can. Note we do not actually create this matching; we just want to reserve the space for this matching to ensure fairness, but really some of these vertices may be moved later on. Then the remaining unmatched vertices in each cluster is either entirely red or entirely blue. We call this amount the *excess* and the color the *excess color*. We label each cluster with both of these.

(E) **Construct red-blue clustering graph.** Next, we would like to construct $H_M = (V_M, E_M)$, our red-blue clustering graph on \mathcal{C}_0 . Let $V_M = \mathcal{C}_0$. In addition, for the within-cluster matchings mentioned in Step (D), let those matches be contained in M . With this start, we will do a matching process to simultaneously construct E_M and the rest of M . Note the unmatched vertices are specifically the excess vertices in each cluster. We will match these with an iterative process given our parameter ℓ :

1. Select a vertex $C_i \in V_M$ with excess at least ℓ to start a new connected component in H_M . Without loss of generality, say its excess color is red.
2. Find a vertex $C_j \in V_M$ whose excess color is blue and whose excess is at least ℓ . Add (C_i, C_j) to E_M .
3. Say without loss of generality that the excess of C_i is less than that of C_j . Then match all the excess in C_i to vertices in the excess of C_j . Now C_j has a smaller excess.
4. If C_j has an excess less than ℓ or C_j is the ℓ th cluster in this component, end this component. Start over at (1) with a new cluster.
5. Otherwise, use C_j as our reference and continue constructing this component at (2).
6. Complete when there are no more clusters with over ℓ excess that are not in a component (or all remaining such clusters have the same excess color).

We would like to construct \mathcal{C}' by merging all clusters in each component. This would be fair if M were a perfect matching, however this is not true yet. In the next step, we handle this.

(F) **Fix unmatched vertices.** We now want to match excess vertices that are unmatched. We do this by bringing vertices from other clusters into the clusters that have unmatched excess, starting with all small unmatched excess. Note that some clusters were never used in Step (E) because they had small excess to start. This means they had many internal red-blue matches. Remove t^2/ℓ^2 of these and put them into clusters in need. For other vertices, we will later describe a process where t/ℓ of the clusters can contribute $108t^2/\ell^2$ vertices to account for unmatched excess. Thus clusters lose at most $108t^2/\ell^2$ vertices, and we account for all unmatched vertices. Call the new clustering \mathcal{C}_1 . Now M is perfect and H_M is unchanged.

(G) **Define \mathcal{C}' .** Finally, we create the clustering \mathcal{C}' by merging the clusters in each component of H_M . Note that Proposition 34 assures \mathcal{C}' is fair. In addition, we will show that cluster sizes in \mathcal{C}_1

are at most $6t$, so \mathcal{C}' has the desired upper bound of $6t\ell$ on cluster size. Finally, we removed at most $\ell + t^2/\ell^2$ vertices from each cluster. This is the desired \mathcal{C} -good clustering.

Further details and the proofs that the above sequence of steps achieve the desired approximation can be found in the next section. While the approximation factor obtained is not as strong as the ones for revenue or value objectives with fairness, we believe cost is a much harder objective with fairness constraints.

E.1 Proof of Theorem 17

This algorithm contains a number of components. We will discuss the claims made by the description step by step. In Step (A), we simply utilize any β -approximation for the unfair approximation. Step (B) is also quite simple. At this point, all that is left is to show how to find \mathcal{C}' , ie, prove Lemma 31 (introduced in Section 6). This occurs in the steps following Step (B). In Step (C), we apply our first changes to the starting clustering from T . We now prove that the cluster sizes can be enforced to be between t and $3t$.

Lemma 35. *Given a clustering \mathcal{C} , we can construct a clustering \mathcal{C}_0 , where each $C \in \mathcal{C}_0$ is a union of clusters in \mathcal{C} and $t \leq |C| < 3t$.*

Proof. We iterate over all clusters in \mathcal{C} whose size are less than t and continually merge them until we create a cluster of size $\geq t$. Note that since the last two clusters we merged were of size $< t$, this cluster is of size $t \leq |C| < 2t$. We then stop this cluster and continue merging the rest of the clusters. At the end, if we are left with a single cluster of size $< t$, we simply merge this with any other cluster, which will then be of size $t \leq |C| < 3t$. \square

Step (D) describes a rather simple process. All we have to do in each cluster is count the amount of each color in each cluster, find which is more, and also compute the difference. No claims are made here.

Step (E) defines a more careful process. We describe this process and its results here.

Lemma 36. *There is an algorithm that, given a clustering \mathcal{C}_0 with $t \leq |C| \leq 3t$ for $C \in \mathcal{C}_0$, can construct a red-blue clustering graph $H_M = (V_M, E_M)$ on \mathcal{C}_0 with underlying matching M such that:*

1. H_M is a forest, and its max component size is ℓ .
2. For every $(C_i, C_j) \in E_M$, there are at least ℓ matches between C_i and C_j in M . In other words, $|M(C_i \cap C_j)| \geq \ell$.
3. For most $C_i \in V_M$, at most ℓ vertices in C_i are unmatched in M . The only exceptions to this rule are (1) exactly one cluster in every ℓ -sized component in H_M , and (2) at most $n/2$ additional clusters.

Proof. We use precisely the process from Step 5. Let $V_M = \mathcal{C}_0$. H_M will look like a bipartite graph with some entirely isolated nodes. We then try to construct components of H_M one-by-one such that (1) the max component size is ℓ , and (2) edges represent at least ℓ matches in M .

Let us show it satisfies the three conditions of the lemma. For condition 1, note that we will always halt component construction once it reaches size ℓ . Thus no component can exceed size ℓ . In addition, for every edge added to the graph, at least one of its endpoints now has small excess and will not be considered later in the program. Thus no cycles can be created, so it is a forest.

For condition 2, consider the construction of any edge $(C_i, C_j) \in E_M$. At this point, we only consider C_i and C_j to be clusters with different-color excess of at least ℓ each. In the next part of the algorithm, we match as much excess as we can between the two clusters. Therefore, there must be at least ℓ underlying matches.

Finally, condition 3 will be achieved by the completion condition. By the completion condition, there are no isolated vertices (besides possibly those leftover of the same excess color) that have over ℓ excess. Whenever we add a cluster to a component, either that cluster matches all of its excess, or the cluster it becomes adjacent to matches all of its excess. Therefore at any time, any component

has at most one cluster with any excess at all. If the component is smaller than ℓ (and is not the final component), then that can only happen when in the final addition, both clusters end up with less than ℓ excess. Therefore, no cluster in this component can have less than ℓ excess. For an ℓ -sized component, by the rule mentioned before, only one cluster can remain with ℓ excess. When the algorithm completes, we are left with a number of large-excess clusters with the same excess color, say red without loss of generality. Assume for contradiction there are more than $n/2$ such clusters, and so there is at least $n\ell/2$. Since we started with half red and half blue vertices, the remaining excess in the rest of the clusters must match up with the large red excess. Thus the remaining at most $n/2$ clusters must have at least $n\ell/2$ blue excess, but this is only achievable if they have large excess left. This is a contradiction. Thus we satisfy condition 3. \square

This concludes Step (E). In Step (F), we will transform the underlying clustering \mathcal{C}_0 such that we can achieve a perfect matching M . This will require removing a small number of vertices from some clusters in \mathcal{C}_0 and putting them in clusters that have unmatched vertices. This process will at most double cluster size.

Lemma 37. *There is an algorithm that, given a clustering \mathcal{C}_0 with $t \leq |C| \leq 3t$ for $C \in \mathcal{C}_0$, finds a clustering \mathcal{C}_1 and an underlying matching M' such that:*

1. *There is a bijection between \mathcal{C}_0 and \mathcal{C}_1 .*
2. *For any cluster $C_0 \in \mathcal{C}_0$ and its corresponding $C_1 \in \mathcal{C}_1$, $|C_0| - |C_1| \leq \ell + 108t^2/\ell^2$. This means that at most ℓ vertices are removed from C_0 in the construction of C_1 .*
3. *For all $C_1 \in \mathcal{C}_1$, $t - \ell - 108t^2/\ell^2 \leq |C_1| \leq 6t$.*
4. *M' is a perfect red-blue matching.*
5. *H_M is a red-blue clustering graph of \mathcal{C}_1 with matching M' , perhaps with additional edges.*

Proof. Use Lemma 36 to find the red-blue clustering graph H_M and its corresponding graph M . Then we know that only one cluster in every ℓ -sized component plus one other cluster can have a larger than ℓ excess. Since cluster sizes are at least t , $|V_M| \geq n/t$. This means that at most $n/(t\ell) + 1 = (n + t\ell)/(t\ell) \leq 2n/(t\ell)$ clusters need more than ℓ vertices. Since the excess is upper bounded by cluster size which is upper bounded by $3t$, this is at most $6n/\ell$ vertices in large excess that need matches.

We will start by removing all small excess vertices from clusters. This removes at most ℓ from any cluster. These vertices will then be placed in clusters with large excess of the right color. If we run out of large excess of the right color that needs matches, since the total amount of red and blue vertices is balanced, that means we can instead transfer the unmatched small excess red vertices to clusters with a small amount of unmatched blue vertices. In either case, this accounts for all the small unmatched excess. Now all we need to account for is at most $6n/\ell$ unmatched vertices in large excess clusters. At this point, note that the large excess should be balanced between red and blue. From now on, we will remove matches from within and between clusters to contribute to this excess. Since this always contributes the same amount of red and blue vertices by breaking matches, we do not have to worry about the balance of colors. We will describe how to distribute these contributions across a large number of clusters.

Consider vertices that correspond to clusters that (ignoring the matching M) started out with at most ℓ excess. So the non-excess portion, which is at least size $t - \ell$, is entirely matched with itself. We will simply remove t^2/ℓ^2 of these matches to contribute.

Otherwise, we will consider vertices that started out with large excess. We must devise a clever way to break matches without breaking too many incident upon a single cluster. For every tree in H_M (since H_M is a forest by Lemma 36), start at the root, and do a breadth-first search over all internal vertices. At any vertex we visit, break ℓ matches between it and its child (recall by Lemma 36 that each edge in H_M represents at least ℓ inter-cluster matches). Thus, each break contributes 2ℓ vertices. We do this for every internal vertex. Since an edge represents at least ℓ matches and the max cluster size is at most $3t$, any vertex can have at most $3t/\ell$ children. Thus the fraction of vertices in H_M that correspond to a contribution of 2ℓ vertices is at least $\ell/(3t)$.

Clearly, the worst case is when all vertices in H_M have large excess, as this means that fewer clusters are ensured to be able to contribute. By Lemma 36, at least $n/2$ of these are a part of completed connected components (ie, of size ℓ or with each cluster having small remaining excess). So consider this case. Since $|V_M| \geq n/(3t)$, then this process yields $n\ell^2/(18t^2)$ vertices. To achieve $6n/\ell$ vertices, we must then run $108t^2/\ell^3$ iterations. If an edge no longer represents ℓ matches because of an earlier iteration, consider it a non-edge for the rest of the process. The only thing left to consider is if a cluster C becomes isolated in H_M during the process. We know C began with at least t vertices, and at most ℓ were removed by removing small excess. So as long as $t > \ell + 108t^2/\ell^2$, we can remove the rest of the $108t^2/\ell^2$ vertices from the non-excess in C (the rest must be non-excess) in the same way as vertices that were isolated in H_M to start. Thus, we can account for the entire set of unmatched vertices without removing more than $108t^2/\ell^2$ vertices from any given cluster.

Now we consider the conditions. Condition 1 is obviously satisfied because we are just modifying clusters in C_0 , not removing them. The second condition is true because of our careful accounting scheme where we only remove $\ell + 108t^2/\ell^2$ vertices per cluster. The same is true for the lower bound in condition 3. When we add them to new clusters, since we only add a vertex to match an unmatched vertex, we at most double cluster size. So the max cluster size is $6t$.

For the fourth condition, note that we explicitly executed this process until all unmatched vertices became matched, and any endpoint in a match was used to create a new match. Thus the new matching, which we call M' , is perfect. It is still red-blue. Finally, note we did not create any matches between clusters. Therefore, no match in M' can violate H_M . Thus condition 5 is met. \square

Finally, we construct our final clustering in Step (G). However, to satisfy the qualities of Lemma 30, we must first argue about the weight loss from each cluster.

Lemma 38. *Consider any clustering \mathcal{C} with cluster sizes between t and $6t$. Say each cluster has a specified r number of red vertices to remove and b number of blue vertices to remove such that $r + b \leq x$ for some x , and r (resp. b) is nonzero only if the number of red (resp. blue) vertices in the cluster is $O(n)$. Then we can remove the desired number of each color while removing at most an $O((x/t)\gamma_t)$ of the weight originally contained within the cluster.*

Proof. Consider some cluster C with parameters r and b . We will focus first on removing red vertices. Let C_r be the red vertex set in C . We create a graph K corresponding to this cluster as follows. Let b_0 be a vertex representing all blue vertices from C , b'_0 be the “complement” vertex to b_0 , and R be a set of vertices r_i corresponding to all red vertices in C . We also add a set of $2r - |C_r| + 2X$ dummy vertices (where X is just some large value that makes it so $2r - |C_r| + X > 0$). $2r - |C_r| + X$ of the dummy vertices will be connected to b_0 with infinite edge weight (denote these δ_i), the other X will be connected to b'_0 with infinite edge weight (denote these δ'_i). This will ensure that b_0 and b'_0 are in the same partitions as their corresponding dummies. Let s_G and s_K be the similarity function in the original graph and new graph respectively.

$$\begin{aligned} s_K(b_0, \delta_i) &= \infty \\ s_K(b'_0, \delta'_i) &= \infty \end{aligned}$$

The blue vertex b_0 is also connected to all r_i with the following weight (where C_b is the set of blue vertices in C):

$$s_K(b_0, r_i) = \sum_{b_j \in C_b} s_G(r_i, b_j) + \frac{1}{2} \sum_{r_j \in R \setminus \{r_i\}} s_G(r_i, r_j)$$

This edge represents the cumulative edge weight between r_i and all blue vertices. The additional summation term, which contains the edge weights between r_i and all other red vertices, is necessary to ensure our bisection cut will also contain the edge weights between two of the removed red vertices.

Next, the edge weights between red vertices must contain the other portion of the corresponding edge weight in the original graph.

$$s_K(r_i, r_j) = \frac{1}{2} s_G(r_i, r_j)$$

Now, we note that there are a total of $2 - |C_r| + 2X + |C_r| = 2r + 2X$ vertices. So a bisection will partition the graph into vertex sets of size $r + X$. Obviously, in any approximation, b_0 must be grouped with all δ_i and b'_0 must be grouped with all δ'_i . This means the b_0 partition must contain $|C_r| - r$ of the R vertices, and the b'_0 partition must contain the other r . These r vertices in the latter partition are the ones we select to move.

Consider any set S of r red vertices in K . Then it is a valid bisection. We now show that the edge weight in the cut for this bisection is exactly the edge weight lost by removing S from K . We can do this algebraically. We start by breaking down the weight of the cut into the weight between the red vertices in S and b_0 , and also the red vertices in S and the red vertices not in S .

$$\begin{aligned}
& s_K(S, V(K) \setminus S) \\
&= \sum_{r_i \in S} s_K(b_0, r_i) + \sum_{r_i \in S, r_j \in R \setminus S} s_K(r_i, r_j) \\
&= \sum_{r_i \in S} \left(\sum_{b_j \in B} s_G(r_i, b_j) + \frac{1}{2} \sum_{r_j \in R \setminus \{r_j\}} s_G(r_i, r_j) \right) \\
&\quad + \sum_{r_i \in S, r_j \in R \setminus S} \frac{1}{2} s_G(r_i, r_j) \\
&= \sum_{r_i \in S} \left(\sum_{b_j \in B} s_G(r_i, b_j) + \frac{1}{2} \sum_{r_j \in R \setminus \{r_j\}} s_G(r_i, r_j) \right. \\
&\quad \left. + \frac{1}{2} \sum_{r_j \in R \setminus S} s_G(r_i, r_j) \right)
\end{aligned}$$

Notice that the two last summations have an overlap. They both contribute half the edge weight between r_i and vertices in $R \setminus S$. Thus, these edges contribute their entire edge weight. All remaining vertices in $S \setminus \{r_i\}$ only contribute half their edge weight. We can then distribute the summation.

$$\begin{aligned}
& s_K(S, V(K) \setminus S) \\
&= \sum_{r_i \in S} \left(\sum_{b_j \in B} s_G(r_i, b_j) + \frac{1}{2} \sum_{r_j \in S \setminus \{r_j\}} s_G(r_i, r_j) \right. \\
&\quad \left. + \sum_{r_j \in R \setminus S} s_G(r_i, r_j) \right) \\
&= \sum_{r_i \in S, b_j \in B} s_G(r_i, b_j) + \frac{1}{2} \sum_{r_i \in S, r_j \in S \setminus \{r_j\}} s_G(r_i, r_j) \\
&\quad + \sum_{r_i \in S, r_j \in R \setminus S} s_G(r_i, r_j)
\end{aligned}$$

In the middle summation, note that every edge $e = (u, v)$ is counted twice when $r_i = u$ and $r_j = v$, and when $r_i = v$ and $r_j = u$. We can then rewrite this as:

$$\begin{aligned}
s_K(S, V(K) \setminus S) &= \sum_{r_i \in S, b_j \in B} s_G(r_i, b_j) \\
&+ \sum_{r_i, r_j \in S} s_G(r_i, r_j) \\
&+ \sum_{r_i \in S, r_j \in R \setminus S} s_G(r_i, r_j)
\end{aligned}$$

When we remove S , we remove the connections between S and blue vertices, the connections within S , and the connections between S and red vertices not in S . This is precisely what this accounts for. Therefore, any bisection on K directly corresponds to removing a vertex set S of r red vertices from C . If we have a γ_t -approximation for minimum weighted bisection, then, this yields a γ_t -approximation for the smallest loss we can achieve from removing r red vertices.

Now we must compare the optimal way to remove r vertices to the total weight in a cluster. Let $\rho = |C_r|$ be the number of red vertices in a cluster. Then the total number of possible cuts to isolate r red vertices is $\binom{\rho}{r}$. Let \mathcal{S} be the set of all possible cuts to isolate r red vertices. Then if we sum over the weight of all possible cuts (where weight here is the weight between the r removed vertices and all vertices, including each other), that will sum over each red-red edge and blue-red edge multiple times. A red-red edge is counted if either of its endpoints is in $S \in \mathcal{S}$, and this happens $2\binom{\rho}{r-1} - \binom{\rho-1}{r-2} \leq 2\binom{\rho}{r-1}$ of the time. A blue-red edge is counted if its red endpoint is in S , which happens $\binom{\rho}{r-1} \leq 2\binom{\rho}{r-1}$. And of course, since no blue-blue edge is covered, each is covered under $2\binom{\rho}{r-1}$ times. Therefore, if we sum over all these cuts, we get at most $2\binom{\rho}{r-1}$ times the weight of all edges in C .

$$\sum_{S \in \mathcal{S}} s(S) \leq 2\binom{\rho}{r-1} s(C)$$

Let OPT be the minimum possible cut. Now since there are $\binom{\rho}{r}$ cuts, we know the lefthand side here is bounded above by $\binom{\rho}{r} s(OPT)$.

$$\binom{\rho}{r} s(OPT) \leq 2\binom{\rho}{r-1} s(C)$$

We can now simplify.

$$s(OPT) \leq \frac{2r}{\rho} s(C)$$

But note we are given $\rho = O(t)$. So if we have a γ_t approximation for the minimum bisection problem, this means we can find a way to remove r vertices such that the removed weight is at most $O(r/t)\gamma_t$. We can do this again to get a bound on the removal of the blue vertices. This yields a total weight removal of $O(x/t)\gamma_t$. \square

Finally, we can prove Lemma 31, which satisfies the conditions of Lemma 30.

Proof. Start by running Lemma 35 on \mathcal{C} to yield \mathcal{C}_0 . Then we can apply Lemma 37 to yield \mathcal{C}_1 with red-blue clustering graph H_M and underlying perfect red-blue matching M' . We create \mathcal{C}' by merging components in H_M into clusters. Since the max component size is ℓ and the max cluster size in \mathcal{C}_1 is $6t$, then the max cluster size in \mathcal{C}' is $6t\ell$. This satisfies condition 2 of being \mathcal{C} -good. In addition, it is fair by Proposition 34.

Finally, we utilize the fact that we only moved at most $\ell + 108t^2\ell^2$ vertices from any cluster, and note that we only move vertices of a certain color if we have $O(n)$ of that color in that cluster. Then by Lemma 38, we know we lost at most $O(\ell\gamma_t/t + t\gamma_t/\ell^2)$ fraction of the weight from any cluster. This satisfies the second condition and therefore \mathcal{C}' is \mathcal{C} -good. \square

Table 5: Impact of different fairlet decomposition on ratio over original average-linkage in percentage (mean \pm std. dev).

Samples	100	200	400	800	1600
CENSUSGENDER, initial	74.12 \pm 2.52	76.16 \pm 3.42	74.15 \pm 1.44	70.17 \pm 1.01	65.02 \pm 0.79
final	92.32 \pm 2.70	95.75 \pm 0.74	95.68 \pm 0.96	96.61 \pm 0.60	97.45 \pm 0.19
CENSUSRACE, initial	65.67 \pm 7.53	65.31 \pm 3.74	61.97 \pm 2.50	59.59 \pm 1.89	56.91 \pm 0.82
final	85.38 \pm 1.68	92.98 \pm 1.89	94.99 \pm 0.52	96.86 \pm 0.85	97.24 \pm 0.63
BANKMARRIAGE, initial	75.19 \pm 2.53	73.58 \pm 1.05	74.03 \pm 1.33	73.68 \pm 0.59	72.94 \pm 0.63
final	93.88 \pm 2.16	96.91 \pm 0.99	96.82 \pm 0.36	97.05 \pm 0.71	97.81 \pm 0.49
BANKAGE, initial	77.48 \pm 1.45	78.28 \pm 1.75	76.40 \pm 1.65	75.95 \pm 0.77	75.33 \pm 0.28
final	91.26 \pm 2.66	95.74 \pm 2.17	96.45 \pm 1.56	97.31 \pm 1.94	97.84 \pm 0.92

F Additional experimental results for revenue

We have conducted experiments on the four datasets for revenue as well. The Table 5 shows the ratio of fair tree built by using average-linkage on different fairlet decompositions. We run Algorithm 1 on the subsamples with Euclidean distances. Then we convert distances into similarity scores using transformation $s(i, j) = \frac{1}{1+d(i, j)}$. We test the performance of the initial random fairlet decomposition and final fairlet decomposition found by Algorithm 1 for revenue objective using the converted similarity scores.

G Additional experimental results for multiple colors

We ran experiments with multiple colors and the results are analogous to those in the paper. We tested both Census and Bank datasets, with age as the protected feature. For both datasets we set 4 ranges of age to get 4 colors and used $\alpha = 1/3$. We ran the fairlet decomposition in [3] and compare the fair hierarchical clustering’s performance to that of average-linkage. The age ranges and the number of data points belonging to each color are reported in Table 6. Colors are named $\{1, 2, 3, 4\}$ descending with regard to the number of points of the color. The vanilla average-linkage has been found to be unfair: if we take the layer of clusters in the tree that is only one layer higher than the leaves, there is always one cluster with $\alpha > \frac{1}{3}$ for the definition of α -capped fairness, showing the tree to be unfair.

Table 6: Age ranges for all four colors for Census and Bank.

Dataset	Color 1	Color 2	Color 3	Color 4
CENSUSMULTICOLOR	(26, 38] : 9796	(38, 48] : 7131	(48, $+\infty$) : 6822	(0, 26] : 6413
BANKMULTICOLOR	(30, 38] : 14845	(38, 48] : 12148	(48, $+\infty$) : 11188	(0, 30] : 7030

As in the main body, in Table 7, we show for each dataset the $\text{ratio}_{\text{value}}$ both at the time of initialization (Initial) and after using the local search algorithm (Final), where $\text{ratio}_{\text{value}}$ is the ratio between the performance of the tree built on top of the fairlets and that of the tree directly built by average-linkage.

Table 7: Impact of Algorithm 1 on $\text{ratio}_{\text{value}}$ in percentage (mean \pm std. dev).

Samples	200	400	800	1600	3200	6400
CENSUSMULTICOLOR, initial	88.55 \pm 0.87	88.74 \pm 0.46	88.45 \pm 0.53	88.68 \pm 0.22	88.56 \pm 0.20	88.46 \pm 0.30
final	99.01 \pm 0.09	99.41 \pm 0.57	99.87 \pm 0.28	99.80 \pm 0.27	100.00 \pm 0.14	99.88 \pm 0.30
BANKMULTICOLOR, initial	90.98 \pm 1.17	91.22 \pm 0.84	91.87 \pm 0.32	91.70 \pm 0.30	91.70 \pm 0.18	91.69 \pm 0.14
final	98.78 \pm 0.22	99.34 \pm 0.32	99.48 \pm 0.16	99.71 \pm 0.16	99.80 \pm 0.08	99.84 \pm 0.05

Table 8 shows the performance of trees built by average-linkage based on different fairlets, for Revenue objective. As in the main body, the similarity score between any two points i, j is $s(i, j) = \frac{1}{1+d(i, j)}$. The entries in the table are mean and standard deviation of ratios between the fair tree performance and the vanilla average-linkage tree performance. This ratio was calculated both at time of initialization (Initial) when the fairlets were randomly found, and after Algorithm 1 terminated (Final).

Table 9 shows the run time of Algorithm 1 with multiple colors.

H Pseudocode for the cost objective

Table 8: Impact of Algorithm 1 on revenue, in percentage (mean \pm std. dev).

Samples	200	400	800	1600	3200
CENSUSMULTICOLOR, initial	75.76 \pm 2.86	73.60 \pm 1.77	69.77 \pm 0.56	66.02 \pm 0.95	61.94 \pm 0.61
final	92.68 \pm 0.97	94.66 \pm 1.66	96.40 \pm 0.61	97.09 \pm 0.60	97.43 \pm 0.77
BANKMULTICOLOR, initial	72.08 \pm 0.98	70.96 \pm 0.69	70.79 \pm 0.72	70.77 \pm 0.49	69.88 \pm 0.53
final	94.99 \pm 0.79	95.87 \pm 2.07	97.19 \pm 0.81	97.93 \pm 0.59	98.43 \pm 0.14

Table 9: Average running time of Algorithm 1 in seconds.

Samples	200	400	800	1600	3200	6400
CENSUSMULTICOLOR	0.43	1.76	7.34	35.22	152.71	803.59
BANKMULTICOLOR	0.43	1.45	6.77	29.64	127.29	586.08

Algorithm 2 Fair hierarchical clustering for cost objective.

Input: Graph G , edge weight $w : E \rightarrow \mathbb{R}$, color $c : V \rightarrow \{\text{red, blue}\}$, parameters t and ℓ

{Step (A)}

$T \leftarrow \text{UNFAIRHC}(G, w)$ {Blackbox unfair clustering that minimizes cost}

{Step (B)}

Let $\mathcal{C} \leftarrow \emptyset$

Do a BFS of T , placing visited cluster C in \mathcal{C} if $|C| \leq t$, and not proceeding to C 's children

{Step (C)}

$\mathcal{C}_0, \mathcal{C}' \leftarrow \emptyset$

for C **in** \mathcal{C} **do**

$\mathcal{C}' \leftarrow \mathcal{C}' \cup C$

if $|\mathcal{C}'| \geq t$ **then**

 Add \mathcal{C}' to \mathcal{C}_0

 Let $\mathcal{C}' \leftarrow \emptyset$

end if

end for

If $|\mathcal{C}'| > 0$, merge \mathcal{C}' into some cluster in \mathcal{C}_0

{Step (D)}

for C **in** \mathcal{C}_0 **do**

 Let $\text{exc}(C) \leftarrow$ majority color in C

 Let $\text{ex}(C) \leftarrow$ difference between majority and minority colors in C

end for

{Step (E)}

$H_M \leftarrow \text{BuildClusteringGraph}(\mathcal{C}_0, \text{ex}, \text{exc})$

{Step (F)}

$fV \leftarrow \text{FixUnmatchedVertices}(\mathcal{C}_0, H_M, \text{ex}, \text{exc})$

{Step (G)}

$\mathcal{C}' \leftarrow \text{ConstructClustering}(\mathcal{C}_0, \text{ex}, \text{exc}, fV)$

return \mathcal{C}'

Algorithm 3 BuildClusteringGraph(\mathcal{C}_0, ex, exc)

$H_M \leftarrow (V_M = \mathcal{C}_0, E_M = \emptyset)$
Let $C_i \in V_M$ be any vertex
Let $\ell \leftarrow n^{1/3} \sqrt{\log n}$
while \exists an unvisited $C_j \in V_M$ such that $exc(C_j) \neq exc(C_i)$ **do**
 Add (C_i, C_j) to E_M
 Swap labels C_i and C_j if $ex(C_j) > ex(C_i)$
 Let $ex(C_i) \leftarrow ex(C_i) - ex(C_j)$
 if $ex(C_i) < \ell$ or $|component(C_i)| \geq \ell$ **then**
 Reassign starting point C_i to an unvisited vertex in V_M
 end if
end while
return H_M

Algorithm 4 FixUnmatchedVertices($\mathcal{C}_0, H_M, ex, exc$)

Let $\ell \leftarrow n^{1/3} \sqrt{\log n}$
for $C \in \mathcal{C}_0 \setminus V_M$ **do**
 Let $fV(C, \text{red}), fV(C, \text{blue}) \leftarrow m^2/\ell^2$
end for
for i from 1 to $108t^2/\ell^3$ **do**
 for each k component in H_M **do**
 for p in a BFS of k **do**
 Let $ch \leftarrow$ some child of p
 $fV(p, exc(p)) \leftarrow fV(p, exc(p)) + \ell$
 $ex(p) \leftarrow ex(p) - \ell$
 $fV(ch, exc(ch)) \leftarrow fV(ch, exc(ch)) + \ell$
 $ex(ch) \leftarrow ex(ch) - \ell$
 if # matches between p and $ch < \ell$ **then**
 Remove (p, ch) from E_M {This creates a new component}
 end if
 end for
 end for
end for
return fV

Algorithm 5 ConstructClustering($\mathcal{C}_0, ex, exc, fV$)

Let $\mathcal{C}', R \leftarrow \emptyset$
for C in \mathcal{C}_0 **do**
 for c in {red, blue} **do**
 Let $f = fV(C, c)$
 Let $C_f = \{v \in C : c(v) = c\}$
 Create the transformed graph L from C_f {Described in the proof of Lemma 38}
 $C' \leftarrow \text{MINWEIGHTBISECTION}(L)$ {Blackbox, returns isolated C_f vertices}
 $C \leftarrow C \setminus C'$
 $R \leftarrow R \cup C'$
 $ex(C) \leftarrow ex(C) - |C'|$
 end for
end for
for $C \in \mathcal{C}_0$ **do**
 Let $S \subset R$ such that $|S| = ex(C)$ with no vertices of color $exc(C)$
 $C = C \cup S$
 $R \leftarrow R \setminus S$
 Add C to \mathcal{C}'
end for
return \mathcal{C}'
