

Hierarchical Clustering: Objective Functions and Algorithms

Vincent Cohen-Addad¹, Varun Kanade^{*2,3}, Frederik Mallmann-Trenn^{4,5}, and Claire Mathieu⁴

¹University of Copenhagen

²University of Oxford

³The Alan Turing Institute

⁴École normale supérieure, CNRS, PSL Research University

⁵Simon Fraser University

Abstract

Hierarchical clustering is a recursive partitioning of a dataset into clusters at an increasingly finer granularity. Motivated by the fact that most work on hierarchical clustering was based on providing algorithms, rather than optimizing a specific objective, Dasgupta (2016) framed similarity-based hierarchical clustering as a combinatorial optimization problem, where a ‘good’ hierarchical clustering is one that minimizes some cost function. He showed that this cost function has certain desirable properties, such as in order to achieve optimal cost disconnected components must be separated first and that in ‘structureless’ graphs, *i.e.*, cliques, all clusterings achieve the same cost.

We take an axiomatic approach to defining ‘good’ objective functions for both similarity and dissimilarity-based hierarchical clustering. We characterize a set of *admissible* objective functions (that includes the one introduced by Dasgupta) that have the property that when the input admits a ‘natural’ ground-truth hierarchical clustering, the ground-truth clustering has an optimal value.

Equipped with a suitable objective function, we analyze the performance of practical algorithms, as well as develop better and faster algorithms for hierarchical clustering. For similarity-based hierarchical clustering, Dasgupta (2016) showed that a simple recursive sparsest-cut based approach achieves an $O(\log^{3/2} n)$ -approximation on worst-case inputs. We give a more refined analysis of the algorithm and show that it in fact achieves an $O(\sqrt{\log n})$ -approximation¹. This improves upon the LP-based $O(\log n)$ -approximation of Roy and Pokutta (2016). For dissimilarity-based hierarchical clustering, we show that the classic average-linkage algorithm gives a factor 2 approximation, and provide a simple and better algorithm that gives a factor 3/2 approximation. This aims at explaining the success of this heuristics in practice. Finally, we consider ‘beyond-worst-case’ scenario through a generalisation of the stochastic block model for hierarchical clustering. We show that Dasgupta’s cost function also has desirable properties for these inputs and we provide a simple algorithm that for graphs generated according to this model yields a $1 + o(1)$ factor approximation.

*This work was supported in part by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

¹Charikar and Chatziafratis (2017) independently proved that the sparsest-cut based approach achieves a $O(\sqrt{\log n})$ approximation.

Contents

1	Introduction	2
1.1	Summary of Our Contributions	4
1.2	Related Work	5
2	Preliminaries	6
2.1	Notation	6
2.2	Ultrametrics	7
3	Quantifying Output Value: An Axiomatic Approach	10
3.1	Admissible Cost Functions	10
3.2	Characterizing Admissible Cost Functions	11
3.2.1	Characterizing g that satisfy conditions of Theorem 3.4	13
3.2.2	Characterizing Objective Functions for Dissimilarity Graphs	15
4	Similarity-Based Inputs: Approximation Algorithms	15
5	Admissible Objective Functions and Algorithms for Random Inputs	17
5.1	A Random Graph Model For Hierarchical Clustering	17
5.2	Objective Functions and Ground-Truth Tree	19
5.3	Algorithm for Clustering in the HSBM	21
6	Dissimilarity-Based Inputs: Approximation Algorithms	26
6.1	Average-Linkage	26
6.2	A Simple and Better Approximation Algorithm for Worst-Case Inputs	27
7	Perfect Ground-Truth Inputs and Beyond	29
7.1	Perfect Ground-Truth Inputs are Easy	30
7.2	A Near-Linear Time Algorithm	33
7.3	Beyond Structured Inputs	34
8	Worst-Case Analysis of Common Heuristics	36

1 Introduction

A *hierarchical clustering* is a recursive partitioning of a dataset into successively smaller clusters. The input is a weighted graph whose edge weights represent pairwise similarities or dissimilarities between datapoints. A hierarchical clustering is represented by a rooted tree where each leaf represents a datapoint and each internal node represents a cluster containing its descendant leaves. Computing a hierarchical clustering is a fundamental problem in data analysis; it is routinely used to analyze, classify, and pre-process large datasets. A hierarchical clustering provides useful information about data that can be used, *e.g.*, to divide a digital image into distinct regions of different granularities, to identify communities in social networks at various societal levels, or to determine the ancestral tree of life. Developing robust and efficient algorithms for computing hierarchical clusterings is of importance in several research areas, such as machine learning, big-data analysis, and bioinformatics.

Compared to flat partition-based clustering (the problem of dividing the dataset into k parts), hierarchical clustering has received significantly less attention from a theory perspective. Partition-based clustering is typically framed as minimizing a well-defined objective such as k -means, k -medians, *etc.* and (approximation) algorithms to optimize these objectives have been a focus of study for at least two decades. On the other hand, hierarchical clustering has rather been studied at a more procedural level in terms of algorithms used in practice. Such algorithms can be broadly classified into two categories, *agglomerative* heuristics which build the candidate cluster tree bottom up, *e.g.*, average-linkage, single-linkage, and complete-linkage, and *divisive* heuristics which build the tree top-down, *e.g.*, bisection k -means, recursive sparsest-cut *etc.* Dasgupta (2016) identified the lack of a well-defined objective function as one of the reasons why the theoretical study of hierarchical clustering has lagged behind that of partition-based clustering.

Defining a Good Objective Function.

What is a ‘good’ output tree for hierarchical clustering? Let us suppose that the edge weights represent similarities (similar datapoints are connected by edges of high weight)². Dasgupta (2016) frames hierarchical clustering as a combinatorial optimization problem, where a good output tree is a tree that minimizes some cost function; but which function should that be? Each (binary) tree node is naturally associated to a cut that splits the cluster of its descendant leaves into the cluster of its left subtree on one side and the cluster of its right subtree on the other, and Dasgupta defines the objective to be the sum, over all tree nodes, of the total weight of edges crossing the cut multiplied by the cardinality of the node’s cluster. In what sense is this good? Dasgupta argues that it has several attractive properties: (1) if the graph is disconnected, *i.e.*, data items in different connected components have nothing to do with one another, then the hierarchical clustering that minimizes the objective function begins by first pulling apart the connected components from one another; (2) when the input is a (unit-weight) clique then no particular structure is favored and all binary trees have the same cost; and (3) the cost function also behaves in a desirable manner for data containing a planted partition. Finally, an attempt to generalize the cost function leads to functions that violate property (2).

In this paper, we take an axiomatic approach to defining a ‘good’ cost function. We remark that in many application, for example in phylogenetics, there exists an unknown ‘ground truth’ hierarchical clustering— the actual ancestral tree of life—from which the similarities are generated (possibly with noise), and the goal is to infer the underlying ground truth tree from the available data. In this sense, a cluster tree is good insofar as it is isomorphic to the (unknown) ground-truth

²This entire discussion can equivalently be phrased in terms of dissimilarities without changing the essence.

cluster tree, and thus a natural condition for a ‘good’ objective function is one such that for inputs that admit a ‘natural’ ground-truth cluster tree, the value of the ground-truth tree is optimal. We provide a formal definition of inputs that admit a ground-truth cluster tree in Section 2.2.

We consider, as potential objective functions, the class of all functions that sum, over all the nodes of the tree, the total weight of edges crossing the associated cut times some function of the cardinalities of the left and right clusters (this includes the class of functions considered by Dasgupta (2016)). In Section 3 we characterize the ‘good’ objective functions in this class and call them *admissible* objective functions. We prove that for any objective function, for any ground-truth input, the ground-truth tree has optimal cost (w.r.t to the objective function) *if and only if* the objective function (1) is symmetric (independent of the left-right order of children), (2) is increasing in the cardinalities of the child clusters, and (3) for (unit-weight) cliques, has the same cost for all binary trees (Theorem 3.4). Dasgupta’s objective function is admissible in terms of the criteria described above.

In Section 5, we consider random graphs that induce a natural clustering. This model can be seen as a noisy version of our notion of ground-truth inputs and a hierarchical stochastic block model. We show that the ground-truth tree has optimal *expected cost* for any admissible objective function. Furthermore, we show that the ground-truth tree has cost at most $(1 + o(1))\text{OPT}$ with high probability for the objective function introduced by Dasgupta (2016).

Algorithmic Results

The objective functions identified in Section 3 allow us to (1) quantitatively compare the performances of algorithms used in practice and (2) design better and faster approximation algorithms.³

Algorithms for Similarity Graphs: Dasgupta (2016) shows that the *recursive ϕ -approximate sparsest cut* algorithm, that recursively splits the input graph using a ϕ -approximation to the sparsest cut problem, outputs a tree whose cost is at most $O(\phi \log n \cdot \text{OPT})$. Roy and Pokutta (2016) recently gave an $O(\log n)$ -approximation by providing a linear programming relaxation for the problem and providing a clever rounding technique. Charikar and Chatziafratis (2017) showed that the recursive ϕ -sparsest cut algorithm of Dasgupta gives an $O(\phi)$ -approximation. In Section 4, we obtain an independent proof showing that the ϕ -approximate sparsest cut algorithm is an $O(\phi)$ -approximation (Theorem 4.1)⁴. Our proof is quite different from the proof of Charikar and Chatziafratis (2017) and relies on a charging argument. Combined with the celebrated result of Arora et al. (2009), this yields an $O(\sqrt{\log n})$ -approximation. The results stated here apply to Dasgupta’s objective function; the approximation algorithms extend to other objective functions, though the ratio depends on the specific function being used. We conclude our analysis of the worst-case setting by showing that all the linkage-based algorithms commonly used in practice can perform rather poorly on worst-case inputs (see Sec. 8).

Algorithms for Dissimilarity Graphs: Many of the algorithms commonly used in practice, *e.g.*, linkage-based methods, assume that the input is provided in terms of pairwise dissimilarity (*e.g.*, points that lie in a metric space). As a result, it is of interest to understand how they fare when compared using admissible objective functions for the dissimilarity setting. When the edge

³For the objective function proposed in his work, Dasgupta (2016) shows that finding a cluster tree that minimizes the cost function is NP-hard. This directly applies to the admissible objective functions for the dissimilarity setting as well. Thus, the focus turns to developing approximation algorithms.

⁴Our analysis shows that the algorithm achieves a 6.75ϕ -approximation and the analysis of Charikar and Chatziafratis (2017) yields a 8ϕ -approximation guarantee. This minor difference is of limited impact since the best approximation guarantee for sparsest-cut is $O(\sqrt{\log n})$.

weights of the input graph represent dissimilarities, the picture is considerably different from an approximation perspective. For the analogue of Dasgupta’s objective function in the dissimilarity setting, we show that the average-linkage algorithm (see Algorithm 3) achieves a 2-approximation (Theorem 6.2). This stands in contrast to other practical heuristic-based algorithms, which may have an approximation guarantee as bad as $\Omega(n^{1/4})$ (Theorem 8.6). Thus, using this objective-function based approach, one can conclude that the average-linkage algorithm is the more robust of the practical algorithms, perhaps explaining its success in practice. We also provide a new, simple, and better algorithm, the locally densest-cut algorithm,⁵ which we show gives a 3/2-approximation (Theorem 6.5). Our results extend to any admissible objective function, though the exact approximation factor depends on the specific choice.

Structured Inputs and Beyond-Worst-Case Analysis: The recent work of Roy and Pokutta (2016) and Charikar and Chatziafratis (2017) have shown that obtaining constant approximation guarantees for worst-case inputs is beyond current techniques (see Section 1.2). Thus, we consider inputs that admit a ‘natural’ ground-truth cluster tree. For such inputs, we show that essentially all the practical algorithms do the right thing, in that they recover the ground-truth cluster tree. Since real-world inputs might exhibit a noisy structure, we consider more general scenarios:

- We consider a natural generalization of the classic stochastic block model that generates random graphs with a hidden ground-truth hierarchical clustering. We provide a simple algorithm based on singular value decomposition (SVD) and agglomerative methods that achieves a $(1 + o(1))$ -approximation for Dasgupta’s objective function (in fact, it recovers the ground-truth tree) with high probability. Interestingly, this algorithm is very similar to approaches used in practice for hierarchical clustering.
- We introduce the notion of a δ -adversarially perturbed ground-truth input, which can be viewed as being obtained from a small perturbation to an input that admits a natural ground truth cluster tree. This approach bears similarity to the stability-based conditions used by Balcan et al. (2008) and Bilu and Linial (2012). We provide an algorithm that achieves a δ -approximation in both the similarity and dissimilarity settings, independent of the objective function used as long as it is admissible according to the criteria used in Section 3.

1.1 Summary of Our Contributions

Our work makes significant progress towards providing a more complete picture of objective-function based hierarchical clustering and understanding the success of the classic heuristics for hierarchical clustering.

- Characterization of ‘good’ objective functions. We prove that for any ground-truth input, the ground-truth tree has strictly optimal cost for an objective function *if and only if*, the objective function (1) is symmetric (independent of the left-right order of children), (2) is monotone in the cardinalities of the child clusters, and (3) for unit-weight cliques, gives the same weight to all binary trees (Theorem 3.4). We refer to such objective functions as *admissible*; according to these criteria Dasgupta’s objective function is admissible.
- Worst-case approximation. First, for similarity-based inputs, we provide a new proof that the recursive ϕ -approximate sparsest cut algorithm is an $O(\phi)$ -approximation (hence an $O(\sqrt{\log n})$ -approximation) (Theorem 4.1) for Dasgupta’s objective function. Second, for

⁵We say that a cut (A, B) is locally dense if moving a vertex from A to B or from B to A does not increase the density of the cut. One could similarly define locally-sparsest-cut.

dissimilarity-based inputs, we show that the classic average-linkage algorithm is a 2-approximation (Theorem 6.2), and provide a new algorithm which we prove is a 3/2-approximation (Theorem 6.5). All those results extend to other cost functions but the approximation ratio is function-dependent.

- Beyond worst-case. First, stochastic models. We consider the *hierarchical stochastic block model* (Definition 5.1). We give a simple algorithm based on SVD and classic agglomerative methods that, with high probability, recovers the ground-truth tree and show that this tree has cost that is $(1 + o(1))\text{OPT}$ with respect to Dasgupta’s objective function (Theorem 5.8). Second, adversarial models. We introduce the notion of δ -perturbed inputs, obtained by a small adversarial perturbation to ground-truth inputs, and give a simple δ -approximation algorithm (Theorem 7.8).
- Perfect inputs, perfect reconstruction. For ground-truth inputs, we note that the algorithms used in practice (the linkage algorithms, the bisection 2-centers, etc.) correctly reconstruct a ground truth tree (Theorems 7.1, 7.3, 7.5). We introduce a simple, faster algorithm that is also optimal on ground-truth inputs (Theorem 7.7).

Techniques: A key observation that is used in several of our results is that *ultrametrics* are the natural notion to model *perfect inputs* for hierarchical clustering. (Although the output of our algorithms (and that of Dasgupta (2016)) are trees and not ultrametrics, it is easy to see that data arising from tree metrics does not exhibit perfect hierarchical structure.) This is, in particular, motivated by the fact that they exhibit a strong relationship with hierarchical clustering heuristics (see Carlsson and Mémoli (2010)). We then use ultrametrics to define perfect inputs (Section 2) and to identify “good” objective functions (Section 3). We also use it to introduce a random graph model for hierarchical clustering (Section 5). Elsewhere, we use techniques such as concentration of measure phenomena and charging schemes for analysing approximation ratios in our results.

1.2 Related Work

The recent paper of Dasgupta (2016) served as the starting point of this work. Dasgupta (2016) defined an objective function for hierarchical clustering and thus formulated the question of constructing a cluster tree as a combinatorial optimization problem. Dasgupta also showed that the resulting problem is NP-hard and that the recursive ϕ -sparsest-cut algorithm achieves an $O(\phi \log n)$ -approximation. Dasgupta’s results have been improved in two subsequent papers. Roy and Pokutta (2016) wrote an integer program for the hierarchical clustering problem using a combinatorial characterization of the ultrametrics induced by Dasgupta’s cost function. They also provide a spreading metric LP and a rounding algorithm based on sphere/region-growing that yields an $O(\log n)$ -approximation. Finally, they show that no polynomial size SDP can achieve a constant factor approximation for the problem and that under the Small Set Expansion (SSE) hypothesis, no polynomial-time algorithm can achieve a constant factor approximation.

Charikar and Chatziafratis (2017) also gave a proof that the problem is hard to approximate within any constant factor under the Small Set Expansion hypothesis. They also proved that the recursive ϕ -sparsest cut algorithm produces a hierarchical clustering with cost at most $O(\phi\text{OPT})$; their techniques appear to be significantly different from ours. Additionally, Charikar and Chatziafratis (2017) introduce a spreading metric SDP relaxation for the hierarchical clustering problem introduced by Dasgupta that has integrality gap $O(\sqrt{\log n})$ and a spreading metric LP relaxation that yields an $O(\log n)$ -approximation to the problem.

On hierarchical clustering more broadly. There is an extensive literature on hierarchical clustering and its applications. It will be impossible to discuss most of it here; for some applications the reader may refer to *e.g.*, (Jardine and Sibson, 1972; Sneath and Sokal, 1962; Felsenstein and Felsenstein, 2004; Castro et al., 2004). Algorithms for hierarchical clustering have received a lot of attention from a practical perspective. For a definition and overview of *agglomerative* algorithms (such as average-linkage, complete-linkage, and single-linkage) see *e.g.*, (Friedman et al., 2001) and for *divisive algorithms* see *e.g.*, Steinbach et al. (2000).

Most previous theoretical work on hierarchical clustering aimed at evaluating the cluster tree output by the linkage algorithms using the traditional objective functions for partition-based clustering, *e.g.*, considering k -median or k -means cost of the clusters induced by the top levels of the tree (see *e.g.*, (Plaxton, 2003; Dasgupta and Long, 2005; Lin et al., 2006)). Previous work also proved that average-linkage can be useful to recover an underlying partition-based clustering when it exists under certain stability conditions (see (Balcan et al., 2008; Balcan and Liang, 2016)). The approach of this paper is different: we aim at associating a cost or a value to each hierarchical clustering and finding the best hierarchical clustering with respect to these objective functions.

In Section 3, we take an axiomatic approach toward *objective functions*. Axiomatic approach toward a *qualitative analysis of algorithms* for clustering where taken before. For example, the celebrated result of Kleinberg (2002) (see also Zadeh and Ben-David (2009)) showed that there is no algorithm satisfying three natural axioms simultaneously. This approach was applied to hierarchical clustering algorithms by Carlsson and Mémoli (2010) who showed that in the case of hierarchical clustering one gets a positive result, unlike the impossibility result of Kleinberg. Their focus was on finding an ultrametric (on the datapoints) that is the closest to the metric (in which the data lies) in terms of the Gromov-Hausdorff distance. Our approach is completely different as we focus on defining objective functions and use these for *quantitative analyses of algorithms*.

Our condition for inputs to have a ground-truth cluster tree, and especially their δ -adversarially perturbed versions, can be to be in the same spirit as that of the stability condition of Bilu and Linial (2012) or Bilu et al. (2013): the input induces a natural clustering to be recovered whose cost is optimal. It bears some similarities with the “strict separation” condition of Balcan et al. (2008), while we do not require the separation to be strict, we do require some additional hierarchical constraints. There are a variety of stability conditions that aim at capturing some of the structure that real-world inputs may exhibit (see *e.g.*, (Awasthi et al., 2012; Balcan et al., 2013, 2008; Ostrovsky et al., 2012)). Some of them induce a condition under which an underlying clustering can be mostly recovered (see *e.g.*, (Bilu and Linial, 2012; Balcan et al., 2009a, 2013) for deterministic conditions and *e.g.*, Arora and Kannan (2001); Brubaker and Vempala (2008); Dasgupta and Schulman (2007); Dasgupta (1999); Balcan et al. (2009b) for probabilistic conditions). Imposing other conditions allows one to bypass hardness-of-approximation results for classical clustering objectives (such as k -means), and design efficient approximation algorithms (see, *e.g.*, (Awasthi et al., 2010; Awasthi and Sheffet, 2012; Kumar and Kannan, 2010)). Eldridge et al. (2016) also investigate the question of understanding hierarchical cluster trees for random graphs generated from graphons. Their goal is quite different from ours—they consider the “single-linkage tree” obtained using the graphon as the ground-truth tree and investigate how a cluster tree that has low *merge distortion* with respect to this be obtained.⁶ This is quite different from the approach taken in our work which is primarily focused on understanding performance with respect to admissible cost functions.

⁶This is a simplistic characterization of their work. However, a more precise characterization would require introducing a lot of terminology from their paper, which is not required in this paper.

2 Preliminaries

2.1 Notation

An undirected weighted graph $G = (V, E, w)$ is defined by a finite set of vertices V , a set of edges $E \subseteq \{\{u, v\} \mid u, v \in V\}$ and a weight function $w : E \rightarrow \mathbb{R}_+$, where \mathbb{R}_+ denotes non-negative real numbers. We will only consider graphs with positive weights in this paper. To simplify notation (and since the graphs are undirected) we let $w(u, v) = w(v, u) = w(\{u, v\})$. When the weights on the edges are not pertinent, we simply denote graphs as $G = (V, E)$. When G is clear from the context, we denote $|V|$ by n and $|E|$ by m . We define $G[U]$ to be the subgraph induced by the nodes of U .

A *cluster tree* or *hierarchical clustering* T for graph G is a rooted binary tree with exactly $|V|$ leaves, each of which is labeled by a distinct vertex $v \in V$.⁷ Given a graph $G = (V, E)$ and a cluster tree T for G , for nodes $u, v \in V$ we denote by $\text{LCA}_T(u, v)$ the lowest common ancestor (furthest from the root) of u and v in T .

For any internal node N of T , we denote the subtree of T rooted at N by T_N .⁸ Moreover, for any node N of T , define $V(N)$ to be the set of leaves of the subtree rooted at N . Additionally, for any two trees T_1, T_2 , define the *union* of T_1, T_2 to be the tree whose root has two children C_1, C_2 such that the subtree rooted at C_1 is T_1 and the subtree rooted at C_2 is T_2 .

Finally, given a weighted graph $G = (V, E, w)$, for any set of vertices $A \subseteq V$, let $w(A) = \sum_{a, b \in A} w(a, b)$ and for any set of edges E_0 , let $w(E_0) = \sum_{e \in E_0} w(e)$. Finally, for any sets of vertices $A, B \subseteq V$, let $w(A, B) = \sum_{a \in A, b \in B} w(a, b)$.

2.2 Ultrametrics

Definition 2.1 (Ultrametric). *A metric space (X, d) is an ultrametric if for every $x, y, z \in X$, $d(x, y) \leq \max\{d(x, z), d(y, z)\}$.*

Similarity Graphs Generated from Ultrametrics

We say that a weighted graph $G = (V, E, w)$ is a *similarity graph generated from an ultrametric*, if there exists an ultrametric (X, d) , such that $V \subseteq X$, and for every $x, y \in V, x \neq y$, $e = \{x, y\}$ exists, and $w(e) = f(d(x, y))$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a non-increasing function.⁹

Dissimilarity Graphs Generated from Ultrametrics

We say that a weighted graph $G = (V, E, w)$ is a *dissimilarity graph generated from an ultrametric*, if there exists an ultrametric (X, d) , such that $V \subseteq X$, and for every $x, y \in V, x \neq y$, $e = \{x, y\}$ exists, and $w(e) = f(d(x, y))$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a non-decreasing function.

Minimal Generating Ultrametric

For a weighted undirected graph $G = (V, E, w)$ generated from an ultrametric (either similarity or dissimilarity), in general there may be several ultrametrics and the corresponding function f

⁷In general, one can look at trees that are not binary. However, it is common practice to use binary trees in the context of hierarchical trees. Also, for results presented in this paper nothing is gained by considering trees that are not binary.

⁸For any tree T , when we refer to a subtree T' (of T) rooted at a node N , we mean the connected subgraph containing all the leaves of T that are descendant of N .

⁹In some cases, we will say that $e = \{x, y\} \notin E$, if $w(e) = 0$. This is fine as long as $f(d(x, y)) = 0$.

mapping distances in the ultrametric to weights on the edges, that generate the same graph. It is useful to introduce the notion of a minimal ultrametric that generates G .

We focus on similarity graphs here; the notion of minimal generating ultrametric for dissimilarity graphs is easily obtained by suitable modifications. Let (X, d) be an ultrametric that generates $G = (V, E, w)$ and f the corresponding function mapping distances to similarities. Then we consider the ultrametric (V, \tilde{d}) defined as follows: (i) $\tilde{d}(u, u) = 0$ and (ii) for $u \neq v$,

$$\tilde{d}(u, v) = \tilde{d}(v, u) = \max_{u', v'} \{d(u', v') \mid f(d(u', v')) = f(d(u, v))\} \quad (1)$$

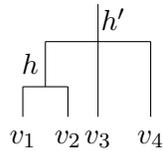
It remains to be seen that (V, \tilde{d}) is indeed an ultrametric. First, notice that by definition, $\tilde{d}(u, v) \geq d(u, v)$ and hence clearly $\tilde{d}(u, v) = 0$ if and only if $u = v$ as d is the distance in an ultrametric. The fact that \tilde{d} is symmetric is immediate from the definition. The only part remaining to check is the so called *isosceles triangles with longer equal sides* conditions—the ultrametric requirement that for any u, v, w , $d(u, v) \leq \max\{d(u, w), d(v, w)\}$ implies that all triangles are isosceles and the two sides that are equal are at least as large as the third side. Let $u, v, w \in V$, and assume without loss of generality that according to the distance d of (V, d) , $d(u, w) = d(v, w) \geq d(u, v)$. From (1) it is clear that $\tilde{d}(u, w) = \tilde{d}(v, w) \geq d(u, w)$. Also, from (1) and the non-increasing nature of f it is clear that if $d(u, v) \leq d(u', v')$, then $\tilde{d}(u, v) \leq \tilde{d}(u', v')$. Thence, (V, \tilde{d}) is an ultrametric. The advantage of considering the minimal ultrametric is the following: if $\mathcal{D} = \{\tilde{d}(u, v) \mid u, v \in V, u \neq v\}$ and $\mathcal{W} = \{w(u, v) \mid u, v \in V, u \neq v\}$, then the restriction of f from $\mathcal{D} \rightarrow \mathcal{W}$ is actually a bijection. This allows the notion of a generating tree to be defined in terms of distances in the ultrametric or weights, without any ambiguity. Applying an analogous definition and reasoning yields a similar notion for the dissimilarity case.

Definition 2.2 (Generating Tree). *Let $G = (V, E, w)$ be a graph generated by a minimal ultrametric (V, d) (either a similarity or dissimilarity graph). Let T be a rooted binary tree with $|V|$ leaves and $|V| - 1$ internal nodes; let \mathcal{N} denote the internal nodes and L the set of leaves of T and let $\sigma : L \rightarrow V$ denote a bijection between the leaves of T and nodes of V . We say that T is a generating tree for G , if there exists a weight function $W : \mathcal{N} \rightarrow \mathbb{R}_+$, such that for $N_1, N_2 \in \mathcal{N}$, if N_1 appears on the path from N_2 to the root, $W(N_1) \leq W(N_2)$. Moreover for every $x, y \in V$, $w(\{x, y\}) = W(LCA_T(\sigma^{-1}(x), \sigma^{-1}(y)))$.*

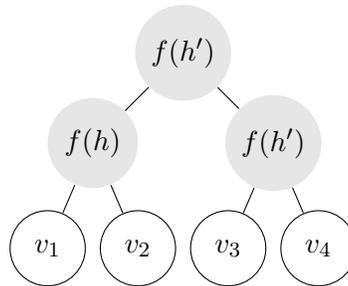
The notion of a generating tree defined above more or less corresponds to what is referred to as a *dendrogram* in the machine learning literature (see *e.g.*, (Carlsson and Mémoli, 2010)). More formally, a dendrogram is a rooted tree (not necessarily binary), where the leaves represent the datapoints. Every internal node in the tree has associated with it a height function h which is the distance between any pairs of datapoints for which it is the least common ancestor. It is a well-known fact that a set of points in an ultrametric can be represented using a dendrogram (see *e.g.*, (Carlsson and Mémoli, 2010)). A dendrogram can easily be modified to obtain a generating tree in the sense of Definition 2.2: an internal node with k children is replaced by an arbitrary binary tree with k leaves and the children of the nodes in the dendrogram are attached to these k leaves. The height h of this node is used to give the weight $W = f(h)$ to all the $k - 1$ internal nodes added when replacing this node. Figure 1 shows this transformation.

Ground-Truth Inputs

Definition 2.3 (Ground-Truth Input). *We say that a graph G is a ground-truth input if it is a similarity or dissimilarity graph generated from an ultrametric. Equivalently, there exists a tree T that is generating for G .*

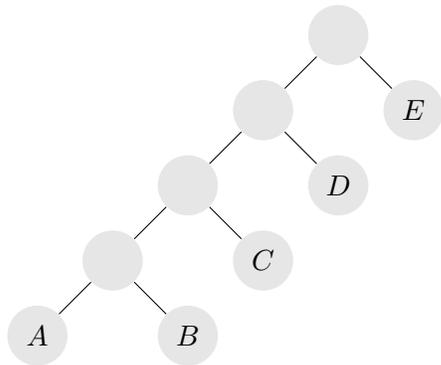


(a) Dendrogram on 4 nodes.

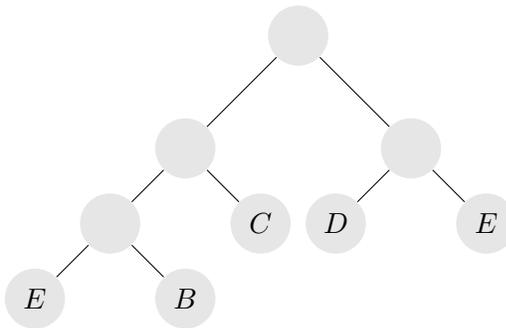


(b) Generating tree equivalent to dendrogram

Figure 1: Dendrogram and equivalent generating tree.



(a)



(b)

Figure 2: (a) Caterpillar tree on 5 nodes with unit-weight edges used to define a tree metric. (b) A candidate cluster tree for the data generated using the tree metric

Motivation. We briefly describe the motivation for defining graphs generated from an ultrametric as ground-truth inputs. We'll focus the discussion on similarity graphs, though essentially the same logic holds for dissimilarity graphs. As described earlier, there is a natural notion of a *generating tree* associated with graphs generated from ultrametrics. This tree itself can be viewed as a cluster tree. The clusters obtained using the generating tree have the property that any two nodes in the same cluster are at least as similar to each other as they are to points outside this cluster; and this holds at every level of granularity. Furthermore, as observed by Carlsson and Mémoli (2010), many practical hierarchical clustering algorithms such as the linkage based algorithms, actually output a dendrogram equipped with a height function, that corresponds to an ultrametric embedding of the data. While their work focuses on algorithms that find embeddings in ultrametrics, our work focuses on finding cluster trees. We remark that these problems are related but also quite different.

Furthermore, our results show that the linkage algorithms (and some other practical algorithms), recover a generating tree when given as input graphs that are generated from an ultrametric. Finally, we remark that relaxing the notion further leads to instances where it is hard to define a 'natural' ground-truth tree. Consider a similarity graph generated by a *tree-metric* rather than an ultrametric, where the tree is the caterpillar graph on 5 nodes (see Fig. 2(a)). Then, it is hard to argue that the tree shown in Fig. 2(b) is not a more suitable cluster tree. For instance, D and E are more similar to each other than D is to B or A . In fact, it is not hard to show that by choosing a suitable function f mapping distances from this tree metric to similarities, Dasgupta's objective function is minimized by the tree shown in Fig. 2(b), rather than the 'generating' tree in Fig. 2(a).

3 Quantifying Output Value: An Axiomatic Approach

3.1 Admissible Cost Functions

Let us focus on the similarity case; in this case we use *cost* and *objective* interchangeably. Let $G = (V, E, w)$ be an undirected weighted graph and let T be a cluster tree for graph G . We want to consider cost functions for cluster trees that capture the quality of the hierarchical clustering produced by T . Following the recent work of Dasgupta (2016), we adopt an approach in which a cost is assigned to each internal node of the tree T that corresponds to the quality of the split at that node.

The Axiom. A natural property we would like the cost function to satisfy is that a cluster tree T has minimum cost if and only if T is a generating tree for G . Indeed, the objective function can then be used to indicate whether a given tree is generating and so, whether it is an underlying ground-truth hierarchical clustering. Hence, the objective function acts as a “guide” for finding the correct hierarchical classification. Note that there may be multiple trees that are generating for the same graph. For example, if $G = (V, E, w)$ is a clique with every edge having the same weight then every tree is a generating tree. In these cases, all the generating tree are *valid* ground-truth hierarchical clusterings.

Following Dasgupta (2016), we restrict the search space for such cost functions. For an internal node N in a clustering tree T , let $A, B \subseteq V$ be the leaves of the subtrees rooted at the left and right child of N respectively. We define the cost Γ of the tree T as the sum of the cost at every internal node N in the tree, and at an individual node N we consider cost functions γ of the form

$$\Gamma(T) = \sum_N \gamma(N), \tag{2}$$

$$\gamma(N) = \left(\sum_{x \in A, y \in B} w(x, y) \right) \cdot g(|A|, |B|) \tag{3}$$

We remark that Dasgupta (2016) defined $g(a, b) = a + b$.

Definition 3.1 (Admissible Cost Function). *We say that a cost function γ of the form (2,3) is admissible if it satisfies the condition that for all similarity graphs $G = (V, E, w)$ generated from a minimal ultrametric (V, d) , a cluster tree T for G achieves the minimum cost if and only if it is a generating tree for G .*

Remark 3.2. *Analogously, for the dissimilarity setting we define admissible value functions to be the functions of the form (2,3) that satisfy: for all dissimilarity graph G generated from a minimal ultrametric (V, d) , a cluster tree T for G achieves the maximum value if and only if it is a generating tree for G .*

Remark 3.3. *The RHS of (3) has linear dependence on the weight of the cut (A, B) in the subgraph of G induced by the vertex set $A \cup B$ as well as on an arbitrary function of the number of leaves in the subtrees of the left and right child of the internal node creating the cut (A, B) . For the purpose of hierarchical clustering this form is fairly natural and indeed includes the specific cost function introduced by Dasgupta (2016). We could define the notion of admissibility for other forms of the cost function similarly and it would be of interest to understand whether they have properties that are desirable from the point of view of hierarchical clustering.*

3.2 Characterizing Admissible Cost Functions

In this section, we give an almost complete characterization of admissible cost functions of the form (3). The following theorem shows that cost functions of this form are admissible if and only if they satisfy three conditions: that all cliques must have the same cost, symmetry and monotonicity.

Theorem 3.4. *Let γ be a cost function of the form (3) and let g be the corresponding function used to define γ . Then γ is admissible if and only if it satisfies the following three conditions.*

1. *Let $G = (V, E, w)$ be a clique, i.e., for every $x, y \in V$, $e = \{x, y\} \in E$ and $w(e) = 1$ for every $e \in E$. Then the cost $\Gamma(T)$ for every cluster tree T of G is identical.*
2. *For every $n_1, n_2 \in \mathbb{N}$, $g(n_1, n_2) = g(n_2, n_1)$.*
3. *For every $n_1, n_2 \in \mathbb{N}$, $g(n_1 + 1, n_2) > g(n_1, n_2)$.*

Proof. We first prove the only if part and then the if part.

Only If Part: Suppose that γ is indeed an admissible cost function. We prove that all three conditions must be satisfied by γ .

1. *All cliques have same cost.* We observe that a clique $G = (V, E, w)$ can be generated from an ultrametric. Indeed, let $X = V$ and let $d(u, v) = d(v, u) = 1$ for every $u, v \in X$ such that $u \neq v$ and $d(u, u) = 0$. Clearly, for $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that is non-increasing and satisfying $f(1) = 1$, (V, d) is a minimal ultrametric generating G .

Let T be any binary rooted tree with leaves labeled by V , i.e., a cluster tree for graph G . For any internal node N of T define $W(N) = 1$ as the weight function. This satisfies the definition of generating tree (Defn. 2.2). Thus, every cluster tree T for G is generating and hence, by the definition of admissibility all of them must be optimal, i.e., they all must have exactly the same cost.

2. $g(n_1, n_2) = g(n_2, n_1)$. This part follows more or less directly from the previous part. Let G be a clique on $n_1 + n_2$ nodes. Let T be any cluster tree for G , with subtrees T_1 and T_2 rooted at the left and right child of the root respectively, such that T_1 contains n_1 leaves and T_2 contains n_2 leaves. The number of edges, and hence the total weight of the edges, crossing the cut induced by the root node of T is $n_1 \cdot n_2$. Let \tilde{T} be a tree obtained by making T_2 be rooted at the left child of the root and T_1 at the right child. Clearly \tilde{T} is also a cluster tree for G and induces the same cut at the root node, hence using the property that all cliques have the same cost, $\Gamma(T) = \Gamma(\tilde{T})$. But $\Gamma(T) = n_1 \cdot n_2 \cdot g(n_1, n_2) + \Gamma(T_1) + \Gamma(T_2)$ and $\Gamma(\tilde{T}) = n_1 \cdot n_2 \cdot g(n_2, n_1) + \Gamma(T_1) + \Gamma(T_2)$. Thence, $g(n_1, n_2) = g(n_2, n_1)$.

3. $g(n_1 + 1, n_2) > g(n_1, n_2)$. Consider a graph on $n_1 + n_2 + 1$ nodes generated from an ultrametric as follows. Let $V_1 = \{v_1, \dots, v_{n_1}\}$, $V_2 = \{v'_1, \dots, v'_{n_2}\}$ and consider the ultrametric $(V_1 \cup V_2 \cup \{v^*\}, d)$ defined by $d(x, y) = 1$ if $x \neq y$ and $x, y \in V_1$ or $x, y \in V_2$, $d(x, y) = 2$ if $x \neq y$ and $x \in V_1, y \in V_2$ or $x \in V_2, y \in V_1$, $d(v^*, x) = d(x, v^*) = 3$ for $x \in V_1 \cup V_2$, and $d(u, u) = 0$ for $u \in V_1 \cup V_2 \cup \{v^*\}$. It can be checked easily by enumeration that this is indeed an ultrametric. Furthermore, if $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is non-increasing and satisfies $f(1) = 2$, $f(2) = 1$ and $f(3) = 0$, i.e., $w(\{u, v\}) = 2$ if u and v are both either in V_1 or V_2 , $w(\{u, v\}) = 1$ if $u \in V_1$ and $v \in V_2$ or the other way around, and $w(\{v^*, u\}) = 0$ for $u \in V_1 \cup V_2$, then $(V_1 \cup V_2, \{v^*\}, d)$ is a minimal ultrametric generating G .

Now consider two possible cluster trees defined as follows: Let T_1 be an arbitrary tree on nodes V_1 , T_2 and arbitrary tree on nodes V_2 . T is obtained by first joining T_1 and T_2 using internal node N and making this the left subtree of the root node ρ and the right subtree of the root node is just the singleton node v^* . T' is obtained by first creating a tree by joining T_1 and the singleton node

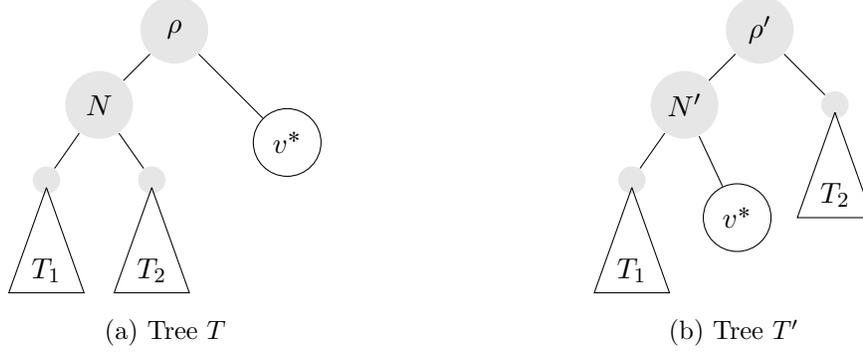


Figure 3: Trees T and T' used to show monotonicity of g .

v^* using internal node N' , this is the left subtree of the root node ρ' and T_2 is the right subtree of the root node. (See Figures 3a and 3b.)

Now it can be checked that T is generating by defining the following weight function. For every internal node M of T_1 , let $W(M) = 1$, similarly for every internal node M of T_2 , let $W(M) = 1$, define $W(N) = 2$ and $W(\rho) = 3$. Now, we claim that T' cannot be a generating tree. This follows from the fact that for a node $u \in V_1, v \in V_2$, the root node $\rho' = \text{LCA}_{T'}(u, v)$, but it is also the case that $\rho' = \text{LCA}_{T'}(v^*, v)$. Thus, it cannot possibly be the case that $W(\rho) = w(\{u, v\})$ and $W(\rho) = w(\{v^*, v\})$ as $w(\{u, v\}) \neq w(\{v^*, v\})$. By definition of admissibility, it follows that $\Gamma(T) < \Gamma(T')$, but $\Gamma(T) = \Gamma(T_1) + \Gamma(T_2) + n_1 \cdot n_2 \cdot g(n_1, n_2)$. The last term arises from the cut at node N ; the root makes no contribution as the cut at the root node ρ has weight 0. On the other hand $\Gamma(T') = \Gamma(T_1) + \Gamma(T_2) + n_1 \cdot n_2 \cdot g(n_1 + 1, n_2)$. There is no cost at the node N' , since the cut has size 0; however, at the root node the cost is now $n_1 \cdot n_2 \cdot g(n_1 + 1, n_2)$ as the left subtree at the root contains $n_1 + 1$ nodes. It follows that $g(n_1 + 1, n_2) > g(n_1, n_2)$.

If Part: For the other direction, we first use the following observation. By condition 2 in the statement of the theorem, every clique on n nodes has the same cost irrespective of the tree used for hierarchical clustering; let $\kappa(n)$ denote said cost. Let $n_1, n_2 \geq 1$, then we have,

$$n_1 \cdot n_2 \cdot g(n_1, n_2) = \kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2) \quad (4)$$

We will complete the proof by induction on $|V|$. The minimum number of nodes required to have a cluster tree with at least one internal node is 2. Suppose $|V| = 2$, then there is a unique (up to interchanging left and right children) cluster tree; this tree is also generating and hence by definition any cost function is admissible. Thus, the base case is covered rather easily.

Now, consider a graph $G = (V, E, w)$ with $|V| = n > 2$. Let T^* be a tree that is generating. Suppose that T is any other tree. Let ρ^* and ρ be the root nodes of the trees respectively. Let V_L^* and V_R^* be the nodes on the left subtree and right subtree of ρ^* ; similarly V_L and V_R in the case of ρ . Let $A = V_L^* \cap V_L$, $B = V_L^* \cap V_R$, $C = V_R^* \cap V_L$, $D = V_R^* \cap V_R$. Let a, b, c and d denote the sizes of A, B, C and D respectively.

We will consider the case when all of $a, b, c, d > 0$; the proof is similar and simpler in case some of them are 0. Let \tilde{T} be a tree with root $\tilde{\rho}$ that has the following structure: Both children of the root are internal nodes, all of A appears as leaves in the left subtree of the left child of the root, B as leaves in the right subtree of the left child of the root, C as leaves in the left subtree of the right child of the root and D as leaves in the right subtree of the right child of the root. We assume that all four subtrees for the sets A, B, C, D are generating and hence by induction optimal. We claim that the cost of \tilde{T} is at least as much as the cost of T^* . To see this note that $V_L^* = A \cup B$. Thus,

the left subtree of ρ^* is optimal for the set V_L^* (by induction), whereas that of $\tilde{\rho}$ may or may not be. Similarly for all the nodes in V_R^* . The only other thing left to account for is the cost at the root. But since ρ^* and $\tilde{\rho}$ induce exactly the same cut on V , the cost at the root is the same. Thus, $\Gamma(\tilde{T}) \geq \Gamma(T^*)$. Furthermore, equality holds if and only if \tilde{T} is also generating for G .

Let W^* denote the weight function for the generating tree T^* such that for all $u, v \in V$, $W^*(\text{LCA}_{T^*}(u, v)) = w(\{u, v\})$. Let ρ_L^* and ρ_R^* denote the left and right children of the root ρ^* of T^* . For all $u_a \in A, u_b \in B$, $w(\{u_a, u_b\}) \geq W^*(\rho_L^*)$. Let

$$x = \frac{1}{ab} \sum_{u_a \in A, u_b \in B} w(\{u_a, u_b\})$$

denote the average weight of the edges going between A and B ; it follows that $x \geq W^*(\rho_L^*)$. Similarly for all $u_c \in C, u_d \in D$, $w(\{u_c, u_d\}) \geq W^*(\rho_R^*)$. Let

$$y = \frac{1}{cd} \sum_{u_c \in C, u_d \in D} w(\{u_c, u_d\})$$

denote the average weight of the edges going between C and D ; it follows that $y \geq W^*(\rho_R^*)$. Finally for every $u \in A \cup B, u' \in C \cup D$, $w(\{u, u'\}) = W^*(\rho^*)$; denote this common value by z . By the definition of generating tree, we know that $x \geq z$ and $y \geq z$.

Now consider the tree T . Let T_L and T_R denote the left and right subtrees of ρ . By induction, it must be that T_L splits A and C as the first cut (or at least that's one possible tree, if multiple cuts exist), similarly T_R first cuts B and D . Both, T and \tilde{T} have subtrees containing only nodes from A, B, C and D . The costs for these subtrees are identical in both cases (by induction). Thus, we have

$$\begin{aligned} \Gamma(T) - \Gamma(\tilde{T}) &= zac \cdot g(a, c) + zbd \cdot g(b, d) + (xab + ycd + z(ad + bc)) \cdot g(a + c, b + d) \\ &\quad - xab \cdot g(a, b) + y \cdot cdg(c, d) - z(a + b)(c + d) \cdot g(a + b, c + d) \\ &= (x - z)ab(g(a + c, b + d) - g(a, b)) + (y - z)cd(g(a + c, b + d) - g(c, d)) \\ &\quad + z((a + c)(b + d) \cdot g(a + c, b + d) + ac \cdot g(a, c) + bd \cdot g(b, d)) \\ &\quad - z((a + b)(c + d) \cdot g(a + b, c + d) + ab \cdot g(a, b) + cd \cdot g(c, d)) \end{aligned}$$

Using (4), we get that the last two expressions above both evaluate to $z(\kappa(a + b + c + d) - \kappa(a) - \kappa(b) - \kappa(c) - \kappa(d))$, but have opposite signs. Thus, we get

$$\Gamma(T) - \Gamma(\tilde{T}) = (x - z)ab(g(a + c, b + d) - g(a, b)) + (y - z)cd(g(a + c, b + d) - g(c, d))$$

It is clear that the above expression is always non-negative and is 0 if and only if $x = z$ and $y = z$. If it is the latter case and it is also the case that $\Gamma(\tilde{T}) = \Gamma(T^*)$, then it must actually be the case that T is a generating tree. \square

3.2.1 Characterizing g that satisfy conditions of Theorem 3.4

Theorem 3.4 give necessary and sufficient conditions on g for cost functions of the form (3) be admissible. However, it leaves open the question of the existence of functions satisfying the criteria and also characterizing the functions g themselves. The fact that such functions exist already follows from the work of Dasgupta (2016), who showed that if $g(n_1, n_2) = n_1 + n_2$, then all cliques have the same cost. Clearly, g is monotone and symmetric and thus satisfies the condition of Theorem 3.4.

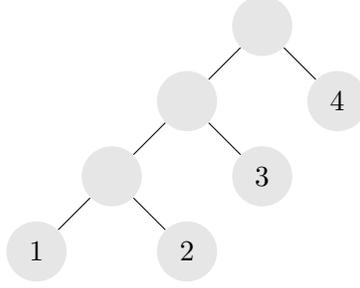


Figure 4: The caterpillar cluster tree for a clique with 4 nodes.

In order to give a more complete characterization, we define g as follows: Suppose $g(\cdot, \cdot)$ is symmetric, we define $g(n, 1)$ for all $n \geq 1$ so that $g(n, 1)/(n + 1)$ is non-decreasing.¹⁰ We consider a particular cluster tree for a clique that is defined using a caterpillar graph, *i.e.*, a cluster tree where the right child of any internal node is a leaf labeled by one of the nodes of G and the left child is another internal node, except at the very bottom. Figure 4 shows a caterpillar cluster tree for a clique on 4 nodes. The cost of the clique on n nodes, say $\kappa(n)$, using this cluster tree is given by

$$\kappa(n) = \sum_{i=0}^{n-1} i \cdot g(i, 1)$$

Now, we enforce the condition that all cliques have the same cost by defining $g(n_1, n_2)$ for $n_1, n_2 > 1$ suitably, in particular,

$$g(n_1, n_2) = \frac{\kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2)}{n_1 \cdot n_2} \quad (5)$$

Thus it only remains to be shown that g is strictly increasing. We show that for $n_2 \leq n_1$, $g(n_1 + 1, n_2) > g(n_1, n_2)$. In order to show this it suffices to show that,

$$n_1(\kappa(n_1 + n_2 + 1) - \kappa(n_1 + 1) - \kappa(n_2)) - (n_1 + 1)(\kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2)) > 0$$

Thus, consider

$$\begin{aligned} & n_1(\kappa(n_1 + n_2 + 1) - \kappa(n_1 + 1) - \kappa(n_2)) - (n_1 + 1)(\kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2)) \\ &= n_1(\kappa(n_1 + n_2 + 1) - \kappa(n_1 + n_2) - \kappa(1) - \kappa(n_1 + 1) + \kappa(n_1) + \kappa(1)) - (\kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2)) \\ &= n_1(n_1 + n_2)g(n_1 + n_2, 1) - n_1^2 g(n_1, 1) - (\kappa(n_1 + n_2) - \kappa(n_1) - \kappa(n_2)) \\ &\geq n_1(n_1 + n_2)g(n_1 + n_2, 1) - n_1^2 g(n_1, 1) - \sum_{i=n_1}^{n_1+n_2-1} i \cdot g(i, 1) \\ &\geq \frac{g(n_1 + n_2, 1)}{n_1 + n_2 + 1} \cdot \left(n_1(n_1 + n_2)(n_1 + n_2 + 1) - n_1^2(n_1 + 1) - \sum_{i=n_1}^{n_1+n_2-1} i(i + 1) \right) > 0 \end{aligned}$$

Above we used the fact that $g(n, 1)/(n + 1)$ is non-decreasing in n and some elementary calculations. This shows that the objective function proposed by Dasgupta (2016) is by no means unique. Only in the last step, do we get an inequality where we use the condition that $g(n, 1)/(n + 1)$ is increasing. Whether this requirement can be relaxed further is also an interesting direction.

¹⁰The function proposed by Dasgupta (2016) is $g(n, 1) = n + 1$, so this ratio is always 1.

3.2.2 Characterizing Objective Functions for Dissimilarity Graphs

When the weights of the edges represent dissimilarities instead of similarities, one can consider objective functions of the same form as (3). As mentioned in Remark 3.2, the difference in this case is that the goal is to maximize the objective function and hence the definition of admissibility now requires that generating trees have a value of the objective that is strictly larger than any tree that is not generating.

The characterization of admissible objective functions as given in Theorem 3.4 for the similarity case continues to hold in the case of dissimilarities. The proof follows in the same manner by appropriately switching the direction of the inequalities when required.

4 Similarity-Based Inputs: Approximation Algorithms

In this section, we analyze the recursive ϕ -sparsest-cut algorithm (see Algorithm 1) that was described previously in Dasgupta (2016). For clarity, we work with the cost function introduced by Dasgupta (2016): The goal is to find a tree T minimizing $\text{cost}(T) = \sum_{N \in T} \text{cost}(N)$ where for each node N of T with children N_1, N_2 , $\text{cost}(N) = w(V(N_1), V(N_2)) \cdot V(N)$. We show that the ϕ -sparsest-cut algorithm achieves a 6.75ϕ -approximation. (Charikar and Chatziafratis (2017) also proved an $O(\phi)$ approximation for Dasgupta’s function.) Our proof also yields an approximation guarantee not just for Dasgupta’s cost function but more generally for any admissible cost function, but the approximation ratio depends on the cost function.

The ϕ -sparsest-cut algorithm (Algorithm 1) constructs a binary tree top-down by recursively finding cuts using a ϕ -approximate sparsest cut algorithm, where the *sparsest-cut* problem asks for a set A minimizing the *sparsity* $w(A, V \setminus A) / (|A| |V \setminus A|)$ of the cut $(A, V \setminus A)$.

Algorithm 1 Recursive ϕ -Sparsest-Cut Algorithm for Hierarchical Clustering

- 1: **Input:** An edge weighted graph $G = (V, E, w)$.
 - 2: $\{A, V \setminus A\} \leftarrow \text{cut with sparsity} \leq \phi \cdot \min_{S \subset V} w(S, V \setminus S) / (|S| |V \setminus S|)$
 - 3: Recurse on $G[A]$ and on $G[V \setminus A]$ to obtain trees T_A and $T_{V \setminus A}$
 - 4: **return** the tree whose root has two children, T_A and $T_{V \setminus A}$.
-

Theorem 4.1. ¹¹ For any graph $G = (V, E)$, and weight function $w : E \rightarrow \mathbb{R}_+$, the ϕ -sparsest-cut algorithm (Algorithm 1) outputs a solution of cost at most $\frac{27}{4}\phi OPT$.

Proof. Let $G = (V, E)$ be the input graph and n denote the total number of vertices of G . Let T denote the tree output by the algorithm and T^* be any arbitrary tree. We will prove that $\text{cost}(T) \leq \frac{27}{4}\phi \text{cost}(T^*)$. ¹²

Recall that for an arbitrary tree T_0 and node N of T_0 , the vertices corresponding to the leaves of the subtree rooted at N is denoted by $V(N)$. Consider the node N_0 of T^* that is the first node reached by the walk from the root that always goes to the child tree with the higher number of leaves, stopping when the subtree of T^* rooted at N_0 contains fewer than $2n/3$ leaves. The *balanced cut* (BC) of T^* is the cut $(V(N_0), V - V(N_0))$. For a given node N with children N_1, N_2 , we say

¹¹For Dasgupta’s function, this was already proved in Charikar and Chatziafratis (2017) with a different constant. The present, independent proof, uses a different method.

¹²The following paragraph bears similarities with the first part of the analysis of (Dasgupta, 2016, Lemma 11) but we obtain a more fine-grained analysis by introducing a charging scheme.

that the cut induced by N is the sum of the weights of the edges between that have one extremity in $V(N_1)$ and the other in $V(N_2)$.

Let $(A \cup C, B \cup D)$ be the cut induced by the root node u of T , where A, B, C, D are such that $(A \cup B, C \cup D)$ is the balanced cut of T^* . Since $(A \cup C, B \cup D)$ is a ϕ -approximate sparsest cut:

$$\frac{w(A \cup C, B \cup D)}{|A \cup C| \cdot |B \cup D|} \leq \phi \frac{w(A \cup B, C \cup D)}{|A \cup B| \cdot |C \cup D|}.$$

By definition of N_0 , $A \cup B$ and $C \cup D$ both have size in $[n/3, 2n/3]$, so the product of their sizes is at least $(n/3)(2n/3) = 2n^2/9$; developing $w(A \cup B, C \cup D)$ into four terms, we obtain

$$\begin{aligned} w(A \cup C, B \cup D) &\leq \phi \frac{9}{2n^2} |A \cup C| |B \cup D| (w(A, C) + w(A, D) + w(B, C) + w(B, D)) \\ &\leq \phi \frac{9}{2} \left[\frac{|B \cup D|}{n} w(A, C) + w(A, D) + w(B, C) + \frac{|A \cup C|}{n} w(B, D) \right], \end{aligned}$$

and so the cost induced by node u of T^* satisfies

$$n \cdot w(A \cup C, B \cup D) \leq \frac{9}{2} \phi |B \cup D| w(A, C) + \frac{9}{2} \phi |A \cup C| w(B, D) + \frac{9}{2} \phi n (w(A, D) + w(B, C)).$$

To account for the cost induced by u , we thus assign a charge of $(9/2)\phi|B \cup D|w(e)$ to each edge e of (A, C) , a charge of $(9/2)\phi|A \cup C|w(e)$ to each edge e of (B, D) , and a charge of $(9/2)\phi n w(e)$ to each edge e of (A, D) or (B, C) .

When we do this for every node u of T , how much does each edge get charged?

Lemma 4.2. *Let $G = (V, E)$ be a graph on n nodes. We consider the above charging scheme for T and T^* . Then, an edge $(v_1, v_2) \in E$ gets charged at most $(9/2)\phi \min((3/2)|V(\text{LCA}_{T^*}(v_1, v_2))|, n)w(e)$ overall, where $\text{LCA}_{T^*}(v_1, v_2)$ denotes the lowest common ancestor of v_1 and v_2 in T^* .*

We temporarily defer the proof and first see how Lemma 4.2 implies the theorem. Observe (as in Dasgupta (2016)) that $\text{cost}(T^*) = \sum_{\{u, v\} \in E} |V(\text{LCA}_{T^*}(u, v))| w(u, v)$. Thanks to Lemma 4.2, when we sum charges assigned because of every node N of T , overall we obtain

$$\text{cost}(T) \leq \frac{9}{2} \phi \sum_{\{v_1, v_2\} \in E} \frac{3}{2} |V(\text{LCA}_{T^*}(v_1, v_2))| w(v_1, v_2) = \frac{27}{4} \phi \text{cost}(T^*).$$

□

Proof of Lemma 4.2. The lemma is proved by induction on the number of nodes of the graph. (The base case is obvious.) For the inductive step, consider the cut $(A \cup C, B \cup D)$ induced by the root node u of T .

- Consider the edges that cross the cut. First, observe that edges of (A, B) or of (C, D) never get charged at all. Second, an edge $e = \{v_1, v_2\}$ of (A, D) or of (B, C) gets charged $(9/2)\phi n w(e)$ when considering the cost induced by node u , and does not get charged when considering any other node of T . In T^* , edge e is separated by the cut $(A \cup B, C \cup D)$ induced by N_0 , so the least common ancestor of v_1 and v_2 is the parent node of N_0 (or above), and by definition of N_0 we have $|V(\text{LCA}_{T^*}(v_1, v_2))| \geq 2n/3$, hence the lemma holds for e .
- An edge $e = \{v_1, v_2\}$ of $G[A] \cup G[C]$ does not get charged when considering the cut induced by node u . Apply Lemma 4.2 to $G[A \cup C]$ for the tree $T_{A \cup C}^*$ defined as the subtree of T^* induced

by the vertices of $A \cup C$ ¹³. By induction, the overall charge to e due to the recursive calls for $G[A \cup C]$ is at most $(9/2)\phi \min((3/2)|V(\text{LCA}_{T_{A \cup C}^*}(v_1, v_2))|, |A \cup C|)w(e)$. By definition of $T_{A \cup C}^*$, we have $|V(\text{LCA}_{T_{A \cup C}^*}(v_1, v_2))| \leq |V(\text{LCA}_{T^*}(v_1, v_2))|$, and $|A \cup C| \leq n$, so the lemma holds for e .

- An edge $\{v_1, v_2\}$ of (A, C) gets a charge of $(9/2)\phi|B \cup D|w(e)$ plus the total charge to e coming from the recursive calls for $G[A \cup C]$ and the tree $T_{A \cup C}^*$. By induction the latter is at most

$$(9/2)\phi \min((3/2)|V(\text{LCA}_{T_{A \cup C}^*}(v_1, v_2))|, |A \cup C|)w(e) \leq (9/2)\phi|A \cup C|w(e).$$

Overall the charge to e is at most $(9/2)\phi n w(e)$. Since the cut induced by node u_0 of T^* separates v_1 from v_2 , we have $|V(\text{LCA}_{T^*}(v_1, v_2))| \geq 2n/3$, hence the lemma holds for e . For edges of (B, D) or of $G[B] \cup G[D]$, a symmetrical argument applies.

□

Remark 4.3. *The recursive ϕ -sparsest-cut algorithm achieves an $O(f_n \phi)$ -approximation for any admissible cost function f , where $f_n = \max_n f(n)/f(\lceil n/3 \rceil)$. Indeed, adapting the definition of the balanced cut as in Dasgupta (2016) and rescaling the charge by a factor of f_n imply the result.*

We complete our study of classical algorithms for hierarchical clustering by showing that the standard agglomerative heuristics can perform poorly (Theorems 8.1, 8.3). Thus, the sparsest-cut-based approach seems to be more reliable in the worst-case. To understand better the success of the agglomerative heuristics, we restrict our attention to ground-truth inputs (Section 7), and random graphs (Section 5), and show that in these contexts these algorithms are efficient.

5 Admissible Objective Functions and Algorithms for Random Inputs

In this section, we initiate a *beyond-worst-case* analysis of the hierarchical clustering problem (see also Section 7.3). We study admissible objective functions in the context of random graphs that have a natural hierarchical structure; for this purpose, we consider a suitable generalization of the stochastic block model to hierarchical clustering.

We show that, for admissible cost functions, an underlying ground-truth cluster tree has optimal expected cost. Additionally, for a subfamily of admissible cost functions (called *smooth*, see Defn. 5.4) which includes the cost function introduced by Dasgupta, we show the following: The cost of the ground-truth cluster tree is with high probability sharply concentrated (up to a factor of $(1 + o(1))$ around its expectation), and so of cost at most $(1 + o(1))\text{OPT}$. This is further evidence that optimising admissible cost functions is an appropriate strategy for hierarchical clustering.

We also provide a simple algorithm based on the SVD based approach of McSherry (2001) followed by a standard agglomerative heuristic yields a hierarchical clustering which is, up to a factor $(1 + o(1))$, optimal with respect to smooth admissible cost functions.

5.1 A Random Graph Model For Hierarchical Clustering

We describe the random graph model for hierarchical clustering, called the hierarchical block model. This model has already been studied earlier, e.g., Lyzinski et al. (2017). However, prior work has

¹³note that $T_{A \cup C}^*$ is not necessarily the optimal tree for $G[A \cup C]$, which is why the lemma was stated in terms of every tree T^* , not just on the optimal tree.

mostly focused on statistical hypothesis testing and exact recovery in some regimes. We will focus on understanding the behaviour of admissible objective functions and algorithms to output cluster trees that have almost optimal cost in terms of the objective function.

We assume that there are k “bottom”-level clusters that are then arranged in a hierarchical fashion. In order to model this we will use a similarity graph on k nodes generated from an ultrametric (see Sec. 2.2). There are n_1, \dots, n_k nodes in each of the k clusters. Each edge is present in the graph with a probability that is a function of the clusters in which their endpoints lie and the underlying graph on k nodes generated from the ultrametric. The formal definition follows.

Definition 5.1 (Hierarchical Stochastic Block Model (HSBM)). *A hierarchical stochastic block model with k bottom-level clusters is defined as follows:*

- Let $\tilde{G}_k = (\tilde{V}_k, \tilde{E}_k, w)$ be a graph generated from an ultrametric (see Sec. 2.2), where $|\tilde{V}_k| = k$ for each $e \in \tilde{E}_k$, $w(e) \in (0, 1)$.¹⁴ Let \tilde{T}_k be a tree on k leaves, let \tilde{N} denote the internal nodes of \tilde{T} and \tilde{L} denote the leaves; let $\tilde{\sigma} : \tilde{L} \rightarrow [k]$ be a bijection. Let \tilde{T} be generating for \tilde{G}_k with weight function $\tilde{W} : \tilde{N} \rightarrow (0, 1)$ (see Defn. 2.2).
- For each $i \in [k]$, let $p_i \in (0, 1]$ be such that $p_i > \tilde{W}(N)$, if N denotes the parent of $\tilde{\sigma}^{-1}(i)$ in \tilde{T} .
- For each $i \in [k]$, there is a fixed constant $f_i \in (0, 1)$; furthermore $\sum_{i=1}^k f_i = 1$.

Then a random graph $G = (V, E)$ on n nodes with sparsity parameter $\alpha_n \in (0, 1]$ is defined as follows: (n_1, \dots, n_k) is drawn from the multinomial distribution with parameters $(n, (f_1, \dots, f_k))$. Each vertex $i \in [n]$ is assigned a label $\psi(i) \in [k]$, so that exactly n_j nodes are assigned the label j for $j \in [k]$. An edge (i, j) is added to the graph with probability $\alpha_n p_{\psi(i)}$ if $\psi(i) = \psi(j)$ and with probability $\alpha_n \tilde{W}(N)$ if $\psi(i) \neq \psi(j)$ and N is the least common ancestor of $\tilde{\sigma}^{-1}(i)$ and $\tilde{\sigma}^{-1}(j)$ in \tilde{T} . The graph $G = (V, E)$ is returned without any labels.

As the definition is rather long and technical, a few remarks are in order.

- Rather than focusing on an arbitrary hierarchy on n nodes, we assume that there are k clusters (which exhibit no further hierarchy) and there is a hierarchy on these k clusters. The model assumes that k is fixed, but in future work, it may be interesting to study models where k itself may be a (modestly growing) function of n . The condition $p_i > \tilde{W}(N)$ (where N is the parent of $\tilde{\sigma}^{-1}(i)$) ensures that nodes in cluster i are strictly more likely to connect to each other than to nodes from any other cluster.
- The graphs generated can be of various sparsity, depending on the parameter α_n . If $\alpha_n \in (0, 1)$ is a fixed constant, we will get dense graphs (with $\Omega(n^2)$ edges), however if $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$, sparser graphs may be achieved. This is similar to the approach taken by Wolfe and Olhede (2013) when considering random graph models generated according to graphons.

We define the *expected graph*, \bar{G} , which is a complete graph where an edge (i, j) has weight $p_{i,j}$ where $p_{i,j}$ is the probability with which it appears in the random graph G . In order to avoid ambiguity, we denote by $\Gamma(T; G)$ and $\Gamma(T; \bar{G})$ the costs of the cluster tree T for the unweighted

¹⁴In addition to \tilde{G}_k being generated from an ultrametric, we make the further assumption that the function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, that maps ultrametric distances to edge weights, has range $(0, 1)$, so that the weight of an edge can be interpreted as a probability of an edge being present. We rule out $w(e) = 0$ as in that case the graph is disconnected and each component can be treated separately.

(random) graph G and weighted graph \bar{G} respectively. Observe that due to linearity (see Eqns. (2) and (3)), for any tree T and any admissible cost function, $\Gamma(T; \bar{G}) = \mathbb{E}[\Gamma(T; G)]$, where the expectation is with respect to the random choices of edges in G (in particular this holds even when conditioning on n_1, \dots, n_k).

Furthermore, note that \bar{G} itself is generated from an ultrametric and the generating trees for \bar{G} are obtained as follows: Let \tilde{T}_k be any generating tree for \tilde{G}_k , let $\hat{T}_1, \hat{T}_2, \dots, \hat{T}_k$ be any binary trees with n_1, \dots, n_k leaves respectively. Let the weight of every internal node of \hat{T}_i be p_i and replace each leaf l in \hat{T}_i by $\tilde{T}_{\hat{\sigma}(l)}$. In particular, this last point allows us to derive Proposition 5.3. We refer to any tree that is generating for the expected graph \bar{G} as a *ground-truth tree* for G .

Remark 5.2. *Although it is technically possible to have $n_i = 0$ for some i under the model, we will assume in the rest of the section that $n_i > 0$ for each i . This avoids getting into the issue of degenerate ground-truth trees; those cases can be handled easily, but add no expository value.*

5.2 Objective Functions and Ground-Truth Tree

In this section, we assume that the graphs represent similarities. This is clearly more natural in the case of unweighted graphs; however, all our results hold in the dissimilarity setting and the proofs are essentially identical.

Proposition 5.3. *Let Γ be an admissible cost function. Let G be a graph generated according to an HSBM (See Defn. 5.1). Let ψ be the (hidden) function mapping the nodes of G to $[k]$ (the bottom-level clusters). Let T be a ground-truth tree for G . Then,*

$$\mathbb{E}[\Gamma(T) \mid \psi] \leq \min_{T'} \mathbb{E}[\Gamma(T') \mid \psi].$$

Moreover, for any tree T' , $\mathbb{E}[\Gamma(T) \mid \psi] = \mathbb{E}[\Gamma(T') \mid \psi]$ if and only if T' is a ground-truth tree.

Proof. As per Remark 5.2, we'll assume that each $n_i > 0$ to avoid degenerate cases. Let \bar{G} be a the expected graph, *i.e.*, \bar{G} is complete and an edge (i, j) has weight p_{ij} , the probability that the edge (i, j) is present in the random graph G generated according to the hierarchical model. Thus, by definition of admissibility $\Gamma(T; \bar{G}) = \min_{T'} \Gamma(T'; \bar{G})$ if and only if T is generating (see Defn. 3.1). As ground-truth trees for G are precisely the generating trees for \bar{G} ; the result follows by observing that for any tree T (not necessarily ground-truth) $\mathbb{E}[\Gamma(T; G) \mid \psi] = \Gamma(T; \bar{G})$, where the expectation is taken only over the random choice of the edges, by linearity of expectation and the definition of the cost function (Eqns. 2 and 3). \square

Definition 5.4. *Let γ be a cost function defined using the function $g(\cdot, \cdot)$ (see Defn. 3.1). We say that the cost function Γ (as defined in Eqn. 2) satisfies the smoothness property if*

$$g_{\max} := \max\{g(n_1, n_2) \mid n_1 + n_2 = n\} = O\left(\frac{\kappa(n)}{n^2}\right),$$

where $\kappa(n)$ is the cost of a unit-weight clique of size n under the cost function Γ .

Fact 5.5. *The cost function introduced by Dasgupta (2016) satisfies the smoothness property.*

Theorem 5.6. *Let $\alpha_n = \omega(\sqrt{\log n/n})$. Let Γ be an admissible cost function satisfying the smoothness property (Defn. 5.4). Let k be a fixed constant and G be a graph generated from an HSBM (as per Defn. 5.1) where the underlying graph \tilde{G}_k has k nodes and the sparsity factor is α_n . Let ψ*

be the (hidden) function mapping the nodes of G to $[k]$ (the bottom-level clusters). For any binary tree T with n leaves labelled by the vertices of G , the following holds with high probability:

$$|\Gamma(T) - \mathbb{E}[\Gamma(T) \mid \psi]| \leq o(\mathbb{E}[\Gamma(T) \mid \psi]).$$

The expectation is taken only over the random choice of edges. In particular if T^* is a ground-truth tree for G , then, with high probability,

$$\Gamma(T^*) \leq (1 + o(1)) \min_{T'} \Gamma(T') = (1 + o(1)) OPT.$$

Proof. Our goal is to show that for any fixed cluster tree T' the cost is sharply concentrated around its expectation with an extremely high probability. We then apply the union bound over all possible cluster trees and obtain that in particular the cost of OPT is sharply concentrated around its expectation. Note that there are at most $2^{c \cdot n \log n}$ possible cluster trees (including labellings of the leaves to vertices of G), where c is a suitably large constant. Thus, it suffices to show that for any cluster tree T' we have

$$\mathbb{P}[|\Gamma(T') - \mathbb{E}[\Gamma(T') \mid \psi]| \geq o(\mathbb{E}[\Gamma(T') \mid \psi])] \leq \exp(-c^* n \log n),$$

where $c^* > c$.

Recall that for a given node N of T' with children N_1, N_2 , we have $\gamma(N) = w(V(N_1), V(N_2)) \cdot g(|V(N_1)|, |V(N_2)|)$ and $\Gamma(T') = \sum_{N \in T'} \gamma(N)$ (see Eqns. (2) and (3)). Let $Y_{i,j} = \mathbf{1}_{(i,j) \in E}$ for all $1 \leq i, j \leq n$ and observe that $\{Y_{i,j} \mid i < j\}$ are independent and $Y_{i,j} = Y_{j,i}$. Furthermore, let $Z_{i,j} = g(|V(\text{child}_1(N^{i,j}))|, |V(\text{child}_2(N^{i,j}))|) \cdot Y_{i,j}$, where $N^{i,j}$ is the node in T' separating nodes i and j and $\text{child}_1(N^{i,j})$ and $\text{child}_2(N^{i,j})$ are the two children of $N^{i,j}$. We can thus write

$$\Gamma(T') = \sum_{N \in T'} g(|V(\text{child}_1(N))|, |V(\text{child}_2(N))|) \sum_{\substack{i \in V(\text{child}_1(N)) \\ j \in V(\text{child}_2(N))}} Y_{i,j} \quad (6)$$

$$= \sum_{N \in T'} \sum_{\substack{i \in V(\text{child}_1(N)) \\ j \in V(\text{child}_2(N))}} Z_{i,j} \quad (7)$$

$$= \sum_{i < j} Z_{i,j}, \quad (8)$$

where we used that every potential edge i, j , $i \neq j$ appears in exactly one cut and that $Z_{i,j} = Z_{j,i}$. Observe that $\sum_{i < j} Z_{i,j}$ is a sum of independent random variables. Assume that the following claim holds.

Claim 5.7. *Let $w_{\min} = \Omega(1)$ be the minimum weight in \tilde{T}_k , the tree generating tree for \tilde{G}_k (see Defn. 5.1), i.e., $w_{\min} = \min_{N \in \tilde{T}_k} \tilde{W}(N)$ and recall that $g_{\max} = \max\{g(n_1, n_2) \mid n_1 + n_2 = n\}$. We have*

1. $\mathbb{E}[\Gamma(T') \mid \psi] \geq \kappa(n) \cdot \alpha_n \cdot w_{\min}$
2. $\sum_{i < j} g(|V(\text{child}_1(N^{i,j}))|, |V(\text{child}_2(N^{i,j}))|)^2 \leq g_{\max} \cdot \kappa(n)$

We defer the proof to later and first finish the proof of Theorem 5.6. We will make use of the slightly generalized version of Hoeffding bounds (see Hoeffding (1963)). For X_1, X_2, \dots, X_m

independent random variables satisfying $a_i \leq X_i \leq b_i$ for $i \in [n]$. Let $X = \sum_{i=1}^m X_i$, then for any $t > 0$

$$\mathbb{P}[|X - \mathbb{E}[X]| \geq t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^m (b_i - a_i)^2}\right). \quad (9)$$

By assumption, there exists a sequence $(y_n)_{n \geq 0}$ such that $\alpha_n = \omega\left(y_n \cdot \sqrt{\frac{\log n}{n}}\right)$ with $y_n = \omega(1)$.

We apply (9) with $t = \mathbb{E}[\Gamma(T') \mid \psi] \cdot \frac{y_n \cdot \sqrt{\frac{\log n}{n}}}{\alpha_n} = o(\mathbb{E}[\Gamma(T') \mid \psi])$ and derive

$$\begin{aligned} & \mathbb{P}\left[|\Gamma(T') - \mathbb{E}[\Gamma(T') \mid \psi]| \leq \mathbb{E}[\Gamma(T') \mid \psi] \cdot \frac{y_n \cdot \sqrt{\frac{\log n}{n}}}{\alpha_n}\right] \geq \\ & \geq 1 - \exp\left(-\frac{2\left(\mathbb{E}[\Gamma(T') \mid \psi] \cdot \frac{y_n \cdot \sqrt{\frac{\log n}{n}}}{\alpha_n}\right)^2}{\sum_{i < j} g(|V(N_1^{i,j})|, |V(N_2^{i,j})|)^2}\right) \\ & \geq 1 - \exp\left(-\frac{2 \cdot \kappa(n) \cdot w_{\min}^2 \cdot y_n^2 \cdot \log n}{g_{\max} \cdot n}\right) \\ & \geq 1 - \exp(-c^* \cdot n \log n), \end{aligned}$$

where the last inequality follows using Claim 5.7 and since $y_n = \omega(1)$. \square

We now turn to the proof of Claim 5.7.

Proof of Claim 5.7. Note that for any two vertices i, j of G , the edge (i, j) exists in G with probability at least $\alpha_n \cdot w_{\min}$. Thus, we have

$$\begin{aligned} \mathbb{E}[\Gamma(T') \mid \psi] &= \sum_{N \in T'} g(|V(\text{child}_1(N))|, |V(\text{child}_2(N))|) \sum_{\substack{i \in V(\text{child}_1(N)) \\ j \in V(\text{child}_2(N))}} \alpha_n \cdot w(i, j) \\ &\geq w_{\min} \cdot \alpha_n \cdot \sum_{N \in T'} g(|V(\text{child}_1(N))|, |V(\text{child}_2(N))|) |V(\text{child}_1(N))| \cdot |V(\text{child}_2(N))| \\ &= w_{\min} \cdot \alpha_n \cdot \kappa(n), \end{aligned}$$

where we made use of Eqn. (5). Furthermore, we have

$$\begin{aligned} & \sum_{i < j} g(|V(\text{child}_1(N^{i,j}))|, |V(\text{child}_2(N^{i,j}))|)^2 = \\ & = \sum_{N \in T'} g(|V(\text{child}_1(N))|, |V(\text{child}_2(N))|)^2 \cdot |V(\text{child}_1(N))| \cdot |V(\text{child}_2(N))| \\ & \leq g_{\max} \sum_{N \in T'} g(|V(\text{child}_1(N))|, |V(\text{child}_1(N))|) \cdot |V(\text{child}_1(N))| \cdot |V(\text{child}_2(N))| \\ & = g_{\max} \cdot \kappa(n), \end{aligned}$$

where we made use of Eqn. (5). \square

5.3 Algorithm for Clustering in the HSBM

In this section, we provide an algorithm for obtaining a hierarchical clustering of a graph generated from an HSBM. The algorithm is quite simple and combines approaches that are used in practice for hierarchical clustering: SVD projections and agglomerative heuristics. See Algorithm 2 for a complete description.

Theorem 5.8. *Let $\alpha_n = \omega(\sqrt{\log n/n})$. Let Γ be an admissible cost function (Defn. 3.1) satisfying the smoothness property (Defn 5.4). Let k be a fixed constant and G be a graph generated from an HSBM (as per Defn. 5.1) where the underlying graph \tilde{G}_k has k nodes and the sparsity factor is α_n . Let T be a ground-truth tree for G . With high probability, Algorithm 2 with parameter k on graph G outputs a tree T' that satisfies $\Gamma(T) \leq (1 + o(1))OPT$.*

Algorithm 2 Agglomerative Algorithm for Recovering Ground-Truth Tree of an HSBM Graph

- 1: **Input:** Graph $G = (V, E)$ generated from an HSBM.
 - 2: **Parameter:** A constant k .
 - 3: Apply (SVD) projection algorithm of (McSherry, 2001, Thm. 12) with parameters $G, k, \delta = |V|^{-2}$, to get $\zeta(1), \dots, \zeta(|V|) \in \mathbb{R}^{|V|}$ for vertices in V , where $\dim(\text{span}(\zeta(1), \dots, \zeta(|V|))) = k$.
 - 4: Run the single-linkage algorithm (Alg. 6) on the points $\{\zeta(1), \dots, \zeta(|V|)\}$ until there are exactly k clusters. Let $\mathcal{C} = \{C_1^\zeta, \dots, C_k^\zeta\}$ be the clusters (of points $\zeta(i)$) obtained. Let $C_i \subseteq V$ denote the set of vertices corresponding to the cluster C_i^ζ .
 - 5: **while** there are at least two clusters in \mathcal{C} **do**
 - 6: Take the pair of clusters C_i, C_j of \mathcal{C} that maximizes $\frac{\text{cut}(C_i, C_j)}{|C_i| \cdot |C_j|}$
 - 7: $\mathcal{C} \leftarrow \mathcal{C} \setminus \{C_i\} \setminus \{C_j\} \cup \{C_i \cup C_j\}$
 - 8: **end while**
 - 9: The sequence of merges in the while-loop (Steps 5 to 8) induces a hierarchical clustering tree on $\{C_1, \dots, C_k\}$, say T'_k with k leaves (represented by C_1, \dots, C_k). Replace each leaf of T'_k by an arbitrary binary tree on $|C_k|$ leaves labelled according to the vertices C_k to obtain T .
 - 10: Repeat the algorithm $k' = 2k \log n$ times. Let $T^1, \dots, T^{k'}$ be the corresponding hierarchical clustering trees.
 - 11: **Output:** Tree T^i (out of the k' candidates) that minimises $\Gamma(T_i)$.
-

Remark 5.9. *In an HSBM, k is a fixed constant. Thus, even if k is not known in advance, one can simply run the Algorithm 2 with all possible different values (constantly many) and return the solution with the minimal cost $\Gamma(T)$.*

Let $G = (V, E)$ be the input graph generated according to an HSBM. Let T be the tree output by Algorithm 2. We divide the proof into two claims that correspond to the outcome of Step 3 and the while-loop (Steps 5 to 8) of Algorithm 2.

We use a result of McSherry (2001) who considers a random graph model with k clusters that is (slightly) more general than the HSBM considered here. The difference is that there is no hierarchical structure on top of the k clusters in his setting; however, his goal is also simply to identify the k clusters and not any hierarchy upon them. The following theorem is derived from (McSherry, 2001, Obs. 11 and Thm.12).

Theorem 5.10 (McSherry (2001)). *Let s be the size of the smallest cluster (of the k clusters) and δ be the confidence parameter. Assume that for all u, v belonging to different clusters with adjacency vectors \mathbf{u}, \mathbf{v} (i.e., u_i is 1 if the edge (u, i) exists in G and 0 otherwise) satisfy*

$$\|\mathbb{E}[\mathbf{u}] - \mathbb{E}[\mathbf{v}]\|_2^2 \geq c \cdot k \cdot (n/s + \log(n/\delta))$$

for a large enough constant c , where $\mathbb{E}[\mathbf{u}]$ is the entry-wise expectation. Then, the algorithm of McSherry (2001, Thm. 12) with parameters G, k, δ projects the columns of the adjacency matrix of G to points $\{\zeta(1), \dots, \zeta(|V|)\}$ in a k -dimensional subspace of $\mathbb{R}^{|V|}$ such that the following holds w.p. at least $1 - \delta$ over the random graph G and with probability $1/k$ over the random bits of the algorithm. There exists $\eta > 0$ such that for any u in the i^{th} cluster and v in the j^{th} cluster:

1. if $i = j$ then $\|\zeta(u) - \zeta(v)\|_2^2 \leq \eta$;
2. if $i \neq j$ then $\|\zeta(u) - \zeta(v)\|_2^2 > 2\eta$,

Recall that $\psi : V \rightarrow [k]$ is the (hidden) labelling assigning each vertex of G to one of the k bottom-level clusters. Let $C_i^* = \{v \in V \mid \psi(v) = i\}$. Recall that $n_i = |V(C_i^*)|$.

The algorithm of (McSherry, 2001, Thm. 12) might fail for two reasons. The first reason is that the random choices by the algorithm yield an incorrect clustering. This happens w.p. at most $1 - 1/k$ and we can simply repeat the algorithm sufficiently many times to be sure that at least once we get the desired result, *i.e.*, the projections satisfy the conclusion of Thm. 5.10. Claims 5.13 and 5.14 show that in this case, Steps 5 to 8 of Alg. 2 produce a tree that has cost close to optimal. The second reason for failure is the randomness in the (random) edge choices. We need to carefully take these random edge choices into account, which are only made once and therefore cannot be repeated. We explicitly describe the bad events below and show that these occur with very low probability. Ultimately, the algorithm simply outputs a tree that has the least cost among all the ones produced (and one of them is guaranteed to have cost $(1 + o(1))\text{OPT}$) with high probability.

In order to apply our algorithm (and therefore also McSherry's algorithm) we need that the graph does not "deviate" too much from its expectation. For this reason we define the following three bad events. Let $\bar{\mathcal{E}}_1$ be the event that there exists i , such that $n_i < f_i n/2$, *i.e.*, at least one of the bottom-level clusters has size that is not representative. Let \mathcal{E}_1 be the complement of $\bar{\mathcal{E}}_1$. If \mathcal{E}_1 holds the term n/s that appears in Thm. 5.10 is a constant. The second bad event is that McSherry's algorithm fails due to the random choice of edges. This happens with probability at most δ (note that the condition in Theorem 5.10 depends on δ), which we set at $\delta = \frac{1}{|V|^2}$. We denote the complement of this event \mathcal{E}_2 . We remark that \mathcal{E}_2 depends only on random choices made in regards to edges (not node labels) and hence holds with probability $1 - \delta$ conditioned on \mathcal{E}_1 .

Furthermore, we define \mathcal{E}_3 as follows. Let C_1^*, \dots, C_k^* be the hidden bottom-level clusters, *i.e.*, $C_i^* = \{v \mid \psi(v) = i\}$; notice that the definition of $\{C_i^*\}$ depends only on the random assignments of labels to vertices and not the random choice of edges. For the partition C_1^*, \dots, C_k^* of V and for any S_1, S_2 , where S_1 and S_2 are disjoint and both sets are unions of some cluster sets from $\{C_1^*, \dots, C_k^*\}$ the following holds:

$$\left| \mathbb{E} \left[\frac{\text{cut}(S_1, S_2)}{|S_1| \cdot |S_2|} \right] - \frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} \right| \leq \frac{1}{\log n} \quad (10)$$

In Claim 5.11 we show that the event $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ occurs with high probability.

In order to prove Theorem 5.8 we establish the following claims.

Claim 5.11. *With high probability, the event $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ holds.*

Proof. We have the following

$$\mathbb{P}[\{\bar{\mathcal{E}}_1\} \cup \{\bar{\mathcal{E}}_2\} \cup \{\bar{\mathcal{E}}_3\}] \leq \mathbb{P}[\{\bar{\mathcal{E}}_1\} \cup \{\bar{\mathcal{E}}_3\}] + \mathbb{P}[\bar{\mathcal{E}}_2 \mid \mathcal{E}_1] \leq \mathbb{P}[\bar{\mathcal{E}}_1] + \mathbb{P}[\bar{\mathcal{E}}_3 \mid \mathcal{E}_1] + \mathbb{P}[\bar{\mathcal{E}}_2 \mid \mathcal{E}_1] \quad (11)$$

In the remainder of the proof we bound each of the terms on the RHS by $1/|V|^2$ which yields the claim.

Consider the first term on the RHS of (11). Observe that $\bar{\mathcal{E}}_1$ depends only on the random choices that assign labels to the vertices according to ψ . Thus, event $\bar{\mathcal{E}}_1$ occurs with a very low probability which is seen by a simple application of the Chernoff-Hoeffding bound, as $\mathbb{E}[n_i] = f_i n$ together with a union bound over all k clusters.

Consider the second term of (11). Let N be some internal node of any generating tree \tilde{T}_k and let $S_1 = V(\text{child}_1(N))$ and $S_2 = V(\text{child}_2(N))$. Note that both S_1 and S_2 are a disjoint union of some of the clusters $\{C_1^*, \dots, C_k^*\}$. Then notice that for any $u \in S_1$ and $v \in S_2$, the probability that the edge (u, v) exists in G is $\alpha_n \tilde{W}(N)$. Thus, conditioned on \mathcal{E}_1 and ψ , we have

$$\mathbb{E} \left[\frac{\text{cut}(S_i, S_j)}{|S_i||S_j|} \right] = \alpha_n \tilde{W}(N)$$

To complete the proof, we will show the following: Note that as the probability of any edge existing is $\Omega(\alpha_n)$, it must be the case that $\mathbb{E} \left[\frac{w(S_1, S_2)}{|S_1||S_2|} \right] = \Omega(\alpha_n)$. Furthermore, since $|S_1|, |S_2| = \Omega(n)$ by the Chernoff-Hoeffding bound for a single pair (S_1, S_2) the above holds with probability $\exp(-cn^2\alpha_n/\log n)$. Thus, even after taking a union bound over all $2^{O(k)}$ possible pairs of sets (S_1, S_2) that can be derived from said partition, Eqn. (12) and thus $\mathcal{E}_3 \mid \mathcal{E}_1$ holds with high probability.

Consider the third term of (11). Recall that $\bar{\mathcal{E}}_2 \mid \mathcal{E}_1$ is the event that that McSherry's algorithm fails due to the random choice of edges. This happens with probability at most δ which we set at $\delta = \frac{1}{|V|^2}$ (we show below that the condition of Thm. 5.10 is satisfied for this value of δ). Putting everything together yields the claim. \square

Observation 5.12. *Let $\alpha_n = \omega(\sqrt{\log n/n})$. Let G be generated by an HSBM. Assume that event \mathcal{E}_1 occurs. Let u, v be two nodes such that $i = \psi(u) \neq \psi(v) = j$. Let \mathbf{u} and \mathbf{v} denote the random variables corresponding to the columns of u and v in the adjacency matrix of G . Then,*

$$\|\mathbb{E}[\mathbf{u} \mid \{\mathcal{E}_1\}] - \mathbb{E}[\mathbf{v} \mid \{\mathcal{E}_1\}]\|_2^2 = \omega(\log n)$$

The expectations are with respect to the random choice of the edges.

Proof. As we know that the event \mathcal{E}_1 occurs, we may conclude that $s = \min_i n_i \geq \frac{n}{2} \min_i f_i$. Thus, $n/s \leq \frac{2}{f_{\min}}$, where $f_{\min} = \min_i f_i$ and hence n/s is bounded by some fixed constant.

Let $q = \tilde{W}(N)$ where N is the LCA $_{\tilde{T}_k}(\tilde{\sigma}^{-1}(i), \tilde{\sigma}^{-1}(j))$ in \tilde{T}_k , the generating tree for \tilde{G}_k used in defining the HSBM. Assuming \mathcal{E}_1 and taking expectations only with respect to the random choice of edges, we have:

$$\begin{aligned} \|\mathbb{E}[\mathbf{u} \mid \mathcal{E}_1] - \mathbb{E}[\mathbf{v} \mid \mathcal{E}_1]\|_2^2 &\geq n_i \alpha_n^2 (p_i - q)^2 / 2 + n_j \alpha_n^2 (p_j - q)^2 / 2 \\ &= \Omega(\alpha_n^2 n) = \omega(\log n) \end{aligned}$$

Above we used that $p_i - q > 0$ and $n_i = \Omega(n)$ for each i . \square

Claim 5.13. *Let $\alpha_n = \omega(\sqrt{\log n/n})$. Let G be generated by an HSBM. Let C_1^*, \dots, C_k^* be the hidden bottom-level clusters, i.e., $C_i^* = \{v \mid \psi(v) = i\}$. Assume that event \mathcal{E}_1 and \mathcal{E}_2 occur. With probability at least $\Omega(1/k)$, the clusters obtained after Step 4 correspond to the assignment ψ , i.e., there exists a permutation $\pi : [k] \rightarrow [k]$, such that $C_j = C_{\pi(j)}^*$.*

Proof. Since we condition on \mathcal{E}_2 , Step 3 can only fail due to the coin flips of the underlying algorithm which happens with probability at most $1 - 1/k$, by Theorem 5.10.

Thus, with probability at least $\Omega(1/k)$ the conclusions of Theorem 5.10 hold. In the rest of the proof we assume that the following holds: There exists $\eta > 0$ such that for any pair of nodes u, v we have

1. if $\psi(u) = \psi(v)$ then $\|\zeta(u) - \zeta(v)\|_2^2 \leq \eta$;
2. if $\psi(u) \neq \psi(v)$ then $\|\zeta(u) - \zeta(v)\|_2^2 > 2\eta$.

Therefore, any linkage algorithm, *e.g.*, single linkage (See Alg. 6), performing merges starting from the set $\{\zeta(1), \dots, \zeta(n)\}$ until there are k clusters will merge clusters at a distance of at most η and hence, the clusters obtained after Step 4 correspond to the assignment ψ . This yields the claim. \square

Claim 5.14. *Let $\alpha_n = \omega(\sqrt{\log n/n})$. Let G be generated according to an HSBM and let T^* be a ground-truth tree for G . Assume that the events $\mathcal{E}_1, \mathcal{E}_2$, and \mathcal{E}_3 occur and that furthermore, the clusters obtained after Step 4 correspond to the assignment ψ , *i.e.*, there exists a permutation $\pi : [k] \rightarrow [k]$ such that for each $v \in C_i$, $\psi(v) = \pi(i)$. Then, the sequence of merges in the while-loop (Steps 5 to 8) followed by Step 9 produces *w.p.* $\Omega(1/k)$ a tree T such that $\Gamma(T) \leq (1 + o(1))OPT$.*

Proof. Steps 5 to 8 based on the true expectations instead of the empirical estimates. Then essentially, the algorithm is performing some linkage algorithm (*i.e.*, average, single, or complete-linkage) on a ground-truth input and hence is guaranteed to recover the generating tree (see Theorem. 7.1).

For simplicity of notation, we will assume that π is the identity permutation, *i.e.*, the algorithm has not only identified the true clusters correctly but also guessed the correct label. This only makes the notation less messy, though the proof is essentially unchanged.

As $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ holds, we have the following: For the partition $\{C_1^*, \dots, C_k^*\}$ of V , where $C_i^* = \{v \mid \psi(v) = i\}$, and for any S_1, S_2 , where S_1 and S_2 are disjoint and are both unions of some of the clusters $\{C_1^*, \dots, C_k^*\}$, the following holds:

$$\left| \mathbb{E} \left[\frac{\text{cut}(S_1, S_2)}{|S_1| \cdot |S_2|} \right] - \frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} \right| \leq \frac{1}{\log n} \quad (12)$$

We restrict our attention to sets S_1, S_2 that are of the form that $S_1 = V(\text{child}_1(N))$ and $S_2 = V(\text{child}_2(N))$ for some internal node N of some generating tree \tilde{T}_k of the graph \tilde{G}_k used to generate G . Then for pairs (S_1, S_2) and (S'_1, S'_2) both of this form, the following holds: there exists $\frac{3}{\log n} > \eta > 0$ such that

1. If $\mathbb{E} \left[\frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} \right] = \mathbb{E} \left[\frac{w(S'_1, S'_2)}{|S'_1| \cdot |S'_2|} \right]$, then $\left| \frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} - \frac{w(S'_1, S'_2)}{|S'_1| \cdot |S'_2|} \right| \leq \eta$
2. If $\mathbb{E} \left[\frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} \right] \neq \mathbb{E} \left[\frac{w(S'_1, S'_2)}{|S'_1| \cdot |S'_2|} \right]$, then $\left| \frac{w(S_1, S_2)}{|S_1| \cdot |S_2|} - \frac{w(S'_1, S'_2)}{|S'_1| \cdot |S'_2|} \right| > 2\eta$

Note that the above conditions are enough to ensure that the algorithm performs the same steps as with perfect inputs, up to an arbitrary choice of tie-breaking rule. Since Theorem 7.1 is true, no matter the tie breaking rule chosen, the proof of the claim follows. \square

We are ready to prove Theorem 5.8.

Proof of Theorem 5.8. Conditioning on $\mathcal{E}_1 \cap \mathcal{E}_2 \mathcal{E}_3$ which occurs w.h.p. we get from Claims 5.13 and 5.14 that w.p. at least $\Omega(1/k)$ the tree T^i (the i^{th} run of the algorithm) obtained in step 9 satisfies $\Gamma(T^i) \leq (1 + o(1))OPT$. It is possible to boost this probability by running Algorithm 2 multiple times. Running it $\Omega(k \log n)$ times and taking the tree with the smallest $\Gamma(T^i)$ yields the result. \square

6 Dissimilarity-Based Inputs: Approximation Algorithms

In this section, we consider general dissimilarity inputs and admissible objective for these inputs. For ease of exposition, we focus on an particular admissible objective function for dissimilarity inputs. Find T maximizing the value function corresponding to Dasgupta's cost function of Section 4: $\text{val}(T) = \sum_{N \in T} \text{val}(N)$ where for each node N of T with children N_1, N_2 , $\text{val}(N) = w(V(N_1), V(N_2)) \cdot V(N)$. This optimization problem is NP-Hard Dasgupta (2016), hence we focus on approximation algorithms.

We show (Theorem 6.2) that average-linkage achieves a 2 approximation for the problem. We then introduce a simple algorithm based on *locally-densest cuts* and show (Theorem 6.5) that it achieves a $3/2 + \varepsilon$ approximation for the problem.

We remark that our proofs show that for any admissible objective function, those algorithms have approximation guarantees, but the approximation guarantee depends on the objective function.

We start with the following elementary upper bound on OPT.

Fact 6.1. *For any graph $G = (V, E)$, and weight function $w : E \rightarrow \mathbb{R}_+$, we have $OPT \leq n \cdot \sum_{e \in E} w(e)$.*

6.1 Average-Linkage

We show that average-linkage is a 2-approximation in the dissimilarity setting.

Theorem 6.2. *For any graph $G = (V, E)$, and weight function $w : E \rightarrow \mathbb{R}_+$, the average-linkage algorithm (Algorithm 3) outputs a solution of value at least $n \sum_{e \in E} w(e)/2 \geq OPT/2$.*

Algorithm 3 Average-Linkage Algorithm for Hierarchical Clustering (dissimilarity setting)

- 1: **Input:** Graph $G = (V, E)$ with edge weights $w : E \mapsto \mathbb{R}_+$
 - 2: Create n singleton trees.
 - 3: **while** there are at least two trees **do**
 - 4: Take trees roots N_1 and N_2 minimizing $\sum_{x \in V(N_1), y \in V(N_2)} w(x, y) / (|V(N_1)| |V(N_2)|)$
 - 5: Create a new tree with root N and children N_1 and N_2
 - 6: **end while**
 - 7: **return** the resulting binary tree T
-

When two trees are chosen at Step 4 of Algorithm 3, we say that they are *merged*. We say that all the trees considered at the beginning of an iteration of the while loop are the trees that are *candidate for the merge* or simply the *candidate trees*.

We first show the following lemma and then prove the theorem.

Lemma 6.3. *Let T be the output tree and A, B be the children of the root. We have,*

$$\frac{w(V(A), V(B))}{|V(A)| \cdot |V(B)|} \geq \frac{w(V(A))}{|V(A)| \cdot (|V(A)| - 1)} + \frac{w(V(B))}{|V(B)| \cdot (|V(B)| - 1)}.$$

Proof. Let $a = |V(A)|(|V(A)| - 1)/2$ and $b = |V(B)|(|V(B)| - 1)/2$. For any node N_0 of T , let $\text{child}_1(N_0)$ and $\text{child}_2(N_0)$ be the two children of N_0 . We first consider the subtree T_A of T rooted at A . We have

$$\begin{cases} w(V(A)) &= \sum_{A_0 \in T_A} w(V(\text{child}_1(A_0)), V(\text{child}_2(A_0))), \\ a &= \sum_{A_0 \in T_A} |V(\text{child}_1(A_0))| \cdot |V(\text{child}_2(A_0))|. \end{cases}$$

By an averaging argument, there exists $A' \in T_A$ with children A_1, A_2 such that

$$\frac{w(V(A_1), V(A_2))}{|V(A_1)| \cdot |V(A_2)|} \geq \frac{w(V(A))}{a}. \quad (13)$$

We now consider the iteration of the while loop at which the algorithm merged the trees A_1 and A_2 . Let A_1, A_2, \dots, A_k and B_1, B_2, \dots, B_ℓ be the trees that were candidate for the merge at that iteration, and such that $V(A_i) \cap V(B) = \emptyset$ and $V(B_i) \cap V(A) = \emptyset$. Observe that the leaves sets of those trees form a partition of the sets $V(A)$ and $V(B)$, so

$$\begin{cases} w(A, B) &= \sum_{i,j} w(V(A_i), V(B_j)), \\ |V(A)| \cdot |V(B)| &= \sum_{i,j} |V(A_i)| \cdot |V(B_j)|. \end{cases}$$

By an averaging argument again, there exists A_i, B_j such that

$$\frac{w(V(A_i), V(B_j))}{|V(A_i)| \cdot |V(B_j)|} \leq \frac{w(V(A), V(B))}{|V(A)| \cdot |V(B)|}. \quad (14)$$

Now, since the algorithm merged A_1, A_2 rather than A_i, B_j , by combining Eq. 13 and 14, we have

$$\frac{w(V(A))}{a} \leq \frac{w(V(A_1), V(A_2))}{|V(A_1)| \cdot |V(A_2)|} \leq \frac{w(V(A_i), V(B_j))}{|V(A_i)| \cdot |V(B_j)|} \leq \frac{w(V(A), V(B))}{|V(A)| \cdot |V(B)|}.$$

Applying the same reasoning to B and taking the sum yields the lemma. \square

Proof of Theorem 6.2. We proceed by induction on the number of the nodes n in the graph. Let A, B be the children of the root of the output tree T . By induction,

$$\text{val}(T) \geq (|V(A)| + |V(B)|) \cdot w(V(A), V(B)) + \frac{|V(A)|}{2} w(V(A)) + \frac{|V(B)|}{2} w(V(B)). \quad (15)$$

Lemma 6.3 implies $(|V(A)| + |V(B)|) \cdot w(V(A), V(B)) \geq |V(B)| w(V(A)) + |V(A)| w(V(B))$. Dividing both sides by 2 and plugging it into (15) yields

$$\text{val}(T) \geq \frac{|V(A)| + |V(B)|}{2} w(V(A), V(B)) + \frac{|V(A)| + |V(B)|}{2} (w(V(A)) + w(V(B))).$$

Observing that $n = |V(A)| + |V(B)|$ and combining $\sum_{e \in E} w(e) = w(V(A), V(B)) + w(V(A)) + w(V(B))$ with Fact 6.1 completes the proof. \square

6.2 A Simple and Better Approximation Algorithm for Worst-Case Inputs

In this section, we introduce a very simple algorithm (Algorithm 5) that achieves a better approximation guarantee. The algorithm follows a divisive approach by recursively computing locally-densest cuts using a local search heuristic (see Algorithm 4). This approach is similar to the recursive-sparsest-cut algorithm of Section 4. Here, instead of trying to solve the densest cut problem (and so being forced to use approximation algorithms), we solve the simpler problem of computing a locally-densest cut. This yields both a very simple local-search-based algorithm and a good approximation guarantee.

We use the notation $A \oplus x$ to mean the set obtained by adding x to A if $x \notin A$, and by removing x from A if $x \in A$. We say that a cut (A, B) is a ε/n -locally-densest cut if for any x ,

$$\frac{w(A \oplus x, B \oplus x)}{|A \oplus x| \cdot |B \oplus x|} \leq \left(1 + \frac{\varepsilon}{n}\right) \frac{w(A, B)}{|A||B|}.$$

The following local search algorithm computes an ε/n -locally-densest cut.

Algorithm 4 Local Search for Densest Cut

- 1: **Input:** Graph $G = (V, E)$ with edge weights $w : E \mapsto \mathbb{R}_+$
 - 2: Let (u, v) be an edge of maximum weight
 - 3: $A \leftarrow \{v\}$, $B \leftarrow V \setminus \{v\}$
 - 4: **while** $\exists x: \frac{w(A \oplus x, B \oplus x)}{|A \oplus x| \cdot |B \oplus x|} > (1 + \varepsilon/n) \frac{w(A, B)}{|A||B|}$ **do**
 - 5: $A \leftarrow A \oplus x$, $B \leftarrow B \oplus x$
 - 6: **end while**
 - 7: **return** (A, B)
-

Theorem 6.4. *Algorithm 4 computes an ε/n -locally-densest cut in time $\tilde{O}(n(n+m)/\varepsilon)$.*

Proof. The proof is straightforward and given for completeness. By definition, the algorithm computes an ε/n -locally densest cut so we only need to argue about the running time. The weight of the cut is initially at least w_{\max} , the weight of the maximum edge weight, and in the end at most nw_{\max} . Since the weight of the cut increases by a factor of $(1 + \varepsilon/n)$ at each iteration, the total number of iterations of the while loop is at most $\log_{1+\varepsilon/n}(nw_{\max}/w_{\max}) = \tilde{O}(n/\varepsilon)$. Each iteration takes time $O(m+n)$, so the running time of the algorithm is $\tilde{O}(n(m+n)/\varepsilon)$. \square

Algorithm 5 Recursive Locally-Densest-Cut for Hierarchical Clustering

- 1: **Input:** Graph $G = (V, E)$, with edge weights $w : E \mapsto \mathbb{R}_+$, $\varepsilon > 0$
 - 2: Compute an ε/n -locally-densest cut (A, B) using Algorithm 4
 - 3: Recurse on $G[A]$ and $G[B]$ to obtain rooted trees T_A and T_B .
 - 4: Return the tree T whose root node has two children, T_A and T_B .
-

Theorem 6.5. *Algorithm 5 returns a tree of value at least*

$$\frac{2n}{3}(1 - \varepsilon) \sum_e w(e) \geq \frac{2}{3}(1 - \varepsilon) OPT,$$

in time $\tilde{O}(n^2(n+m)/\varepsilon)$.

The proof relies on the following lemma.

Lemma 6.6. *Let (A, B) be an ε/n -locally-densest cut. Then,*

$$(|A| + |B|)w(A, B) \geq 2(1 - \varepsilon)(|B|w(A) + |A|w(B)).$$

Proof. Let $v \in A$. By definition of the algorithm,

$$(1 + \varepsilon/n) \frac{w(A, B)}{|A||B|} \geq \frac{w(A \setminus \{v\}, B \cup \{v\})}{(|A| - 1)(|B| + 1)}.$$

Rearranging,

$$\frac{(|A| - 1)(|B| + 1)}{|A||B|} (1 + \varepsilon/n)w(A, B) \geq w(A \setminus \{v\}, B \cup \{v\}) = w(A, B) + w(v, A) - w(v, B)$$

Summing over all vertices of A , we obtain

$$|A| \frac{(|A| - 1)(|B| + 1)}{|A||B|} (1 + \varepsilon/n)w(A, B) \geq |A|w(A, B) + 2w(A) - w(A, B).$$

Rearranging and simplifying,

$$(|A| - 1)(|B| + 1) \frac{\varepsilon}{n} w(A, B) + (|A| - 1)(1 + \varepsilon/n)w(A, B) \geq 2|B|w(A).$$

Since $|B| + 1 \leq n$, this gives

$$|A|w(A, B) \geq 2(1 - \varepsilon)|B|w(A).$$

Proceeding similarly with B and summing the two inequalities yields the lemma. \square

Proof of Theorem 6.5. We first show the approximation guarantee. We proceed by induction on the number of vertices. The base case is trivial. By inductive hypothesis,

$$\text{val}(T) \geq nw(A, B) + \frac{2}{3} \cdot (1 - \varepsilon)(|A|w(A) + |B|w(B)),$$

where $n = |A| + |B|$. Lemma 6.6 implies

$$nw(A, B) = (|A| + |B|)w(A, B) \geq 2(1 - \varepsilon)(|B|w(A) + |A|w(B)).$$

Hence,

$$\frac{|A| + |B|}{3} w(A, B) \geq \frac{2}{3} (1 - \varepsilon)(|B|w(A) + |A|w(B)).$$

Therefore,

$$\text{val}(T) \geq \frac{2n}{3} (1 - \varepsilon)(w(A, B) + w(A) + w(B)) = (1 - \varepsilon) \frac{2n}{3} \sum_e w(e).$$

To analyze the running time, observe that by Theorem 6.4, a recursive call on a graph $G' = (V', E')$ takes time $\tilde{O}(|V'|(|V'| + |E'|)/\varepsilon)$ and that the depth of the recursion is $O(n)$. \square

Remark 6.7. *The average-linkage and the recursive locally-densest-cut algorithms achieve an $O(g_n)$ - and $O(h_n)$ -approximation respectively, for any admissible cost function f , where $g_n = \max_n f(n)/f(\lceil n/2 \rceil)$. $h_n = \max_n f(n)/f(\lceil 2n/3 \rceil)$. An almost identical proof yields the result.*

Remark 6.8. *In Section 8, we show that other commonly used algorithms, such as complete-linkage, single-linkage, or bisection 2-Center, can perform arbitrarily badly (see Theorem 8.6). Hence average-linkage is more robust in that sense.*

7 Perfect Ground-Truth Inputs and Beyond

In this section, we focus on ground-truth inputs. We state that when the input is a perfect ground-truth input, commonly used algorithms (single linkage, average linkage, and complete linkage; as well as some divisive algorithms – the bisection k -Center and sparsest-cut algorithms) yield a tree of optimal cost, hence (by Definition 3.1) a ground-truth tree. Some of those results are folklore (and straightforward when there are no ties), but we have been unable to pin down a reference, so we include them for completeness (Section 7.1). We also introduce a faster optimal algorithm for “strict” ground-truth inputs (Section 7.2). The proofs present no difficulty. The meat of this section is Subsection 7.3, where we go beyond ground-truth inputs; we introduce δ -adversarially-perturbed ground-truth inputs and design a simple, more robust algorithm that, *for any* admissible objective function, is a δ -approximation.

7.1 Perfect Ground-Truth Inputs are Easy

Algorithm 6 Linkage Algorithm for Hierarchical Clustering (similarity setting)

- 1: **Input:** A graph $G = (V, E)$ with edge weights $w : E \mapsto \mathbb{R}_+$
 - 2: Create n singleton trees. Root labels: $\mathcal{C} = \{\{v_1\}, \dots, \{v_n\}\}$
 - 3: Define $\text{dist} : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}_+$: $\text{dist}(C_1, C_2) = \begin{cases} \frac{1}{|C_1||C_2|} \sum_{x \in C_1, y \in C_2} w((x, y)) & \text{Average Linkage} \\ \min_{x \in C_1, y \in C_2} w((x, y)) & \text{Single Linkage} \\ \max_{x \in C_1, y \in C_2} w((x, y)) & \text{Complete Linkage} \end{cases}$
 - 4: **while** there are at least two trees **do**
 - 5: Take the two trees with root labels C_1, C_2 such that $\text{dist}(C_1, C_2)$ is maximum
 - 6: Create a new tree by making those two tree children of a new root node labeled $C_1 \cup C_2$
 - 7: Remove C_1, C_2 from \mathcal{C} , add $C_1 \cup C_2$ to \mathcal{C} , and update dist
 - 8: **end while**
 - 9: **return** the resulting binary tree T
-

In the following, we refer to the *tie breaking* rule of Algorithm 6 as the rule followed by the algorithm for deciding which of C_i, C_j or C_k, C_ℓ to merge, when $\max_{C_1, C_2 \in \mathcal{C}} \text{dist}(C_1, C_2) = \text{dist}(C_i, C_j) = \text{dist}(C_k, C_\ell)$.

Theorem 7.1.¹⁵ *Assume that the input is a (dissimilarity or similarity) ground-truth input. Then, for any admissible objective function, the agglomerative heuristics average-linkage, single-linkage, and complete-linkage (see Algorithm 6) return an optimal solution. This holds no matter the tie breaking rule of Algorithm 6.*

Proof. We focus on the similarity setting; the proof for the dissimilarity setting is almost identical. We define the *candidate trees* after t iterations of the while loop to be sets of trees in \mathcal{C} at that time. The theorem follows from the following statement, which we will prove by induction on t : If $C^t = \{C_1, \dots, C_k\}$ denotes the set of clusters after t iterations, then there exists a generating tree T^t for G , such that the candidate trees are subtrees of T^t .

For the base case, initially each candidate tree contains exactly one vertex and the statement holds. For the general case, let C_1, C_2 be the two trees that constitute the t^{th} iteration. By induction, there exists a generating tree T^{t-1} for G , and associated weights W^{t-1} (according to

¹⁵This Theorem may be folklore, at least when there are no ties, but we have been unable to find a reference.

Definition 2.2) such that C_1 and C_2 are subtrees of T^{t-1} , rooted at nodes N^1 and N^2 of T^{t-1} respectively.

To define T^t , we start from T^{t-1} . Consider the path $P = \{N^1, N_1, N_2, \dots, N_k, N^2\}$ joining N^1 to N^2 in T^{t-1} and let $N_r = \text{LCA}_{T^{t-1}}(N^1, N^2)$. If N_r is the parent of N_1 and N_2 , then $T^t = T^{t-1}$, else do the following transformation: remove the subtrees rooted at N^1 and at N^2 ; create a new node N^* as second child of N_k , and let N_1 and N_2 be its children. This defines T^t . To define W^t , extend W^{t-1} by setting $W^t(N^*) = W(N_r)$.

Claim 7.2. For any $N_i, N_j \in P$, $W^{t-1}(N_i) = W^{t-1}(N_j)$.

Thanks to the inductive hypothesis, with Claim 7.2 it is easy to verify that W^t certifies that T^t is generating for G . \square

Proof of Claim 7.2. Fix a node N_i on the path from N_r to N^1 (the argument for nodes on the path from N_r to N^2 is similar). By induction $W^{t-1}(N_i) \geq W^{t-1}(N_r)$. We show that since the linkage algorithms merge the trees C_1 and C_2 , we also have $W^{t-1}(N_i) \leq W^{t-1}(N_r)$ and so $W^{t-1}(N_i) = W^{t-1}(N_r)$, hence the claim. Let $w_0 = W^{t-1}(N_r)$.

By induction, for all $u \in C_1, v \in C_2$, $w(u, v) = w_0$, and thus $\text{dist}(C_1, C_2) = w_0$ in the execution of all the algorithms. Fix a candidate tree $C' \in \mathcal{C}^t$, $C' \neq C_1, C_2$ and $C' \subseteq V(N_i)$. Since \mathcal{C} is a partition of the vertices of the graph and since candidate trees are subtrees of T^{t-1} , such a cluster exists. Thus, for $u \in C_1, v \in C'$ $w(u, v) = W^{t-1}(\text{LCA}_{T^{t-1}}(u, v)) = W^{t-1}(N_i) \geq w_0$ since N_i is a descendant of N_r .

It is easy to check that by their definitions, for any of the linkage algorithms, we thus have that $\text{dist}(C_1, C') \geq w_0 = \text{dist}(C_1, C_2)$. But since the algorithms merge the clusters at maximum distance, it follows that $\text{dist}(C_1, C') \leq \text{dist}(C_1, C_2) = w_0$ and therefore, $W^{t-1}(N_i) \leq W^{t-1}(N_r)$ and so, $W^{t-1}(N_i) = W^{t-1}(N_r)$ and the claim follows. This is true no matter the tie breaking chosen for the linkage algorithms. \square

Divisive Heuristics. In this section, we focus on two well-known divisive heuristics: (1) the bisection 2-Center which uses a partition-based clustering objective (the k -Center objective) to divide the input into two (non necessarily equal-size) parts (see Algorithm 7), and (2) the recursive sparsest-cut algorithm, which can be implemented efficiently for ground-truth inputs (Lemma 7.6).

Algorithm 7 Bisection 2-Center (similarity setting)

- 1: **Input:** A graph $G = (V, E)$ and a weight function $w : E \mapsto \mathbb{R}_+$
 - 2: Find $\{u, v\} \subseteq V$ that maximizes $\min_x \max_{y \in \{u, v\}} w(x, y)$
 - 3: $A \leftarrow \{x \mid w(x, u) \geq \max_{y \in \{u, v\}} w(x, y)\}$
 - 4: $B \leftarrow V \setminus A$.
 - 5: Apply Bisection 2-Center on $G[A]$ and $G[B]$ to obtain trees T_A, T_B respectively
 - 6: **return** The union tree of T_A, T_B .
-

Loosely speaking, we show that this algorithm computes an optimal solution if the optimal solution is unique. More precisely, for any similarity graph G , we say that a tree T is *strictly generating* for G if there exists a weight function W such that for any nodes N_1, N_2 , if N_1 appears on the path from N_2 to the root, then $W(N_1) < W(N_2)$ and for every $x, y \in V$, $w(x, y) = W(\text{LCA}_T(x, y))$. In this case we say that the input is a strict ground-truth input. In the context of dissimilarity, an analogous notion can be defined and we obtain a similar result.

Theorem 7.3. ¹⁶ *For any admissible objective function, the bisection 2-Center algorithm returns an optimal solution for any similarity or dissimilarity graph G that is a strict ground-truth input.*

Proof. We proceed by induction on the number of nodes in the graph. Consider a strictly generating tree T and the corresponding weight function W . Consider the root node N_r of T and let N_1, N_2 be the children of the root. Let (α, β) be the cut induced by the root node of T (i.e., $\alpha = V(N_1)$, $\beta = V(N_2)$). Define w_0 to be the weight of an edge between $u \in \alpha$ and $v \in \beta$ for any u, v (recall that since T is strictly generating all the edges between α and β are of same weight). We show that the bisection 2-Centers algorithm divides the graph into α and β . Applying the inductive hypothesis on both subgraphs yields the result.

Suppose that the algorithm locates the two centers in β . Then, $\min_x \max_{y \in \{u, v\}} w(x, y) = w_0$ since the vertices of α are connected by an edge of weight w_0 to the centers. Thus, the value of the clustering is w_0 . Now, consider a clustering consisting of a center c_0 in α and a center c_1 in β . Then, for each vertex u , we have $\max_{c \in \{c_0, c_1\}} w(u, c) \geq \min(W(N_1), W(N_2)) > W(N_r) = w_0$ since T and W are strictly generating; Hence a strictly better clustering value. Therefore, the algorithm locates $x \in \alpha$ and $y \in \beta$. Finally, it is easy to see that the partitioning induced by the centers yields parts $A = \alpha$ and $B = \beta$. \square

Remark 7.4. *To extend our result to (non-strict) ground-truth inputs, one could consider the following variant of the algorithm (which bears similarities with the popular elbow method for partition-based clustering): Compute a k -Center clustering for all $k \in \{1, \dots, n\}$ and partition the graph according to the k -Center clustering of the smallest $k > 1$ for which the value of the clustering increases. Mimicking the proof of Theorem 7.3, one can show that the tree output by the algorithm is generating.*

We now turn to the recursive sparsest-cut algorithm (i.e., the recursive ϕ -sparsest-cut algorithm of Section 4, for $\phi = 1$). The recursive sparsest-cut consists in recursively partitioning the graph according to a sparsest cut of the graph. We show (1) that this algorithm yields a tree of optimal cost and (2) that computing a sparsest cut of a similarity graph generated from an ultrametric can be done in linear time. Finally, we observe that the analogous algorithm for the dissimilarity setting consists in recursively partitioning the graph according to the densest cut of the graph and achieves similar guarantees (and similarly the densest cut of a dissimilarity graph generated from an ultrametric can be computed in linear time).

Theorem 7.5. ¹⁷ *For any admissible objective function, the recursive sparsest-cut (respectively densest-cut) algorithm computes a tree of optimal cost if the input is a similarity (respectively dissimilarity) ground-truth input.*

Proof. The proof, by induction, has no difficulty and it may be easier to recreate it than to read it.

Let T be a generating tree and W be the associated weight function. Let N_r be the root of T , N_1, N_2 the children of N_r , and $(\alpha = V(N_1), \beta = V(N_2))$ the induced root cut. Since T is strictly generating, all the edges between α and β are of same weight w , which is therefore also the sparsity of (α, β) . For every edge (u, v) of the graph, $w(u, v) = W(\text{LCA}_T(u, v)) \geq w$, so every cut has sparsity at least w , so (α, β) has minimum sparsity.

Now, consider the tree T^* computed by the algorithm, and let (γ, δ) denote the sparsest-cut used by the algorithm at the root (in case of ties it might not be different from (α, β)). By induction the algorithm on $G[\gamma]$ and $G[\delta]$ gives two generating trees T_γ and T_δ with associated weight functions

¹⁶This Theorem may be folklore, but we have been unable to find a reference.

¹⁷This Theorem may be folklore, at least when there are no ties, but we have been unable to find a reference.

W_γ and W_δ . To argue that T^* is generating, we define W^* as follows, where N_r^* denotes the root of T^* .

$$W^*(N) = \begin{cases} W_\gamma(N) & \text{if } N \in T_\gamma \\ W_\delta(N) & \text{if } N \in T_\delta \\ w & \text{if } N = N_r^* \end{cases}$$

By induction $w(u, v) = W(\text{LCA}_T(u, v))$ if either both $u, v \in \gamma$, or both $u, v \in \delta$. For any $u \in \gamma, v \in \delta$, we have $w(u, v) = w = W(N_r^*) = W(\text{LCA}_T(u, v))$. Finally, since $w \leq w(u, v)$ for any u, v , we have $W(N_r^*) = w \leq W(N)$, for any $N \in T^*$, and therefore T^* is generating. \square

We then show how to compute a sparsest-cut of a graph that is a ground-truth input.

Lemma 7.6. *If the input graph is a ground-truth input then the sparsest cut is computed in $O(n)$ time by the following algorithm: pick an arbitrary vertex u , let w_{\min} be the minimum weight of edges adjacent to u , and partition V into $A = \{x \mid w(u, x) > w_{\min}\}$ and $B = V \setminus A$.*

Proof. Let $w_{\min} = w(u, v)$. We show that $w(A, B)/(|A||B|) = w_{\min}$ and since w_{\min} is the minimum edge weight of the graph, that the cut (A, B) only contains edges of weight w_{\min} . Fix a generating tree T . Consider the path from u to the root of T and let N_0 be the first node on the (bottom-up) path such that $W(N_0) = w_{\min}$. For any vertex $x \in A$, we have that $w(u, x) > w_{\min}$. Hence by definition, we have that N_0 is an ancestor of $\text{LCA}_T(u, x)$. Therefore, for any other node y such that $w(u, y) = w_{\min}$, we have $\text{LCA}_T(u, y) = \text{LCA}_T(x, y)$ and so, $w(x, y) = W(\text{LCA}_T(x, y)) = W(\text{LCA}_T(u, y)) = w_{\min}$. It follows that all the edges in the cut (A, B) are of weight w_{\min} and so, the cut is a sparsest cut. \square

7.2 A Near-Linear Time Algorithm

In this section, we propose a simple, optimal, algorithm for computing a generating tree of a ground-truth input. For any graph G , the running time of this algorithm is $O(n^2)$, and $\tilde{O}(n)$ if there exists a tree T that is *strictly generating* for the input. For completeness we recall that for any graph G , we say that a tree T is strictly generating for G if there exists a weight function W such that for any nodes N_1, N_2 , if N_1 appears on the path from N_2 to the root, then $W(N_1) < W(N_2)$ and for every $x, y \in V$, $w(x, y) = W(\text{LCA}_T(x, y))$. In this case we say that the inputs is a *strict ground-truth input*.

The algorithm is described for the similarity setting but could be adapted to the dissimilarity case to achieve the same performances.

Algorithm 8 Fast and Simple Algorithm for Hierarchical Clustering on Perfect Data (similarity setting)

- 1: **Input:** A graph $G = (V, E)$ and a weight function $w : E \mapsto \mathbb{R}_+$
 - 2: $p \leftarrow$ random vertex of V
 - 3: Let $w_1 > \dots > w_k$ be the edge weights of the edges that have p as an endpoint
 - 4: Let $B_i = \{v \mid w(p, v) = w_i\}$, for $1 \leq i \leq k$.
 - 5: Apply the algorithm recursively on each $G[B_i]$ and obtain a collection of trees T_1, \dots, T_k
 - 6: Define T_0^* as a tree with p as a single vertex
 - 7: For any $1 \leq i \leq k$, define T_i^* to be the union of T_{i-1}^* and T_i
 - 8: Return T_k^*
-

Theorem 7.7. *For any admissible objective function, Algorithm 8 computes a tree of optimal cost in time $O(n \log^2 n)$ with high probability if the input is a strict ground-truth input or in time $O(n^2)$ if the input is a (non-necessarily strict) ground-truth input.*

Proof. We proceed by induction on the number of vertices in the graph. Let p be the first pivot chosen by the algorithm and let B_1, \dots, B_k be the sets defined by p at Step 4 of the algorithm, with $w(p, u) > w(p, v)$, for any $u \in B_i, v \in B_{i+1}$.

We show that for any $u \in B_i, v \in B_j, j > i$, we have $w(u, v) = w(p, v)$. Consider a generating tree T and define $N_1 = \text{LCA}_T(p, u)$ and $N_2 = \text{LCA}_T(p, v)$. Since T, h, σ is generating and $w(p, u) > w(p, v)$, we have that N_2 is an ancestor of N_1 , by Definition 2.2. Therefore, $\text{LCA}_T(u, v) = N_2$, and so $w(u, v) = W(N_2) = w(p, v)$. Therefore, combining the inductive hypothesis on any $G[B_i]$ and by Definition 2.2 the tree output by the algorithm is generating.

A bound of $O(n^2)$ for the running time follows directly from the definition of the algorithm. We now argue that the running time is $O(n \log^2 n)$ with high probability if the input is strongly generated from a tree T . First, it is easy to see that a given recursive call on a subgraph with n_0 vertices takes $O(n_0)$ time. Now, observe that if at each recursive call the pivot partitions the n_0 vertices of its subgraph into buckets of size at most $2n_0/3$, then applying the master theorem implies a total running time of $O(n \log n)$. Unfortunately, there are trees where picking an arbitrary vertex as a pivot yields a single bucket of size $n - 1$.

Thus, consider the node N of T that is the first node reached by the walk from the root that always goes to the child tree with the higher number of leaves, stopping when the subtree of T rooted at N contains fewer than $2n/3$ but at least $n/3$ leaves. Since T is strongly generating we have that the partition into B_1, \dots, B_k induced by any vertex $v \in V(N)$ is such that any B_i contains less than $2n/3$ vertices. Indeed, for any u such that $\text{LCA}_T(u, v)$ is an ancestor of N and $x \in V(N)$, we have that $w(u, v) < w(x, v)$, and so u and x belong to different parts of the partition B_1, \dots, B_k .

Since the number of vertices in $V(N)$ is at least $n/3$, the probability of picking one of them is at least $1/3$. Therefore, since the pivots are chosen independently, after $c \log n$ recursive calls, the probability of not picking a vertex of $V(N)$ as a pivot is $O(1/n^c)$. Taking the union bound yields the theorem. \square

7.3 Beyond Structured Inputs

Since real-world inputs might sometimes differ from our definition of ground-truth inputs introduced in the Section 2, we introduce the notion of δ -adversarially-perturbed ground-truth inputs. This notion aims at accounting for noise in the data. We then design a simple and arguably more reliable algorithm (a robust variant of Algorithm 8) that achieves a δ -approximation for δ -adversarially-perturbed ground-truth inputs in $O(n(n + m))$ time. An interesting property of this algorithm is that its approximation guarantee is the same for any admissible objective function.

We first introduce the definition of δ -adversarially-perturbed ground-truth inputs. For any real $\delta \geq 1$, we say that a weighted graph $G = (V, E, w)$ is a δ -adversarially-perturbed ground-truth input if there exists an ultrametric (X, d) , such that $V \subseteq X$, and for every $x, y \in V, x \neq y, e = \{x, y\}$ exists, and $f(d(x, y)) \leq w(e) \leq \delta f(d(x, y))$, where $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a non-increasing function. This defines δ -adversarially-perturbed ground-truth inputs for similarity graphs and an analogous definition applies for dissimilarity graphs.

We now introduce a robust, simple version of Algorithm 8 that returns a δ -approximation if the input is a δ -adversarially-perturbed ground-truth inputs. Algorithm 8 was partitioning the input graph based on a single, random vertex. In this slightly more robust version, the partition is built iteratively: Vertices are added to the current part if there exists at least one vertex in the current

part or in the parts that were built before with which they share an edge of high enough weight (see Algorithm 9 for a complete description).

Algorithm 9 Robust and Simple Algorithm for Hierarchical Clustering on δ -adversarially-perturbed ground-truth inputs (similarity setting)

```

1: Input: A graph  $G = (V, E)$  and a weight function  $w : E \mapsto \mathbb{R}_+$ , a parameter  $\delta$ 
2:  $p \leftarrow$  arbitrary vertex of  $V$ 
3:  $i \leftarrow 0$ 
4:  $\tilde{V}_i \leftarrow \{p\}$ 
5: while  $\tilde{V}_i \neq V$  do
6:   Let  $p_1 \in \tilde{V}_i, p_2 \in V \setminus \tilde{V}_i$  s.t.  $(p_1, p_2)$  is an edge of maximum weight in the cut  $(\tilde{V}_i, V \setminus \tilde{V}_i)$ 
7:    $w_i \leftarrow w(p_1, p_2)$ 
8:    $B_i \leftarrow \{u \mid w(p_1, u) = w_i\}$ 
9:   while  $\exists u \in V \setminus (\tilde{V}_i \cup B_i)$  s.t.  $\exists v \in B_i \cup \tilde{V}_i, w(u, v) \geq w_i$  do
10:     $B_i \leftarrow B_i \cup \{u\}$ .
11:   end while
12:    $\tilde{V}_{i+1} \leftarrow \tilde{V}_i \cup B_i$ 
13:    $i \leftarrow i + 1$ 
14: end while
15: Let  $B_1, \dots, B_k$  be the sets obtained
16: Apply the algorithm recursively on each  $G[B_i]$  and obtain a collection of trees  $T_1, \dots, T_k$ 
17: Define  $T_0^*$  as a tree with  $p$  as a single vertex
18: For any  $1 \leq i \leq k$ , define  $T_i^*$  to be the union of  $T_{i-1}^*$  and  $T_i$ 
19: Return  $T_k^*$ 

```

Theorem 7.8. *For any admissible objective function, Algorithm 9 returns a δ -approximation if the input is a δ -adversarially-perturbed ground-truth input.*

To prove the theorem we introduce the following lemma whose proof is temporarily deferred. The lemma states that the tree built by the algorithm is almost generating (up to a factor of δ in the edge weights).

Lemma 7.9. *Let T be a tree output by Algorithm 9, let \mathcal{N} be the set of internal nodes of T . For any node N with children N_1, N_2 there exists a function $\omega : \mathcal{N} \mapsto \mathbb{R}_+$, such that for any $u \in V(N_1), v \in V(N_2)$, $\omega(N) \leq w(u, v) \leq \delta\omega(N)$. Moreover, for any nodes N, N' , if N' is an ancestor of N , we have that $\omega(N) \geq \omega(N')$.*

Assuming Lemma 7.9, the proof of Theorem 7.8 is as follows.

Proof of Theorem 7.8. Let $G = (V, E)$, $w : E \mapsto \mathbb{R}_+$ be the input graph and T^* be a tree of optimal cost. By Lemma 7.9, the tree T output by the algorithm is such that for any node N with children N_1, N_2 there exists a real $\omega(N)$, such that for any $u \in V(N_1), v \in V(N_2)$, $\omega \leq w(u, v) \leq \delta\omega$. Thus, consider the slightly different input graph $G' = (V, E, w')$, where $w' : E \mapsto \mathbb{R}_+$ is defined as follows. For any edge (u, v) , define $w'(u, v) = \omega(\text{LCA}_T(u, v))$. Since by Lemma 7.9, for any nodes N, N' of T , if N' is an ancestor of N , we have that $\omega(N) \geq \omega(N')$ and by definition 2.2, T is generating for G' . Thus, for any admissible cost function, we have that for G' , $\text{cost}_{G'}(T) \leq \text{cost}_{G'}(T^*)$.

Finally, observe that for any edge e , we have $w'(e) \leq w(e) \leq \delta w'(e)$. It follows that $\text{cost}_G(T) \leq \delta \text{cost}_{G'}(T)$ for any admissible cost function and $\text{cost}_{G'}(T^*) \leq \text{cost}_G(T^*)$. Therefore, $\text{cost}_G(T) \leq \delta \text{cost}_G(T^*) = \delta \text{OPT}$. \square

Proof of Lemma 7.9. We proceed by induction on the number of vertices in the graph (the base case is trivial). Consider the first recursive call of the algorithm. We show the following claim.

Claim 7.10. *For any $1 \leq i \leq k$, for any $y \in \tilde{V}_i$, $x \in B_i$, $w_i \geq w(x, y) \geq w_i/\delta$. Additionally, for any $x, y \in B_i$, $w(x, y) \geq w_i/\delta$.*

We first argue that Claim 7.10 implies the lemma. Let T be the tree output by the algorithm. Consider the nodes on the path from p to the root of T ; Let N_i denote the node whose subtree is the union of T_{i-1}^* and T_i . By definition, $V(T_{i-1}^*) = \tilde{V}_i$ and $V(T_i) = B_i$. Applying Claim 7.10 and observing that $w_i > w_{i+1}$ implies that the lemma holds for all the nodes on the path. Finally, since for any edge $\{u, v\}$, for $u, v \in B_i$, we also have $w(u, v) \geq w_i/\delta$, combining with the inductive hypothesis on B_i implies the lemma for all the nodes of the subtree T_i . \square

Proof of Claim 7.10. Let (X, d) and f be a pair of ultrametric and function that is generating for G . Fix $i \in \{1, \dots, k\}$. For any vertex $x \in B_i$, let $\sigma(x)$ denote a vertex y that is in \tilde{V}_i or inserted to B_i before x and such that $w(y, x) = w_i$. For any vertex v , let $\sigma^i(x)$ denotes the vertex obtained by applying σ i times to x (i.e., $\sigma^2(x) = \sigma(\sigma(x))$). By definition of the algorithm that for any $x \in B_i$, $\exists s \geq 1$, such that $\sigma^s(x) \in \tilde{V}_i$.

Fix $x \in B_i$. For any $y \in \tilde{V}_i$, we have that $w(y, x) \leq w_i$ since otherwise, the algorithm would have added x before.

Now, let $y \in \tilde{V}_i$ or y inserted to B_i prior to x . We aim at showing that $w(y, x) \geq w_i/\delta$. Observe that since X, d is an ultrametric, $d(x, y) \leq \max(d(x, \sigma(x)), d(\sigma(x), y))$.

We now “follow” σ by applying the function σ to $\sigma(x)$ and repeating until we reach $\sigma(x)^\ell = z \in \tilde{V}_i$. Combining with the definition of an ultrametric, it follows that

$$d(x, y) \leq \max(d(x, \sigma(x)), d(\sigma(x), \sigma^2(x)), \dots, d(\sigma(x)^{\ell-1}, z), d(z, y)).$$

If y was in \tilde{V}_i , we define $\hat{y} = y$. Otherwise y is also in B_i (and so was added to B_i before x). We then proceed similarly than for x and “follow” σ . In this case, let $\hat{y} = \sigma^k(y) \in \tilde{V}_i$. Applying the definition of an ultrametric again, we obtain

$$d(x, y) \leq \max(d(x, \sigma(x)), d(\sigma(x), \sigma^2(x)), \dots, d(\sigma^{\ell-1}, z), d(z, \hat{y}), d(y, \sigma(y)), \dots, d(\sigma^{k-1}(y), \hat{y})).$$

Assume for now that $d(z, \hat{y})$ is not greater than the others. Applying the definition of a δ -adversarially-perturbed input, we have that

$$\delta w(x, y) \geq \min(\dots, w(\sigma^a(x), \sigma^{a+1}(x)), \dots, w(\sigma^b(y), \sigma^{b+1}(y)), \dots).$$

Following the definition of σ , we have $w(v, \sigma(v)) \geq w_i, \forall v$. Therefore, we conclude $\delta w(x, y) \geq w_i$.

We thus turn to the case where $d(z, \hat{y})$ is greater than the others. Since both $z, \hat{y} \in \tilde{V}_i$, we have that they belong to some B_{j_0}, B_{j_1} , where $j_0, j_1 < i$. We consider the minimum j such that a pair at distance at least $d(z, \hat{y})$ was added to \tilde{V}_j . Consider such a pair $u, v \in \tilde{V}_j$ satisfying $d(u, v) \geq d(z, \hat{y})$ and suppose w.l.o.g that $v \in B_{j-1}$ (we could have either $u \in B_{j-1}$ or $u \in \tilde{V}_{j-1}$). Again, we follow the path $\sigma(v), \sigma(\sigma(v)), \dots$, until we reach $\sigma^{r_1}(v) \in \tilde{V}_{j-1}$ and similarly for u : $\sigma^{r_2}(u) \in \tilde{V}_{j-1}$. Applying the definition of an ultrametric this yields that

$$d(u, v) \leq \max(\dots, d(\sigma^a(u), \sigma^{a+1}(u)), \dots, d(\sigma^b(v), \sigma^{b+1}(v)), \dots, d(\sigma^{r_1}(v), \sigma^{r_2}(u))). \quad (16)$$

Now the difference is that \tilde{V}_{j-1} does not contain any pair at distance at least $d(z, \hat{y})$. Therefore, we have $d(\sigma^{r_1}(v), \sigma^{r_2}(u)) < d(z, \hat{y})$. Moreover, recall that by definition of u, v , $d(z, \hat{y}) \leq d(u, v)$.

Thus, $d(\sigma^{r_1}(v), \sigma^{r_2}(u))$ is not the maximum in Equation 16 since it is smaller than the left-hand side. Simplifying Equation 16 yields

$$d(x, y) < d(z, \hat{y}) \leq d(u, v) \leq \max(\dots, d(\sigma^a(u), \sigma^{a+1}(u)), \dots, d(\sigma^b(v), \sigma^{b+1}(v)), \dots).$$

By definition of a δ -adversarially-perturbed input, we obtain $\delta w(x, y) \geq \min_{\ell} w(\sigma^{\ell}(b), \sigma^{\ell+1}(b)) \geq w_j$. Now, it is easy to see that for $j < i$, $w_i < w_j$ and therefore $\delta w(x, y) \geq w_i$.

We conclude that for any $y \in \tilde{V}_i$, $x \in B_i$, $w_i \geq w(x, y) \geq w_i/\delta$ and for $x, y \in B_i$, we have that $w(x, y) \geq w_i/\delta$, as claimed. \square

8 Worst-Case Analysis of Common Heuristics

The results presented in this section shows that for both the similarity and dissimilarity settings, some of the widely-used heuristics may perform badly. The proofs are not difficult nor particularly interesting, but the results stand in sharp contrast to structured inputs and help motivate our study of inputs beyond worst case.

Similarity Graphs. We show that for very simple input graphs (*i.e.*, unweighted trees), the linkage algorithms (adapted to the similarity setting, see Algorithm 6) may perform badly.

Theorem 8.1. *There exists an infinite family of inputs on which the single-linkage and complete-linkage algorithms output a solution of cost $\Omega(nOPT/\log n)$.*

Proof. The family of inputs consists of the graphs that represent paths of length $n > 2$. More formally, Let G_n be a graph on n vertices such that $V = \{v_1, \dots, v_n\}$ and that has the following edge weights. Let $w(v_{i-1}, v_i) = w(v_i, v_{i+1}) = 1$, for all $1 < i < n$ and for any i, j , $j \notin \{i-1, i, i+1\}$, define $w(v_i, v_j) = 0$.

Claim 8.2. $OPT(G_n) = O(n \log n)$.

Proof. Consider the tree T^* that recursively divides the path into two subpaths of equal length. We aim at showing that $\text{cost}(T^*) = O(n \log n)$. The cost induced by the root node is n (since there is only one edge joining the two subpaths). The cost induced by each child of the root is $n/2$ since there is again only one edge joining the two sub-subpaths and now only $n/2$ leaves in the two subtrees. A direct induction shows that for a descendant at distance i from the root, the cost induced by this node is $n/2^i$. Since the number of children at distance i is 2^i , we obtain that the total cost induced by all the children at distance i is n . Since the tree divides the graph into two subgraph of equal size, there are at most $O(\log n)$ levels and therefore, the total cost of T (and so $OPT(G_n)$) is $O(n \log n)$. \square

Complete-Linkage. We show that the complete-linkage algorithm could perform a sequence of merges that would induce a tree of cost $\Omega(n^2)$. At start, each cluster contains a single vertex and so, the algorithm could merge any two clusters $\{v_i\}, \{v_{i+1}\}$ with $1 \leq i < n$ since their distance are all 1 (and it is the maximum edge weight in the graph). Suppose w.l.o.g that the algorithm merges v_1, v_2 . This yields a cluster C_1 such that the maximum distance between vertices of C_1 and v_3 is 1. Thus, assume w.l.o.g that the second merge of the algorithm is C_1, v_3 . Again, this yields a cluster C_2 whose maximum distance to v_4 is also 1. A direct induction shows that the algorithm output a tree whose root induces the cut $(V \setminus \{v_n\}, v_n)$ and one of the child induces the cut $(V \setminus \{v_{n-1}, v_n\}, v_{n-1})$ and so on. We now argue that this tree \hat{T} has cost $\Omega(n^2)$. Indeed, for any $1 < i \leq n$, we have $V(\text{LCA}_{\hat{T}}(v_{i-1}, v_i)) = i$. Thus the cost is at least $\sum_{i=2}^n i = \Omega(n^2)$.

Single-Linkage. We now turn to the case of the single-linkage algorithm. Recall that the algorithm merges the two candidate clusters C_i, C_j that minimize $w(u, v)$ for $u \in C_i, v \in C_j$.

At start, each cluster contains a single vertex and so, the algorithm could merge any two clusters $\{v_i\}, \{v_j\}$ for $j \notin \{i-1, i, i+1\}$ since the edge weight is 0 (and it is the minimum edge weight). Suppose w.l.o.g that the algorithm merges v_1, v_3 . This yields a cluster C_1 such that the distance between vertices of C_1 and any v_i for $i \equiv 1 \pmod 2$ is 0. Thus, assume w.l.o.g that the second merge of the algorithm is C_1, v_5 . A direct induction shows that w.l.o.g the output tree contains a node \tilde{N} such that $V(\tilde{N})$ contains all the vertices of odd indices. Now observe that the cost of the tree is at least $|V(\tilde{N})| \cdot w(V(\tilde{N}), V \setminus V(\tilde{N})) = \Omega(n^2)$.

Thus, by Claim 8.2 the single-linkage and complete-linkage algorithms can output a solution of cost $\Omega(n \text{OPT} / \log n)$. \square

Theorem 8.3. *There exists an infinite family of inputs on which the average-linkage algorithm output a solution of cost $\Omega(n^{1/3} \text{OPT})$.*

Proof. For any $n = 2^i$ for some integer i , we define a tree $T_n = (V, E)$ as follows. Let $k = n^{1/3}$. Let $P = (u_1, \dots, u_k)$ be a path of length k (i.e., for each $1 \leq i < k$, we have an edge between u_i and u_{i+1}). For each u_i , we define a collection $P_i = \{P_1^i = (V_1^i, E_1^i), \dots, P_k^i = (V_k^i, E_k^i)\}$ of k paths of length k and for each P_j^i we connect one of its extremities to u_i . Define $V_i = \{u_i\} \cup_j V_j^i$.

Claim 8.4. $\text{OPT}(T_n) \leq 3n^{4/3}$

Proof. Consider the following non-binary solution tree T^* : Let the root have children N_1, \dots, N_k such that $V(N_i) = V_i$ and for each child N_i let it have children N_i^j such that $V(N_i^j) = V_j^i$. Finally, for each N_i^j let the subtree rooted at N_i^j be any tree.

We now analyze the cost of T^* . Observe that for each edge e in the path P , we have $|V(\text{LCA}_{T^*}(e))| = n$. Moreover, for each edge e connecting a path P_j^i to u_i , we have $|V(\text{LCA}_{T^*}(e))| = k^2 = n^{2/3}$. Finally, for each edge e whose both endpoints are in a path P_j^i , we have that $|V(\text{LCA}_{T^*}(e))| \leq k = n^{1/3}$.

We now sum up over all edges to obtain the overall cost of T_n . There are $k = n^{1/3}$ edges in P ; They incur a cost of $nk = n^{4/3}$. There are k^2 edges joining a vertex u_i to a path P_j^i ; They incur a cost of $k^2 \cdot n^{2/3} = n^{4/3}$. Finally, there are k^3 edges whose both endpoints are in a path P_j^i ; They incur a cost of $k^3 \cdot n^{1/3} \leq n^{4/3}$. Thus, the total cost of this tree is at most $3n^{4/3} \geq \text{OPT}(T_n)$. \square

We now argue that there exists a sequence of merges done by the average-linkage algorithm that yield a solution of cost at least $n^{5/3}$.

Claim 8.5. *There exists a sequence of merges and an integer t such that the candidate trees at time t have leaves sets $\{\{u_1, \dots, u_k\}\} \cup_{i,j} \{V_j^i\}$.*

Equipped with this claim, we can finish the proof of the proposition. Since there is no edge between V_j^i and $V_{j'}^{i'}$ for $i' \neq i$ or $j' \neq j$ the distance between those trees in the algorithm will always be 0. However, the distance between the tree \tilde{T} that has leaves set $\{u_1, \dots, u_k\}$ and any other tree is positive (since there is one edge joining those two sets of vertices in T_n). Thus, the algorithm will merge \tilde{T} with some tree whose vertex set is exactly V_j^i for some i, j . For the same reasons, the resulting cluster will be merged to a cluster whose vertex set is exactly $V_{j'}^{i'}$, and so on. Hence, after $n/2k = k^2/2$ such merges, the tree \tilde{T} has a leaves set of size $k \cdot k^2/2 = n/2$. However, the number of edges from this cluster to the other candidate clusters is $k^2/2$ (since the other remaining clusters corresponds to vertex sets V_j^i for some i, j). For each such edge e we have $|V(\text{LCA}_T(e))| \geq n/2$.

Since there are $k^2/2$ of them, the resulting tree has cost $\Omega(n^{5/3})$. Combining with Claim 8.4 yields the theorem. \square

We thus turn to the proof of Claim 8.5.

Proof of Claim 8.5. Given a graph G a set of candidate trees \mathcal{C} , define G/\mathcal{C} to be the graph resulting from the contraction of all the edges whose both endpoints belong to the same cluster. We show a slightly stronger claim. We show that for any graph G and candidate trees \mathcal{V} such that

1. All the candidate clusters in \mathcal{V} have the same size; and
2. There exists a bijection ϕ between vertices $v \in T_n$ and vertices in G/\mathcal{C} ;

There exists a sequence of merges and an integer t such that the candidate trees at time t have leaves sets $\{\{\phi(u_1), \dots, \phi(u_k)\}\} \cup_{i,j} \{\phi(V_j^i)\}$ where $\phi(V_j^i) = \{\phi(v) \mid v \in V_j^i\}$.

This slightly stronger statement yields the claim by observing that T_n and the candidate trees at the start of the algorithm satisfies the conditions of the statement.

We proceed by induction on the number of vertices of the graph. Let $V_j^i = \{v_j^i(1), \dots, v_j^i(k)\}$ such that $(v_j^i(\ell), v_j^i(\ell + 1)) \in E_j^i$ for any $1 \leq \ell < k$, and $(v_j^i(k), u_i) \in E$.

We argue that the algorithm could perform a sequence of merges that results in the following set \mathcal{C} of candidate trees. \mathcal{C} contains candidate trees $U^i = \phi(u_{2i-1}) \cup \phi(u_{2i})$ for $1 \leq i < k/2$, and for each i, j , candidate trees $v_{i,j,\ell} = \phi(v_j^i(2\ell - 1) \cup \phi(v_j^i(2\ell)))$, for $1 \leq \ell < k/2$. Let s_0 be the number of vertices in each candidate tree.

At first, all the trees contain a single vertex and so, for each adjacent vertices of the graph the distance between their corresponding trees in the algorithm is $1/s_0$. For any non-adjacent pair of vertices, the corresponding trees are at distance 0. Thus, w.l.o.g assume the algorithm first merges u_1, u_2 . Then, the distance between the newly created tree U^1 and any other candidate tree C is 0 if there is no edge between u_1 and u_2 and C or $1/(2s_0)$ if there is one (since U^1 contains now two vertices). For the other candidate trees the distance is unchanged. Thus, the algorithm could merge vertices u_3, u_4 . Now, observe that the distance between U^2 and U^1 is at most $1/(4s_0)$. Thus, it is possible to repeat the argument and assume that the algorithm merges the candidate trees corresponding to u_5, u_6 . Repeating this argument $k/2$ times yields that after $k/2$ merges, the algorithm has generated the candidates trees $U_1, \dots, U_{k/2-1}$. The other candidate trees still contain a single vertex. Thus, the algorithm is now forced to merge candidate trees that contains single vertices that are adjacent (since their distance is $1/s_0$ and any other distance is $< 1/s_0$). Assume, w.l.o.g, that the algorithm merges $v_1^1(1), v_1^1(2)$. Again, applying a similar reasoning to each $v_1^1(2\ell - 1), v_1^1(2\ell)$ yields the set of candidate clusters $v_{1,1,1}, \dots, v_{1,1,k/2-1}$. Applying this argument to all sets V_j^i yields that the algorithm could perform a sequence of merges that results in the set \mathcal{C} of candidate clusters described above.

Now, all the clusters have size $2s_0$ and there exists a bijection between vertices of G/\mathcal{C} and $T_{n/2}$. Therefore, combining with the induction hypothesis yields the claim. \square

Dissimilarity Graphs. We now show that single-linkage, complete-linkage, and bisection 2-Center might return a solution that is arbitrarily bad compared to OPT in some cases. Hence, since average-linkage achieves a 2-approximation in the worst-case it seems that it is more robust than the other algorithms used in practice.

Theorem 8.6. *For each of the single-linkage, complete-linkage, and bisection 2-Center algorithms, there exists a family of inputs for which the algorithm outputs a solution of value $O(\text{OPT}/n)$.*

Proof. We define the family of inputs as follow. For any $n > 2$, the graph G_n consists of n vertices $V = \{v_1, \dots, v_{n-1}, u\}$ and the edge weights are the following: For any $i, j \in \{1, \dots, n-1\}$, $w(v_i, v_j) = 1$, for any $1 < i \leq n-1$, $w(v_i, u) = 1$, and $w(v_1, u) = W$ for some fixed $W \geq n^3$. Consider the tree T^* whose root induces a cut $(V \setminus \{u\}, \{u\})$. Then, the value of this tree (and so OPT) is at least nW , since $|V(\text{LCA}_{T^*}(v_1, u))| = n$.

Single-Linkage. At start, all the clusters are at distance 1 from each other except v_1 and u that are at distance W . Thus, suppose that the first merge generates a candidate tree C_1 whose leaves set is $\{v_1, v_2\}$. Now, since $w(v_2, u) = 1$, we have that all the clusters are at distance 1 from each other. Therefore, the next merge could possibly generate the cluster C_2 with leaves sets $\{u, v_1, v_2\}$. Assume w.l.o.g that this is the case and let T be the tree output by the algorithm. We obtain $|V(\text{LCA}_T(u, v_1))| = 3$ and so, since for any v_i, v_j , $|V(v_i, v_j)| \leq n$, $\text{val}(T) \leq n^2 + 3W \leq 4W$, since $W > W^3$. Hence, $\text{val}(T) = O(\text{val}(T^*)/n)$.

Complete-Linkage. Again, at first all the clusters are at distance 1 from each other except v_1 and u that are at distance W . Since the algorithm merges the two clusters that are at maximum distance, it merges u and v_1 . Again, let T be the tree output by the algorithm. We have $\text{val}(T) \leq n^2 + 2W \leq 3W$, since $W > W^3$. Hence, $\text{val}(T) = O(\text{val}(T^*)/n)$.

Bisection 2-Center. It is easy to see that for any location of the two centers, the cost of the clustering is 1. Thus, suppose that the algorithm locates centers in v_2, v_3 and that the induced partitioning is $\{v_1, v_2, u\}, V \setminus \{v_1, v_2, u\}$. It follows that $|V(\text{LCA}_T(u, v_1))| \leq 3$ and so, $\text{val}(T) \leq n^2 + 3W \leq 4W$, since $W > W^3$. Again, $\text{val}(T) = O(\text{val}(T^*)/n)$. \square

Proposition 8.7. *For any input \mathcal{I} lying in a metric space, for any solution tree T for \mathcal{I} , we have $\text{val}(T) = O(\text{OPT})$.*

Proof. Consider a solution tree T and the node u_0 of T_S^* that is the first node reached by the walk from the root that always goes to the child tree with the higher number of leaves, stopping when the subtree of T_S^* rooted at u_0 contains fewer than $2n/3$ leaves. Let $A = V(u_0), B = V \setminus V(u_0)$. Note that the number of edges in $G[A]$ is at most $a = \binom{|A|}{2}$, the number of edges in $G[B]$ is at most $b = \binom{|B|}{2}$, whereas the number of edges in the cut (A, B) is $|A| \cdot |B|$. Recall that $n/3 \leq |A|, |B| \leq 2n/3$ and so $a, b = \Theta(|A| \cdot |B|)$. Finally observe that for each edge $(u, v) \in G[A]$, we have $w(u, v) \leq w(u, x) + w(x, v)$ for any $x \in B$. Thus, since $a + b = \Theta(|A| \cdot |B|)$, by a simple counting argument, we deduce $\text{val}(T) = \Omega(n \sum_e w(e))$ and by Fact 6.1, $\Omega(\text{OPT})$. \square

Acknowledgments

The authors are grateful to Sanjoy Dasgupta for sharing thoughtful comments at various stages of this project.

References

- Sanjeev Arora and Ravi Kannan. Learning mixtures of arbitrary gaussians. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 247–257, 2001. doi: 10.1145/380752.380808. URL <http://doi.acm.org/10.1145/380752.380808>.
- Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009. doi: 10.1145/1502793.1502794. URL <http://doi.acm.org/10.1145/1502793.1502794>.

- Pranjal Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012, and 16th International Workshop, RANDOM 2012, Cambridge, MA, USA, August 15-17, 2012. Proceedings*, pages 37–49, 2012. doi: 10.1007/978-3-642-32512-0_4. URL http://dx.doi.org/10.1007/978-3-642-32512-0_4.
- Pranjal Awasthi, Avrim Blum, and Or Sheffet. Stability yields a PTAS for k-median and k-means clustering. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 309–318, 2010. doi: 10.1109/FOCS.2010.36. URL <http://dx.doi.org/10.1109/FOCS.2010.36>.
- Pranjal Awasthi, Avrim Blum, and Or Sheffet. Center-based clustering under perturbation stability. *Inf. Process. Lett.*, 112(1-2):49–54, 2012. doi: 10.1016/j.ipl.2011.10.006. URL <http://dx.doi.org/10.1016/j.ipl.2011.10.006>.
- Maria-Florina Balcan and Yingyu Liang. Clustering under perturbation resilience. *SIAM J. Comput.*, 45(1):102–155, 2016. doi: 10.1137/140981575. URL <http://dx.doi.org/10.1137/140981575>.
- Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. A discriminative framework for clustering via similarity functions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, pages 671–680, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-047-0. doi: 10.1145/1374376.1374474. URL <http://doi.acm.org/10.1145/1374376.1374474>.
- Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Approximate clustering without the approximation. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 1068–1077, 2009a. URL <http://dl.acm.org/citation.cfm?id=1496770.1496886>.
- Maria-Florina Balcan, Heiko Röglin, and Shang-Hua Teng. Agnostic clustering. In *Algorithmic Learning Theory, 20th International Conference, ALT 2009, Porto, Portugal, October 3-5, 2009. Proceedings*, pages 384–398, 2009b. doi: 10.1007/978-3-642-04414-4_31. URL http://dx.doi.org/10.1007/978-3-642-04414-4_31.
- Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Clustering under approximation stability. *J. ACM*, 60(2):8, 2013. doi: 10.1145/2450142.2450144. URL <http://doi.acm.org/10.1145/2450142.2450144>.
- Yonatan Bilu and Nathan Linial. Are stable instances easy? *Combinatorics, Probability & Computing*, 21(5):643–660, 2012. doi: 10.1017/S0963548312000193. URL <http://dx.doi.org/10.1017/S0963548312000193>.
- Yonatan Bilu, Amit Daniely, Nati Linial, and Michael E. Saks. On the practically interesting instances of MAXCUT. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, pages 526–537, 2013. doi: 10.4230/LIPIcs.STACS.2013.526. URL <http://dx.doi.org/10.4230/LIPIcs.STACS.2013.526>.
- S. Charles Brubaker and Santosh Vempala. Isotropic PCA and affine-invariant clustering. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 551–560, 2008. doi: 10.1109/FOCS.2008.48. URL <http://dx.doi.org/10.1109/FOCS.2008.48>.

- Gunnar Carlsson and Facundo Mémoli. Characterization, stability and convergence of hierarchical clustering methods. *Journal of Machine Learning Research*, 11:1425–1470, 2010.
- Rui M Castro, Mark J Coates, and Robert D Nowak. Likelihood based hierarchical clustering. *IEEE Transactions on signal processing*, 52(8):2308–2321, 2004.
- Moses Charikar and Vaggos Chatziafratis. Approximate hierarchical clustering via sparsest cut and spreading metrics. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 841–854, 2017. doi: 10.1137/1.9781611974782.53. URL <http://dx.doi.org/10.1137/1.9781611974782.53>.
- Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 634–644, 1999. doi: 10.1109/SFFCS.1999.814639. URL <http://dx.doi.org/10.1109/SFFCS.1999.814639>.
- Sanjoy Dasgupta. A cost function for similarity-based hierarchical clustering. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, pages 118–127, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4132-5. doi: 10.1145/2897518.2897527. URL <http://doi.acm.org/10.1145/2897518.2897527>.
- Sanjoy Dasgupta and Philip M Long. Performance guarantees for hierarchical clustering. *Journal of Computer and System Sciences*, 70(4):555–569, 2005.
- Sanjoy Dasgupta and Leonard J. Schulman. A probabilistic analysis of EM for mixtures of separated, spherical gaussians. *Journal of Machine Learning Research*, 8:203–226, 2007. URL <http://www.jmlr.org/papers/v8/dasgupta07a.html>.
- Justin Eldridge, Mikhail Belkin, and Yusu Wang. Graphons, mergeons, and so on! In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2307–2315, 2016. URL <http://papers.nips.cc/paper/6089-graphons-mergeons-and-so-on>.
- Joseph Felsenstein and Joseph Felsenstein. *Inferring phylogenies*, volume 2. Sinauer Associates Sunderland, 2004.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. ISSN 01621459. URL <http://www.jstor.org/stable/2282952>.
- N. Jardine and R. Sibson. *Mathematical Taxonomy*. Wiley series in probability and mathematical statistiscs. John Wiley & Sons, 1972. URL <https://books.google.dk/books?id=uMYZogEACAAJ>.
- Jon Kleinberg. An impossibility theorem for clustering. In *NIPS*, volume 15, pages 463–470, 2002.
- Amit Kumar and Ravindran Kannan. Clustering with spectral norm and the k-means algorithm. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 299–308, 2010. doi: 10.1109/FOCS.2010.35. URL <http://dx.doi.org/10.1109/FOCS.2010.35>.

- Guolong Lin, Chandrashekhar Nagarajan, Rajmohan Rajaraman, and David P Williamson. A general approach for incremental approximation and hierarchical clustering. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1147–1156. Society for Industrial and Applied Mathematics, 2006.
- Vince Lyzinski, Minh Tang, Avanti Athreya, Youngser Park, and Carey E Priebe. Community detection and classification in hierarchical stochastic blockmodels. *IEEE Transactions on Network Science and Engineering*, 4(1):13–26, 2017.
- Frank McSherry. Spectral partitioning of random graphs. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 529–537, 2001. doi: 10.1109/SFCS.2001.959929. URL <http://dx.doi.org/10.1109/SFCS.2001.959929>.
- Rafail Ostrovsky, Yuval Rabani, Leonard J. Schulman, and Chaitanya Swamy. The effectiveness of Lloyd-type methods for the k-means problem. *J. ACM*, 59(6):28, 2012. doi: 10.1145/2395116.2395117. URL <http://doi.acm.org/10.1145/2395116.2395117>.
- C Greg Plaxton. Approximation algorithms for hierarchical location problems. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 40–49. ACM, 2003.
- Aurko Roy and Sebastian Pokutta. Hierarchical clustering via spreading metrics. In *Advances In Neural Information Processing Systems*, pages 2316–2324, 2016.
- Peter HA Sneath and Robert R Sokal. Numerical taxonomy. *Nature*, 193(4818):855–860, 1962.
- Michael Steinbach, George Karypis, and Vipin Kumar. A comparison of document clustering techniques. In *In KDD Workshop on Text Mining*, 2000.
- Patrick J Wolfe and Sofia C Olhede. Nonparametric graphon estimation. *arXiv preprint arXiv:1309.5936*, 2013.
- Reza Bosagh Zadeh and Shai Ben-David. A uniqueness theorem for clustering. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 639–646. AUAI Press, 2009.