
Hierarchical Clustering with Structural Constraints

Vaggos Chatziafratis^{*1} Rad Niazadeh^{*1} Moses Charikar¹

Abstract

Hierarchical clustering is a popular unsupervised data analysis method. For many real-world applications, we would like to exploit prior information about the data that imposes constraints on the clustering hierarchy, and is not captured by the set of features available to the algorithm. This gives rise to the problem of *hierarchical clustering with structural constraints*. Structural constraints pose major challenges for bottom-up approaches like average/single linkage and even though they can be naturally incorporated into top-down divisive algorithms, no formal guarantees exist on the quality of their output. In this paper, we provide provable approximation guarantees for two simple top-down algorithms, using a recently introduced optimization viewpoint of hierarchical clustering with pairwise similarity information (Dasgupta, 2016). We show how to find good solutions even in the presence of conflicting prior information, by formulating a *constraint-based regularization* of the objective. Furthermore, we explore a variation of this objective for dissimilarity information (Cohen-Addad et al., 2018) and improve upon current techniques. Finally, we demonstrate our approach on a real dataset for the taxonomy application.

1. Introduction

Hierarchical clustering (HC) is a widely used data analysis tool, ubiquitous in information retrieval, data mining, and machine learning (see a survey by (Berkhin, 2006)). This clustering technique represents a given dataset as a binary tree; each leaf represents an individual data point and each internal node represents a cluster on the leaves of its descendants. HC has become the most popular method for

gene expression data analysis (Eisen et al., 1998), and also has been used in the analysis of social networks (Leskovec et al., 2014; Mann et al., 2008), bioinformatics (Diez et al., 2015), image and text classification (Steinbach et al., 2000), and even in analysis of financial markets (Tumminello et al., 2010). It is attractive because it provides richer information at all levels of granularity simultaneously, compared to more traditional *flat* clustering approaches like k -means.

Recently, (Dasgupta, 2016) formulated HC as a combinatorial optimization problem, giving a principled way to compare the performance of different HC algorithms. This optimization viewpoint has since received a lot of attention (Roy & Pokutta, 2016; Charikar & Chatziafratis, 2017; Cohen-Addad et al., 2017; Moseley & Wang, 2017; Cohen-Addad et al., 2018) that has led not only to the development of new algorithms but also to theoretical justifications for the observed success of popular algorithms (e.g. average-linkage).

However, in real applications of clustering, the user often has *background knowledge* about the data that may not be captured by the input to the clustering algorithm. There is a rich body of work on constrained (flat) clustering formulations that take into account such user input in the form of “cannot link” and “must link” constraints (Wagstaff & Cardie, 2000; Wagstaff et al., 2001; Bilenko et al., 2004; Rangapuram & Hein, 2012). Very recently, “semi-supervised” versions of HC that incorporate additional constraints have been studied (Vikram & Dasgupta, 2016), where the natural form of such constraints is triplet (or “must link before”) constraints $ab|c$ ¹: these require that valid solutions contain a sub-cluster with a, b together and c previously separated from them.² Such triplet constraints, as we show later, can encode more general structural constraints in the form of rooted subtrees. Surprisingly, such simple triplet constraints already pose significant challenges for bottom-up linkage methods. (Figure 1).

Our work is motivated by applying the optimization lens to study the interaction of hierarchical clustering algorithms with structural constraints. Constraints can be fairly natu-

¹Department of Computer Science, Stanford University, Stanford, CA, USA. Correspondence to: Moses Charikar <moses@cs.stanford.edu>, Vaggos Chatziafratis <vaggos@stanford.edu>, Rad Niazadeh <rad@cs.stanford.edu>.

¹Hierarchies on data imply that all datapoints are linked at the highest level and all are separated at the lowest level, hence “cannot link” and “must link” constraints are not directly meaningful.

²For a concrete example from taxonomy of species, a triplet constraint may look like (TUNA, SALMON|LION).




Figure 1. (Left) Example of a triplet constraint $uv|w$ and more general rooted tree constraints on 4 points u, v, w, z . (Right) Example with only two constraints $ab|c, a'b'|c'$ demonstrating that popular distance-based linkage algorithms may fail to produce valid HC. Here they get stuck after 3 merging steps (green edges).

rally incorporated into top-down (i.e. divisive) algorithms for hierarchical clustering; but can we establish guarantees on the quality of the solution they produce? Another issue is that incorporating constraints from multiple experts may lead to a conflicting set of constraints; can the optimization viewpoint of hierarchical clustering still help us obtain good solutions even in the presence of infeasible constraints? Finally, different objective functions for HC have been studied in the literature; do algorithms designed for these objectives behave similarly in the presence of constraints? To the best of our knowledge, this is the first work to propose a unified approach for constrained HC through the lens of optimization and to give provable approximation guarantees for a collection of fast and simple top-down algorithms that have been used for unconstrained HC in practice (e.g. community detection in social networks (Mann et al., 2008)).

Background on Optimization View of HC. (Dasgupta, 2016) introduced a natural optimization framework for HC. Given a weighted graph $G(V, E, w)$ and pairwise similarities $w_{ij} \geq 0$ between the n data points $i, j \in V$, the goal is to find a hierarchical tree T^* such that

$$T^* = \arg \min_{\text{all trees } T} \sum_{(i,j) \in E} w_{ij} \cdot |T_{ij}| \quad (1)$$

where T_{ij} is the subtree rooted at the lowest common ancestor of i, j in T and $|T_{ij}|$ is the number of leaves it contains.³ We denote (1) as *similarity-HC*. For applications where the geometry of the data is given by dissimilarities, again denoted by $\{w_{ij}\}_{(i,j) \in E}$, (Cohen-Addad et al., 2018) proposed an analogous approach, where the goal is to find a hierarchical tree T^* such that

$$T^* = \arg \max_{\text{all trees } T} \sum_{(i,j) \in E} w_{ij} \cdot |T_{ij}| \quad (2)$$

We denote (2) as *dissimilarity-HC*. A comprehensive list of desirable properties of the aforementioned objectives can be found in (Dasgupta, 2016; Cohen-Addad et al., 2018). In

³Observe that in HC, all edges get cut eventually. Therefore it is better to postpone cutting “heavy” edges to when the clusters become small, i.e. as far down the tree as possible.

particular, if there is an underlying ground-truth hierarchical structure in the data, then T^* can recover the ground-truth. Also, both objectives are NP-hard to optimize, so the focus is on approximation algorithms.

Our Results. *i)* We design algorithms that take into account both the geometry of the data, in the form of similarities, and the structural constraints imposed by the users. Our algorithms emerge as the natural extensions of Dasgupta’s original recursive sparsest cut algorithm and the recursive balanced cut suggested in (Charikar & Chatziafratis, 2017). We generalize previous analyses to handle constraints and we prove an $O(k\alpha_n)$ -approximation guarantee⁴, thus surprisingly matching the best approximation guarantee of the *unconstrained* HC problem for constantly many constraints.

ii) In the case of infeasible constraints, we extend the similarity-HC optimization framework, and we measure the quality of a possible tree T by a *constraint-based regularized* objective. The regularization naturally favors solutions with as few constraint violations as possible and as far down the tree as possible (similar to the motivation behind similarity-HC objective). For this problem, we provide a top-down $O(k\alpha_n)$ -approximation algorithm by drawing an interesting connection to an instance of the hypergraph sparsest cut problem.

iii) We then change gears and study the dissimilarity-HC objective. Surprisingly, we show that known top-down techniques do not cope well with constraints, drawing a contrast with the situation for similarity-HC. Specifically, the (*locally*) *densest cut heuristic* performs poorly even if there is only one triplet constraint, blowing up its approximation factor to $O(n)$. Moreover, we improve upon the state-of-the-art in (Cohen-Addad et al., 2018), by showing a simple randomized partitioning is a $\frac{2}{3}$ -approximation algorithm. We also give a deterministic *local-search* algorithm with the same worst-case guarantee. Furthermore, we show that our randomized algorithm is robust under constraints, mainly because of its “exploration” behavior. In fact, besides the number of constraints, we propose an inherent notion of *dependency measure* among constraints to capture this behavior quantitatively. This helps us not only to explain why “non-exploring” algorithms may perform poorly, but also gives tight guarantees for our randomized algorithm.

Experimental results. We run experiments on the Zoo dataset (Lichman, 2013) to demonstrate our approach and the performance of our algorithms for a taxonomy application. Due to lack of space, we present these results in the full online version of our paper (Chatziafratis et al., 2018).

Constrained HC work-flow in Practice. Throughout this paper, we develop different tools to handle user-defined

⁴For n data points, $\alpha_n = O(\sqrt{\log n})$ is the best approximation factor for the sparsest cut and k is the number of constraints.

structural constraints for hierarchical clustering. Here we describe a recipe on how to use our framework in practice.

(1) *Preprocessing constraints to form triplets.* User-defined structural constraints as rooted binary subtrees are convenient for the user and hence for the usability of our algorithm. The following proposition (whose proof is in the supplement) allows us to focus on studying HC with just triplet constraints.

Proposition 1. *Given constraints as a rooted binary subtree T on k data points ($k \geq 3$), there is linear time algorithm that returns an equivalent set of at most k triplet constraints.*

(2) *Detecting feasibility.* The next step is to see if the set of triplet constraints is consistent, i.e. whether there exists a HC satisfying all the constraints. For this, we use a simple linear time algorithm called BUILD (Aho et al., 1981).

(3) *Hard constraints vs. regularization.* BUILD can create a hierarchical decomposition that satisfies triplet constraints, but ignores the *geometry* of the data, whereas our goal here is to consider both simultaneously. Moreover, in the case that the constraints are infeasible, we aim to output a clustering that minimizes the cost of violating constraints combined with the cost of the clustering itself.

• *Feasible instance:* to output a feasible HC, we propose using *Constrained Recursive Sparsest Cut* (CRSC) or *Constrained Recursive Balanced Cut* (CRBC): two simple top-down algorithms which are natural adaptations of recursive sparsest cut (Mann et al., 2008; Dasgupta, 2016) or recursive balanced cut (Charikar & Chatziafratis, 2017) to respect constraints (Section 2).

• *Infeasible instance:* in this case, we turn our attention to a regularized version of HC, where the cost of violating constraints is added to the tree cost. We then propose an adaptation of CRSC, namely *Hypergraph Recursive Sparsest Cut* (HRSC) for the regularized problem (Section 3).

See the supplementary materials for a real-world application examples, demonstrating our proposed HC work-flow.

Further related work. Similar to (Vikram & Dasgupta, 2016), constraints in the form of triplet queries have been used in an (adaptive) active learning framework by (Tamuz et al., 2011; Emamjomeh-Zadeh & Kempe, 2018), showing that approximately $O(n \log n)$ triplet queries are enough to learn an underlying HC. Other forms of user interaction in order to improve the quality of the produced clusterings have been used in (Balcan & Blum, 2008; Awasthi et al., 2014) where they prove that interactive feedback in the form of cluster split/merge requests can lead to significant improvements. Robust algorithms for HC in the presence of noise were studied in (Balcan et al., 2014) and a variety of sufficient conditions on the similarity function that would allow linkage-style methods to produce good clusters was

explored in (Balcan et al., 2008). On a different setting, the notion of triplets has been used as a measure of *distance* between hierarchical decomposition trees on the same data points (Brodal et al., 2013). More technically distant analogs of how to use relations among triplet points have recently been proposed in (Kleindessner & von Luxburg, 2017) for defining kernel functions corresponding to high-dimensional embeddings.

2. Constrained Sparsest (Balanced) Cut

Given an instance of the constrained hierarchical clustering, our proposed CRSC algorithm uses a blackbox α_n -approximation algorithm for the sparsest cut problem (the best-known approximation factor for this problem is $O(\sqrt{\log n})$ due to (Arora et al., 2009)). Moreover, it also maintains the feasibility of the solution in a top-down approach by recursive partitioning of what we call the *supergraph* G' . Informally speaking, the supergraph is a simple data structure to track the progress of the algorithm and the resolved constraints.

More formally, for every constraint $ab|c$ we merge the nodes a and b into a *supernode* $\{a, b\}$ while maintaining the edges in G (now connecting to their corresponding supernodes). Note that G' may have parallel edges, but this can easily be handled by grouping edges together and replacing them with the sum of their weights. We repeatedly continue this *merging* procedure until there are no more constraints. Observe that any feasible solution needs to start splitting the original graph G by using a cut that is also present in G' . When cutting the graph $G' = (G_1, G_2)$, if a constraint $ab|c$ is *resolved*,⁵ then we can safely *unpack* the supernode $\{a, b\}$ into two nodes again (unless there is another constraint $ab|c'$ in which case we should keep the supernode). By continuing and recursively finding approximate sparsest cuts on the supergraph G_1 and G_2 , we can find a feasible hierarchical decomposition of G respecting all triplet constraints. Next, we show the approximation guarantees for our algorithm.

Algorithm 1 CRSC

- 1: Given G and the triplet constraints $ab|c$, run BUILD to create the supergraph G' .
 - 2: Use a blackbox access to an α_n -approximation oracle for the sparsest cut problem, i.e. $\arg \min_{S \subseteq V} \frac{w_{G'}(S, \bar{S})}{|S| \cdot |\bar{S}|}$. ($|S|$ = total size of supernodes in S)
 - 3: Given the output cut (S, \bar{S}) , separate the graph G' into two pieces $G_1(S, E_1)$ and $G_2(V \setminus S, E_2)$.
 - 4: Recursively compute a HC T_1 for G_1 using only G_1 's active constraints. Similarly compute T_2 for G_2 .
 - 5: Output $T = (T_1, T_2)$.
-

Analysis of CRSC Algorithm. The main result of this section is the following theorem:

⁵A constraint $ab|c$ is *resolved*, if c gets separated from a, b .

Theorem 1. Given a weighted graph $G(V, E, w)$ with k triplet constraints $ab|c$ for $a, b, c \in V$, the CRSC algorithm outputs a HC respecting all triplet constraints and is $O(k\alpha_n)$ -approximation for the HC-similarity objective (1).

Notations and Definitions. We slightly abuse notation by having OPT denote the optimum hierarchical decomposition or its optimum value as measured by (1). Similarly for CRSC. For $t \in [n]$, $\text{OPT}(t)$ denotes the maximal clusters in OPT of size at most t . Note that $\text{OPT}(t)$ induces a partitioning of V . We use $\text{OPT}(t)$ to denote edges cut by $\text{OPT}(t)$ (i.e. edges with endpoints in different clusters in $\text{OPT}(t)$) or their total weight; the meaning will be clear from context. For convenience, we define $\text{OPT}(0) = \sum_{(i,j) \in E} w_{ij}$. For a cluster A created by CRSC, a constraint $ab|c$ is *active* if $a, b, c \in A$, otherwise $ab|c$ is resolved and can be discarded.

Overview of the Analysis. There are three main ingredients: The first is to view a HC of n datapoints as a collection of partitions, one for each level $t = n - 1, \dots, 1$, as in (Charikar & Chatziafratis, 2017). For a level t , the partition consists of maximal clusters of size at most t . The total cost incurred by OPT is then a combination of costs incurred at each level of this partition. This is useful for comparing our CRSC cost with OPT . The second idea is in handling constraints and it is the main obstacle where previous analyses (Charikar & Chatziafratis, 2017; Cohen-Addad et al., 2018) break down: constraints inevitably limit the possible cuts that are feasible at any level, and since the set of active constraints⁶ differ for CRSC and OPT , a direct comparison between them is impossible. If we have no constraints, we can charge the cost of partitioning a cluster A to lower levels of the OPT decomposition. However, when we have triplet constraints, the partition induced by the lower levels of OPT in a cluster A will not be feasible in general (Figure 2). The natural way to overcome this obstacle is merging pieces of this partition so as to respect constraints and using higher levels of OPT , but it still may be impossible to compare CRSC with OPT if all pieces are merged. We overcome this difficulty by an indirect comparison between the CRSC cost and lower levels $\frac{r}{6k_A}$ of OPT , where k_A is the number of active constraints in A . Finally, after a cluster-by-cluster analysis bounding the CRSC cost for each cluster, we exploit disjointness of clusters of the same level in the CRSC partition allowing us to combine their costs.

Proof of Theorem 1. We start by borrowing the following facts from (Charikar & Chatziafratis, 2017), modified slightly for the purpose of our analysis (proofs are provided in the supplementary materials).

Fact 1 (Decomposition of OPT). The total cost paid by OPT can be decomposed into costs of the different levels in

⁶All constraints are *active* in the beginning of CRSC.

the OPT partition, i.e. $\text{OPT} = \sum_{t=0}^n w(\text{OPT}(t))$.

Fact 2 (OPT at scaled levels). Let $k \leq \frac{n}{6}$ be the number of constraints. Then, $\text{OPT} \geq \frac{1}{6k} \cdot \sum_{t=0}^n w(\text{OPT}(\lfloor \frac{t}{6k} \rfloor))$.




Figure 2. The main obstacle in the constrained HC problem is that our algorithm has different *active constraints* compared to OPT . Both $ab|c, de|f$ constraints are resolved by the cut $\text{OPT}(t)$.

Given the above facts, we look at any cluster A of size r produced by the algorithm. Here is the main technical lemma that allows us to bound the cost of CRSC for partitioning A .

Lemma 1. Suppose CRSC partitions a cluster A ($|A| = r$) in two clusters (B_1, B_2) (w.l.o.g. $|B_1| = s, |B_2| = r - s, s \leq \lfloor \frac{r}{2} \rfloor \leq r - s$). Let the size $r \geq 6k$ and let $l = 6k_A$, where k_A denotes the number of active constraints for A . Then: $r \cdot w(B_1, B_2) \leq 4\alpha_n \cdot s \cdot w(\text{OPT}(\lfloor \frac{r}{l} \rfloor) \cap A)$.

Proof. The cost incurred by CRSC for partitioning A is $r \cdot w(B_1, B_2)$. Now consider $\text{OPT}(\lfloor \frac{r}{l} \rfloor)$. This induces a partitioning of A into pieces $\{A_i\}_{i \in [m]}$, where by design $|A_i| = \gamma_i |A|$, $\gamma_i \leq \frac{1}{l}, \forall i \in [m]$. Now, consider the cuts $\{(A_i, A \setminus A_i)\}$. Even though all m cuts are allowed for OPT , for CRSC some of them are forbidden: for example, in Figure 2, the constraints $ab|c, de|f$ render 4 out of the 6 cuts infeasible. But how many of them can become infeasible with k_A active constraints? Since every constraint is involved in at most 2 cuts, we may have at most $2k_A$ infeasible cuts. Let $F \subseteq [m]$ denote the index set of feasible cuts, i.e. if $i \in F$, the cut $(A_i, A \setminus A_i)$ is feasible for CRSC. To cut A , we use an α_n -approximation of sparsest cut, whose sparsity is upper bounded by *any feasible cut*:

$$\begin{aligned} \frac{w(B_1, B_2)}{s(r-s)} &\leq \alpha_n \cdot \text{SP.CUT}(A) \leq \alpha_n \min_{i \in F} \frac{w(A_i, A \setminus A_i)}{|A_i||A \setminus A_i|} \\ &\leq \alpha_n \frac{\sum_{i \in F} w(A_i, A \setminus A_i)}{\sum_{i \in F} |A_i||A \setminus A_i|} \end{aligned}$$

where for the last inequality we used the standard fact that $\min_i \frac{\mu_i}{\nu_i} \leq \frac{\sum_i \mu_i}{\sum_i \nu_i}$ for $\mu_i \geq 0$ and $\nu_i > 0$. We also have the following series of inequalities:

$$\begin{aligned} \alpha_n \frac{\sum_{i \in F} w(A_i, A \setminus A_i)}{\sum_{i \in F} |A_i||A \setminus A_i|} &\leq \alpha_n \frac{2w(\text{OPT}(\lfloor \frac{r}{l} \rfloor) \cap A)}{r^2 \sum_{i \in F} \gamma_i (1 - \gamma_i)} \\ &\leq 4\alpha_n \frac{w(\text{OPT}(\lfloor \frac{r}{l} \rfloor) \cap A)}{r^2} \end{aligned}$$

where the first inequality holds because we double count some (potentially all) edges of $\text{OPT}(\lfloor \frac{r}{7} \rfloor) \cap A$ (these are the edges cut by $\text{OPT}(\lfloor \frac{r}{7} \rfloor)$ that are also present in cluster A , i.e. they have both endpoints in A) and the second inequality holds because $\gamma_i \leq \frac{1}{6k} \implies 1 - \gamma_i \geq \frac{6k-1}{6k}$ and

$$\begin{aligned} \sum_{i \in F} \gamma_i (1 - \gamma_i) &\geq \sum_{i=1}^m \gamma_i (1 - \gamma_i) - 2 \sum_{i \in [m] \setminus F} \frac{1}{6k} \\ &\geq \frac{6k-1}{6k} \sum_{i=1}^m \gamma_i - \frac{2k}{6k} = \frac{4k-1}{6k} \geq 1/2 \end{aligned}$$

Finally, we are ready to prove the lemma by combining the above inequalities ($\frac{r-s}{r} \leq 1$):

$$\begin{aligned} r \cdot w(B_1, B_2) &= r \cdot s(r-s) \cdot \frac{w(B_1, B_2)}{s(r-s)} \\ &\leq r \cdot s(r-s) \cdot 4\alpha_n \frac{w(\text{OPT}(\lfloor \frac{r}{7} \rfloor) \cap A)}{r^2} \\ &\leq 4\alpha_n \cdot s \cdot w(\text{OPT}(\lfloor \frac{r}{7} \rfloor) \cap A). \quad \square \end{aligned}$$

It is clear that we exploited the charging to lower levels of OPT , since otherwise if all pieces in A were merged, the denominator with the $|A_i|$'s would become 0. The next lemma lets us combine the costs incurred by CRSC for different clusters A (proof is in the supplementary materials)

Lemma 2 (Combining the costs of clusters in CRSC). *The total CRSC cost for partitioning all clusters A into (B_1, B_2) (with $|A| = r_A, |B_1| = s_A$) is bounded by:*

$$\begin{aligned} (1) \quad \sum_{A: |A| \geq 6k} r_A \cdot w(B_1, B_2) &\leq O(\alpha_n) \cdot \sum_{t=0}^n w(\text{OPT}(\lfloor \frac{t}{6k} \rfloor)) \\ (2) \quad \sum_{A: |A| < 6k} r_A w(B_1, B_2) &\leq 6k \cdot \text{OPT} \end{aligned}$$

Combining [Fact 2](#) and [Lemma 2](#) finishes the proof. \square

Remark 1. In the supplementary material, we prove how one can use *balanced cut*, i.e. finding a cut S such that

$$\arg \min_{S \subseteq V: |S| \geq n/3, |\bar{S}| \geq n/3} w_{G'}(S, \bar{S}) \quad (3)$$

instead of sparsest cut, and using approximation algorithms for this problem achieves the same approximation factor as in [Theorem 1](#), but with better running time.

Theorem 2 (The divisive algorithm using balanced cut). *Given a weighted graph $G(V, E, w)$ with k triplet constraints $ab|c$ for $a, b, c \in V$, the constrained recursive balanced cut algorithm CRBC (same as CRSC, but using balanced cut instead of sparsest cut) outputs a HC respecting all triplet constraints and achieves an $O(k\alpha_n)$ -approximation for Dasgupta's HC objective. Moreover, the running time is almost linear time.*

3. Constraints and Regularization

Previously, we assumed that constraints were feasible. However, in many practical applications, users/experts may disagree, hence our algorithm may receive conflicting constraints as input. Here we want to explore how to still output a satisfying HC that is a good in terms of objective (1) (similarity-HC) and also respects the constraints as much as possible. To this end, we propose a *regularized* version of Dasgupta's objective, where the regularizer measures quantitatively the degree by which constraints get violated.

Informally, the idea is to penalize a constraint more if it is violated at top levels of the decomposition compared to lower levels. We also allow having different violation *weights* for different constraints (potentially depending on the expertise of the users providing the constraints). More concretely, inspired by the Dasgupta's original objective function, we consider the following optimization problem:

$$\min_{T \in \mathcal{T}} \left(\sum_{(i,j) \in E} w_{ij} |T_{ij}| + \lambda \cdot \sum_{ab|c \in \mathcal{K}} c_{ab|c} |T_{ab}| \cdot \mathbf{1}\{ab|c \text{ is violated}\} \right), \quad (4)$$

where \mathcal{T} is the set of all possible binary HC trees for the given data points, \mathcal{K} is the set of the k triplet constraints, T_{ab} is the size of the subtree rooted at the least common ancestor of a, b , and $c_{ab|c}$ is defined as the *base cost* of violating triplet constraint $ab|c$. Note that the regularization parameter $\lambda \geq 0$ allows us to interpolate between satisfying the constraints or respecting the geometry of the data.

Hypergraph Recursive Sparsest Cut In order to design approximation algorithms for the regularized objective, we draw an interesting connection to a different problem, which we call *3-Hypergraph Hierarchical Clustering (3HHC)*. An instance of this problem consists of a hypergraph $G^{\mathcal{H}} = (V, E, E^{\mathcal{H}})$ with edges E , and hyperedges of size 3, $E^{\mathcal{H}}$, together with similarity weights for edges, $\{w_{ij}\}_{(i,j) \in E}$, and similarity weights for 3-hyperedges,⁷ $\{w_{ij|k}\}_{(i,j,k) \in E^{\mathcal{H}}}$. We now think of HC on the hypergraph $G^{\mathcal{H}}$, where for every binary tree T we define the cost to be the natural extension of Dasgupta's objective:

$$\sum_{(i,j) \in E} w_{ij} |T_{ij}| + \sum_{(i,j,k) \in E^{\mathcal{H}}} w_{ij|k}^T |T_{ijk}| \quad (5)$$

where $w_{ij|k}^T$ is either equal to $w_{ij|k}, w_{jk|i}$ or $w_{ki|j}$, and T_{ijk} is either the subtree rooted at $\text{LCA}(i, j)$,⁸ $\text{LCA}(i, k)$ or $\text{LCA}(k, j)$, all depending on how T cuts the 3-hyperedge

⁷We have 3 different weights corresponding to the 3 possible ways of partitioning $\{i, j, k\}$ in two parts: $w_{ij|k}, w_{jk|i}$ and $w_{ki|j}$.

⁸ $\text{LCA}(i, j)$ denotes the lowest common ancestor of $i, j \in T$.

$\{i, j, k\}$. The goal is to find a hierarchical clustering of this hypergraph, so as to minimize the cost (5) of the tree.

Reduction from Regularization to 3HHC. Given an instance of HC with constraints (with their costs of violations) and a parameter λ , we create a hypergraph $G^{\mathcal{H}}$ so that the total cost of any binary clustering tree in the 3HHC problem (5) corresponds to the regularized objective of the same tree as in (4). $G^{\mathcal{H}}$ has exactly the same set of vertices, (normal) edges and (normal) edge weights as in the original instance of the HC problem. Moreover, for every constraint $ab|c$ (with cost $c_{ab|c}$) it has a hyperedge $\{a, b, c\}$, to which we assign three weights $w_{ab|c} = 0, w_{ac|b} = w_{bc|a} = \lambda \cdot c_{ab|c}$. Therefore, we ensure that any divisive algorithm for the 3HHC problem avoids the cost $|T_{abc}| \cdot \lambda \cdot c_{ab|c}$ only if it chops $\{a, b, c\}$ into $\{a, b\}$ and $\{c\}$ at some level, which matches the regularized objective.

Reduction from 3HHC to Hypergraph Sparsest Cut. A natural generalization of the sparsest cut problem for our hypergraphs, which we call *Hyper Sparsest Cut (HSC)*, is the following problem:

$$\arg \min_{S \subseteq V} \left(\frac{w(S, \bar{S}) + \sum_{(i,j,k) \in E^{\mathcal{H}}} w_{ijk}^S}{|S||\bar{S}|} \right),$$

where $w(S, \bar{S})$ is the weight of the cut (S, \bar{S}) and w_{ijk}^S is either equal to $w_{ij|k}, w_{jk|i}$ or $w_{ki|j}$, depending on how (S, \bar{S}) chops the hyperedge $\{i, j, k\}$. Now, similar to (Charikar & Chatziafratis, 2017; Dasgupta, 2016), we can recursively run a blackbox approximation algorithm for HSC to solve 3HHC. The main result of this section is the following technical proposition, whose proof is analogous to that of Theorem 1 (provided in the supplementary materials).

Proposition 2. *Given the hypergraph $G^{\mathcal{H}}$ with k hyperedges, and given access an α_n -approximation algorithm for HSC, the Recursive Hypergraph Sparsest Cut (R-HSC) algorithm achieves an $O(k\alpha_n)$ -approximation.*

Reduction from HSC back to Sparsest Cut. We now show how to get an α_n -approximation oracle for our instance of the HSC problem by a general reduction to sparsest cut. Our reduction is simple: given a hypergraph $G^{\mathcal{H}}$ and all the weights, create an instance of sparsest cut with the same vertices, (normal) edges and (normal) edge weights. Moreover, for every 3-hyperedge $\{a, b, c\}$ consider adding a *triangle* to the graph, i.e. three weighted edges connecting $\{a, b, c\}$, where:

$$\begin{aligned} w'_{ab} &= \frac{w_{bc|a} + w_{ac|b} - w_{ab|c}}{2} = \lambda \cdot c_{ab|c}, \\ w'_{ac} &= \frac{w_{bc|a} + w_{ab|c} - w_{ac|b}}{2} = 0, \\ w'_{bc} &= \frac{w_{ac|b} + w_{ab|c} - w_{bc|a}}{2} = 0. \end{aligned}$$

This construction can be seen in Figure 3. The important observation is that $w'_{ab} + w'_{ac} = w_{bc|a}, w'_{ab} + w'_{bc} = w_{ac|b}$ and $w'_{bc} + w'_{ac} = w_{ab|c}$, which are exactly the weights associated with the corresponding splits of the 3-hyperedge $\{a, b, c\}$. So, correctness of the reduction⁹ follows as the weight of each cut is preserved between the hypergraph and the graph after adding the triangles. For a discussion on extending this gadget more generally, see the supplement.



Figure 3. Transforming a 3-hyperedge to a triangle.

4. Variations on a Theme

In this section we study dissimilarity-HC, and we look into the problem of designing approximation algorithms for both unconstrained and constrained hierarchical clustering. In (Cohen-Addad et al., 2017), they show that average linkage is a $\frac{1}{2}$ -approximation for this problem and they propose a top-down approach based on locally densest cut achieving a $(\frac{2}{3} - \epsilon)$ -approximation in time $\tilde{O}\left(\frac{n^2(n+m)}{\epsilon}\right)$. Notably, when ϵ gets small the running time blows up.

Here, we prove that the most natural randomized algorithm for this problem, i.e. recursive random cutting, is a $\frac{2}{3}$ -approximation with expected running time $O(n \log n)$. We further derandomize this algorithm to get a simple deterministic local-search style $\frac{2}{3}$ -approximation algorithm.

If we also have structural constraints for the dissimilarity-HC, we show that the existing approaches fail. In fact we show that they lead to an $\Omega(n)$ -approximation factor due to the lack of “exploration” (e.g. recursive densest cut). We then show that recursive random cutting is robust to adding user constraints, and indeed it preserves a constant approximation factor when there are, roughly speaking, constantly many user constraints.

Randomized $\frac{2}{3}$ -approximation. Consider the most natural randomized algorithm for hierarchical clustering, i.e. recursively partition each cluster into two, where each point in the current cluster independently flips an unbiased coin and based on the outcome, it is put in one of the two parts.

Theorem 3. *Recursive-Random-Cutting is a $\frac{2}{3}$ -approximation for maximizing dissimilarity-HC objective.*

Proof sketch. An alternative view of Dasgupta’s objective is to divide the reward of the clustering tree between all possible triples $\{i, j, k\}$, where $(i, j) \in E$ and k is another point (possibly equal to i or j). Now, in any hierarchical clustering tree, if at the moment right before i and j become separated the vertex k has still been in the same cluster as

⁹Since all weights in the final graph are non-negative, standard techniques for Sparsest Cut can be used.

$\{i, j\}$, then this triple contributes w_{ij} to the objective function. We claim this event happens with probability exactly $\frac{2}{3}$. To see this, consider an infinite independent sequence of coin flips for i, j , and k . Without loss of generality, condition on i 's sequence to be all heads. The aforementioned event happens only if j 's first tails in its sequence happens no later than k 's first tails in its sequence. This happens with probability $\sum_{i \geq 1} \frac{1}{2} (\frac{1}{4})^{i-1} = \frac{2}{3}$. Therefore, the algorithm gets the total reward $\frac{2n}{3} \sum_{(i,j) \in E} w_{ij}$ in expectation. Moreover, the total reward of any hierarchical clustering is upper-bounded by $n \sum_{(i,j) \in E} w_{ij}$, which completes the proof of the $\frac{2}{3}$ -approximation. \square

Remark 2. This algorithm runs in time $O(n \log n)$ in expectation, due to the fact that the binary clustering tree has expected depth $O(\log n)$ (see for example (Cormen et al., 2009)) and at each level we only perform n operations.

We now derandomize the recursive random cutting algorithm using the *method of conditional expectations*. At every recursion, we go over the points in the current cluster one by one, and decide whether to put them in the “left” partition or “right” partition for the next recursion. Once we make a decision for a point, we fix that point and go to the next one. Roughly speaking, these local improvements can be done in polynomial time, which will result in a simple local-search style deterministic algorithm.

Theorem 4. *There is a deterministic local-search style $\frac{2}{3}$ -approximation algorithm for maximizing dissimilarity-HC objective that runs in time $O(n^2(n+m))$.*

Maximizing the Objective with User Constraints

From a practical point of view, one can think of many settings in which the output of the hierarchical clustering algorithm should satisfy user-defined hard constraints. Now, combining the new perspective of maximizing Dasgupta’s objective with this practical consideration raises a natural question: which algorithms are robust to adding user constraints, in the sense that a simple variation of these algorithms still achieve a decent approximation factor?

• **Failure of “Non-exploring” Approaches.** Surprisingly enough, there are convincing reasons that adapting existing algorithms for maximizing Dasgupta’s objective (e.g. those proposed in (Cohen-Addad et al., 2018)) to handle user constraints is either challenging or hopeless. First, bottom-up algorithms, e.g. average-linkage, fail to output a feasible outcome if they only consider each constraint separately and not all the constraints jointly (as we saw in Figure 1). Second, maybe more surprisingly, the natural extension of (locally) Recursive-Densest-Cut¹⁰ algorithm proposed in (Cohen-Addad et al., 2018) to handle user constraints performs poorly in the worst-case, even when we

¹⁰While a locally densest cut can be found in poly-time, densest cut is NP-hard, making our negative result stronger.

have only one constraint. Recursive-Densest-Cut proceeds by repeatedly picking the cut that has maximum density, i.e. $\arg \max_{S \subseteq V} \frac{w(S, \bar{S})}{|S| \cdot |\bar{S}|}$ and making two clusters. To handle the user constraints, we run it recursively on the supergraph generated by the constraints, similar to the approach in Section 2. Note that once the algorithm resolves a triplet constraint, it also breaks its corresponding supernode.

Now consider the following example in Figure 4, in which there is just one triplet constraint $ab|c$. The weight W should be thought of as large and ϵ as small. By choosing appropriate weights on the edges of the clique K_n , we can fool the algorithm into cutting the dense parts in the clique, without ever resolving the $ab|c$ constraint until it is too late. The algorithm gets a gain of $O(n^3 + W)$ whereas OPT gets $\Omega(nW)$ by starting with the removal of the edge (b, c) and then removing (a, b) , thus enjoying a gain of $\approx nW$.




Figure 4. $\Omega(n)$ -approximation lower bound instance for the constrained Recursive-Densest-Cut algorithm.

• **Constrained Recursive Random Cutting.** The example in Figure 4, although a bit pathological, suggests that a meaningful algorithm for this problem should explore cutting low-weight edges that might lead to resolving constraints, maybe randomly, with the hope of unlocking rewarding edges that were hidden before this exploration.

Formally, our approach is showing that the natural extension of recursive random cutting for the constrained problem, i.e. by running it on the supergraph generated by constraints and unpacking supernodes as we resolve the constraints (in a similar fashion to CSC), achieves a constant factor approximation when the constraints have bounded *dependency*. In the remaining of this section, we define an appropriate notion of dependency between the constraints, under the name of *dependency measure* and analyze the approximation factor of constrained recursive random cutting (Constrained-RRC) based on this notion.

Suppose we are given an instance of hierarchical clustering with triplet constraints $\{c_1, \dots, c_k\}$, where $c_i = x^i | y^i z^i, \forall i \in [k]$. For any triplet constraint c_i , let’s call the pair $\{y^i, z^i\}$ the *base*, and z^i the *key* of the constraint. We first partition our constraints into equivalence classes $\mathcal{C}_1, \dots, \mathcal{C}_N$, where $\mathcal{C}_i \subseteq \{c_1, \dots, c_k\}$. For every i, j , the constraints c_i and c_j belong to the same class \mathcal{C} if they share




Figure 5. Description of a class \mathcal{C} with base $\{x, y\}$.




Figure 6. Classes $\{\mathcal{C}_i, \mathcal{C}_j\}$, and two situations for having $\mathcal{C}_i \rightarrow \mathcal{C}_j$.

the same base (see Figure 5).

Definition 1 (Dependency digraph). The Dependency digraph is a directed graph with vertex set $\{\mathcal{C}_1, \dots, \mathcal{C}_L\}$. For every i, j , there is a directed edge $\mathcal{C}_i \rightarrow \mathcal{C}_j$ if $\exists c = x|yz, c' = x'|y'z'$, such that $c \in \mathcal{C}_i, c' \in \mathcal{C}_j$, and either $\{x, z\} = \{y', z'\}$ or $\{x, y\} = \{y', z'\}$ (see Figure 6).

The dependency digraph captures how groups of constraints impact each other. Formally, the existence of the edge $\mathcal{C}_i \rightarrow \mathcal{C}_j$ implies that all the constraints in \mathcal{C}_j should be resolved before one can separate the two endpoints of the (common) base edge of the constraints in \mathcal{C}_i .

Remark 3. If the constraints $\{c_1, \dots, c_k\}$ are feasible, i.e. there exists a hierarchical clustering that can respect all the constraints, the dependency digraph is clearly acyclic.

Definition 2 (Layered dependency subgraph). Given any class \mathcal{C} , the *layered dependency subgraph* of \mathcal{C} is the induced subgraph in the dependency digraph by all the classes that are reachable from \mathcal{C} . Moreover, the vertex set of this subgraph can be partitioned into layers $\{\mathcal{I}_0, \mathcal{I}_1, \dots, \mathcal{I}_L\}$, where L is the maximum length of any directed path leaving \mathcal{C} and \mathcal{I}_l is a subset of classes where the length of the longest path from \mathcal{C} to each of them is exactly equal to l (see Figure 7).

We are now ready to define a crisp quantity for every dependency graph. This will later help us give a more meaningful and refined *beyond-worst-case* guarantee for the approximation factor of the Constrained-RRC algorithm.

Definition 3 (Dependency measure). Given any class \mathcal{C} ,




Figure 7. Layered dependency subgraph of class \mathcal{C} .

the *dependency measure* of \mathcal{C} is defined as

$$\text{DM}(\mathcal{C}) \triangleq \prod_{l=0}^L (1 + \sum_{\mathcal{C}' \in \mathcal{I}_l} |\mathcal{C}'|),$$

where $\mathcal{I}_0, \dots, \mathcal{I}_L$ are the layers of the dependency subgraph of \mathcal{C} , as in Definition 2. Moreover, the dependency measure of a set of constraints $\text{DMC}(\{c_1, \dots, c_k\})$ is defined as $\max_{\mathcal{C}} \text{DM}(\mathcal{C})$, where the maximum is taken over all the classes generated by $\{c_1, \dots, c_k\}$.

Intuitively speaking, the notion of the dependency measure quantitatively expresses how “deeply” the base of a constraint is *protected* by the other constraints, i.e. how many constraints need to be resolved first before the base of a particular constraint is unpacked and the Constrained-RRC algorithm can enjoy its weight. This intuition is formalized through the following theorem, whose proof is deferred to the supplementary materials.

Theorem 5. *The constrained recursive random cutting (Constrained-RRC) algorithm is an α -approximation algorithm for maximizing dissimilarity-HC objective objective given a set of feasible constraints $\{c_1, \dots, c_k\}$, where*

$$\alpha = \frac{2(1 - k/n)}{3 \cdot \text{DMC}(\{c_1, \dots, c_k\})} \leq \frac{2(1 - k/n)}{3 \cdot \max_{\mathcal{C}} \text{DM}(\mathcal{C})}$$

Corollary 1. *Constrained-RRC is an $O(1)$ -approximation for maximizing dissimilarity-HC objective, given feasible constraints of constant dependency measure.*

5. Conclusion

We studied the problem of hierarchical clustering when we have structural constraints on the feasible hierarchies. We followed the optimization viewpoint that was recently developed in (Dasgupta, 2016; Cohen-Addad et al., 2018) and we analyzed two natural top-down algorithms giving provable approximation guarantees. In the case where the constraints are infeasible, we proposed and analyzed a regularized version of the HC objective by using the hypergraph version of the sparsest cut. Finally, we explored a variation of Dasgupta’s objective and improved upon previous techniques, both in the unconstrained and in the constrained setting.

Acknowledgements

Vaggos Chatziafratis was partially supported by ONR grant N00014-17-1-2562. Rad Niazadeh was supported by Stanford Motwani fellowship. Moses Charikar was supported by NSF grant CCF-1617577 and a Simons Investigator Award. We would also like to thank Leo Keselman, Aditi Raghunathan and Yang Yuan for providing comments on an earlier draft of the paper. We also thank the anonymous reviewers for their helpful comments and suggestions.

References

- Aho, A. V., Sagiv, Y., Szymanski, T. G., and Ullman, J. D. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM Journal on Computing*, 10(3):405–421, 1981.
- Arora, S., Rao, S., and Vazirani, U. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):5, 2009.
- Awasthi, P., Balcan, M., and Voevodski, K. Local algorithms for interactive clustering. In *International Conference on Machine Learning*, pp. 550–558, 2014.
- Balcan, M.-F. and Blum, A. Clustering with interactive feedback. In *International Conference on Algorithmic Learning Theory*, pp. 316–328. Springer, 2008.
- Balcan, M.-F., Blum, A., and Vempala, S. A discriminative framework for clustering via similarity functions. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pp. 671–680. ACM, 2008.
- Balcan, M.-F., Liang, Y., and Gupta, P. Robust hierarchical clustering. *The Journal of Machine Learning Research*, 15(1):3831–3871, 2014.
- Berkhin, P. A survey of clustering data mining techniques. In *Grouping multidimensional data*, pp. 25–71. Springer, 2006.
- Bilenko, M., Basu, S., and Mooney, R. J. Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 11. ACM, 2004.
- Brodal, G. S., Fagerberg, R., Mailund, T., Pedersen, C. N., and Sand, A. Efficient algorithms for computing the triplet and quartet distance between trees of arbitrary degree. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1814–1832. Society for Industrial and Applied Mathematics, 2013.
- Charikar, M. and Chatziafratis, V. Approximate hierarchical clustering via sparsest cut and spreading metrics. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 841–854. Society for Industrial and Applied Mathematics, 2017.
- Chatziafratis, V., Niazadeh, R., and Charikar, M. Hierarchical clustering with structural constraints. *arXiv preprint arXiv:1805.09476*, 2018. URL <https://arxiv.org/abs/1805.09476>.
- Cohen-Addad, V., Kanade, V., and Mallmann-Trenn, F. Hierarchical clustering beyond the worst-case. In *Advances in Neural Information Processing Systems*, pp. 6202–6210, 2017.
- Cohen-Addad, V., Kanade, V., Mallmann-Trenn, F., and Mathieu, C. Hierarchical clustering: Objective functions and algorithms. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 378–397. SIAM, 2018.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009. ISBN 0262033844, 9780262033848.
- Dasgupta, S. A cost function for similarity-based hierarchical clustering. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 118–127. ACM, 2016.
- Diez, I., Bonifazi, P., Escudero, I., Mateos, B., Muñoz, M. A., Stramaglia, S., and Cortes, J. M. A novel brain partition highlights the modular skeleton shared by structure and function. *Scientific reports*, 5:10532, 2015.
- Eisen, M. B., Spellman, P. T., Brown, P. O., and Botstein, D. Cluster analysis and display of genome-wide expression patterns. *Proceedings of the National Academy of Sciences*, 95(25):14863–14868, 1998.
- Emamjomeh-Zadeh, E. and Kempe, D. Adaptive hierarchical clustering using ordinal queries. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 415–429. SIAM, 2018.
- Kleindessner, M. and von Luxburg, U. Kernel functions based on triplet comparisons. In *Advances in Neural Information Processing Systems*, pp. 6810–6820, 2017.
- Leskovec, J., Rajaraman, A., and Ullman, J. D. *Mining of massive datasets*. Cambridge university press, 2014.
- Lichman, M. Uci machine learning repository, zoo dataset, 2013. URL <http://archive.ics.uci.edu/ml/datasets/zoo>.
- Mann, C. F., Matula, D. W., and Olinick, E. V. The use of sparsest cuts to reveal the hierarchical community structure of social networks. *Social Networks*, 30(3):223–234, 2008.

- Moseley, B. and Wang, J. Approximation bounds for hierarchical clustering: Average linkage, bisecting k-means, and local search. In *Advances in Neural Information Processing Systems*, pp. 3097–3106, 2017.
- Patané, J. S., Martins, J., and Setubal, J. C. Phylogenomics. In *Comparative Genomics*, pp. 103–187. Springer, 2018.
- Rangapuram, S. S. and Hein, M. Constrained 1-spectral clustering. In *Artificial Intelligence and Statistics*, pp. 1143–1151, 2012.
- Roy, A. and Pokutta, S. Hierarchical clustering via spreading metrics. In *Advances in Neural Information Processing Systems*, pp. 2316–2324, 2016.
- Steinbach, M., Karypis, G., Kumar, V., et al. A comparison of document clustering techniques. In *KDD workshop on text mining*, volume 400, pp. 525–526. Boston, 2000.
- Tamuz, O., Liu, C., Belongie, S., Shamir, O., and Kalai, A. T. Adaptively learning the crowd kernel. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pp. 673–680. Omnipress, 2011.
- Tumminello, M., Lillo, F., and Mantegna, R. N. Correlation, hierarchies, and networks in financial markets. *Journal of Economic Behavior & Organization*, 75(1):40–58, 2010.
- Vikram, S. and Dasgupta, S. Interactive bayesian hierarchical clustering. In *International Conference on Machine Learning*, pp. 2081–2090, 2016.
- Wagstaff, K. and Cardie, C. Clustering with instance-level constraints. *AAAI/IAAI*, 1097:577–584, 2000.
- Wagstaff, K., Cardie, C., Rogers, S., Schrödl, S., et al. Constrained k-means clustering with background knowledge. In *ICML*, volume 1, pp. 577–584, 2001.