

# Hierarchical Correlation Clustering and Tree Preserving Embedding

Morteza Haghiri Chehreghani  
 Chalmers University of Technology  
 Gothenburg, Sweden  
 morteza.chehreghani@chalmers.se

Mostafa Haghiri Chehreghani  
 Amirkabir University of Technology (Tehran Polytechnic)  
 Tehran, Iran  
 mostafa.chehreghani@gmail.com

## Abstract

We propose a hierarchical correlation clustering method that extends the well-known correlation clustering to produce hierarchical clusters applicable to both positive and negative pairwise dissimilarities. Then, in the following, we study unsupervised representation learning with such hierarchical correlation clustering. For this purpose, we first investigate embedding the respective hierarchy to be used for tree preserving embedding and feature extraction. Thereafter, we study the extension of minimax distance measures to correlation clustering, as another representation learning paradigm. Finally, we demonstrate the performance of our methods on several datasets.

## 1. Introduction

Data clustering plays an essential role in unsupervised learning and exploratory data analysis. It is used in a variety of applications including web mining, network analysis, image segmentation, bioinformatics, user analytics and knowledge management. Its goal is to partition the data into groups in a way that the objects in the same cluster are more similar according to some criterion, compared to the objects in different clusters.

Many clustering methods partition the data into  $K$  flat clusters for example,  $K$ -means [55], spectral clustering [62,68] and correlation clustering [8]. In many applications, however, the clusters are preferred to be presented at different levels, encompassing both high-level and detailed information. Hierarchical clustering is useful to produce such structures, usually encoded by a *dendrogram*. A dendrogram is a tree data structure where each node corresponds to a cluster, with the leaf nodes (those at the bottom of the tree) containing only one object. Higher-level clusters are formed by aggregating lower-level clusters and the inter-cluster dissimilarity between them.

Hierarchical clustering can be performed either in an agglomerative (i.e., bottom-up) or in a divisive (i.e., top-down) manner [56]. Agglomerative methods are often computationally

more efficient, making them more popular in practice [64]. In both approaches, the clusters are aggregated or split based on various criteria, such as *single*, *average*, *centroid*, *complete* and *Ward*. Several studies aim to improve these methods. The works in [49,52] focus on the statistical significance of hierarchical clustering. [24,25,65] formulate this problem as an optimization problem and propose approximate solutions. [82] considers multiple dissimilarities for a pair of clusters, and [11,17] suggest merging multiple clusters at each step instead of one. [6] employs global information to eliminate the influence of noisy similarities, and [19] proposes to apply agglomerative methods to small subsets of the data instead of individual data objects. [33,38] augment agglomerative methods with probabilistic models, and finally, [23,60] propose efficient but approximate methods for hierarchical clustering.

On the other hand, most clustering methods, either flat or hierarchical, assume non-negative pairwise (dis)similarities. However, in several practical applications, pairwise similarities can be any real number, positive or negative. For example, it could be preferable for a user or oracle to indicate whether two objects are similar (considered a positive relation) or dissimilar (considered a negative relation), rather than solely providing a positive (non-negative) pairwise similarity, even if the two objects are dissimilar. The former approach yields more precise information because, in the latter scenario, the dissimilarity between two objects (i.e., zero similarity) could be confused with a lack of available information. Some relevant applications for this setting include image segmentation with higher order correlation information [47,48], webpage segmentation [12], community detection over graphs [67], social media mining [73], analysis of connections over web [43], dealing with attraction/rejection data [26], automated label generation from clicks [3] and entity resolution [7,34].

Hence, a specialized clustering model known as *correlation clustering* has been developed to work with such data. This model was first introduced on the graphs with only +1 or -1 pairwise similarities [7,8], and then was generalized to the graphs with arbitrary positive or negative

edge weights [5, 13, 26]. The original model obtains the number of clusters automatically. The variant in [20] limits the number of clusters to fixed  $K$  clusters. Semidefinite programming (SDP) relaxation provides tight approximation bounds in particular for maximizing the agreements [13, 58], although it is computationally inefficient in practice [74]. Then, [15, 74] provide efficient greedy algorithms based on local search and Frank-Wolfe optimization with a fast convergence rate.

However, all of these methods produce flat correlation clusters. In this paper, we first propose a *Hierarchical Correlation Clustering* (HCC) method that handles both positive and negative pairwise (dis)similarities and produces clusters at different levels (Section 3). To the best of our knowledge, this work is one of the first extensions of the well-known correlation clustering to hierarchical clustering.<sup>1</sup> A hierarchical correlation clustering, also called HCC, is developed in [76]. This method offers a 0.4767 approximation, but lacking experimental evaluation. Furthermore, unlike our method, the method in [76] does not follow the generic agglomerative clustering procedure.

We note that unlike flat clustering, hierarchical clustering yields an ordering among the objects, i.e., objects that join earlier in the hierarchy are closer to each other than those that join at later steps. This implies that hierarchical clustering induces a new (dis)similarity measure between the objects, connected to the way the objects join each other to form clusters at different levels. Thereby, in the following, we consider two representation learning methods related to hierarchical clustering and study their adaptation to hierarchical correlation clustering. This enables us to not only use HCC for producing hierarchical clusters, but also to employ it for computing a suitable similarity/distance measure as an intermediate data processing step.

One way to perform representation learning from hierarchical clustering is to compute an embedding that corresponds to the respective hierarchy. *Tree preserving embedding* [69, 70] is a method that achieves this for the special case of single linkage method. Later, [21] develops tree preserving embedding for various standard agglomerative clustering methods. We then adapt these works (in particular the later work [21]) to develop a tree preserving embedding for HCC dendrograms (Section 4) where the embedded features can be used as a set of new features for an arbitrary downstream task. In this way, we can investigate HCC for the purpose of computing relevant features for a probabilistic method such as Gaussian Mixture Model (GMM), instead of solely using HCC for the purpose of hierarchical clustering. This enables us to apply a method like GMM for

<sup>1</sup>We note the so-called hierarchical correlation clustering methods proposed in [2, 36, 53] are irrelevant to the well-studied correlation clustering problem [7, 8]; they study for example the correlation coefficients for high-dimensional data.

clustering the pairwise similarities that can be positive or negative numbers, a task that was not possible before.

Another representation learning paradigm that we study is called *minimax* dissimilarity, a graph-based method that is tightly connected to hierarchical clustering. It provides a sophisticated way to infer transitive relations and extract manifolds and elongated clusters in an unsupervised way [14, 32, 46, 54]. Thereby, for the first time, we study minimax dissimilarities on the graphs with positive and negative (dis)similarities, i.e., with correlation clustering (Section 5). We show that using minimax dissimilarities with correlation clustering not only helps for extracting elongated patterns, but also yields a significant reduction in the computational complexity, i.e., from NP-hardness to a polynomial runtime.

We finally perform several experiments on various datasets to demonstrate the effectiveness of our methods in different settings (Section 6).

## 2. Notations and Definitions

A dataset is characterized by a set of  $n$  objects with indices  $\mathbf{O} = \{1, \dots, n\}$  and a pairwise similarity or dissimilarity matrix. An  $n \times n$  matrix  $\mathbf{S} \in \mathbb{R}^{n \times n}$  represents the pairwise similarities between the objects, whereas, the pairwise dissimilarities are shown by matrix  $\mathbf{D} \in \mathbb{R}^{n \times n}$ . Both of similarities and dissimilarities can be positive or negative. This property allows us to convert the pairwise similarities to dissimilarities by a simple transformation such as  $\mathbf{D} = -\mathbf{S}$ , i.e., the pairwise dissimilarities are obtained by the negation of the similarities and vice versa.<sup>2</sup> The objects and the pairwise (dis)similarities are represented by graph  $\mathcal{G} = (\mathbf{O}, \mathbf{S})$  or  $\mathcal{G} = (\mathbf{O}, \mathbf{D})$ .

A cluster is represented by a set, e.g.,  $\mathbf{v}$ , which includes the objects belong to that. The function  $dis(\mathbf{u}, \mathbf{v})$  denotes the inter-cluster dissimilarity between clusters  $\mathbf{u}$  and  $\mathbf{v}$  that can be defined according to different criteria. A hierarchical clustering solution can be represented by a dendrogram  $T$  defined as a rooted ordered tree such that, i) each node  $\mathbf{v}$  in  $T$  includes a non-empty subset of the objects corresponding to a cluster, i.e.,  $\mathbf{v} \subseteq \mathbf{O}, |\mathbf{v}| \geq 1, \forall \mathbf{v} \in T$ , with the leaf nodes including distinct single objects, and ii) the overlapping clusters are ordered, i.e.,  $\forall \mathbf{u}, \mathbf{v} \in T$ , if  $\mathbf{u} \cap \mathbf{v} \neq \emptyset$ , then either  $\mathbf{u} \subseteq \mathbf{v}$  or  $\mathbf{v} \subseteq \mathbf{u}$ . The latter condition implies that between every two overlapping nodes an ancestor-descendant relation holds, i.e.,  $\mathbf{u} \subseteq \mathbf{v}$  indicates  $\mathbf{v}$  is an ancestor of  $\mathbf{u}$ , and  $\mathbf{u}$  is a descendant of  $\mathbf{v}$ .

The clusters at the lowest level, called *leaf* clusters/node, are the individual distinct objects, i.e.,  $\mathbf{v}$  is a leaf cluster if

<sup>2</sup>Such a *nonparametric* transformation resolves the issues related to obtaining a proper similarity measure from pairwise dissimilarities. For example, with kernels, e.g., RBF kernels, finding the optimal parameter(s) is often crucial and nontrivial, and the optimal parameters occur inside a very narrow range [61, 80]. Moreover, the methods we develop in this paper are unaffected by the choice of the transformation  $\mathbf{D}$ ; for example in Algorithm 1, we only use the pairwise similarities  $\mathbf{S}$ .

and only if  $|\mathbf{v}| = 1$ . A cluster at a higher level contains the union of the objects of its children. The root of a dendrogram is defined as the cluster at the highest level which has the maximum size, i.e., all other clusters are its descendants.  $linkage(\mathbf{v})$ ,  $\mathbf{v} \in T$  returns the dissimilarity between the children of  $\mathbf{v}$  based on the criterion used to compute the dendrogram (i.e.,  $dis(c_l, c_r)$  where  $c_l$  and  $c_r$  indicate the two child clusters of  $\mathbf{v}$ ). For simplicity of explanation, w.l.g., we assume every non-leaf cluster has two children. The level of cluster  $\mathbf{v}$ , i.e.,  $level(\mathbf{v})$ , is determined by

$$level(\mathbf{v}) = \max(level(c_l), level(c_r)) + 1. \quad (1)$$

For the leaf clusters,  $level()$  and  $dis()$  return 0. Every connected subtree of  $T$  whose leaf clusters contain only individual objects from  $\mathbf{O}$  constitutes a dendrogram on this subset of objects. We require that every common node present in both  $T$  and the subtree must have the same child nodes or clusters. We use  $\mathcal{T}_T$  to refer to the set of all (sub)dendrograms obtained in this way from  $T$ .

### 3. Hierarchical Correlation Clustering

Agglomerative methods begin with each object in a separate cluster, and then at each round, combine the two clusters that have a *minimal* dissimilarity according to a criterion (defined by the  $dis(\cdot, \cdot)$  function) until only one cluster remains. For example, the *single* linkage (SL) criterion [71] defines the dissimilarity between two clusters as the dissimilarity between their nearest members ( $dis(\mathbf{u}, \mathbf{v}) = \min_{i \in \mathbf{u}, j \in \mathbf{v}} D_{i,j}$ ), whereas, *complete* linkage (CL) [50] uses the dissimilarity between their farthest members ( $dis(\mathbf{u}, \mathbf{v}) = \max_{i \in \mathbf{u}, j \in \mathbf{v}} D_{i,j}$ ). On the other hand, the *average* linkage (AL) criterion [72] considers the average of the inter-cluster dissimilarities as the dissimilarity between the two clusters ( $dis(\mathbf{u}, \mathbf{v}) = \sum_{i \in \mathbf{u}, j \in \mathbf{v}} \frac{D_{i,j}}{|\mathbf{u}||\mathbf{v}|}$ ). These methods can be shown to be shift-invariant, as mentioned in Proposition 1 [18].

**Proposition 1** *Single linkage, complete linkage and average linkage methods are invariant w.r.t. the shift of the pairwise dissimilarities by an arbitrary real number  $\alpha$ .*

Thus, we can still use these methods even with possibly negative pairwise dissimilarities as shifting the pairwise dissimilarities (by a large enough constant) to make them non-negative does not change the solution.

However, clustering the data consisting of positive and negative dissimilarities is usually conducted by *correlation clustering*. Thus, despite the applicability of single linkage, average linkage and complete linkage methods, we propose a novel hierarchical clustering consistent with the standard correlation clustering, called Hierarchical Correlation Clustering (HCC). This method is thus adapted to positive/negative pairwise (dis)similarities, and as our experi-

ments confirm, it outperforms the other methods (i.e., SL, CL, and AL) when applied to such data.

The cost function for flat (standard) correlation clustering accounts for disagreements (i.e., negative similarities inside clusters and positive similarities between clusters) and is written by [20]

$$R^{CC}(\mathbf{v}_1, \dots, \mathbf{v}_K; \mathbf{S}) = \frac{1}{2} \sum_{k=1}^K \sum_{i,j \in \mathbf{v}_k} (|\mathbf{S}_{ij}| - \mathbf{S}_{ij}) + \frac{1}{2} \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} (|\mathbf{S}_{ij}| + \mathbf{S}_{ij}), \quad (2)$$

where  $K$  is the number of clusters and  $\mathbf{v}_k$ 's indicate the different clusters.

We may rewrite the cost function as

$$R^{CC}(\mathbf{v}_1, \dots, \mathbf{v}_K; \mathbf{S}) = - \underbrace{\frac{1}{2} \sum_{k=1}^K \sum_{k'=1}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} \mathbf{S}_{ij}}_{constant} + \underbrace{\frac{1}{2} \sum_{k=1}^K \sum_{i,j \in \mathbf{v}_k} |\mathbf{S}_{ij}| + \frac{1}{2} \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} |\mathbf{S}_{ij}|}_{constant} + \frac{1}{2} \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} \mathbf{S}_{ij} + \frac{1}{2} \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} \mathbf{S}_{ij}. \quad (3)$$

We then have

$$R^{CC}(\mathbf{v}_1, \dots, \mathbf{v}_K; \mathbf{S}) = constant + \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} \mathbf{S}_{ij} \equiv constant - \sum_{k=1}^K \sum_{\substack{k'=1, \\ k' \neq k}}^K \sum_{i \in \mathbf{v}_k} \sum_{j \in \mathbf{v}_{k'}} D_{ij}. \quad (4)$$

Therefore, correlation clustering aims to minimize the inter-cluster similarities, and in other words, it maximizes the inter-cluster dissimilarities. This formulation in Eq. 4 inspires us for a *consistent* way of defining a new inter-cluster dissimilarity function for hierarchical (agglomerative) correlation clustering of positive and negative (dis)similarities. At each step, we merge the two clusters that have a minimal dissimilarity (or a maximal similarity), where we define the dissimilarity between the two clusters

$\mathbf{u}$  and  $\mathbf{v}$  as

$$dis^{CC}(\mathbf{u}, \mathbf{v}) = \sum_{i \in \mathbf{u}} \sum_{j \in \mathbf{v}} \mathbf{D}_{ij} = - \sum_{i \in \mathbf{u}} \sum_{j \in \mathbf{v}} \mathbf{S}_{ij}. \quad (5)$$

We emphasize that HCC is consistent with the generic agglomerative clustering framework applied with, for example, *single* linkage, *average* linkage, *complete* linkage and other criteria. The only difference is the definition of the inter-cluster dissimilarity function where with HCC we use  $dis^{CC}(\cdot, \cdot)$  defined in Eq. 5 (inspired from the cost function of flat correlation clustering). Other than this, the algorithmic procedure is consistent. Algorithm 1 in Appendix A describes the pseudocode of the HCC algorithm.

## 4. Feature Extraction from HCC

As mentioned, HCC represents the relations between objects according to the way they join to form the hierarchy. In this section, we use this intuition and develop a data representation consistent with HCC. For this purpose, we adapt the methods in [69, 70] and in particular [21] to our setting. Hereby, we first introduce distance functions over HCC, and then, investigate the embedding of such a distance function. This procedure leads to obtaining a set of features from HCC for each object which then can be used in the downstream task.

### 4.1. Distance functions over HCC

Given dendrogram  $T$ , each cluster  $\mathbf{v} \in T$  represents the root of a dendrogram  $T' \in \mathcal{T}_T$ .  $T'$  admits the properties of its root cluster, i.e.,  $level(T') = \max_{\mathbf{u} \in T'} level(\mathbf{u}) = level(\mathbf{v})$  and  $linkage(T') = \max_{\mathbf{u} \in T'} linkage(\mathbf{u}) = linkage(\mathbf{v})$ , since the root cluster has the maximum linkage and level among the clusters in  $T'$ . Hence, in this way, we define functions such as  $level()$  and  $linkage()$  for the dendrograms as well.

The  $linkage()$  function may seem to be a natural choice for defining a distance function over a HCC dendrogram. Specifically, one can define the dendrogram-based distance function  $\mathbf{X}_{ij}$  over dendrogram  $T$  between  $i, j \in \mathbf{O}$  as

$$\mathbf{X}_{ij} = \min linkage(T') \quad \text{s.t.} \quad i, j \in T', \text{ and } T' \in \mathcal{T}_T. \quad (6)$$

This choice corresponds to the linkage of the smallest cluster that includes both  $i$  and  $j$ . This in particular makes sense for the single linkage dendrogram, and it would be consistent with the tree-preserving embedding in [69, 70]. If the original dissimilarity matrix  $\mathbf{D}$  contains negative values, then using Proposition 1, one can sufficiently shift the pairwise dissimilarities to make all of them non-negative, without changing the structure of the dendrogram and the order of the clusters. Therefore, the conditions for a valid distance function including non-negativity still hold.

However, for the HCC dendrogram, the linkage function might not fulfill the conditions for a distance function. For example, consider a set of  $n$  objects where all the pairwise similarities are +1, i.e., the dissimilarities are thus  $-1$ . Then, the linkage function will always return negative values which would violate the non-negativity condition of a valid distance function. On the other hand, the HCC linkage  $dis^{CC}(\cdot, \cdot)$  is *not shift-invariant* (similar to the standard flat correlation clustering [18]) and we cannot use the shift trick in Proposition 1. Let  $\mathbf{D}^\alpha$  shows the shifted pairwise dissimilarities, i.e.,  $\mathbf{D}_{i,j}^\alpha = \mathbf{D}_{ij} + \alpha$ . Then,  $dis^{CC}(\mathbf{u}, \mathbf{v})$  between two clusters  $\mathbf{u}$  and  $\mathbf{v}$  based on  $\mathbf{D}^\alpha$  is given by

$$\begin{aligned} dis^{CC}(\mathbf{u}, \mathbf{v}) &= \sum_{i \in \mathbf{u}, j \in \mathbf{v}} \mathbf{D}_{i,j}^\alpha = \sum_{i \in \mathbf{u}, j \in \mathbf{v}} (\mathbf{D}_{i,j} + \alpha) \\ &= \sum_{i \in \mathbf{u}, j \in \mathbf{v}} \mathbf{D}_{i,j} + \alpha |\mathbf{u}| |\mathbf{v}|. \end{aligned} \quad (7)$$

With  $\alpha > 0$ , this shift would induce a bias for the HCC linkage to choose imbalanced clusters. In other words,  $dis^{CC}(\cdot, \cdot)$  is not shift-invariant and we cannot shift the pairwise dissimilarities in  $\mathbf{D}$  to make them nonnegative.

Therefore, we consider another choice, i.e., the  $level()$  function used in [21]. It is nonnegative and satisfies the desired conditions. Then,  $\mathbf{X}_{ij}$  is now computed by

$$\mathbf{X}_{ij} = \min level(T') \quad \text{s.t.} \quad i, j \in T' \text{ and } T' \in \mathcal{T}_T. \quad (8)$$

Intuitively, Eq. 8 selects the level of the smallest cluster/dendrogram that contains both  $i$  and  $j$ . The lower the level at which the two objects join, the greater the similarity or proximity between them, indicating a closer relationship in the hierarchical clustering structure. In other words, a higher level in the dendrogram signifies a later fusion of the two objects, suggesting that they share fewer common characteristics compared to objects fused at lower levels.

### 4.2. Embedding the HCC-based distances

After applying the distance function in Eq. 8, we obtain an  $n \times n$  matrix representing pairwise HCC-based distances among objects. It is usually preferred to obtain vector-based representations for objects rather than pairwise distances. Models like Gaussian Mixture Models (GMMs) which involve mixture density estimation (see, e.g., [75]), can only be applied to vectors. Additionally, working with vector-based data simplifies feature selection. Hence, it is desired to compute an embedding of the objects into a new space, so that their pairwise squared Euclidean distances in the new space match their pairwise distances obtained from the dendrogram.

The matrix of pairwise distances  $\mathbf{X}$  obtained via Eq. 8 induces an *ultrametric* [21, 51]. The primary distinction between a *metric* and an *ultrametric* is that the *addition* operation in the triangle inequality for a metric is replaced by a *maximum* operation, i.e., with ultrametric we have

$$\forall i, j, k : \mathbf{X}_{ij} \leq \max(\mathbf{X}_{ik}, \mathbf{X}_{kj}). \quad (9)$$

The connection between ultrametric and trees is well-established in mathematics [42,57]. Here we instantiate it to our setting via making the argument in [21] more accurate.

It is evident that when  $\mathbf{X}_{ij} \leq \mathbf{X}_{ik}$ , the inequality in Eq. 9 is satisfied. Conversely, if  $\mathbf{X}_{ij} > \mathbf{X}_{ik}$ , it implies that objects  $i$  and  $k$  are included in the same cluster (shown by  $c_{i,k}$ ) before  $i$  and  $j$  join (to form cluster  $c_{i,j}$ ). The bottom-up hierarchical clustering process then continues until  $c_{i,j,k}$  is formed, encompassing all three objects  $i, j, k$ . Notice that  $i$  and  $k$  join  $j$  simultaneously via  $c_{i,k}$ . In this case, according to Eqs. 8 and 9, and the relationships illustrated in Figure 1, we conclude

$$\mathbf{X}_{ij} = \mathbf{X}_{kj} \leq \max(\mathbf{X}_{ik}, \mathbf{X}_{kj}). \quad (10)$$




Figure 1. Illustration of ultrametric property of  $\mathbf{X}$ .

Ultrametric matrices, in turn, exhibit positive definiteness [31, 78], and such positive definite matrices result in inducing an Euclidean embedding [66]. Thereby, after ensuring the existence of such an embedding, we can employ a proper method to compute it. Specifically, we use the method proposed in [83] known as *multidimensional scaling* [59]. This method proposes first centering  $\mathbf{X}$  to obtain a Mercer kernel and then performing an eigenvalue decomposition. It works as follows.

1. We center  $\mathbf{X}$  by

$$\mathbf{B} \leftarrow -\frac{1}{2}(\mathbf{I}_n - \frac{1}{n}\mathbf{L}_n)\mathbf{X}(\mathbf{I}_n - \frac{1}{n}\mathbf{L}_n), \quad (11)$$

where  $\mathbf{I}_n$  is an identity matrix of size  $n \times n$  and  $\mathbf{L}_n$  represents an  $n \times n$  matrix filled entirely with ones. With this centering, the sum of both the rows and columns in matrix  $\mathbf{B}$  becomes zero.

2. With applying the transformation in step 1,  $\mathbf{B}$  becomes a positive semidefinite matrix, i.e., all the eigenvalues are nonnegative. Thus, we decompose  $\mathbf{B}$  into its eigenbasis:

$$\mathbf{B} = \mathbf{Y}\mathbf{Z}\mathbf{Y}^T, \quad (12)$$

where  $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$  represents the eigenvectors  $\mathbf{y}_i$  and  $\mathbf{Z} = \text{diag}(z_1, \dots, z_n)$  is a diagonal matrix of eigenvalues  $z_1 \geq \dots \geq z_l \geq z_{l+1} = 0 = \dots = z_n$ .

3. We calculate the  $n \times l$  matrix  $\mathbf{A}$  as the matrix of new data features:

$$\mathbf{A} = \mathbf{Y}_l(\mathbf{Z}_l)^{1/2}, \quad (13)$$

with  $\mathbf{Y}_l = (\mathbf{y}_1, \dots, \mathbf{y}_l)$  and  $\mathbf{Z}_l = \text{diag}(z_1, \dots, z_l)$ , where  $l$  specifies the dimensionality of the new vectors.

In the embedded space, the new dimensions are arranged based on their corresponding eigenvalues. One can opt to select only the most significant ones, rather than utilizing all of them. Therefore, computing such an embedding also offers the benefit of feature selection.

We also note that many clustering methods can be written in matrix factorization form via for example spectral  $K$ -means [28]. This induces an embedding and hence a set of relevant features. However, for general positive/negative similarity matrices no exact embedding might be feasible due to violating positive semidefiniteness. The method that we described here provides a solution to this challenge, i.e., enables us to extract features when the base pairwise similarities are positive and negative.

## 5. Correlation Clustering and Minimax Dissimilarities

Finally, we study minimax dissimilarities for correlation clustering, a graph-based method that corresponds to constructing a hierarchical clustering.

Given graph  $\mathcal{G}(\mathbf{O}, \mathbf{D})$ , the minimax (MM) dissimilarity between  $i$  and  $j$  is defined as

$$\mathbf{D}_{ij}^{MM} = \min_{p \in \mathcal{P}_{ij}(\mathcal{G})} \max_{1 \leq l \leq |p|-1} \mathbf{D}_{p(l)p(l+1)}, \quad (14)$$

where  $\mathcal{P}_{ij}(\mathcal{G})$  is the set of all paths between  $i$  and  $j$  over  $\mathcal{G}(\mathbf{O}, \mathbf{D})$ . Each path  $p$  is specified by a sequence of object indices, i.e.,  $p(l)$  indicates the  $l^{\text{th}}$  object on the path.

Minimax dissimilarities enable a clustering algorithm to capture the inherent patterns and manifolds in an unsupervised and nonparametric way by extracting the transitive connections [16,32]. For example, if object  $i$  is similar to object  $j$ ,  $j$  is similar to  $k$ , and  $k$  is similar to  $l$ , then the minimax dissimilarity between  $i$  and  $l$  will be small, even though their direct dissimilarity might be large. The reason is that minimax dissimilarity finds the connectivity path  $i \rightarrow j \rightarrow \dots \rightarrow k \rightarrow l$  and connects  $i$  and  $l$  via this path. This property is helpful in finding elongated clusters and manifolds of arbitrary shapes in an unsupervised way.

Minimax dissimilarities have been so far solely used with nonnegative pairwise dissimilarities. In the case of possible negative dissimilarities, we may use a trick similar to Proposition 1. As shown in Lemma 1, minimax paths are invariant w.r.t. the shift of the pairwise dissimilarities.

**Lemma 1** Consider graphs  $\mathcal{G}(\mathbf{O}, \mathbf{D})$  and  $\mathcal{G}^\alpha(\mathbf{O}, \mathbf{D}^\alpha)$ , where the pairwise dissimilarities (edge weights) in

$\mathcal{G}^\alpha(\mathbf{O}, \mathbf{D}^\alpha)$  are shifted by constant  $\alpha$ , i.e.,  $\mathbf{D}_{i,j}^\alpha = \mathbf{D}_{i,j} + \alpha$ . Then, the minimax paths between every pair of objects  $i$  and  $j$  are identical on graphs  $\mathcal{G}(\mathbf{O}, \mathbf{D})$  and  $\mathcal{G}^\alpha(\mathbf{O}, \mathbf{D}^\alpha)$ .

All the proofs are in Appendix B. Hence, given a dissimilarity matrix  $\mathbf{D}$ , one can subtract  $\alpha := \min(\mathbf{D})$  from all the elements to obtain  $\mathbf{D}^\alpha$ . Then, the minimax dissimilarities can be computed from  $\mathcal{G}^\alpha(\mathbf{O}, \mathbf{D}^\alpha)$ . After computing the minimax dissimilarities from  $\mathcal{G}^\alpha$ , we may add  $\alpha$  to all the pairwise minimax dissimilarities. We can obtain the minimax similarities  $\mathbf{S}_{ij}^{MM}$  via  $\mathbf{S}_{ij}^{MM} = -\mathbf{D}_{ij}^{MM}$ , if needed.

We demonstrate that for correlation clustering there exists a simpler method to calculate the minimax dissimilarities intended for use in correlation clustering. According to Theorem 1, performing correlation clustering on minimax dissimilarities can be achieved in polynomial time via computing the connected components of the unweighted graph  $\mathcal{G}(\mathbf{O}, \mathbf{S}')$ , where the similarity matrix  $\mathbf{S}'$  is obtained by  $\mathbf{S}'_{ij} = 1$  if  $\mathbf{S}_{ij} > 0$ , and  $\mathbf{S}'_{ij} = 0$  otherwise.

**Theorem 1** *The optimal clusters of the correlation clustering on graph  $\mathcal{G}(\mathbf{O}, \mathbf{S}^{MM})$  are equal to the connected components of graph  $\mathcal{G}(\mathbf{O}, \mathbf{S}')$ .*

As mentioned, correlation clustering on an arbitrary similarity matrix  $\mathbf{S}$  is NP-hard [8, 26]. Therefore, using minimax (dis)similarities with correlation clustering not only helps for extracting elongated complex patterns, but also yields a significant reduction in the computational complexity, i.e., from NP-hardness to a polynomial runtime.

Among the approximate algorithms proposed for correlation clustering on complete graphs with discrete weights, the method in [5] provides a 3-factor approximation. With a randomly selected unclustered object in the graph, this method greedily finds the object’s positive neighbors (those with similarity +1) to form a new cluster. Then, it repeats this procedure for the remaining objects. One can conclude that in the optimal solution of correlation clustering on graph  $\mathcal{G}(\mathbf{O}, \mathbf{S}^{MM})$ , only the positive neighbors of an object will be in the same cluster as the object is, i.e., interestingly the 3-factor approximation algorithm in [5] becomes exact when applied to  $\mathcal{G}(\mathbf{O}, \mathbf{S}^{MM})$  (Theorem 2).

**Theorem 2** *Assume the edge weights of graph  $\mathcal{G}(\mathbf{O}, \mathbf{S})$  are either +1 or -1. Then, the approximate algorithm in [5] is exact when applied to the minimax similarities, i.e., to graph  $\mathcal{G}(\mathbf{O}, \mathbf{S}^{MM})$ .*

## 6. Experiments

In this section, we describe our experimental results on various datasets. We compare our methods with single linkage (SL), complete linkage (CL) and average linkage (AL).<sup>3</sup>

<sup>3</sup>Some criteria, e.g. *centroid*, *median* and *Ward* compute a representative for each cluster and then compute the inter-cluster dissimilarities by

As mentioned, there are several improvements over these basic methods. However, such contributions are orthogonal to our contribution. Moreover, it is unclear how such improvements can be extended to the dissimilarities that can be both positive and negative. Thus, we limit our baselines to these methods which as mentioned in Proposition 1, are applicable to such data.

In our studies, we have access to the true labels. Therefore, consistent with several previous studies on hierarchical clustering, e.g. [4, 9, 29], we evaluate the results according to the following criteria: i) Normalized Mutual Information (MI) [79] that measures the mutual information between the true and the estimated clustering solutions, and ii) Normalized Rand score (Rand) [41] that obtains the similarity between the two solutions. We do not use the labels to infer the clustering solution, they are only used for evaluation. Therefore, we are still in the unsupervised setting where the ground-truth labels play the role of an external evaluator.

In Appendix C, we describe additional experimental results, in particular on datasets from other domains.

### 6.1. HCC on UCI data

We first investigate the hierarchical correlation clustering on the following six UCI datasets [45]. (i) *Breast Tissue*: includes electrical impedance measurements of freshly excised 106 tissue samples from the breast. The number of clusters is 6. (ii) *Cardiotocography*: contains 2126 measurements of fetal heart rate and uterine contraction features on cardiotocograms in 10 clusters. (iii) *Image Segmentation*: contains 2310 samples from images of 7 outdoor clusters. (iv) *ISOLET*: 7797 samples consisting of spoken attributes of different letters (26 clusters). (v) *Leaf*: 340 images of leaf specimens originating from 40 different plant species (clusters) each described by 16 attributes. (vi) *One-Hundred Plant*: 1600 samples of leafs each described by 64 features, from in total 100 types (clusters). The ground-truth labels are shown by  $\mathbf{c}^*$ , i.e.,  $\mathbf{c}_i^*$  shows the true label for object  $i$ . We assume an oracle reveals the pairwise similarities  $\mathbf{S}$  according to the (flip) noise parameter  $\eta$ . If  $\mathbf{c}_i^* = \mathbf{c}_j^*$  then  $\mathbf{S}_{i,j} = \mathcal{U}(0, 1)$  with probability  $1 - \eta$  and  $\mathbf{S}_{i,j} = \mathcal{U}(-1, 0)$  with probability  $\eta$ . If  $\mathbf{c}_i^* \neq \mathbf{c}_j^*$  then  $\mathbf{S}_{i,j} = \mathcal{U}(-1, 0)$  with probability  $1 - \eta$  and  $\mathbf{S}_{i,j} = \mathcal{U}(0, 1)$  with probability  $\eta$ . The function  $\mathcal{U}(\cdot, \cdot)$  returns a uniform random number within the specified range. For each  $\eta$  we repeat the experiments 20 times and report the average results. This setup provides a systematic approach to study the robustness of various methods to noise.

Figure 2 shows the results for different datasets as a function of the noise level  $\eta$  w.r.t. MI. Rand scores shown in Figure 5 in Appendix C exhibit a consistent behavior. We

the distances between the representatives. Computing such representatives might not be feasible for possibly negative pairwise dissimilarities. Thus, we do not consider them.




Figure 2. MI score of different hierarchical clustering methods applied to UCI datasets, where x-axis shows the parameter  $\eta$ .

Table 1. Performance of different tree preserving embedding methods on UCI datasets applied with GMM.

	<i>Breast Tissue</i>		<i>Cardiotocography</i>		<i>Image Segm.</i>		<i>ISOLET</i>		<i>Leaf</i>		<i>One-Hun. Plant</i>	
method	MI	Rand	MI	Rand	MI	Rand	MI	Rand	MI	Rand	MI	Rand
SL	0.006	0.008	0.007	0.007	0.008	0.001	0.009	0.016	0.008	0.003	0.015	0.020
SL+GMM	0.093	0.077	0.120	0.135	0.239	0.250	0.192	0.174	0.155	0.161	0.083	0.077
CL	0.227	0.166	0.077	0.056	0.187	0.125	0.057	0.017	0.081	0.038	0.029	0.008
CL+GMM	0.251	0.171	0.081	0.060	0.201	0.154	0.061	0.043	0.140	0.129	0.054	0.049
AL	0.542	0.519	0.391	0.479	0.518	0.495	0.257	0.165	0.181	0.106	0.066	0.023
AL+GMM	0.550	0.513	0.422	0.463	0.522	0.501	0.240	0.179	0.152	0.143	0.081	0.065
HCC	0.903	0.900	<b>0.987</b>	<b>0.994</b>	0.945	0.943	0.938	<b>0.918</b>	0.429	0.373	0.159	0.104
HCC+GMM	<b>0.914</b>	<b>0.911</b>	0.979	0.974	<b>0.960</b>	<b>0.966</b>	<b>0.941</b>	0.917	<b>0.462</b>	<b>0.401</b>	<b>0.183</b>	<b>0.217</b>

observe that among different methods, HCC performs significantly better and produces more robust clusters w.r.t. the noise parameter  $\eta$ . The results on *Leaf* and *One-Hundred Plant* are worse with all the methods. The reason is that these datasets are complex, having many clusters (respectively, 40 and 100 clusters) and fairly a small number of objects per cluster.

## 6.2. Tree preserving embedding on UCI data

In the following, we investigate tree preserving embedding and feature extraction. After computing the embeddings from different hierarchical clustering methods, we apply Gaussian Mixture Model (GMM) to the extracted features and evaluate the final clustering using the ground-truth solution. This kind of embedding enables us to apply methods like GMM to positive and negative pairwise similarities, a task that was not possible before. Since the extracted features appear in the form of vectors, thus, the final clustering method is not limited to GMM, and in principle, any numerical clustering can benefit from this embedding.

Table 1 shows the results for different UCI datasets. We observe that the embeddings obtained by HCC (i.e., ‘HCC+GMM’) yield significantly better results compared to the embeddings from other methods (i.e., ‘SL+GMM’, ‘CL+GMM’ and ‘AL+GMM’). It is worth noting that the results of embeddings (e.g., ‘HCC+GMM’) typically surpass the results obtained from the hierarchical clustering alone (e.g., ‘HCC’). This observation supports the idea that employing HCC to compute an embedding (for extracting new features) for a clustering method such as GMM may

be advantageous, yielding superior results compared to using HCC exclusively for clustering purposes. This verifies why tree preserving embedding can be effective in general.

## 6.3. Experiments on Fashion-MNIST

Next, we investigate HCC and tree-preserving embedding on two randomly selected subsets of Fashion-MNIST dataset [81]. Fashion MNIST consists of  $28 \times 28$  images of Zalando’s articles. Each subset consists of 5,000 samples/objects, where we compute the pairwise cosine similarities between them and then apply different methods. Table 2 shows the performance of different methods on these datasets. We observe that, consistent with the previous experiments, both ‘HCC’ and ‘HCC+GMM’ yield improving the results compared to the baselines. Furthermore, employing HCC to compute intermediate features for GMM (i.e., ‘HCC+GMM’) achieves higher scores compared to using ‘HCC’ alone for generating final clusters.

## 6.4. Correlation clustering with minimax dissimilarities

Finally, we investigate the use of minimax dissimilarities with correlation clustering. As mentioned, minimax dissimilarities are especially useful for extracting elongated and complex patterns in an unsupervised way. Thereby, we apply ‘*minimax + correlation clustering*’ to a number of datasets with visually elongated and arbitrarily shaped clusters, and compare the results with ‘*correlation clustering*’ alone. The datasets are shown in Figure 3. For each dataset, we simply construct the K-nearest neighbor graph (using

Table 2. Performance of different methods on MNIST and Fashion MNIST datasets. The embeddings by HCC yield better results.

method	Fashion MNIST 1		Fashion MNIST 2	
	MI	Rand	MI	Rand
SL	0.322	0.206	0.241	0.196
SL+GMM	0.411	0.335	0.384	0.340
CL	0.403	0.293	0.546	0.379
CL+GMM	0.478	0.426	0.574	0.362
AL	0.464	0.468	0.602	0.534
AL+GMM	<b>0.608</b>	0.551	0.647	0.553
HCC	0.499	0.475	0.666	<b>0.557</b>
HCC+GMM	0.581	<b>0.586</b>	<b>0.693</b>	0.548

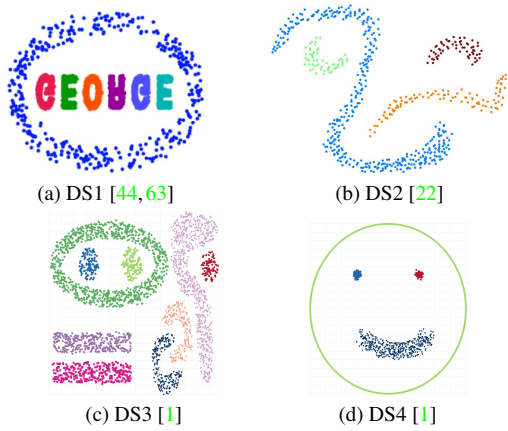


Figure 3. The datasets with arbitrarily shaped clusters, where ‘minimax + correlation clustering’ achieves perfect clustering.

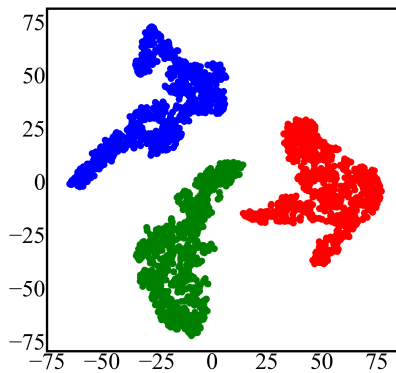


Figure 4. Embeddings of vehicle motion trajectories computed via dynamic time wrapping and t-SNE [27, 39].

the ordinary Euclidean distance with a typical  $K$  such as 3). The edges in the nearest neighborhood receive a positive weight (e.g., +1), and a negative weight (e.g., -1) otherwise. We then apply either ‘minimax + correlation clustering’ or ‘correlation clustering’ to the resultant graph. We observe that for all the datasets, ‘minimax + correlation clustering’

yields the perfect clustering, i.e., MI and Rand are equal to 1. Whereas with ‘correlation clustering’ alone these scores are very low. We acknowledge the existence of other clustering methods, such as DBSCAN [30], capable of achieving perfect clustering on these datasets. However, these methods often involve *crucial hyperparameters*, the tuning of which can be challenging in unsupervised learning. To our knowledge, in addition to computational and theoretical benefits, ‘minimax + correlation clustering’ stands out as the sole method capable of achieving perfect results on datasets with elongated clusters of arbitrary shapes, and fulfills the following two promises: i) *it eliminates the need to fix critical hyperparameters, a task often intricate in unsupervised learning*, and ii) *it automatically determines the correct number of clusters without requiring prior fixing*.

In the following, we consider the interesting application of clustering vehicle motion trajectories for ensuring safety in self-driving [39]. In this application, 1,024 trajectories consisting of drive-by left, drive-by right and cut-in types are prepared, where some of them are collected in real-world and some others are generated using Recurrent Auto-Encoder GAN [27]. Next, dynamic time wrapping [10] and t-SNE [77] are employed to map the temporal data onto a two-dimensional space, as illustrated in Figure 4 [27, 39]. We then apply ‘minimax + correlation clustering’ to this data and compare it with ‘correlation clustering’. Similar to the datasets in Figure 3, ‘minimax + correlation clustering’ achieves perfect clustering with  $MI = Rand = 1$  and accurately determines the correct number of clusters.

## 7. Conclusion

We proposed a new hierarchical clustering method, called HCC, suitable when the (dis)similarities can be positive or negative. This method is consistent with the general algorithmic procedure for agglomerative clustering and only differs in the way the inter-cluster dissimilarity function is defined. We then considered embedding the HCC dendrograms, which provides extracting useful features to apply for example GMM for clustering positive and negative similarities. In the following, we studied the use of minimax dissimilarities with correlation clustering and showed that it yields reduction in the computational complexity, in addition to a possibility for extracting elongated manifolds. Finally, we demonstrated the effectiveness of the methods on several datasets in different settings.

## Acknowledgement

The work of Morteza Haghiri Chehreghani was partially supported by the Swedish Research Council VR (grant number 2023-04809) and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.



## References

- [1] Clustering datasets. <https://github.com/milaan9/Clustering-Datasets/>. 8
- [2] Elke Aichert, Christian Böhm, Peer Kröger, and Arthur Zimek. Mining hierarchies of correlation clusters. In *18th International Conference on Scientific and Statistical Database Management, SSDBM*, pages 119–128, 2006. 2
- [3] Rakesh Agrawal, Alan Halverson, Krishnaram Kenthapadi, Nina Mishra, and Panayiotis Tsaparas. Generating labels from clicks. In *International Conference on Web Search and Web Data Mining, WSDM*, pages 172–181, 2009. 1
- [4] Julien Ah-Pine. An efficient and effective generic agglomerative hierarchical clustering approach. *J. Mach. Learn. Res.*, 19:42:1–42:43, 2018. 6
- [5] Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5):23:1–23:27, 2008. 2, 6, 13
- [6] Maria-Florina Balcan, Yingyu Liang, and Pramod Gupta. Robust hierarchical clustering. *J. Mach. Learn. Res.*, 15(1):3831–3871, 2014. 1
- [7] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. In *43rd Symposium on Foundations of Computer Science FOCS*, 2002. 1, 2
- [8] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004. 1, 2, 6
- [9] MohammadHossein Bateni, Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Raimondas Kiveris, Silvio Lattanzi, and Vahab S. Mirrokni. Affinity clustering: Hierarchical clustering at scale. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6864–6874, 2017. 6
- [10] R. Bellman and R. Kalaba. On adaptive control processes. *Automatic Control, IRE Transactions on*, 4(2):1–9, 1959. 8
- [11] Michel Bruynooghe. Méthodes nouvelles en classification automatique de données taxinomiques nombreuses. *Statistique et analyse des données*, 2(3):24–42, 1977. 1
- [12] Deepayan Chakrabarti, Ravi Kumar, and Kunal Punera. A graph-theoretic approach to webpage segmentation. In *International Conference on World Wide Web, WWW*, pages 377–386, 2008. 1
- [13] Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. In *44th Symposium on Foundations of Computer Science (FOCS)*, pages 524–533, 2003. 2
- [14] Morteza Haghir Chehreghani. Classification with minimax distance measures. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 1784–1790, 2017. 2
- [15] Morteza Haghir Chehreghani. Clustering by shift. In *IEEE International Conference on Data Mining (ICDM)*, pages 793–798, 2017. 2
- [16] Morteza Haghir Chehreghani. Unsupervised representation learning with minimax distance measures. *Mach. Learn.*, 109(11):2063–2097, 2020. 5
- [17] Morteza Haghir Chehreghani. Reliable agglomerative clustering. In *International Joint Conference on Neural Networks, IJCNN*, pages 1–8. IEEE, 2021. 1
- [18] Morteza Haghir Chehreghani. Shift of pairwise similarities for data clustering. *Mach. Learn.*, 112(6):2025–2051, 2023. 3, 4, 14
- [19] Morteza Haghir Chehreghani, Hassan Abolhassani, and Mostafa Haghir Chehreghani. Improving density-based methods for hierarchical clustering of web pages. *Data Knowl. Eng.*, 67(1):30–50, 2008. 1
- [20] Morteza Haghir Chehreghani, Alberto Giovanni Busetto, and Joachim M. Buhmann. Information theoretic model validation for spectral clustering. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012. 2, 3
- [21] Morteza Haghir Chehreghani and Mostafa Haghir Chehreghani. Learning representations from dendrograms. *Mach. Learn.*, 109(9-10):1779–1802, 2020. 2, 4, 5
- [22] Dongdong Cheng, Qingsheng Zhu, Jinlong Huang, Quanzhang Wu, and Lijun Yang. A novel cluster validity index based on local cores. *IEEE Trans. Neural Networks Learn. Syst.*, 30(4):985–999, 2019. 8
- [23] Michael Cochez and Hao Mou. Twister tries: Approximate hierarchical agglomerative clustering for average distance in linear time. In *International Conference on Management of Data (ACM SIGMOD)*, pages 505–517, 2015. 1
- [24] Vincent Cohen-Addad, Varun Kanada, Frederik Mallmann-Trenn, and Claire Mathieu. Hierarchical clustering: Objective functions and algorithms. In *SODA*, pages 378–397, 2018. 1
- [25] Vincent Cohen-Addad, Varun Kanade, and Frederik Mallmann-Trenn. Hierarchical clustering beyond the worst-case. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6201–6209. 2017. 1
- [26] Erik D. Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immerlica. Correlation clustering in general weighted graphs. *Theor. Comput. Sci.*, 361(2-3):172–187, 2006. 1, 2, 6
- [27] Andreas Demetriou, Henrik Alfvåg, Sadeq Rahrovani, and Morteza Haghir Chehreghani. A deep learning framework for generation and analysis of driving scenario trajectories. *SN Comput. Sci.*, 4(3):251, 2023. 8
- [28] Inderjit S. Dhillon, Yuqiang Guan, and Brian Kulis. Kernel k-means: spectral clustering and normalized cuts. In *Tenth ACM International Conference on Knowledge Discovery and Data Mining, SIGKDD*, pages 551–556, 2004. 5
- [29] Laxman Dhulipala, David Eisenstat, Jakub Lacki, Vahab S. Mirrokni, and Jessica Shi. Hierarchical agglomerative graph clustering in nearly-linear time. In *International Conference on Machine Learning (ICML)*, 2021. 6
- [30] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231, 1996. 8
- [31] Miroslav Fiedler. Ultrametric sets in euclidean point spaces. *ELA. The Electronic Journal of Linear Algebra*, 3:23–30, 1998. 5

- [32] Bernd Fischer and Joachim M. Buhmann. Path-based clustering for grouping of smooth curves and texture segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(4):513–518, 2003. 2, 5
- [33] Chris Fraley and Adrian E Raftery. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97:611–631, 2002. 1
- [34] Lise Getoor and Ashwin Machanavajjhala. Entity resolution: Theory, practice & open challenges. *Proc. VLDB Endow.*, 5(12), 2012. 1
- [35] C. Gower and G. J. S. Ross. Minimum spanning trees and single linkage cluster analysis. *Journal of the Royal Statistical Society. Series C*, 18(1), 1969. 12
- [36] Yi Gu and Chaoli Wang. A study of hierarchical correlation clustering for scientific volume data. In *ISVC*, pages 437–446, 2010. 2
- [37] Haghiri Chehreghani, Morteza. *Information-theoretic validation of clustering algorithms*. PhD thesis, 2013. 14
- [38] Nicholas A. Heard. Iterative reclassification in agglomerative clustering. *Journal of Computational and Graphical Statistics*, 20(4):920–936, 2012. 1
- [39] Fazeleh Sadat Hoseini, Sadegh Rahrovani, and Morteza Haghiri Chehreghani. Vehicle motion trajectories clustering via embedding transitive relations. In *24th IEEE International Intelligent Transportation Systems Conference, ITSC*, pages 1314–1321. IEEE, 2021. 8
- [40] T.C. Hu. The maximum capacity route problem. *Operations Research*, 9:898–900, 1961. 12
- [41] L. Hubert and P. Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985. 6
- [42] Bruce Hughes. Trees and ultrametric spaces: a categorical equivalence. *Advances in Mathematics*, 189(1):148–191, 2004. 5
- [43] Dmitri V. Kalashnikov, Zhaoqi Chen, Sharad Mehrotra, and Rabia Nuray-Turan. Web people search via connection analysis. *IEEE Trans. Knowl. Data Eng.*, 20(11):1550–1565, 2008. 1
- [44] George Karypis. Cluto - a clustering toolkit, 2002. Retrieved from the University of Minnesota Digital Conservancy. 8
- [45] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The UCI machine learning repository. <http://archive.ics.uci.edu/datasets>. 6
- [46] Kye-Hyeon Kim and Seungjin Choi. Neighbor search with global geometry: a minimax message passing algorithm. In *Twenty-Fourth International Conference on Machine Learning, ICML*, pages 401–408, 2007. 2
- [47] Sungwoong Kim, Sebastian Nowozin, Pushmeet Kohli, and Chang Dong Yoo. Higher-order correlation clustering for image segmentation. In *Advances in Neural Information Processing Systems 24 (NIPS)*, pages 1530–1538, 2011. 1
- [48] Sungwoong Kim, Chang Dong Yoo, Sebastian Nowozin, and Pushmeet Kohli. Image segmentation using higher-order correlation clustering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(9):1761–1774, 2014. 1
- [49] Patrick K. Kimes, Yufeng Liu, David Neil Hayes, and James Stephen Marron. Statistical significance for hierarchical clustering. *Biometrics*, 73(3):811–821, 2017. 1
- [50] G. N. Lance and W. T. Williams. A general theory of classificatory sorting strategies. *The Computer Journal*, 9(4):373–380, 1967. 3
- [51] Bruno Leclerc. Description combinatoire des ultramétriques. *Mathématiques et Sciences Humaines*, 73:5–37, 1981. 4
- [52] Mark A. Levenstien, Yaning Yang, and Jurg Ott. Statistical significance for hierarchical clustering in genetic association and microarray expression studies. *BMC Bioinformatics*, 4(1), 2003. 1
- [53] Tom Liebmann, Gunther H. Weber, and Gerik Scheuermann. Hierarchical correlation clustering in multiple 2d scalar fields. *Comput. Graph. Forum*, 37(3):1–12, 2018. 2
- [54] Anna V. Little, Mauro Maggioni, and James M. Murphy. Path-based spectral clustering: Guarantees, robustness to outliers, and fast algorithms. *J. Mach. Learn. Res.*, 21:6:1–6:66, 2020. 2
- [55] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *5th Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967. 1
- [56] Oded Maimon and Lior Rokach. *Data Mining and Knowledge Discovery Handbook*. 2005. 1
- [57] Álvaro Martínez-Pérez and Manuel A. Morón. Uniformly continuous maps between ends of  $\mathbb{R}$ -trees. *Mathematische Zeitschrift*, 263(3):583–606, 2008. 5
- [58] Claire Mathieu and Warren Schudy. Correlation clustering with noisy input. In *SODA '10*, pages 712–728, 2010. 2
- [59] A. Mead. Review of the development of multidimensional scaling methods. *Journal of the Royal Statistical Society: Series D*, 41(1):27–39, 1992. 5
- [60] Daniel Mullner. Modern hierarchical, agglomerative clustering algorithms. *CoRR*, abs/1109.2378, 2011. 1
- [61] Boaz Nadler and Meirav Galun. Fundamental limitations of spectral clustering. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1017–1024, 2007. 2
- [62] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems (NIPS)*, pages 849–856, 2001. 1
- [63] Dehua Peng, Zhipeng Gui, Dehe Wang, Yuncheng Ma, Zichen Huang, Yu Zhou, and Huayi Wu. Clustering by measuring local direction centrality for data with heterogeneous density and weak connectivity. *Nature Communications*, 13(1), 2022. 8
- [64] J. Podani. *Introduction to the exploration of multivariate biological data*. Backhuys Publishers, 2000. 1
- [65] Aurko Roy and Sebastian Pokutta. Hierarchical clustering via spreading metrics. *J. Mach. Learn. Res.*, 18(1):3077–3111, 2017. 1
- [66] I. J. Schoenberg. On certain metric spaces arising from euclidean spaces by a change of metric and their imbedding in hilbert space. *Annals of Mathematics*, 38(4):787–793, 1937. 5
- [67] Jessica Shi, Laxman Dhulipala, David Eisenstat, Jakob Lacki, and Vahab S. Mirrokni. Scalable community detection via parallel correlation clustering. *Proc. VLDB Endow.*, 14(11):2305–2313, 2021. 1

- [68] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000. [1](#)
- [69] A. Shieh, T. B. Hashimoto, and E. M. Airoldi. Tree preserving embedding. In *International Conference on Machine Learning (ICML)*, pages 753–760, 2011. [2](#), [4](#)
- [70] A. D. Shieh, T. B. Hashimoto, and E. M. Airoldi. Tree preserving embedding. *Proceedings of the National Academy of Sciences*, 108(41):16916–16921, 2011. [2](#), [4](#)
- [71] Peter Henry Andrews Sneath. The application of computers to taxonomy. *Journal of General Microbiology*, 17:201–226, 1957. [3](#)
- [72] R. R. Sokal and C. D. Michener. A statistical method for evaluating systematic relationships. *Uni. of Kansas Science Bulletin*, 38:1409–1438, 1958. [3](#)
- [73] Jiliang Tang, Yi Chang, Charu C. Aggarwal, and Huan Liu. A survey of signed network mining in social media. *ACM Comput. Surv.*, 49(3):42:1–42:37, 2016. [1](#)
- [74] Erik Thiel, Morteza Haghiri Chehreghani, and Devdatt P. Dubhashi. A non-convex optimization approach to correlation clustering. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, pages 5159–5166, 2019. [2](#)
- [75] D. M. Titterton, etc., A. F. M. Smith, and U. E. Makov. *Statistical analysis of finite mixture distribution*. Probability & Mathematical Statistics S. John Wiley & Sons, 1985. [4](#)
- [76] D. Vainstein, V. Chatziafratis, G. Citovsky, A. Rajagopalan, M. Mahdian, and Y. Azar. Hierarchical clustering via sketches and hierarchical correlation clustering. In *AISTATS*, pages 559–567, 2021. [2](#)
- [77] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. [8](#)
- [78] Richard S. Varga and Reinhard Nabben. On symmetric ultrametric matrices. In *Numerical Linear Algebra*, 1993. [5](#)
- [79] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.*, 11:2837–2854, 2010. [6](#)
- [80] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. [2](#)
- [81] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. [7](#)
- [82] Pelin Yildirim and Derya Birant. K-linkage: A new agglomerative approach for hierarchical clustering. 17:77–88, 2017. [1](#)
- [83] Gale Young and A. Householder. Discussion of a set of points in terms of their mutual distances. *Psychometrika*, 3(1):19–22, 1938. [5](#)