

Contrastive Multi-view Hyperbolic Hierarchical Clustering

Fangfei Lin^{1,2*}, Bing Bai², Kun Bai², Yazhou Ren¹, Peng Zhao¹ and Zenglin Xu^{3,4†}

¹University of Electronic Science and Technology of China, Chengdu, China

²Tencent Security Big Data Lab, Tencent Inc., China

³Harbin Institute of Technology, Shenzhen, China

⁴Department of Network Intelligence, Peng Cheng National Lab, Shenzhen, China

phoebe.lin1108@gmail.com, {icebai, kunbai}@tencent.com, yazhou.ren@uestc.edu.cn, zhaop211@gmail.com, zenglin@gmail.com

Abstract

Hierarchical clustering recursively partitions data at an increasingly finer granularity. In real-world applications, multi-view data have become increasingly important. This raises a less investigated problem, i.e., multi-view hierarchical clustering, to better understand the hierarchical structure of multi-view data. To this end, we propose a novel neural network-based model, namely Contrastive Multi-view Hyperbolic Hierarchical Clustering (CMHHC). It consists of three components, i.e., multi-view alignment learning, aligned feature similarity learning, and continuous hyperbolic hierarchical clustering. First, we align sample-level representations across multiple views in a contrastive way to capture the view-invariance information. Next, we utilize both the manifold and Euclidean similarities to improve the metric property. Then, we embed the representations into a hyperbolic space and optimize the hyperbolic embeddings via a continuous relaxation of hierarchical clustering loss. Finally, a binary clustering tree is decoded from optimized hyperbolic embeddings. Experimental results on five real-world datasets demonstrate the effectiveness of the proposed method and its components.

1 Introduction

Clustering is one of the fundamental problems in data analysis, which aims to categorize unlabeled data points into clusters. Existing clustering methods can be divided into partitional clustering and hierarchical clustering (HC) [Jain *et al.*, 1999]. The difference is that partitional clustering produces only one partition, while hierarchical clustering produces a nested series of partitions. Compared with partitional clustering, hierarchical clustering reveals more information about fine-grained similarity relationships and structures of data.

Hierarchical clustering has gained intensive attention. There are usually two types of hierarchical methods, i.e.,

the agglomerative methods and continuous ones. Agglomerative methods include Single-Linkage, Complete-Linkage, Average-Linkage, Ward-Linkage, etc.; continuous methods include Ultrametric Fitting (UFit) [Chierchia and Perret, 2020] and Hyperbolic Hierarchical Clustering (HyphC) [Chami *et al.*, 2020], etc. Existing HC methods are usually limited to data from a single source. However, with the advances of data acquisition in real-world applications, data of an underlying signal is often collected from heterogeneous sources or feature subsets [Li *et al.*, 2018]. For example, images can be described by the local binary pattern (LBP) descriptor and the scale-invariant feature transform (SIFT) descriptor; websites can be represented by text, pictures, and other structured metadata. Data from different views may contain complementary and consensus information. Hence it would be greatly beneficial to utilize multi-view data to fertilize hierarchical clustering.

Compared with partitional clustering and single-view hierarchical clustering, multi-view hierarchical clustering is less investigated. It requires a finer-grained understanding of both the consistency and differences among multiple views. To this end, we propose a novel Contrastive Multi-view Hyperbolic Hierarchical Clustering (CMHHC) model, as depicted in Figure 1. The proposed model consists of three components, i.e., multi-view alignment learning, aligned feature similarity learning, and continuous hyperbolic hierarchical clustering. First, to encourage consistency among multiple views and to suppress view-specific noise, we align the representations by contrastive learning, whose intuition is that the same instance from different views should be mapped closer while different instances should be mapped separately. Next, to improve the metric property among data instances, we exploit more reasonable similarity of aligned features measured both on manifolds and in the Euclidean space, upon which we mine hard positives and hard negatives for unsupervised metric learning. Besides, we also assign autoencoders to each view to regularize the model and prevent model collapse. Then, we embed the representations with good similarities into the hyperbolic space and optimize the hyperbolic embeddings via the continuous relaxation of Dasgupta’s discrete objective for HC [Chami *et al.*, 2020]. Finally, the tree decoding algorithm helps decode the binary clustering tree from optimized hyperbolic embedding with low distortion.

To the best of our knowledge, the only relevant work on

*This work was done when Fangfei Lin was an intern at Tencent.

†Corresponding Author

hierarchical clustering for multi-view data is the multi-view hierarchical clustering (MHC) model proposed by Zheng *et al.* [2020]. MHC clusters multi-view data by alternating the cosine distance integration and the nearest neighbor agglomeration. Compared with MHC and naïve methods like multi-view concatenation followed by single-view hierarchical agglomerative clustering, our method enjoys several advantages. Firstly, compared with simply concatenating multiple views in naïve methods or the averaging operation in MHC, CMHHC incorporates a contrastive multi-view alignment process, which can better utilize the complementary and consensus information among different views to learn meaningful view-invariance features of instances. Secondly, compared with the shallow HC framework, deep representation learning and similarity learning are applied in our model to match complex real-world multi-view datasets and obtain more discriminative representations. Thirdly, compared with heuristic agglomerative clustering, our method is oriented to gradient-based clustering process by optimizing multi-view representation learning and hyperbolic hierarchical clustering loss functions. These loss functions provide explicit evidence for measuring the quality of tree structure, which is crucial to achieving better HC performance.

The contributions of this work are summarized as follows:

- To our knowledge, we propose the first deep neural network model for multi-view hierarchical clustering, which can capture the aligned and discriminative representations across multiple views and perform hierarchical clustering at diverse levels of granularity.
- To learn representative and discriminative multi-view embeddings, we exploit a contrastive representation learning module (to align representations across multiple views) and an aligned feature similarity learning module (to consider the manifold similarity). These embeddings are helpful to the downstream task of similarity-based clustering.
- We validate our framework with five multi-view datasets and demonstrate that CMHHC outperforms existing HC and MHC algorithms in terms of the Dendrogram Purity (DP) measurement.

2 Related Work

We review related work from three perspectives: hierarchical clustering, multi-view clustering, and hyperbolic geometry.

2.1 Hierarchical Clustering

Hierarchical clustering raises a recursive way for partitioning a dataset into successively finer clusters. Classical heuristics, like Single Linkage, Complete Linkage, Average Linkage, and Ward Linkage, are often the methods of choice for small datasets but may not scale well to large datasets as the running time scales cubically with the sample size [Kobren *et al.*, 2017]. Another problem with these heuristics is the lack of good objective functions, so there is no solid theoretical basis to support HC algorithms. To overcome the challenge, Dasgupta [2016] defined a proper discrete hierarchical clustering loss on all possible hierarchies. Re-

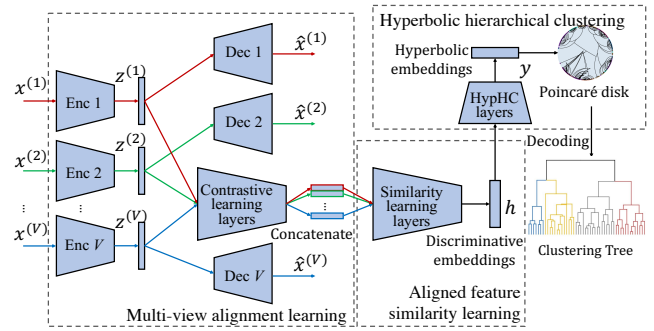


Figure 1: Overview of CMHHC. V different autoencoders are assigned for V different views. Contrastive learning layers are to align sample-level representations across multiple views. Similarity learning layers are to learn better metric property over aligned features. Then, HypHC layers are to optimize the tree-like embeddings in hyperbolic space. Finally, optimal hyperbolic embeddings are decoding into discrete clustering trees.

cently, gradient-based HC has gained increasing research attention, like feature-based gHHC [Monath *et al.*, 2019], and similarity-based UFit [Chierchia and Perret, 2020] and HypHC [Chami *et al.*, 2020].

This paper deals with multi-view hierarchical clustering. The most relevant work to this paper is MHC [Zheng *et al.*, 2020], which performs the cosine distance integration and the nearest neighbor agglomeration alternately. However, this shallow HC framework ignores extracting meaningful and consistent information from multiple views, probably resulting in degenerated clustering performance.

2.2 Multi-view Clustering

Existing multi-view clustering (MVC) methods mainly focus on partitional clustering. Traditional multi-view clustering includes four types, i.e., multi-view subspace clustering [Li *et al.*, 2019a], multi-view spectral clustering [Kang *et al.*, 2020], multi-view matrix factorization-based clustering [Cai *et al.*, 2013], and canonical correlation analysis (CCA)-based clustering [Chaudhuri *et al.*, 2009]. However, many MVC methods do not meet the requirements for complex non-linear situations. Thus, deep MVC approaches have been attached recently [Trosten *et al.*, 2021; Xu *et al.*, 2021; Xu *et al.*, 2022]. For example, Andrew *et al.* [2013] and Wang *et al.* [2015] proposed deep versions of CCA, termed as deep CCA and deep canonically correlated autoencoders respectively. Also, End-to-end Adversarial-attention network for Multi-modal Clustering (EAMC) [Zhou and Shen, 2020] leveraged adversarial learning and attention mechanism to achieve the separation and compactness of cluster structure.

Despite many works towards partitional clustering, there is minimal research towards multi-view hierarchical clustering, which we address in this paper.

2.3 Hyperbolic Geometry

Hyperbolic geometry is a non-Euclidean geometry with a constant negative curvature, which drops the parallel line postulate of the postulates of Euclidean geometry [Sala *et al.*, 2018]. Since the surface area lying in hyperbolic space grows

exponentially with its radius, hyperbolic space can be seen as a continuous version of trees whose number of leaf nodes also increases exponentially with the depth [Monath *et al.*, 2019]. Hyperbolic geometry has been adopted to fertilize research related to tree structures recently [Chami *et al.*, 2020; Nickel and Kiela, 2017; Yan *et al.*, 2021].

3 Method

Given a set of N data points including K clusters and V different views $\{\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^V\}_{i=1}^N$, where $\mathbf{x}_i^v \in \mathbb{R}^{D_v}$ denotes the i -th D_v -dimensional instance from the v -th view, we aim to produce a nested series of partitions, where the more similar multi-view samples are grouped earlier and have lower lowest common ancestors (LCAs). To this end, we establish a multi-view hierarchical clustering framework named CMHHC. We first introduce the network architecture, and then we define the loss functions and introduce the optimization process.

3.1 Network Architecture

The proposed network architecture consists of a multi-view alignment learning module, an aligned feature similarity learning module, and a hyperbolic hierarchical clustering module, which is shown in Figure 1. We introduce the details of each component as follows.

Multi-view Alignment Learning

Data from different sources tend to contain complementary and consensus information, so how to extract comprehensive representations and suppress the view-specific noise is the major task in multi-view representation learning. Therefore, we design the multi-view alignment learning module to map the data of each view into a low-dimensional aligned space. First of all, we assign a deep autoencoder [Hinton and Salakhutdinov, 2006] to each view. The reconstruction from input to output not only helps each view keep the view-specific information to prevent model collapse, but also fertilizes subsequent similarity computing through representation learning. Specifically, considering the v -th view $\mathbf{X}^{(v)}$, the corresponding encoder is represented as $\mathbf{Z}^{(v)} = E^{(v)}(\mathbf{X}^{(v)}; \theta_e^{(v)})$, and the decoder is represented as $\hat{\mathbf{X}}^{(v)} = D^{(v)}(\mathbf{Z}^{(v)}; \theta_d^{(v)})$, where $\theta_e^{(v)}$ and $\theta_d^{(v)}$ denote the autoencoder network parameters, and $\mathbf{Z}^{(v)} \in \mathbb{R}^{N \times D_{ac}}$ and $\hat{\mathbf{X}}^{(v)} \in \mathbb{R}^{N \times D_v}$ denote learned D_{ac} -dimensional latent features and the reconstructed output, respectively.

For the still-detached multiple views, inspired by recent works with contrastive learning [Xu *et al.*, 2022], we propose to achieve the consistency of high-level semantic features across views in a contrastive way, i.e., corresponding instances of different views should be mapped closer, while different instances should be mapped more separately. To be specific, we define a cross-view positive pair as two views describing the same object, and a negative pair as two different objects from the same view or two arbitrary views. We encode the latent features to aligned features denoted as $\mathbf{H}^{(v)} = f_{con}(\mathbf{Z}^{(v)}; \theta_c)$, where θ_c denotes the parameters of contrastive learning encoder $f_{con}(\cdot)$, and $\mathbf{H}^{(v)} \in \mathbb{R}^{N \times D_h}$.

Then, we compute the cosine similarity [Chen *et al.*, 2020] of two representations $\mathbf{h}_i^{(v_1)}$ and $\mathbf{h}_j^{(v_2)}$ ($v_1 \neq v_2$):

$$d_{i,j}^{(v_1)(v_2)} = \frac{\mathbf{h}_i^{(v_1)T} \mathbf{h}_j^{(v_2)}}{\|\mathbf{h}_i^{(v_1)}\| \cdot \|\mathbf{h}_j^{(v_2)}\|}. \quad (1)$$

To achieve consistency of all views, we expect the similarity scores of positive pairs $(\mathbf{h}_i^{(v_1)}, \mathbf{h}_i^{(v_2)})$ to be larger, and those of negative pairs $(\mathbf{h}_i^{(v_1)}, \mathbf{h}_j^{(v_2)})$ and $(\mathbf{h}_i^{(v)}, \mathbf{h}_j^{(v)})$ to be smaller.

Aligned Feature Similarity Learning

After the multi-view alignment learning, comprehensive representations of multi-view data are obtained, and view-specific noise is suppressed. However, the aligned hidden representations do not explicitly guarantee an appropriate similarity measurement that preserves desired distance structure between pairs of multi-view instances, which is crucial to similarity-based hierarchical clustering. Therefore, we devise an aligned feature similarity learning module to optimize the metric property of similarity used in hierarchy learning.

Intuitively, samples may be mapped to some manifold embedded in a high dimensional Euclidean space, so it would be beneficial to utilize the distance on manifolds to improve the similarity measurement [Isken *et al.*, 2018]. This module performs unsupervised metric learning by mining positives and negatives pairs with both the manifold and Euclidean similarities. Similar with Isken *et al.* [2018], we measure the Euclidean similarity of any sample pair via the mapping function $w_{i,j}^e(\mathbf{h}_i, \mathbf{h}_j) = \max(0, \mathbf{h}_i^T \mathbf{h}_j)^3$ and refer to $\text{ENN}_k(\mathbf{h}_i)$ as Euclidean k nearest neighbors (k NN) set of \mathbf{h}_i , where $\mathbf{h}_i \in \mathbf{H}$ is the concatenate aligned representation.

In terms of the manifold similarity, the Euclidean affinity graph needs to be calculated as preparations. The elements of the affinity matrix A are weighted as $a_{ij} = w_{i,j}^e(\mathbf{h}_i, \mathbf{h}_j)$ when \mathbf{h}_i and \mathbf{h}_j are both the Euclidean k NN nodes to each other, or else $a_{ij} = 0$. Following the spirit of advanced random walk model [Isken *et al.*, 2017], we can get the convergence solution \mathbf{r}_i efficiently, where an element $r_i(j)$ denotes the ‘‘best’’ walker from the i -th node to the j -th node. Therefore, the manifold similarity function can be defined as $w_{i,j}^m(\mathbf{h}_i, \mathbf{h}_j) = \mathbf{r}_i(j)$. Similarly, we denote $\text{MNN}_k(\mathbf{h}_i)$ as the manifold k NN set of \mathbf{h}_i .

To this end, we try to consider every data point in the dataset as an anchor in turn so that the hard mining strategy keeps feasible under the premise of acceptable computability. Given an anchor \mathbf{h}_i from \mathbf{H} , we select k_{pos} nearest feature vectors on manifolds, which are not that close in the Euclidean space, as good positives. By $\text{ENN}_{k_{\text{pos}}}(\mathbf{h}_i)$ and $\text{MNN}_{k_{\text{pos}}}(\mathbf{h}_i)$, the hard positive set is descendingly sorted by the manifold similarity as:

$$S_{\text{pos}}(\mathbf{h}_i) = \text{MNN}_{k_{\text{pos}}}(\mathbf{h}_i) - \text{ENN}_{k_{\text{pos}}}(\mathbf{h}_i), \quad (2)$$

where k_{pos} decides how hard the selected positives are, which is a completely detached value from k . However, to keep gained pseudo-label information with little noise, good negatives are expected to be not only relatively far from the anchor on manifolds but also in the Euclidean space. So the hard negative set is in the descending order according to the Euclidean

similarity, denoted as:

$$S_{\text{neg}}(\mathbf{h}_i) = S_{\text{all}}(\mathbf{H}) - (\text{ENN}_{k_{\text{neg}}}(\mathbf{h}_i) + \text{MNN}_{k_{\text{neg}}}(\mathbf{h}_i)), \quad (3)$$

where $S_{\text{all}}(\mathbf{H})$ is the set of all feature vectors, and k_{neg} is the value of the nearest neighbors for negatives, separated from k . Intuitively, a larger value of k_{pos} leads to harder positives for considering those with relatively lower confidence, and tolerating the intra-cluster variability. Similarly, a smaller value of k_{neg} means harder negatives for distinguishing the anchors from more easily-confused negatives.

After obtaining hard tuples as the supervision of similarity learning for multi-view input, we are able to get clustering-friendly embeddings on top of the aligned feature space via an encoder $\mathbf{E} = f_{\text{mom}}(\mathbf{H}; \theta_m)$, where $\mathbf{E} \in \mathbb{R}^{N \times D_e}$ is the discriminative embeddings of D_e dimensions and θ_m is the parameter of the encoder $f_{\text{mom}}(\cdot)$.

Hyperbolic Hierarchical Clustering

To reach the goal of hierarchical clustering, we adopt the continuous optimizing process of Hyperbolic Hierarchical Clustering [Chami *et al.*, 2020] as the guidance of this module, which is stacked on the learned common-view embeddings. By means of an embedder $f_{\text{hc}}(\cdot)$ parameterized by θ_{hc} , we embed the learned similarity graph from common-view space to a specific hyperbolic space, i.e., the Poincaré Model $\mathbb{B}_h = \{\|\mathbf{y}\|_2 \leq 1, \mathbf{y} \in \mathbb{R}^h\}$ whose curvature is constant -1 . The distance of the geodesic between any two points $\mathbf{y}_i, \mathbf{y}_j \in \mathbb{B}_h$ in hyperbolic space is $d_{\mathbb{B}}(\mathbf{y}_i, \mathbf{y}_j) = \cosh^{-1}(1 + \frac{2\|\mathbf{y}_i - \mathbf{y}_j\|_2^2}{(1 - \|\mathbf{y}_i\|_2^2)(1 - \|\mathbf{y}_j\|_2^2)})$. With this prior knowledge, we can find the optimal hyperbolic embeddings \mathbf{Y}^* pushed to the boundary of the ball via a relaxed form of improved Dasgupta’s cost function. Moreover, the hyperbolic LCA of two hyperbolic embeddings \mathbf{y}_i and \mathbf{y}_j is the point on their geodesic nearest to the origin of the ball. It can be represented as $\mathbf{y}_i \vee \mathbf{y}_j := \arg \min_{\mathbf{y}_o \in \mathbf{y}_i \rightsquigarrow \mathbf{y}_j} d(\mathbf{o}, \mathbf{y})$. The LCA of two hyperbolic embeddings is an analogy to that of two leaf nodes t_i and t_j in a discrete tree, where the tree-like LCA is the node closest to the root node on the two points’ shortest path [Chami *et al.*, 2020]. Based on this, the best hyperbolic embeddings can be decoded back to the original tree via getting merged iteratively along the direction of the ball radius from boundary to origin.

3.2 Loss Functions and Optimization Process

This section introduces the loss functions for CMHHC, including multi-view representation learning loss and hierarchical clustering loss, and discusses the optimization process.

Multi-view Representation Learning Loss

In our model, we jointly align multi-view representations and learn the reliable similarities in common-view space by integrating the autoencoder reconstruction loss \mathcal{L}_r , the multi-view contrastive loss \mathcal{L}_c and the positive weighted triplet metric loss \mathcal{L}_m , so the objective for multi-view representation learning loss \mathcal{L}_{mv} is defined as:

$$\mathcal{L}_{\text{mv}} = \mathcal{L}_r + \mathcal{L}_c + \mathcal{L}_m. \quad (4)$$

First of all, $\mathcal{L}_r^{(v)}$ is the v -th view loss function of reconstructing $\mathbf{x}_i^{(v)}$, so the complete reconstruction loss for all views is:

$$\mathcal{L}_r = \sum_{v=1}^V \mathcal{L}_r^{(v)} = \sum_{v=1}^V \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{x}_i^{(v)} - D^{(v)}(E^{(v)}(\mathbf{x}_i^{(v)})) \right\|_2^2. \quad (5)$$

As for the second term, assuming that alignment between every two views ensures alignment among all views, we introduce total multi-view mode of contrastive loss as follows:

$$\mathcal{L}_c = \sum_{v_1=1}^V \sum_{v_2=1, v_2 \neq v_1}^V \mathcal{L}_c^{(v_1)(v_2)}, \quad (6)$$

where the contrastive loss function between reference view v_1 and contrast view v_2 is:

$$\mathcal{L}_c^{(v_1)(v_2)} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{d_{i,i}^{(v_1)(v_2)}/\tau}}{\sum_{j=1, j \neq i}^N e^{d_{i,j}^{(v_1)(v_1)}/\tau} + \sum_{j=1}^N e^{d_{i,j}^{(v_1)(v_2)}/\tau}}, \quad (7)$$

where τ is the temperature hyperparameter for multi-view contrastive loss.

The third term is for similarity learning. Hard tuples includes an anchor $\mathbf{h}_i \in \mathbf{H}$, a hard positive sample $\mathbf{h}_i^{\text{pos}} \in S_{\text{pos}}(\mathbf{h}_i)$ and a hard negative sample $\mathbf{h}_i^{\text{neg}} \in S_{\text{neg}}(\mathbf{h}_i)$. Both the single $\mathbf{h}_i^{\text{pos}}$ and the single $\mathbf{h}_i^{\text{neg}}$ are randomly selected from corresponding sets. Then we calculate the embeddings of the tuple: $\mathbf{e}_i = f_{\text{mom}}(\mathbf{h}_i; \theta_m)$, $\mathbf{e}_i^{\text{pos}} = f_{\text{mom}}(\mathbf{h}_i^{\text{pos}}; \theta_m)$ and $\mathbf{e}_i^{\text{neg}} = f_{\text{mom}}(\mathbf{h}_i^{\text{neg}}; \theta_m)$, so that we can measure the similarities of the embeddings via the weighted triplet loss:

$$\mathcal{L}_m = \frac{1}{N} \sum_{i=1}^N w^m(\mathbf{e}_i, \mathbf{e}_i^{\text{pos}}) [m + \|\mathbf{e}_i - \mathbf{e}_i^{\text{pos}}\|_2^2 - \|\mathbf{e}_i - \mathbf{e}_i^{\text{neg}}\|_2^2]_+, \quad (8)$$

where $w^m(\mathbf{e}_i, \mathbf{e}_i^{\text{pos}})$ represents the degree of contribution of every tuple. Weighting the standard triplet loss by the similarity between the anchor and the positive on manifolds relieves the pressure from the tuples with too hard positives.

It is worth mentioning that \mathcal{L}_{mv} in Eq (4) is optimized by mini-batch training so that the method can scale up to large multi-view datasets.

Hierarchical Clustering Loss

A “good” hierarchical tree means that more similar data instances should be merged earlier. Dasgupta [2016] first proposed an explicit hierarchical clustering loss function:

$$\mathcal{L}_{\text{Dasgupta}}(T; w_{ij}) = \sum_{ij} w_{ij}^e |\text{leaves}(T[i \vee j])|, \quad (9)$$

where w_{ij}^e is the Euclidean similarity between i and j and $T[i \vee j]$ is the subtree rooted at the LCA of the i and j nodes, and $\text{leaves}(T[i \vee j])$ denotes the set of descendant leaves of internal node $T[i \vee j]$.

Here, we adopt the differentiable relaxation of constrained Dasguta’s objective [Chami *et al.*, 2020]. Our pairwise similarity graph is learned from the multi-view representation learning process so that the unified similarity measurement narrows the gap between representation learning and clustering. The hyperbolic hierarchical clustering loss for our model

is defined as:

$$\mathcal{L}_{\text{hc}}(\mathbf{Y}; w^e, \tau_c) = \sum_{i,j,k} (w_{i,j}^e + w_{i,k}^e + w_{j,k}^e - w_{i,j,k}^{\text{hyp}}(\mathbf{Y}; w, \tau_c) + \sum_{i,j} (w_{i,j}^e)). \quad (10)$$

We compute the hierarchy of any embedding triplet through the similarities of all embedding pairs among three samples:

$$w_{i,j,k}^{\text{hyp}}(\mathbf{Y}; w^e, \tau_c) = (w_{i,j}^e, w_{i,k}^e, w_{j,k}^e) \cdot \sigma_{\tau_c}(d_{\mathbb{B}}(o, \mathbf{y}_i \vee \mathbf{y}_j), d_{\mathbb{B}}(o, \mathbf{y}_i \vee \mathbf{y}_k), d_{\mathbb{B}}(o, \mathbf{y}_j \vee \mathbf{y}_k))^T, \quad (11)$$

where $\sigma_{\tau_c}(\cdot)$ is the softmax function $\sigma_{\tau_c}(d)_i = e^{d_i/\tau_c} / \sum_j e^{d_j/\tau_c}$ scaled by the temperature parameter τ_c for the hyperbolic hierarchical clustering loss. Eq. (10) theoretically asks for similarities among all tuples of the dataset, which takes a high time complexity of $O(N^3)$. However, we could sample N^2 -order triplets by obtaining all possible node pairs and then choosing the third node randomly from the rest [Chami *et al.*, 2020]. With mini-batch training and sampling tricks, HC can scale to large datasets with acceptable computation complexity.

Then the optimal hyperbolic embeddings are denoted as:

$$\mathbf{Y}^* = \arg \min_{\mathbf{Y}} \mathcal{L}_{\text{hc}}(\mathbf{Y}; w^e, \tau_c). \quad (12)$$

Finally, the nature of negative curvature and bent tree-like geodesics in hyperbolic space allows us to decode the binary tree T in the original space by grouping two similar embeddings whose hyperbolic LCA is the farthest to the origin:

$$T = \text{dec}(\mathbf{Y}^*), \quad (13)$$

where $\text{dec}(\cdot)$ is the decoding function [Chami *et al.*, 2020].

Optimization Process

The entire training process of our framework has two steps: (A) Optimizing the multi-view representation learning loss with Eq. (4), and (B) Optimizing the hyperbolic hierarchical clustering loss with Eq. (10).

4 Experiments

4.1 Experimental Setup

Datasets

We conduct our experiments on the following five real-world multi-view datasets.

- MNIST-USPS [Peng *et al.*, 2019] is a two-view dataset with 5000 hand-written digital (0-9) images. The MNIST view is in 28×28 size, and the USPS view is in 16×16 size.
- BDGP [Li *et al.*, 2019b] contains 2500 images of *Drosophila* embryos divided into five categories with two extracted features. One view is 1750-dim visual features, and the other view is 79-dim textual features.
- Caltech101-7 [Dueck and Frey, 2007] is established with 5 diverse feature descriptors, including 40-dim wavelet moments (WM), 254-dim CENTRIST, 1,984-dim HOG, 512-dim GIST, and 928-dim LBP features, with 1400 RGB images sampled from 7 categories.

- COIL-20 contains object images of 20 categories. Following Trosten *et al.* [2021], we establish a variant of COIL-20, where 480 grayscale images of 128×128 pixel size are depicted from 3 different random angles.
- Multi-Fashion [Xu *et al.*, 2022] is a three-view dataset with 10,000 28×28 images of different fashionable designs, where different views of each sample are different products from the same category.

Baseline Methods

We demonstrate the effects of our CMHHC by comparing with the following three kinds of hierarchical clustering methods. Note that for single-view methods, we concatenate all views into a single view to provide complete information [Peng *et al.*, 2019].

Firstly, we compare CMHHC with conventional linkage-based discrete single-view hierarchical agglomerative clustering (HAC) methods, including Single-linkage, Complete-linkage, Average-linkage, and Ward-linkage algorithms. Secondly, we compare CMHHC with the most representative similarity-based continuous single-view hierarchical clustering methods, i.e., UFit [Chierchia and Perret, 2020], and HypHC [Chami *et al.*, 2020]. For UFit, we adopt the best loss function “Closest+Size”. As for HypHC, we directly use the continuous Dasgupta’s objective. We utilized the corresponding open-source versions for both of the methods and followed the default parameters in the codes provided by the authors. Lastly, to the best of our knowledge, MHC [Zheng *et al.*, 2020] is the only proposed multi-view hierarchical clustering method. We implemented it with Python as there was no open-source implementation available.

Hierarchical Clustering Metrics

Unlike partitional clustering, binary clustering trees, i.e., the results of hierarchical clustering, can provide diverse-granularity cluster information when the trees are truncated at different altitudes. Hence, general clustering metrics, such as clustering accuracy (ACC) and Normalized Mutual Information (NMI), is not able to display the characteristics of clustering hierarchies comprehensively. To this end, following previous literature [Kobren *et al.*, 2017; Monath *et al.*, 2019], we validate hierarchical clustering performance via the Dendrogram Purity (DP) measurement. To sum up, DP measurement is equivalent to the average purity over the leaf nodes of LCA of all available data point pairs in the same ground-truth clusters. Clustering tree with higher DP value contains purer subtrees and keeps a more consistent structure with ground-truth flat partitions.

Implementation Details

Our entire CMHHC model is implemented with PyTorch. We first pretrain V autoencoders for 200 epochs and contrastive learning encoder for 10, 50, 50, 50 and 100 epochs on BDGP, MNIST-USPS, Caltech101-7, COIL-20, and Multi-Fashion respectively, and then finetune the whole multi-view representation learning process for 50 epochs, and finally train the hyperbolic hierarchical clustering loss for 50 epochs. The batch size is set to 256 for representation learning and 512 for hierarchical clustering, using the Adam and hyperbolic-matched Riemannian optimizer [Kochurov *et al.*, 2020] re-

| Method | MNIST-USPS | BDGP | Caltech101-7 | COIL-20 | Multi-Fashion |
|------------------------|---------------|---------------|---------------|---------------|---------------|
| HAC (Single-linkage) | 29.81% | 61.88% | 23.67% | 72.56% | 27.89% |
| HAC (Complete-linkage) | 54.36% | 56.57% | 30.19% | 69.95% | 48.72% |
| HAC (Average-linkage) | 69.67% | 45.91% | 30.90% | 73.14% | 65.70% |
| HAC (Ward-linkage) | 80.38% | 58.61% | 35.69% | 80.81% | 72.33% |
| UFit | 21.67% | 69.20% | 19.00% | 55.41% | 25.94% |
| HyperHC | 32.99% | 31.21% | 22.46% | 28.50% | 25.65% |
| MHC | 78.27% | 89.14% | 45.22% | 66.50% | 54.81% |
| CMHHC (Ours) | 94.49% | 91.53% | 66.52% | 84.89% | 96.25% |

Table 1: Dendrogram Purity (DP) results of baselines and CMHHC. We concatenate the features of multi-view data for single-view baselines, including HAC, UFit, and HyperHC.

spectively. The learning rate is set to $5e^{-4}$ for Adam, and a search over $[5e^{-4}, 1e^{-3}]$ for Riemannian Adam of different datasets. We empirically set $\tau = 0.5$ for all datasets, while $\tau_c = 5e^{-2}$ for BDGP, MNIST-USPS, and Multi-Fashion and $\tau_c = 1e^{-1}$ for Caltech101-7 and COIL-20. We run the model 5 times and report the results with the lowest value of L_c . In addition, we create an adjacency graph with 50 Euclidean nearest neighbors to compute manifold similarities. We make the general rule that the k_{pos} value equals $N/2K$ and the k_{neg} value equals N/K , making the selected tuples hard and reliable.

4.2 Experimental Results

Performance Comparison with Baselines

Experimental DP results are reported in Table 1. The results illustrate that our unified CMHHC model outperforms comparing baseline methods. As it shows, our model gains a significant growth in DP by 11.47% on BDGP, 1.61% on MNIST-USPS, 20.00% on Caltech101-7, 4.08% on COIL-20 and 23.92% on Multi-Fashion over the second-best method. The underlying reason is that CMHHC captures much more meaningful multi-view aligned embeddings instead of concatenating all views roughly without making full use of the complementary and consensus information. Our deep model greatly exceeds the level of the only multi-view hierarchical clustering work MHC, especially on Caltech101-7, COIL-20, and large-scale Multi-Fashion. This result can be attributed to the alignment and discrimination of the multi-view similarity graph learning for hyperbolic hierarchical clustering. Additionally, the performance gap between our model and deep continuous UFit and HypHC reflects the limitations of fixing input graphs without an effective similarity learning process.

Ablation Study

We conduct an ablation study to evaluate the effects of the components in the multi-view representation learning module. To be specific, we refer to the CMHHC without autoencoders for reconstruction of multiple views as CMHHC_{-AE}, without contrastive learning submodel as CMHHC_{-Con}, and without similarity learning as CMHHC_{-Sim}. We train CMHHC_{-AE}, CMHHC_{-Con} and CMHHC_{-Sim} after removing the corresponding network layers. Table 2 shows the experimental results on 5 datasets. The results show that the proposed components contribute to the final hierarchical clustering performance in almost all cases.

| Ablation | MNIST-USPS | BDGP | Caltech101-7 | COIL-20 | Multi-Fashion |
|-----------------------|---------------|---------------|---------------|---------------|---------------|
| CMHHC | 94.49% | 91.53% | 66.52% | 84.89% | 96.25% |
| CMHHC _{-AE} | 92.92% | 86.19% | 68.50% | 27.16% | 89.06% |
| CMHHC _{-Con} | 43.10% | 24.51% | 33.02% | 14.73% | 43.57% |
| CMHHC _{-Sim} | 89.78% | 90.10% | 19.74% | 53.40% | 44.65% |

Table 2: Ablation study results.




Figure 2: Visualization of a truncated subtree from the decoded tree on the MNIST-USPS dataset. The right part represents the sampled subtree structure of LCA #9986, where MNIST images are framed in blue, and USPS images are framed in orange. We observe that more similar pairs will have lower LCAs. For example, images belonging to the same category (like #3497 and #4354 of digit 5) are grouped together first, i.e., have the lowest LCA, while less similar images (like #3497 of digit 5 and #140 of digit 7) are merged at the highest LCA in the subtree.

Case Study

We qualitatively evaluate a truncated subtree structure learned via our method for the hierarchical tree. We plot the sampled MNIST-USPS subtrees of the final clustering tree in Figure 2. As shown, the similarity between two nodes is getting more substantial from the root to the leaves, indicating that the hierarchical tree can reveal fine-grained similarity relationships and ground-truth flat partitions for the multi-view data.

5 Conclusion

This paper proposed a novel multi-view hierarchical clustering framework based on deep neural networks. Employing multiple autoencoders, contrastive multi-view alignment learning, and unsupervised similarity learning, we capture the invariance information across views and learn the meaningful metric property for similarity-based continuous hierarchical clustering. Our method aims at providing a clustering tree with high interpretability oriented towards multi-view data, highlighting the importance of representations' alignment and discrimination, and indicating the potential of gradient-based hyperbolic hierarchical clustering. Extensive experiments illustrate CMHHC is capable of clustering multi-view data at diverse levels of granularity.

Acknowledgments

This work was partially supported by the National Key Research and Development Program of China (No. 2018AAA0100204), and a key program of fundamental research from Shenzhen Science and Technology Innovation Commission (No. JCYJ20200109113403826).

References

- [Andrew *et al.*, 2013] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *ICML*, pages 1247–1255, 2013.
- [Cai *et al.*, 2013] Xiao Cai, Feiping Nie, and Heng Huang. Multi-view k-means clustering on big data. In *IJCAI*, pages 2598–2604, 2013.
- [Chami *et al.*, 2020] Ines Chami, Albert Gu, Vaggos Chatzifratris, and Christopher Ré. From trees to continuous embeddings and back: Hyperbolic hierarchical clustering. In *NeurIPS*, pages 15065–15076, 2020.
- [Chaudhuri *et al.*, 2009] Kamalika Chaudhuri, Sham M Kakade, Karen Livescu, and Karthik Sridharan. Multi-view clustering via canonical correlation analysis. In *ICML*, pages 129–136, 2009.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.
- [Chierchia and Perret, 2020] Giovanni Chierchia and Benjamin Perret. Ultrametric fitting by gradient descent. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124004, 2020.
- [Dasgupta, 2016] Sanjoy Dasgupta. A cost function for similarity-based hierarchical clustering. In *STOC*, pages 118–127, 2016.
- [Dueck and Frey, 2007] Delbert Dueck and Brendan J Frey. Non-metric affinity propagation for unsupervised image categorization. In *ICCV*, pages 1–8, 2007.
- [Hinton and Salakhutdinov, 2006] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [Iscen *et al.*, 2017] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *CVPR*, pages 2077–2086, 2017.
- [Iscen *et al.*, 2018] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Mining on manifolds: Metric learning without labels. In *CVPR*, pages 7642–7651, 2018.
- [Jain *et al.*, 1999] Anil K Jain, M Narasimha Murty, and Patrick J Flynn. Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323, 1999.
- [Kang *et al.*, 2020] Zhao Kang, Guoxin Shi, Shudong Huang, Wenyu Chen, Xiaorong Pu, Joey Tianyi Zhou, and Zenglin Xu. Multi-graph fusion for multi-view spectral clustering. *Knowledge-Based Systems*, 189:105102, 2020.
- [Kobren *et al.*, 2017] Ari Kobren, Nicholas Monath, Akshay Krishnamurthy, and Andrew McCallum. A hierarchical algorithm for extreme clustering. In *SIGKDD*, pages 255–264, 2017.
- [Kochurov *et al.*, 2020] Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in PyTorch. *arXiv preprint arXiv:2005.02819*, 2020.
- [Li *et al.*, 2018] Yingming Li, Ming Yang, and Zhongfei Zhang. A survey of multi-view representation learning. *TKDE*, 31(10):1863–1883, 2018.
- [Li *et al.*, 2019a] Ruihuang Li, Changqing Zhang, Huazhu Fu, Xi Peng, Tianyi Zhou, and Qinghua Hu. Reciprocal multi-layer subspace learning for multi-view clustering. In *ICCV*, pages 8172–8180, 2019.
- [Li *et al.*, 2019b] Zhaoyang Li, Qianqian Wang, Zhiqiang Tao, Quanxue Gao, and Zhaohua Yang. Deep adversarial multi-view clustering network. In *IJCAI*, pages 2952–2958, 2019.
- [Monath *et al.*, 2019] Nicholas Monath, Manzil Zaheer, Daniel Silva, Andrew McCallum, and Amr Ahmed. Gradient-based hierarchical clustering using continuous representations of trees in hyperbolic space. In *SIGKDD*, pages 714–722, 2019.
- [Nickel and Kiela, 2017] Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *NeurIPS*, pages 6338–6347, 2017.
- [Peng *et al.*, 2019] Xi Peng, Zhenyu Huang, Jiancheng Lv, Hongyuan Zhu, and Joey Tianyi Zhou. Comic: Multi-view clustering without parameter selection. In *ICML*, pages 5092–5101, 2019.
- [Sala *et al.*, 2018] Frederic Sala, Chris De Sa, Albert Gu, and Christopher Ré. Representation tradeoffs for hyperbolic embeddings. In *ICML*, pages 4460–4469, 2018.
- [Trosten *et al.*, 2021] Daniel J Trosten, Sigurd Lokse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering representation alignment for multi-view clustering. In *CVPR*, pages 1255–1265, 2021.
- [Wang *et al.*, 2015] Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In *ICML*, pages 1083–1092, 2015.
- [Xu *et al.*, 2021] Jie Xu, Yazhou Ren, Huayi Tang, Xiaorong Pu, Xiaofeng Zhu, Ming Zeng, and Lifang He. Multi-VAE: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering. In *ICCV*, pages 9234–9243, 2021.
- [Xu *et al.*, 2022] Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiaofeng Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. In *CVPR*, 2022.
- [Yan *et al.*, 2021] Jiexi Yan, Lei Luo, Cheng Deng, and Heng Huang. Unsupervised hyperbolic metric learning. In *CVPR*, pages 12465–12474, 2021.
- [Zheng *et al.*, 2020] Qinghai Zheng, Jihua Zhu, and Shuangxun Ma. Multi-view hierarchical clustering. *arXiv preprint arXiv:2010.07573*, 2020.
- [Zhou and Shen, 2020] Runwu Zhou and Yi-Dong Shen. End-to-end adversarial-attention network for multi-modal clustering. In *CVPR*, pages 14619–14628, 2020.