

Hyperbolic-Constraint Point Cloud Reconstruction from Single RGB-D Images

Wenrui Li¹, Zhe Yang², Wei Han¹, Hengyu Man¹, Xingtao Wang^{13*} and Xiaopeng Fan¹³⁴

¹Harbin Institute of Technology

²University of Electronic Science and Technology of China

³Harbin Institute of Technology Suzhou Research Institute

⁴Peng Cheng Laboratory

liwr618@163.com; xiazhe200006@gmail.com; 2021111641@stu.hit.edu.cn;

manhengyu@hotmail.com; xtwang@hit.edu.cn; fxp@hit.edu.cn

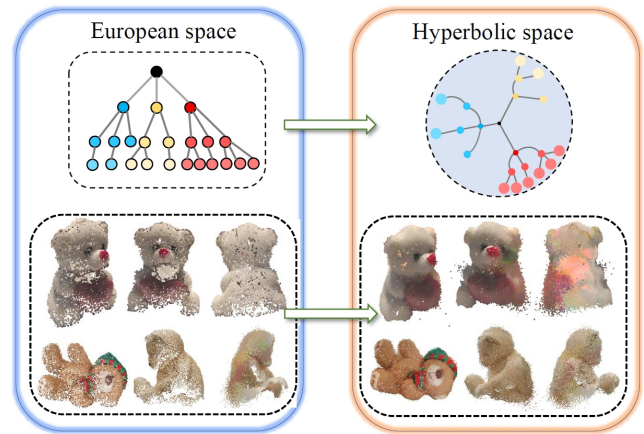
Abstract

Reconstructing desired objects and scenes has long been a primary goal in 3D computer vision. Single-view point cloud reconstruction has become a popular technique due to its low cost and accurate results. However, single-view reconstruction methods often rely on expensive CAD models and complex geometric priors. Effectively utilizing prior knowledge about the data remains a challenge. In this paper, we introduce hyperbolic space to 3D point cloud reconstruction, enabling the model to represent and understand complex hierarchical structures in point clouds with low distortion. We build upon previous methods by proposing a hyperbolic Chamfer distance and a regularized triplet loss to enhance the relationship between partial and complete point clouds. Additionally, we design adaptive boundary conditions to improve the model’s understanding and reconstruction of 3D structures. Our model outperforms most existing models, and ablation studies demonstrate the significance of our model and its components. Experimental results show that our method significantly improves feature extraction capabilities. Our model achieves outstanding performance in 3D reconstruction tasks.

Introduction

Point cloud reconstruction is now a leading technique in 3D computer vision, valued for its precise positional information and detailed reconstruction capabilities. Various research efforts have achieved significant progress in multiple areas, including text, images, and videos (Chen et al. 2023; Cheng et al. 2023; Chou, Bahat, and Heide 2023). The affordability and convenience of single-view reconstruction make it a valuable tool across many domains. Advanced algorithms and computational power enable the reconstruction of a 3D model from a single image. This method is especially useful in situations where obtaining multi-angle images is challenging. However, single-view point cloud reconstruction methods frequently depend on costly CAD models (Sun et al. 2018) and specific priors (Chen et al. 2021), which limit their generalization and scalability.

To address these limitations, researchers are continuously exploring new methods and architectures to improve the effectiveness and application range of single-view point cloud



Visualization Comparison

Figure 1: This figure illustrates the efficient embedding of a tree structure from Euclidean space into hyperbolic space. The left side depicts a multi-level tree structure in Euclidean space, while the right side shows the projection of hyperbolic space. Nodes of different colors represent various categories. In hyperbolic space, the node distribution more accurately reflects their hierarchical relationships.

reconstruction. MCC (Wu et al. 2023) developed a simple yet effective large-scale learning framework utilizing vanilla transformers. They directly trained the positions of 3D points, extending the reconstruction targets to various objects and scenes. NU-MCC (Lionar et al. 2023) enhanced MCC’s approach by introducing a lightweight transformer variant to reduce computational costs and employing a flexible hybrid representation to improve geometric and texture details. Notably, both works reproject the RGB-D image back to visible points and use some of these points as query points to encode image features through transformers. This innovative method uses partial points to reflect the whole. However, previous methods failed to capture the relationship between local and overall structures, as they could not represent the hierarchical, tree-like nature of the objects forming the point clouds (Sarkar 2012a).

Point cloud data exhibits a rich hierarchical structure, where simple geometric shapes such as disks, squares, and

*Corresponding author.

triangles serve as fundamental building blocks for more complex forms. These basic shapes are regarded as "universal ancestors" and can be found across various objects. As illustrated in Fig. 1, this hierarchical structure can be effectively represented as a tree. The lower levels of the tree represent commonly used basic geometric shapes, while the leaves correspond to the specific combinations of these shapes in particular objects. For instance, a simple table comprises legs, a tabletop, and a support structure. The legs connect to and support the tabletop via the support structure, forming a complete table. Essentially, this hierarchy takes the form of a tree, where the legs and support structure serve as the root nodes, with the tabletop as the leaf node. Clearly, understanding this internal hierarchical structure is essential for accurately capturing the 3D object. However, when dealing with data featuring complex structures such as hierarchies or tree-like forms, Euclidean space encounters challenges (Fang et al. 2023), (Chami et al. 2021), (Guo, Guo, and Yu 2022), (Hong et al. 2023). Embedding such data in Euclidean space can result in significant distortions (Fang et al. 2023). In hyperbolic space, distance calculations better preserve the intrinsic structure of the data, more accurately representing the true relationships between data points. This limitation arises because the volume in Euclidean space grows polynomially with its radius, rather than exponentially, thereby restricting its capacity to represent tree-like data with an exponential number of leaves (Montanaro, Valsesia, and Magli 2022).

Hyperbolic space offers a natural solution to this issue. We redefined the hyperbolic chamfer distance in hyperbolic space and introduced a regularization loss and a triplet loss to emphasize the relationship between local point clouds and the overall point cloud. Local features can be embedded in lower-dimensional regions of hyperbolic space, while global features occupy higher-dimensional regions. This approach maximizes mutual information between parts and the whole, thereby enhancing the effectiveness of feature extraction. The main contributions can be summarised as follows:

- We introduce hyperbolic space into 3D point cloud reconstruction for the first time, facilitating the representation and understanding of complex hierarchical structures within 3D objects with minimal distortion. This method more accurately captures the dynamic relationship between parts and the whole, offering new perspectives and techniques for 3D reconstruction.
- We propose the hyperbolic Chamfer distance and regularized triplet loss, which effectively enhance the relationship between partial and complete point clouds. Additionally, we design adaptive boundary conditions to enhance the model's reconstruction of 3D structures.
- Our model demonstrates the superiorities among current mainstream models. Ablation studies highlight the significance of our model and its individual components.

Related Work

Single-View Point Cloud Reconstruction

A single image provides limited information and is subject to occlusions and viewpoint constraints. Single-view 3D re-

construction is a complex and challenging research area. Traditionally, this task relies on rich supervision information, such as using CAD models (Sun et al. 2018), point clouds (Mescheder et al. 2019) (Henrich et al. 2024), or meshes as training targets to predict 3D shapes from a single image. These methods have shown impressive reconstruction results on simple synthetic datasets like ShapeNet (Chang et al. 2015) but are usually limited to specific categories of objects, with limited applicability and generality of the reconstruction. With the advancement of technology, researchers have started exploring weakly supervised and category-agnostic 3D reconstruction methods, aiming to learn 3D representations from a broader and more diverse range of real-world data (Reizenstein et al. 2021) (Yu et al. 2023). This includes using RGB-D images or multi-view images to train models, allowing them to capture richer geometric details and visual features. For example, the MCC (Wu et al. 2023) project learns compact 3D appearance and geometry from large-scale multi-view data, demonstrating generalization capabilities in natural environments. Additionally, the recent NU-MCC (Lionar et al. 2023) achieves more efficient 3D reconstruction by introducing the Repulsive Unsigned Distance Field and a neighborhood decoder.

Local and Global Point Cloud Processing

In the field of point cloud data processing and analysis, effectively integrating local and global information is crucial for understanding and utilizing 3D geometric data (Rao, Lu, and Zhou 2020). Early methods, such as PointNet (Qi et al. 2017a), modeled directly on raw 3D coordinate sets, independently processed each point through shared multi-layer perceptrons (MLP), demonstrating the need for point permutation invariance in point cloud processing. However, PointNet failed to fully leverage local structural information of the point cloud, which has been key to the success of convolutional neural networks (CNNs) in 2D image processing. To address this limitation, PointNet++ (Qi et al. 2017b) introduced a hierarchical feature learning strategy, enhancing model performance by progressively extracting local features at different scales.

Subsequent research, such as PointCNN (Li et al. 2018), DGCNN (Phan et al. 2018), and PointConv (Wu, Qi, and Fuxin 2019), treated point clouds as graph structures and explored the importance of local structures by processing points and their neighborhood relationships within the graph convolution framework. These methods effectively captured local relationships between points through dynamically computed k-nearest neighbor graphs or other graph structures, thereby improving feature representation and model generalization. In terms of global feature learning, methods like PointGLR (Rao, Lu, and Zhou 2023) have attempted to map local features into a representation space with global properties (e.g., hypersphere), achieving an effective fusion of global and local features.

Hyperbolic Geometry

Hyperbolic space, due to its inherent property of encoding the intrinsic hierarchical structure of data, has become

a significant direction for constructing structured representations. Hyperbolic space is a Riemannian manifold with constant negative curvature, making it particularly suitable for handling data with prominent hierarchical structures, such as tree-like structures and complex networks. Inspired by Sarkar’s work (Sarkar 2012b), which demonstrated that tree structures could be embedded in hyperbolic space with minimal distortion, several studies have since explored how to achieve effective representation learning in this non-Euclidean manifold. In the field of machine learning, especially in natural language processing and graph data analysis (Khurlov et al. 2020; Ying et al. 2018), embeddings in hyperbolic space have been proven to effectively capture the hierarchical structure and complexity of data. For example, hyperbolic neural networks (Ganea, Becigneul, and Hoffmann 2018) and Poincaré embeddings leverage the advantages of hyperbolic geometry to improve model performance and interpretability. Hyperbolic embeddings also show potential in encoding the hierarchical structure of visual data, such as zero-shot classification (Hong et al. 2023) and audio-visual question answering (Yang, Li, and Cheng 2024).

Recently, hyperbolic space has been explored for processing point cloud data (Lin et al. 2023) (Montanaro, Valsesia, and Magli 2022), particularly in expressing the hierarchical attributes of 3D objects. 3D objects naturally exhibit hierarchical structures, with parts of varying sizes forming the entire object, where smaller parts may be shared among different categories of objects, while larger parts are more category-specific. In point cloud object reconstruction, hyperbolic space is more suitable than traditional Euclidean space for handling objects with pronounced hierarchical or branching structures due to its unique negative curvature properties. When reconstructing models with complex branches, hyperbolic space can naturally express the progressive unfolding and branching characteristics of these structures. Additionally, the geometric structure of hyperbolic space supports exponential spatial expansion, allowing for the efficient embedding and representation of highly hierarchical data structures in lower dimensions, avoiding the data sparsity and computational inefficiency issues that arise in Euclidean space with increasing dimensions. This characteristic enables hyperbolic space to precisely capture and represent the intrinsic hierarchy of objects while maintaining computational manageability and efficiency.

Method

The overview of HcPCR are illustrated in Fig. 2. Our task is to reconstruct the complete point cloud of an object from a given RGB-D image. For each RGB-D image, we denote the RGB image as I and the depth map as D . Firstly, we process the image I into 16×16 patches, and then obtain the global features f_w using a pre-trained Vision Transformer (ViT) (Dosovitskiy et al. 2021). For the local features f_p , we perform unprojection of the depth map D into the 3D world to obtain a partial point cloud of size $N \times 3$, where $N = H \times W$. Due to objective factors like sensor uncertainties, the obtained partial point cloud is redundant and noisy. Therefore, we pool it and encode the local features P using ViT. Finally, we employ an anchor point prediction op-

eration similar to NU-MCC to further encode the features (E_N). For specific details about the encoder, please refer to NU-MCC. All losses in NU-MCC are represented as \mathcal{L}_N .

Hyperbolic Network

Hyperbolic space is a curved space with constant negative curvature ($k < 0$). To transfer feature vectors from Euclidean space to hyperbolic space, we need to perform two stages: hyperbolic mapping and tangent space mapping.

1) **Hyperbolic Mapping.** In this paper, we use the Poincaré ball model (Nickel and Kiela 2017) to represent hyperbolic space. We define an n -dimensional Poincaré ball with curvature k as follows:

$$\mathbb{P}_k^n = \{x \in \mathbb{R}^n : \|x\| < 1/|k|\}, \quad (1)$$

where $\|\cdot\|$ denotes the Frobenius norm.

To project a point $x \in \mathbb{R}^n$ from Euclidean space \mathbb{R}^n into an n -dimensional Poincaré ball \mathbb{P}_k^n with curvature k , we perform the following hyperbolic mapping:

$$x_P = \Gamma_{\mathbb{P}}(x) = \begin{cases} x, & \text{if } \|x\| \leq 1/|k| \\ \frac{1-\varepsilon}{|k|} \frac{x}{\|x\|}, & \text{else} \end{cases} \quad (2)$$

where x_P is the projected point in the Poincaré sphere and ε is a scaling factor adjusted according to the curvature k .

2) **Tangent Space Mapping.** Hyperbolic spaces differ from vector spaces in that they do not themselves support operations such as unusual addition. In order to facilitate such operations in hyperbolic spaces, we must employ Möbius addition (Khurlov et al. 2020):

$$z_P \oplus_k x_P = \frac{(1 + 2|k|\langle z_P, x_P \rangle + |k|\|x_P\|^2)z_P}{1 + 2|k|\langle z_P, x_P \rangle + |k|^2\|z_P\|^2\|x_P\|^2} + \frac{(1 - |k|\|z_P\|^2)x_P}{1 + 2|k|\langle z_P, x_P \rangle + |k|^2\|z_P\|^2\|x_P\|^2}, \quad (3)$$

the hyperbolic tangent space is a local Euclidean space that approximates the hyperbolic space. To facilitate the calculation of loss, we project points from hyperbolic space into the hyperbolic tangent space. Given a point $z_P \in \mathbb{P}_k^n$ in the Poincaré sphere, we can generate a tangent plane $T_z \mathbb{P}_k^n$:

$$x_{Tg} = \frac{2}{\sqrt{|k|\lambda_k(z_P)}} \tanh^{-1} \left(\sqrt{|k|} \| -z_P \oplus_k x_P \| \right) \cdot \frac{-z_P \oplus_k x_P}{\| -z_P \oplus_k x_P \|}, \quad (4)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product and $x_{Tg} \in T_z \mathbb{P}_k^n$. To facilitate understanding and analysis, we set $z_P = 0_P$.

Hyperbolic Distance

Point cloud distance is a method used to measure the dissimilarity between two point cloud sets. In practical applications, due to the unordered nature of point clouds, shape-level distance typically relies on statistical analysis of point-to-point distances. Among these, Chamfer Distance (CD) is

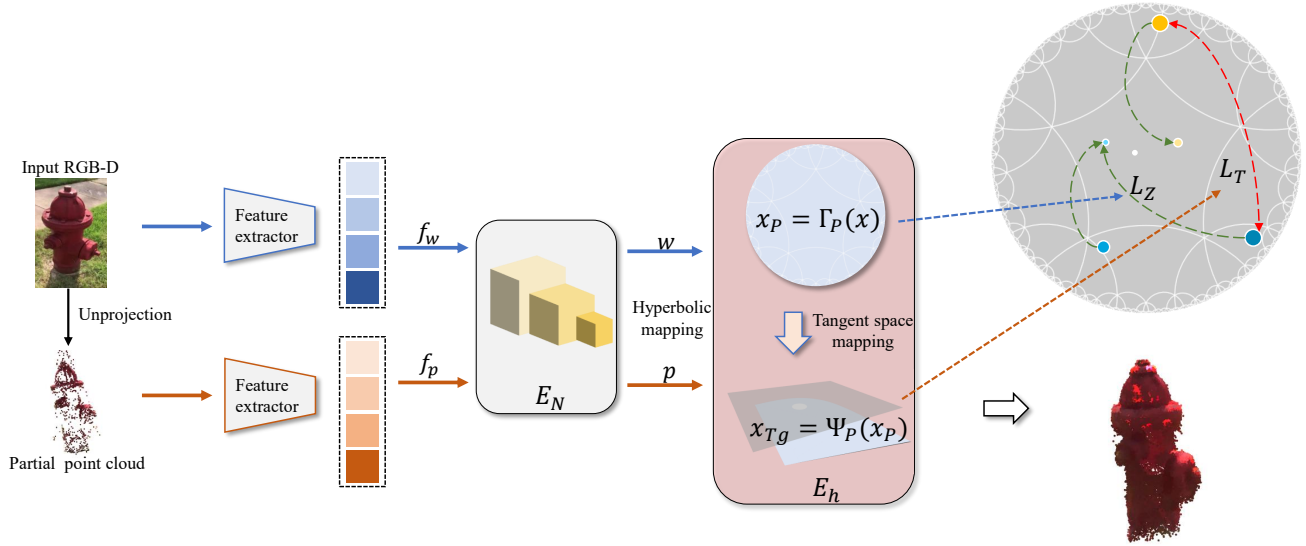


Figure 2: Architecture of the HcPCR. To enhance the understanding of the hierarchical structure of the point clouds, we embed both partial and complete features into a hyperbolic space. The lower-level features (partial point cloud) are embedded into the lower-dimensional central region, while the higher-level features (complete point cloud) are distributed in the higher-dimensional peripheral region. Additionally, a triplet loss is utilized to push apart point clouds of different categories and to pull together point clouds with similar attributes.

widely used due to its relatively efficient computation (Guo et al. 2021) (Wu et al. 2021). CD evaluates the dissimilarity between two point clouds by calculating the average distance from each point in one point cloud to the nearest point in the other point cloud. Generally, the Chamfer distance of point clouds is defined as follows:

$$D(x_i, y_i) = \frac{1}{|x_i|} \sum_j \min_k d(x_{ij}, y_{ik}) + \frac{1}{|y_i|} \sum_k \min_j d(x_{ij}, y_{ik}). \quad (5)$$

In Euclidean space, we typically define the function d as:

$$d(x_{ij}, y_{ik}) = \begin{cases} \|x_{ij} - y_{ik}\| & \text{as L1-CD;} \\ \|x_{ij} - y_{ik}\|^2 & \text{as L2-CD,} \end{cases} \quad (6)$$

where $\|\cdot\|$ denotes the Euclidean ℓ_2 norm of the vector.

Although Chamfer Distance (CD) performs well in computing point cloud distances, the gradient of the CD distance assigns higher weights to points with larger distances, making it sensitive to outliers (Lin et al. 2023). Hyperbolic space offers greater flexibility than Euclidean space in measuring distances between points, potentially correcting matching errors in CD. Therefore, we introduce a new distance metric, the Hyperbolic Chamfer Distance (HyperCD):

$$d(\mathbf{x}, \mathbf{y}) = \frac{2}{\sqrt{k}} \operatorname{arctanh} \left(\sqrt{k} \|\mathbf{x} \oplus_k \mathbf{y}\| \right). \quad (7)$$

It is worth mentioning that when $k \rightarrow 0$, the hyperbolic distance can be simply expressed as the Euclidean distance:

$$\lim_{k \rightarrow 0} d(\mathbf{x}, \mathbf{y}) = 2\|\mathbf{x} - \mathbf{y}\|. \quad (8)$$

Hyperbolic Point Cloud Loss

During the training of HcPCR, we defined a loss function \mathcal{L} that consists of three different components: the NU-MCC loss \mathcal{L}_N , the regularization loss \mathcal{L}_Z , and the triplet loss \mathcal{L}_T :

$$\mathcal{L} = \mathcal{L}_N + \mathcal{L}_Z + \mathcal{L}_T. \quad (9)$$

Regularisation Loss. As introduced in the previous section, we need to project w and p from Euclidean space to hyperbolic space. We use a hyperbolic mapping (Eq. (2)) to map Euclidean feature vectors into hyperbolic space, and then obtain the hyperbolic embeddings of the entire point cloud W through tangent space mapping (Eq. (4)). For the partial point cloud P , we repeat the same process.

To achieve part-whole hierarchy, we define the following regularization loss:

$$\mathcal{L}_Z = \max \left(0, -\Gamma_{\mathbb{P}}(W^+) + \Gamma_{\mathbb{P}}(P^+) + \frac{\gamma}{N} \right), \quad (10)$$

where W^+ and P^+ are the hyperbolic representations of the whole and part of the same point cloud.

The regularizer \mathcal{L}_Z in Eq. (10) introduces a part-whole hierarchy by encouraging partial embeddings to be closer to the center of the Poincaré ball, while the overall embeddings are closer to the edge. Inspired by (Montanaro, Valsesia, and Magli 2022), we employed a variable margin γ/N that depends on the number of points N in a part P . This means that simpler shapes, composed of fewer points, will be farther from the representation of the whole object and have a lower hyperbolic norm (i.e., closer to the center). In contrast,

Method	Distance	Acc↓	Com↓	Pre↑	Recall↑	L_1 ↓	F1↑
MCC	Euclidean	0.172	0.144	0.689	0.727	0.316	0.739
NU-MCC(base)	Euclidean	0.121	0.146	0.792	0.840	0.267	0.802
baseline+ L_Z	Euclidean	0.149	0.152	0.763	0.813	0.301	0.786
baseline+ L_T	Euclidean	0.119	0.134	0.817	0.835	0.253	0.814
baseline+ L_Z+L_T	Euclidean	0.127	0.139	0.728	0.853	0.266	0.806
baseline+ L_Z	Hyperbolic	0.112	0.108	0.827	0.853	0.220	0.836
baseline+ L_T	Hyperbolic	0.107	0.101	0.832	0.868	0.208	0.853
baseline+ L_Z+L_T (ours)	Hyperbolic	0.104	0.094	0.848	0.884	0.198	0.861

Table 1: The performance comparison on CO3D-v2 datasets.

embeddings of larger parts will gradually approach the edge of the Poincaré ball based on their size.

Relying on a fixed margin defined by hyperparameters might not effectively handle the complex structure of point cloud data, thereby affecting the quality of overall and partial feature representations. To address this issue, we introduced adaptive margins by incorporating whole-part features into the hyperbolic manifold. Specifically, we first concatenate the features p and w into a vector l , and then replace the margin with a linear layer (MLP). The detailed method is as follows:

$$\gamma = \gamma_0 \cdot \text{sigmoid}(\text{MLP}(l)). \quad (11)$$

where $\text{MLP}(\cdot)$ denotes a fully connected layer, and γ_0 represents the initial margin.

Triplet Loss. To ensure that parts from different classes maintain a larger geodesic distance, we aim to pull parts and wholes from the same class closer together while mapping parts from different classes further apart. We define the following triplet loss:

$$\mathcal{L}_T = \max(0, d(\Psi_{\mathbb{P}}(\mathcal{W}^+), \Psi_{\mathbb{P}}(\mathcal{P}^+)) - d(\Psi_{\mathbb{P}}(\mathcal{W}^+), \Psi_{\mathbb{P}}(\mathcal{P}^-)) + \varepsilon) \quad (12)$$

where ε is a hyperparameter used to control the degree of separation between positive and negative samples, while \mathcal{P}^- is the embedding of a part from a different category.

The triplet loss \mathcal{L}_T in Eq. (12) encourages the correct clustering of objects and parts in hyperbolic space. Specifically, parts and wholes from the same class are brought closer together, while parts from different classes are mapped farther apart. This ensures that parts from different classes maintain greater geodesic distances. δ is a hyperparameter that controls the separation degree between positive and negative samples.

Experiments

Dataset

We conducted object reconstruction experiments primarily on the CO3D-v2 dataset (Reizenstein et al. 2021). To evaluate the performance of our HcPCR model, we compared it with the latest baseline models considered the best in this

field. For consistency, we followed the dataset setup as described in (Wu et al. 2023).

CO3D-v2 is a dataset used in computer vision research for 3D object recognition and reconstruction. It is an extended version of the original CO3D dataset, featuring more videos and object categories. The dataset consists of videos captured mainly in natural environments, providing multi-angle views. It includes approximately 37,000 short videos covering 51 object categories. The videos include object segmentation masks, and 3D point clouds and depth maps automatically extracted using the COLMAP software. We followed the NU-MCC setup, using 10 categories for evaluation and the remaining 41 for training.

Evaluation Metrics

To fair comparison, we adopted the evaluation method proposed by (Wu et al. 2023; Lionar et al. 2023; Murez et al. 2020). The evaluation metrics include the L1 distances for accuracy (Acc) and completeness (Comp), their sum as the Chamfer distance (CD), precision (Prec), recall, and the F-score, which is the harmonic mean of precision and recall. For detailed formulas refer to (Murez et al. 2020).

Implementation Details

We adopted the feature extraction method described in NU-MCC (Lionar et al. 2023). The input image resolution is 224×224. We used ViT to encode P and I , generating 768-dimensional feature vectors. Additionally, we followed the data augmentation strategy described in [MCC], applying random scaling and random rotation along each axis. The HcPCR model was trained on an Nvidia A100 GPU. For HcPCR, we set the curvature k to -0.14, α to 2.0, γ_0 to 1000, and ε to 4. The initial learning rate was set to 0.0001. We used the Adam optimizer and trained for 100 epochs.

Results Comparison

We compared the HcPCR model with baseline models, including MCC and NU-MCC, on the CO3D-V2 dataset. Table 1 demonstrates that the HcPCR model outperforms all baseline models across every aspect. Specifically, the HcPCR model shows superior performance in all domains

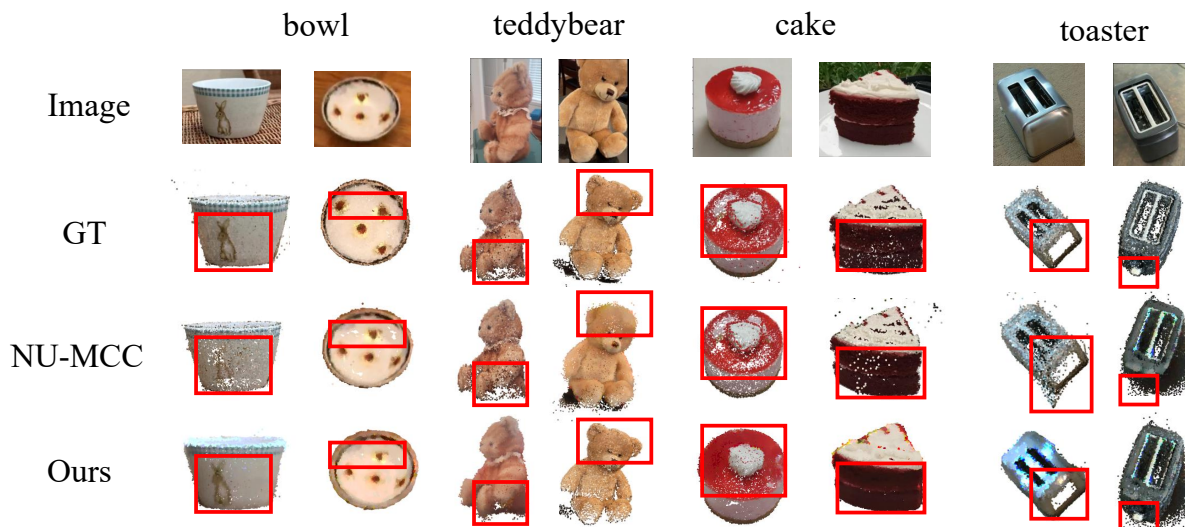


Figure 3: Visualisation comparison on CO3D-v2 validation set. We have selected three categories. Each category contained two samples for evaluation.

compared to NU-MCC. Our method achieves average improvements of 5.9%, 5.2%, and 5.6% in F1 score, completeness, and accuracy, respectively. These results indicate that the HcPCR has superior capability in point cloud reconstruction, attributed to its ability to hierarchically explore 3D structures by leveraging the properties of hyperbolic space.

To investigate the necessity of regularization loss and triplet loss in hyperbolic space, we trained the NU-MCC model using different loss functions. When using Euclidean distance, adding triplet loss effectively enhances the model’s expressiveness. However, merely adding regularization loss leads to an overall decline in model performance. This occurs because the regularization term was originally designed to enhance repulsion between complete and partial point clouds, aiming for a hierarchical embedding of the point cloud structure. However, since Euclidean space cannot effectively represent the tree-like structure of point clouds, this operation instead distances or even severs the original connection between partial and complete point clouds. This further demonstrates the inherent advantage of hyperbolic space in handling tree-structured data. Conversely, when using hyperbolic distance as the metric, the addition of both regularization and triplet losses significantly improves model performance. As shown in Table ??, the model achieves its best performance when using the full loss function in hyperbolic space.

Ablation Study

Evaluating the Effectiveness of Different Margins γ . To validate the superiority of adaptive edges, we compared different preset margin values with the adaptive margins component. The experimental results are shown in left of Fig. 4. It is evident that using different preset values without the adaptive margins significantly affects model performance. This observation highlights the necessity of the adaptive




Figure 4: Ablation Study of different margin and curvature.

margins component.

Evaluate the Effectiveness of the Model Under Different Curvatures k . We conducted experiments under different curvatures to understand the impact of hyperbolic geometry on the performance of the model with hyperbolic embeddings. The model’s performance was tested under various curvatures, and the results are shown in the right of Fig. 4. The best performance was achieved at a curvature of -0.14. This indicates that this value provides the optimal hyperbolic embedding effect.

Ablation Study for δ -Hyperbolicity. While processing point cloud datasets, we observed that their intrinsic geometric properties often display significant non-Euclidean characteristics due to their inherently complex hierarchical structures. In Table 2, we evaluated the point cloud embeddings extracted by a basic encoder and a fully enhanced hyperbolic network encoder to assess their effectiveness in revealing these hierarchical features. To this end, we employed the δ -hyperbolicity measure to evaluate the extent of hierarchical structures, a method previously shown to effectively detect hierarchical characteristics (Khrukov et al. 2020). The δ values computed in the experiments evaluated the overall similarity of sample features to an ideal tree-like structure.

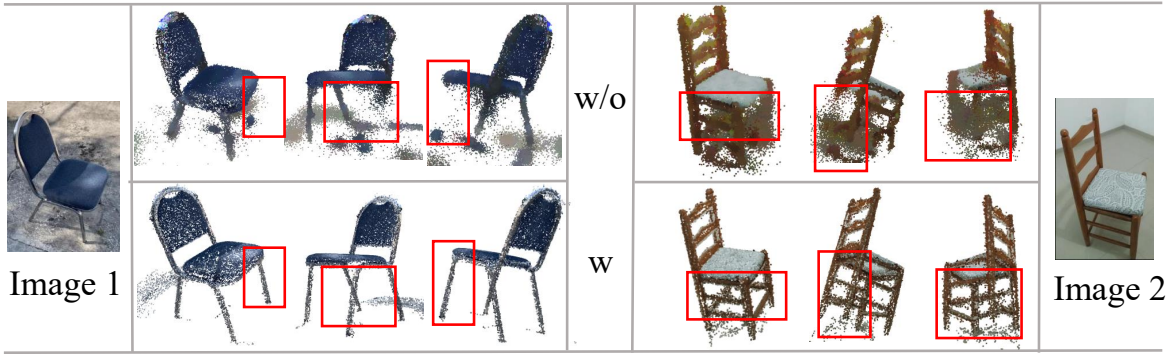


Figure 5: Visualisation results of ablation study with (w) or without (w/o) the full loss function.




Figure 6: The HcPCR embedding is used to represent the CO3D-v2 dataset on the Poincaré disk. Each point within the disk corresponds to a sample, with different colors representing different categories.

Visualization on the CO3D-v2 Dataset. Fig. 3 shows the visual comparison of the reconstructed structures by different models. This figure demonstrates that HcPCR significantly enhances the reconstruction of the overall contours. Compared to the baseline, the geometric appearance of the reconstructed objects is much clearer. This improvement could be attributed to the hyperbolic constraint, which preserves the details of certain point clouds from the original data. Additionally, our reconstructed 3D structures are noticeably more complete.

Visualization for Ablation Study on Loss Function. To further validate the effectiveness of the hyperbolic geometry constraint, we present visual results using the full loss function and using only the initial L_Z loss, in Fig. 5. It is evident that, with the addition of the hyperbolic geometry constraint, the model generates more refined and detailed textures. No-

Encoder	δ
E_N	0.326
$E_N + E_h$	0.294

Table 2: The δ -hyperbolicity value measure the non-Euclidean geometric properties of the data. A lower δ value indicates the data exhibits more pronounced hyperbolicity.

tably, for geometrically prominent parts of the object, the reconstructed shapes are more streamlined and complete. Additionally, our model demonstrates stronger generalization ability. Under the hyperbolic geometry constraint, the reconstruction of occluded object structures (highlighted in the red box) in the images still performs remarkably well.

The Hierarchical Structure of Point Cloud Data in Hyperbolic Space. We used UMAP (McInnes, Healy, and Melville 2020) to project high-dimensional point clouds onto a two-dimensional Poincaré disk, illustrating how hyperbolic space explores the hierarchical structure of point cloud data, as shown in Fig. 6. We visualized 10 categories, each containing 3-4 partial point clouds and one complete point cloud. The number of points in the partial clouds gradually increases, causing these partial objects to move closer to the edge of the disk. As more points are added and the partial objects gradually combine into a complete object, their hyperbolic norm increases, ultimately pushing the entire object towards the disk’s edge. This process highlights the impact of hyperbolic space on hierarchical embedding.

Conclusion

We propose a hyperbolic neural network for reconstructing point cloud objects. To enhance the hierarchical structure of point clouds, we introduce regularization and a triplet loss in hyperbolic space. Additionally, we design an adaptive margins component within the regularization loss to mitigate the heterogeneity caused by fixed margin values. Furthermore, we conduct an ablation study to verify the effectiveness of our model. Finally, we visualize how hierarchical embeddings of tree-like data are achieved in hyperbolic space and compare the reconstruction results against other models.

Acknowledgments

This work was supported in part by the National Key R&D Program of China (2021YFF0900500), the National Natural Science Foundation of China (NSFC) under grants 62441202, U22B2035, 20240222, and the Fundamental Research Funds for the Central Universities under grants HIT.DZJJ.2024025.

References

- Chami, I.; Gu, A.; Nguyen, D. P.; and Re, C. 2021. HoroPCA: Hyperbolic Dimensionality Reduction via Horospherical Projections. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 1419–1429. PMLR.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; Xiao, J.; Yi, L.; and Yu, F. 2015. ShapeNet: An Information-Rich 3D Model Repository. arXiv:1512.03012.
- Chen, H.; Liu, S.; Chen, W.; Li, H.; and Hill, R. 2021. Equivariant Point Network for 3D Point Cloud Analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14514–14523.
- Chen, R.; Chen, Y.; Jiao, N.; and Jia, K. 2023. Fantasia3D: Disentangling Geometry and Appearance for High-quality Text-to-3D Content Creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 22246–22256.
- Chen, Z.; Pu, B.; Zhao, L.; He, J.; and Liang, P. 2025. Divide and augment: Supervised domain adaptation via sample-wise feature fusion. *Information Fusion*, 115: 102757.
- Chen, Z.; Zhang, J.; Lai, Z.; Chen, J.; Liu, Z.; and Li, J. 2022. Geometry-aware guided loss for deep crack recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4703–4712.
- Cheng, Y.-C.; Lee, H.-Y.; Tulyakov, S.; Schwing, A. G.; and Gui, L.-Y. 2023. SDFusion: Multimodal 3D Shape Completion, Reconstruction, and Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4456–4465.
- Chou, G.; Bahat, Y.; and Heide, F. 2023. Diffusion-SDF: Conditional Generative Modeling of Signed Distance Functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2262–2272.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929.
- Fang, P.; Harandi, M.; Le, T.; and Phung, D. 2023. Hyperbolic geometry in computer vision: A survey. *arXiv preprint arXiv:2304.10764*.
- Ganea, O.; Becigneul, G.; and Hofmann, T. 2018. Hyperbolic Neural Networks. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Guo, Y.; Guo, H.; and Yu, S. X. 2022. CO-SNE: Dimensionality Reduction and Visualization for Hyperbolic Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 21–30.
- Guo, Y.; Wang, H.; Hu, Q.; Liu, H.; Liu, L.; and Bennamoun, M. 2021. Deep Learning for 3D Point Clouds: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12): 4338–4364.
- Henrich, P.; Gyenes, B.; Scheikl, P. M.; Neumann, G.; and Mathis-Ullrich, F. 2024. Registered and Segmented Deformable Object Reconstruction From a Single View Point Cloud. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 3129–3138.
- Hong, J.; Hayder, Z.; Han, J.; Fang, P.; Harandi, M.; and Petersson, L. 2023. Hyperbolic Audio-visual Zero-shot Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 7873–7883.
- Khrulkov, V.; Mirvakhobova, L.; Ustinova, E.; Oseledets, I.; and Lempitsky, V. 2020. Hyperbolic Image Embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, W.; Ma, Z.; Deng, L.-J.; Fan, X.; and Tian, Y. 2023a. Neuron-Based Spiking Transmission and Reasoning Network for Robust Image-Text Retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(7): 3516–3528.
- Li, W.; Ma, Z.; Deng, L.-J.; Wang, P.; Shi, J.; and Fan, X. 2023b. Reservoir Computing Transformer for Image-Text Retrieval. In *Proceedings of the 31st ACM International Conference on Multimedia*, MM '23, 5605–5613. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701085.
- Li, W.; Xiong, R.; and Fan, X. 2024. Multi-layer Probabilistic Association Reasoning Network for Image-Text Retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 1–1.
- Li, W.; Zhao, X.-L.; Ma, Z.; Wang, X.; Fan, X.; and Tian, Y. 2023c. Motion-Decoupled Spiking Transformer for Audio-Visual Zero-Shot Learning. MM '23, 3994–4002. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701085.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. PointCNN: Convolution On X-Transformed Points. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Li, Z.; Liao, J.; Tang, C.; Zhang, H.; Li, Y.; Bian, Y.; Sheng, X.; Feng, X.; Li, Y.; Gao, C.; et al. 2024. USTC-TD: A Test Dataset and Benchmark for Image and Video Coding in 2020s. *arXiv preprint arXiv:2409.08481*.
- Lin, F.; Yue, Y.; Hou, S.; Yu, X.; Xu, Y.; Yamada, K. D.; and Zhang, Z. 2023. Hyperbolic Chamfer Distance for Point Cloud Completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 14595–14606.

- Lionar, S.; Xu, X.; Lin, M.; and Lee, G. H. 2023. NU-MCC: Multiview Compressive Coding with Neighborhood Decoder and Repulsive UDF. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 63011–63022. Curran Associates, Inc.
- McInnes, L.; Healy, J.; and Melville, J. 2020. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv:1802.03426.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Montanaro, A.; Valsesia, D.; and Magli, E. 2022. Rethinking the compositionality of point clouds through regularization in the hyperbolic space. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 33741–33753. Curran Associates, Inc.
- Murez, Z.; van As, T.; Bartolozzi, J.; Sinha, A.; Badrinarayanan, V.; and Rabinovich, A. 2020. Atlas: End-to-End 3D Scene Reconstruction from Posed Images. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J.-M., eds., *Computer Vision – ECCV 2020*, 414–431. Cham: Springer International Publishing. ISBN 978-3-030-58571-6.
- Nickel, M.; and Kiela, D. 2017. Poincaré Embeddings for Learning Hierarchical Representations. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Phan, A. V.; Nguyen, M. L.; Nguyen, Y. L. H.; and Bui, L. T. 2018. DGCNN: A convolutional neural network over large-scale labeled graphs. *Neural Networks*, 108: 533–543.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Rao, Y.; Lu, J.; and Zhou, J. 2020. Global-Local Bidirectional Reasoning for Unsupervised Representation Learning of 3D Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rao, Y.; Lu, J.; and Zhou, J. 2023. PointGLR: Unsupervised Structural Representation Learning of 3D Point Clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2193–2207.
- Reizenstein, J.; Shapovalov, R.; Henzler, P.; Sbordone, L.; Labatut, P.; and Novotny, D. 2021. Common Objects in 3D: Large-Scale Learning and Evaluation of Real-Life 3D Category Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10901–10911.
- Sarkar, R. 2012a. Low Distortion Delaunay Embedding of Trees in Hyperbolic Plane. In van Kreveld, M.; and Speckmann, B., eds., *Graph Drawing*, 355–366. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-25878-7.
- Sarkar, R. 2012b. Low Distortion Delaunay Embedding of Trees in Hyperbolic Plane. In van Kreveld, M.; and Speckmann, B., eds., *Graph Drawing*, 355–366. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Sun, X.; Wu, J.; Zhang, X.; Zhang, Z.; Zhang, C.; Xue, T.; Tenenbaum, J. B.; and Freeman, W. T. 2018. Pix3D: Dataset and Methods for Single-Image 3D Shape Modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tang, C.; Sheng, X.; Li, Z.; Zhang, H.; Li, L.; and Liu, D. 2024. Offline and Online Optical Flow Enhancement for Deep Video Compression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5118–5126.
- Wu, C.-Y.; Johnson, J.; Malik, J.; Feichtenhofer, C.; and Gkioxari, G. 2023. Multiview Compressive Coding for 3D Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9065–9075.
- Wu, T.; Pan, L.; Zhang, J.; Wang, T.; Liu, Z.; and Lin, D. 2021. Density-aware Chamfer Distance as a Comprehensive Metric for Point Cloud Completion. arXiv:2111.12702.
- Wu, W.; Qi, Z.; and Fuxin, L. 2019. PointConv: Deep Convolutional Networks on 3D Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yang, Z.; Li, W.; and Cheng, G. 2024. SHMamba: Structured Hyperbolic State Space Model for Audio-Visual Question Answering. arXiv:2406.09833.
- Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. KDD '18. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355520.
- Yu, X.; Xu, M.; Zhang, Y.; Liu, H.; Ye, C.; Wu, Y.; Yan, Z.; Zhu, C.; Xiong, Z.; Liang, T.; Chen, G.; Cui, S.; and Han, X. 2023. MVIImgNet: A Large-Scale Dataset of Multi-View Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9150–9161.