

# Seg-HGNN: Unsupervised and Light-Weight Image Segmentation with Hyperbolic Graph Neural Networks

Debjyoti Mondal  
d.mondal@samsung.com

Rahul Mishra  
mishra.rl@samsung.com

Chandan Pandey  
chandan.p@samsung.com

Samsung R&D Institute India  
Bangalore

---

## Abstract

Image analysis in the euclidean space through linear hyperspaces is well studied. However, in the quest for more effective image representations, we turn to hyperbolic manifolds. They provide a compelling alternative to capture complex hierarchical relationships in images with remarkably small dimensionality. To demonstrate hyperbolic embeddings' competence, we introduce a light-weight hyperbolic graph neural network for image segmentation, encompassing patch-level features in a very small embedding size. Our solution, Seg-HGNN, surpasses the current best unsupervised method by 2.5%, 4% on VOC-07, VOC-12 for localization, and by 0.8%, 1.3% on CUB-200, ECSSD for segmentation, respectively. With less than 7.5k trainable parameters, Seg-HGNN delivers effective and fast ( $\approx 2$  images/second) results on very standard GPUs like the GTX1650. This empirical evaluation presents compelling evidence of the efficacy and potential of hyperbolic representations for vision tasks.

## 1 Introduction

Image segmentation and object localization are crucial tasks with diverse applications. Accurately tracing objects in images and pinpointing their spatial coordinates are essential in fields such as robotics, medical imaging, and augmented reality.

Traditional methods, primarily in the Euclidean space, have made notable progress in these tasks [1, 2, 3]. However, with the increasing complexity and volume of visual data, novel approaches are needed for efficiency, scalability, and richer insights.

Images naturally have latent local hierarchies. Hyperbolic geometry, with its capacity to capture hierarchical and tree-like structures, is particularly well-suited for complex, interconnected visual data [4, 5]. Although hyperbolic operations are expensive, recent advancements [6] have reduced compute needs to a large extent.

In the era of data-driven techniques and large-scale image analysis, our method prioritizes reducing computational and memory demands while being unsupervised and in the

hyperbolic realm. With minimal dimensionality and a test-time training approach, our solution is ideal for real-time applications on resource-constrained edge devices. Our work paves the way for more efficient and accessible image analysis solutions, extending beyond segmentation and object localization.

Through this work, we make the following contributions:

1. **Hyperbolic GNNs for Image Analysis:** We present Seg-HGNN, a novel framework for image segmentation that captures latent structures within images, while being constrained on a hyperbolic manifold. Leveraging Hyperbolic Graph Neural Networks (HGNNs), we show how this non-Euclidean framework improves segmentation and object localization.
2. **Richer image representations in low dimensions:** We show a method of achieving locally aware and semantically rich image representations in low embedding sizes and benchmark their performance.

## 2 Related Work

**Image segmentation.** It is the process of dividing an image into multiple distinct regions or segments, each of which corresponds to a meaningful object. Deep Convolutional Neural Networks (CNNs) [63, 69] have shown promising performance for this task and are the go-to backbone of recent methods. The coming of encoder-decoder models [9, 24, 68] have given way to numerous state-of-the-art variants for segmentation tasks.

While transformers [46] revolutionized language processing, Vision Transformers (ViT) [9] were shown to be at-par with CNNs [25, 43]. However, these models need a huge amount of training data. Caron et al. [10] trained ViTs using self-Distillation with NO labels (DINO), and it was seen that the generated attention maps corresponded to semantic segments in the image. Works of Melas-Kyriazi et al. [54], Shi and Malik [40] explore classical graph theory, using deep features for localization and segmentation. Seg-HGNN takes these ideas into the hyperbolic manifold, while being lightweight and unsupervised.

**Graph Convolutional Networks. (GCNs)** GCNs [61] are networks that take advantage of structured data, and have achieved remarkable results in complex tasks like drug discovery [63] and protein analysis [29]. DGCNet [54] utilizes a dual GCN framework to model image feature context in both coordinate and feature spaces, and merges them back together. Hu et al. [26] introduces a class-wise dynamic GCN module to cluster similar pixels together and dynamically aggregate features. DeepCut [10] constructs a graph with the pair-wise affinities between local image features and performs correlation clustering. Seg-HGNN also tries to use this expressiveness of GCNs.

**Hyperbolic spaces.** Known for embedding hierarchies and tree-like structures with minimal distortion in low dimensions [65], they have inspired hyperbolic variants of neural network blocks [23, 41]. Khrulkov et al. [30] shows the presence of hierarchies in image datasets, which brought along early success in hyperbolic computer vision for few-shot, zero-shot, and unsupervised learning [22, 67, 61]. Exploiting the intrinsic negative curvature of hyperbolic manifolds, they emerge as a powerful framework for geometric deep learning. However, many rely on tangent spaces for aggregation and message passing [12, 62]. Dai et al. [18] presents a fully hyperbolic GNN by introducing a Lorentz linear transform.

**Image segmentation in hyperbolic spaces.** Atigh et al. [9] formulates dataset-level hierarchy and does pixel-level classification in the hyperbolic space in a supervised-setup. They highlight the richness of embeddings, offering boundary information and uncertainty measures in low dimensions.

## 3 Preliminaries

### 3.1 Hyperbolic Space

A Riemannian manifold  $(\mathcal{M}, g)$  is a smooth, connected space where each point  $x \in \mathcal{M}$ , has a tangent space  $\mathcal{T}_x\mathcal{M}$ , that behaves like it is Euclidean. Hyperbolic spaces are a specific type of Riemannian manifold that exhibit constant negative curvature [6]. They are studied using a few isometric models[10]. We choose the Lorentz model for its numerical stability [66].

The Lorentz model for an  $n$ -dimensional hyperbolic space is defined by the manifold  $\mathcal{L} = \{x = [x_0, x_1, \dots, x_n] \in \mathbb{R}^{n+1} : \langle x, x \rangle_{\mathcal{L}} = -1, x_0 > 0\}$ , with the metric tensor  $g = \text{diag}([-1, \mathbf{1}_n^T])$ . We describe a few needed operations below.

**Inner product.** The Lorentz inner product is defined as

$$\langle x, y \rangle_{\mathcal{L}} = x^T g y = -x_0 y_0 + \sum_{i=1}^n x_i y_i \quad (1)$$

**Exponential and logarithmic maps.** An exponential map  $\exp_x(v)$ , projects a vector  $v \in \mathcal{T}_x\mathcal{M}$ , onto the manifold  $\mathcal{M}$ . A logarithmic map,  $\log_x(y)$ , is the inverse operation, which maps a point  $y \in \mathcal{L}$  to the tangent space of  $x$ , that is  $\mathcal{T}_x\mathcal{M}$ . These functions complement each other by satisfying  $\log_x(\exp_x(v)) = v$ . For  $x, y \in \mathcal{L}$ , and  $v \in \mathcal{T}_x\mathcal{L}$ , they are defined as

$$\exp_x(v) = \cosh(\|v\|_{\mathcal{L}})x + \sinh(\|v\|_{\mathcal{L}}) \frac{v}{\|v\|_{\mathcal{L}}} \quad (2)$$

$$\log_x(v) = \frac{\text{arccosh}(-\langle x, v \rangle_{\mathcal{L}})}{\sqrt{\langle x, v \rangle_{\mathcal{L}}^2 - 1}} (v + \langle x, v \rangle_{\mathcal{L}} x) \quad (3)$$

where  $\|v\|_{\mathcal{L}} = \sqrt{\langle v, v \rangle_{\mathcal{L}}}$ .

**Isometric bijections.** The Poincaré ball  $\mathcal{B}$  and the Klein model  $\mathcal{K}$  are two other models of hyperbolic spaces. There are bijective relations connecting all these model. For a point  $x = [x_0, x_1, \dots, x_n] \in \mathcal{L}$ , its corresponding point  $b = [b_0, b_1, \dots, b_{n-1}] \in \mathcal{B}$ , is obtained as

$$p_{\mathcal{L} \rightarrow \mathcal{B}}(x) = \frac{[x_1, \dots, x_n]}{x_0 + 1}, p_{\mathcal{B} \rightarrow \mathcal{L}}(b) = \frac{[1 + \|b\|^2, 2b]}{1 - \|b\|^2} \quad (4)$$

Similarly for  $x = [x_0, x_1, \dots, x_n] \in \mathcal{L}$  and  $k = [k_0, k_1, \dots, k_{n-1}] \in \mathcal{K}$ , we have

$$p_{\mathcal{L} \rightarrow \mathcal{K}}(x) = \frac{[x_1, \dots, x_n]}{x_0}, p_{\mathcal{K} \rightarrow \mathcal{L}}(k) = \frac{[1, k]}{\sqrt{1 - \|k\|^2}} \quad (5)$$

### 3.2 Hyperbolic Graph Convolutional Networks (GCNs)

GCNs [51] introduce the convolution operation on structured, graph data. A graph  $\mathcal{G}$  is made of a set of vertices  $\mathcal{V}$  and edges  $\mathcal{E}$ . Each node  $i$  has an associated embedding  $h_i^0$  and a set of neighbours,  $\mathcal{N}_i$ . The convolution operation for a single layer is done with three underlying sub-operations, namely, feature transformation, message-passing, and update. Constructing a GCN in the hyperbolic space presents unique challenges due to the need to uphold the constraints of hyperbolicity across all layers. Several strategies have addressed these challenges in notable works [12, 52]. They use eq.(3) and eq.(2) to perform transformation and message passing in the tangent space and then project it back to the manifold.

**Hyperbolic Feature Transformation.** Using the feature transformation defined as a matrix-vector multiplication in the Euclidean space, as in Kipf and Welling [51], breaks the hyperbolic constraint when applied to hyperbolic node representations. To avoid this, and to make the aggregation and message-passing steps take less compute, we adopt the Lorentz linear transform introduced in Dai et al. [18], which is defined as

$$\begin{aligned} y &= Wx \\ \text{s.t. } W &= \begin{bmatrix} 1 & 0^\top \\ 0 & \tilde{W} \end{bmatrix}, \tilde{W}^\top \tilde{W} = I \end{aligned} \quad (6)$$

where they show that  $x, y \in \mathcal{L}$  and  $\tilde{W}$  is constrained to be on the Stiefel manifold [24].

**Hyperbolic Neighborhood Aggregation.** We resort to the Einstein midpoint [45] for aggregating the node representations. However, since the Einstein midpoint is defined in the Klein model, we use the bijections defined in eq.(5).

Dai et al. [18] also tells that applying a non-linear activation on the Poincaré ball model does not break hyperbolicity, that is,  $\forall b \in \mathcal{B}$ , we get  $\sigma(b) \in \mathcal{B}$ .

In summary, a graph layer  $l$  in the hyperbolic space would look like

$$\begin{aligned} \tilde{k}_i^l &= p_{\mathcal{L} \rightarrow \mathcal{K}}(W^l h_i^{l-1}) \\ m_i^l &= p_{\mathcal{K} \rightarrow \mathcal{L}}\left(\sum_{j \in \mathcal{N}_i} w_{ij} \gamma_j \tilde{k}_j^l / \sum_{j \in \mathcal{N}_i} \gamma_j\right) \\ h_i^l &= p_{\mathcal{B} \rightarrow \mathcal{L}}(\sigma(p_{\mathcal{L} \rightarrow \mathcal{B}}(m_i^l))) \end{aligned} \quad (7)$$

where  $h_i \in \mathcal{L}$ ,  $W$  is the Lorentz linear transformation matrix,  $w_{ij}$  is the weight of the edge between nodes  $i$  and  $j$ , and  $\mathcal{N}_i = \mathcal{N}_i \cup \{i\}$  is the neighborhood of node  $i$ .  $\gamma_i = \frac{1}{\sqrt{1 - \|k_i\|^2}}$  is called the Lorentz factor, and is used during the calculation of the Einstein midpoint.

### 3.3 Segmentation as a Graph Clustering problem

The image is first divided into patches and each one is a vertex in graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . We model the task of segmentation as partitioning these patches into  $k$  disjoint clusters  $C_1, C_2, \dots, C_k$  such that  $\cup_i C_i = \mathcal{V}$ . We use the Normalized Cuts [40] objective, and for a partition  $\mathcal{P}$  of graph  $\mathcal{G}$ , the Normalized-Cut cost is given by

$$\text{Ncut}(\mathcal{P}) = \frac{\text{cut}(\mathcal{P})}{\text{assoc}(\mathcal{P})} + \frac{\text{cut}(\bar{\mathcal{P}})}{\text{assoc}(\bar{\mathcal{P}})} \quad (8)$$

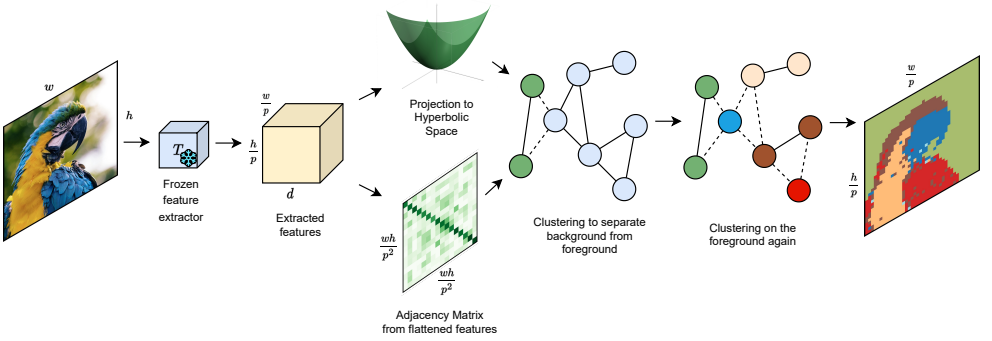


Figure 1: **Method overview.** We extract patch-level features from a frozen model  $T$ , flattening them to  $f$ . Edge weights are obtained from the Gram matrix (eqn.(9)) of  $f$ . We project  $f$  to the Lorentz manifold to obtain initial node embeddings. We optimize an unsupervised graph partitioning loss to obtain  $k$  clusters.

where  $\text{cut}(\cdot)$  is the weight of all removed edges between the partitions,  $\text{assoc}(\cdot)$  is the total weight of partition  $\mathcal{P}$ , and  $\bar{\mathcal{P}}$  is the other partition. In practice, the algorithm recursively bisects the graph using the eigen vectors of the graph Laplacian matrix to form a binary tree structure, representing a hierarchical clustering of the graph. The partitioning process continues until a stopping criterion is met.

## 4 Method

### 4.1 Patch-level features

We use a vision transformer network [19]  $T$ , to get the patch-level features. For an image with size  $m \times n$ , we get  $\frac{mn}{p^2}$  patches, where  $p$  is the patch size of transformer  $T$ . Amir et al. [2], and Melas-Kyriazi et al. [34] have shown that the *key* matrix from the last layer of  $T$  exhibits superior performance in various tasks. Therefore, the image features are represented as  $f \in \mathbb{R}^{\frac{mn}{p^2} \times d}$ , containing the *key*-values of embedding size  $d$  for each patch.

### 4.2 Obtaining Hyperbolic features

The obtained features  $f$  are in the Euclidean space and need to be exported to the Lorentz space  $\mathcal{L}$ . Let  $o_{\mathcal{L}} := [1, 0, \dots, 0]$  be called the origin in  $\mathcal{L}$ . From the definition in eq.(1), we see that  $\langle o_{\mathcal{L}}, [0, f_i] \rangle_{\mathcal{L}} = 0$ . This allows us to think of  $[0, f_i]$  as a vector in the tangent space of  $o_{\mathcal{L}}$ . We can then just use the exponential map defined in eq.(2) to get representations on  $\mathcal{L}$ .

### 4.3 Edge weights and Clustering loss

The edge weights are obtained from the correlation matrix of the transformer features,  $ff^T$ . An additional normalizing factor comes from the idea of normalized graph Laplacian [6, 53]. For an edge and its weight  $\tilde{w}_{ij}$ , we divide it by the square root of the degrees of the associated nodes. Also, as normalized-cut works with positive weights only, we threshold these values

at 0.

$$w_{ij} = \frac{\max(0, \tilde{w}_{ij})}{\sqrt{|\tilde{\mathcal{N}}_i| |\tilde{\mathcal{N}}_j|}} = \frac{\max(0, \tilde{w}_{ij})}{\frac{mm}{p^2}} \quad (9)$$

After passing through a one-layer hyperbolic GCN, we use a fully-connected layer followed by a softmax( $\cdot$ ), to get the probabilities of cluster assignment for each patch, as a matrix  $\mathcal{S}$ .

$$\begin{aligned} \mathcal{F} &= \text{HyperbolicGCN}(f^{\mathcal{L}}, w) \\ \mathcal{F}' &= \log_{o_{\mathcal{L}}}(\mathcal{F}) \\ \mathcal{F}'' &= \text{FullyConnected}(\mathcal{F}') \\ \mathcal{S} &= \text{softmax}(\mathcal{F}'') \end{aligned} \quad (10)$$

We use the relaxed normalization-cut proposed in Bianchi et al. [8], to get the clustering loss.

$$loss_{\text{N-cut}} = -\frac{\text{tr}(\mathcal{S}^T w \mathcal{S})}{\text{tr}(\mathcal{S}^T D \mathcal{S})} + \left\| \frac{\mathcal{S}^T \mathcal{S}}{\|\mathcal{S}^T \mathcal{S}\|_F} - \frac{I_k}{\sqrt{k}} \right\|_F \quad (11)$$

where  $\text{tr}(\cdot)$  is the trace function,  $D = \text{diag}(\sum_j w_{ij})$ ,  $k$  is the number of clusters we want to segment into, and  $I$  is the identity matrix.

## 4.4 Optimization on Hyperbolic manifolds

We keep the transformer  $T$  frozen.  $\text{HyperbolicGCN}(\cdot)$  and the  $\text{FullyConnected}(\cdot)$  layers contain the only trainable parameters. The operations and parameters in the latter layer are euclidean, and can be learned by any standard gradient descent optimizer. The transformation matrix  $W$  in eq.(6) is bounded by the orthogonality constraint of submatrix  $\tilde{W}$ . The set of matrices with orthonormal columns form another Riemannian manifold called the Stiefel manifold [9]. We use Riemannian stochastic gradient descent optimizer to learn  $\tilde{W}$ .

## 4.5 Segmentation and Localization

The cluster assignment probabilities ( $\mathcal{S}$ ) and the cluster count ( $k$ ) facilitate segmentation and localization. For object localization, we set  $k$  as 2 and color-map the patches to assigned clusters. We draw a bounding box around clusters with an area greater than that of 4 patches. Following Aflalo et al. [10], the cluster which appears on more than 2 edges is called the background. The method used for Object Segmentation is identical to Localization, except the inclusion of bounding boxes. For Semantic Part Segmentation, we use a top-down recursive approach. After separating the foreground patches from the image with  $k = 2$ , we perform another round of clustering with  $k = 4$  on the foreground patches only. This also reduces the bias of clustering towards large clusters.

# 5 Experiments and Results

## 5.1 Training Details

We use DINO [11] trained ViT-S with a patch size of 8 as the feature extractor  $T$ . These features are projected to the 16-dimensional Lorentz space, unless stated otherwise. The proposed  $\text{HyperbolicGCN}(\cdot)$  has one graph layer, with both input and output dimension size

Method	Object Localization		Object Segmentation		
	VOC-07	VOC-12	CUB-200	DUTS	ECSSD
Selective Search [14]	18.8	20.9	-	-	-
EdgeBoxes [54]	31.1	31.6	-	-	-
DINO-[CLS] [10]	45.8	46.2	-	-	-
LOST [12]	61.9	64.0	-	-	-
OneGAN [9]	-	-	55.5	-	-
Voynov et al. [47]	-	-	68.3	49.8	-
Spectral Methods [54]	62.7	66.4	76.9	51.4	73.3
TokenCut [50]	68.8	72.1	-	57.6	71.2
DeepCut [11]	69.8	72.2	78.2	<b>59.5</b>	74.6
Seg-HGNN (ours)	<b>72.3</b>	<b>76.1</b>	<b>79.0</b>	57.6	<b>75.9</b>

Table 1: Results comparing Object Localization and Object Segmentation performance.

of 16. The FullyConnected( $\cdot$ ) layer follows with a stack of two linear layers with hidden-layer dimension 32 and output size of  $K = 2$  or 4, for tasks in Section 4.5.

We use a test-time training paradigm, where patches from the image to be segmented are clustered by the hyperbolic graph framework. Our model has almost 7.3k parameters on the euclidean, and 256 parameters on the stiefel manifold. This totals to **7.5k trainable parameters**, compared to 30k parameters in the current state-of-the-art [11].

With the relaxed normalization-cut loss, we train our model for 10 epochs for localization and object segmentation, and for 100 epochs for semantic part segmentation. We use a learning rate of 0.01 for the euclidean parameters and 0.1 for the stiefel parameters. More implementation details and performance analysis are provided in the supplementary material.

## 5.2 Results

In this sub-section, we assess Seg-HGNN’s performance on popular benchmarks for localization and segmentation. We compare it against unsupervised and some supervised methods.

**Localization and Object Segmentation.** In Table 1, we report the object localization performance of our unsupervised approach on PASCAL VOC 2007 [20] and 2012 [21] datasets. Here, we use the Correct Localization (CorLoc) metric, which is the percentage share of images where intersection-over-union with the ground truth bounding box is greater than 0.5.

Table 1 also compares Seg-HGNN’s segmentation performance on three datasets : CUB (Caltech-UCSD Birds-200-2011) [48], a widely used dataset containing images of birds for single object segmentation and semantic part segmentation with 5794 test images, DUTS [49] with 5019 test images, and ECSSD [52] with 1000 images. We report our results using mean Intersection-over-Union (mIOU).

**Semantic Part Segmentation.** In Table 2 we present the performance of Seg-HGNN on the CUB dataset for semantic part segmentation. We use the Normalized Mutual Information (NMI), and the Adjusted Rank Index (ARI) scores here. Seg-HGNN is at par with all other unsupervised methods, including some supervised methods [16, 27, 28]. Figure 3 shows a few samples and corresponding projections of the segmented part embeddings. These are direct projections from the Poincaré ball model. We can see that Seg-HGNN tries to put semantically similar parts together.




Figure 2: A few samples comparing the quality of segmentation achieved by Seg-HGNN. Here, the second row shows the predicted masks and the third row has the ground-truth.

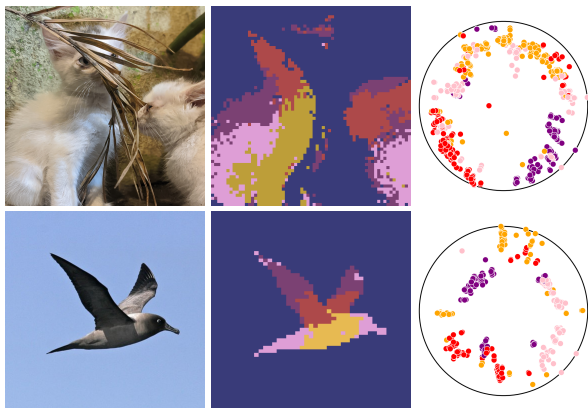


Figure 3: Samples showing the quality of semantic segmentation. The first image shows that even for multiple objects, the underlying semantic behind clustering remains intact. For example, the heads of both the cats are clustered together in red. The projected hyperbolic embeddings also show the quality of clustering.

Method	NMI	ARI
SCOPS [28]	24.4	7.1
Huang and Li [27]	26.1	13.2
Choudhury et al. [14]	43.5	19.6
DFE [17]	25.9	12.4
Amir et al. [4]	38.9	16.1
DeepCut [10]	<b>43.9</b>	20.2
Seg-HGNN (ours)	42.3	<b>20.8</b>

Table 2: **Semantic Segmentation results** over the CUB-200 dataset. Here, the first three methods use ground truth masks for supervision.



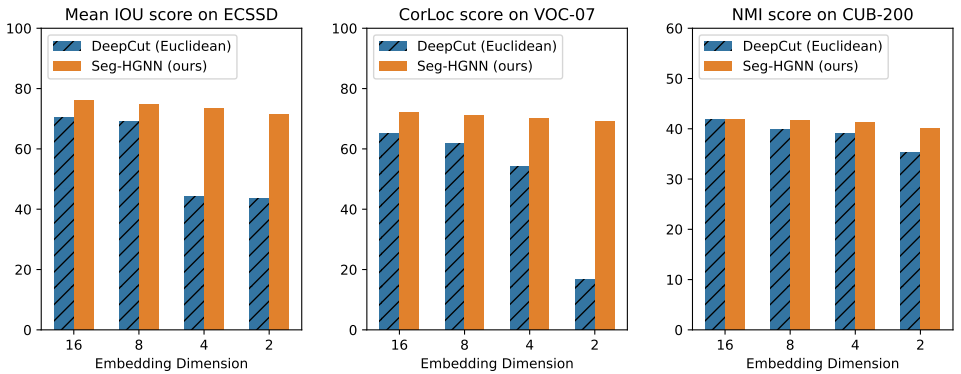


Figure 4: Low-dimensional performance of hyperbolic embeddings.

### 5.3 Effectiveness of low-dimensional embeddings

Through Figure 4, we present the effectiveness of hyperbolic embeddings in a low-dimensional setting. While performance remains similar at higher dimensions, there’s a drastic drop in Euclidean scores. Hyperbolic embeddings can hold information for sizes as low as  $d = 2$ . The results emphasize that low dimensional hyperbolic embeddings might be a go to solution when explainability, reduced complexity and memory footprint are of concern.

### 5.4 Comparing used resources

Hyperbolic operations are said to be compute-demanding. Since even with  $d = 2$  we get reliable performance, Table 3 suggests faster inference speeds and low resource-usage.

Model	V-RAM (GB) ↓				Infer Rate (img/sec) ↑			
	d = 2	4	8	16	2	4	8	16
Seg-HGNN	1.2	1.3	1.5	1.83	2.1	1.9	1.85	1.54
DeepCut	1.81 (d = 64)				1.44 (d = 64)			

Table 3: Resource usage and inference rates on ECSSD, including pre and post-processing.

### 5.5 Effect of adding Seg-HGNN and Conclusion

We benchmark the performance of a few unsupervised methods, both with and without Seg-HGNN. Table 4 shows that adding Seg-HGNN improves scores for both tasks. We conclude that Seg-HGNN and hyperbolic manifolds effectively represent complex hierarchical structures even in minimal dimensions.

Method		VOC-07	ECSSD
MoCo-V3 [15]	without Seg-HGNN	46.19	36.7
	with Seg-HGNN	<b>61.2</b>	<b>53.3</b>
DINO [16]	without Seg-HGNN	45.8	51.2
	with Seg-HGNN	<b>72.3</b>	<b>75.9</b>

Table 4: The effect of adding Seg-HGNN on top of pretrained methods.

## References

- [1] Amit Aflalo, Shai Bagon, Tamar Kashti, and Yonina Eldar. Deepcut: Unsupervised segmentation using graph neural networks clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 32–41, 2023.
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021.
- [3] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne Van Noord, and Pascal Mettes. Hyperbolic image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4453–4462, 2022.
- [4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [5] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(11), 2006.
- [6] Ricardo Bendetti and Carlo Petronio. *Lectures on hyperbolic geometry*. Springer Science and Business Media, 2012.
- [7] Yaniv Benny and Lior Wolf. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 514–530. Springer, 2020.
- [8] Filippo Maria Bianchi, Daniele Grattarola, and Cesare Alippi. Spectral clustering with graph neural networks for graph pooling. In *International conference on machine learning*, pages 874–883. PMLR, 2020.
- [9] William M Boothby. *An introduction to differentiable manifolds and Riemannian geometry*. Academic press, 1986.
- [10] James W Cannon, William J Floyd, Richard Kenyon, and Walter R Parry. *Hyperbolic geometry*, 1997. Flavors of Geometry.
- [11] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [12] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. *Advances in neural information processing systems*, 32, 2019.
- [13] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.

- [14] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [15] X Chen, S Xie, and K He. An empirical study of training self-supervised vision transformers. in 2021 ieee. In *CVF International Conference on Computer Vision (ICCV)*, pages 9620–9629, 2021.
- [16] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Unsupervised part discovery from contrastive reconstruction. *Advances in Neural Information Processing Systems*, 34:28104–28118, 2021.
- [17] Edo Collins, Radhakrishna Achanta, and Sabine Susstrunk. Deep feature factorization for concept discovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–352, 2018.
- [18] Jindou Dai, Yuwei Wu, Zhi Gao, and Yunde Jia. A hyperbolic-to-hyperbolic graph convolutional network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 154–163, 2021.
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [20] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>, .
- [21] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>, .
- [22] Pengfei Fang, Mehrtash Harandi, and Lars Petersson. Kernel methods in hyperbolic spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10665–10674, 2021.
- [23] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in neural information processing systems*, 31, 2018.
- [24] Bin Gao, Nguyen Thanh Son, P-A Absil, and Tatjana Stykel. Riemannian optimization on the symplectic stiefel manifold. *SIAM Journal on Optimization*, 31(2):1546–1575, 2021.
- [25] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 574–584, 2022.
- [26] Hanzhe Hu, Deyi Ji, Weihao Gan, Shuai Bai, Wei Wu, and Junjie Yan. Class-wise dynamic graph convolution for semantic segmentation. In *European Conference on Computer Vision*, pages 1–17. Springer, 2020.

- [27] Zixuan Huang and Yin Li. Interpretable and accurate fine-grained recognition via region grouping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8662–8672, 2020.
- [28] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 869–878, 2019.
- [29] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [30] Valentin Khruikov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6418–6428, 2020.
- [31] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- [32] Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *Advances in neural information processing systems*, 32, 2019.
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [34] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8364–8375, 2022.
- [35] Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems*, 30, 2017.
- [36] Maximilian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International conference on machine learning*, pages 3779–3788. PMLR, 2018.
- [37] Jiwoong Park, Junho Cho, Hyung Jin Chang, and Jin Young Choi. Unsupervised hyperbolic representation learning via message passing auto-encoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5516–5526, 2021.
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.

- [39] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
- [40] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- [41] Ryohei Shimizu, Yusuke Mukuta, and Tatsuya Harada. Hyperbolic neural networks++. *arXiv preprint arXiv:2006.08210*, 2020.
- [42] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. *arXiv preprint arXiv:2109.14279*, 2021.
- [43] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7262–7272, 2021.
- [44] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104: 154–171, 2013.
- [45] Abraham A Ungar. *Analytic Hyperbolic Geometry*. World Scientific, 2005. doi: 10.1142/5914. URL <https://www.worldscientific.com/doi/abs/10.1142/5914>.
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [47] Andrey Voynov, Stanislav Morozov, and Artem Babenko. Object segmentation without labels with large-scale generative models. In *International Conference on Machine Learning*, pages 10596–10606. PMLR, 2021.
- [48] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [49] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 136–145, 2017.
- [50] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14543–14553, 2022.
- [51] Jiexi Yan, Lei Luo, Cheng Deng, and Heng Huang. Unsupervised hyperbolic metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12465–12474, 2021.

- [52] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1155–1162, 2013.
- [53] Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018.
- [54] Li Zhang, Xiangtai Li, Anurag Arnab, Kuiyuan Yang, Yunhai Tong, and Philip HS Torr. Dual graph convolutional network for semantic segmentation. *arXiv preprint arXiv:1909.06121*, 2019.
- [55] Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *Proceedings of the 20th International conference on Machine learning (ICML-03)*, pages 912–919, 2003.
- [56] C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 391–405. Springer, 2014.