

Is Hyperbolic Space All You Need for Medical Anomaly Detection?

Alvaro Gonzalez-Jimenez^{2,3}[0000-0002-1337-9430], Simone Lionetti²[0000-0001-7305-8957], Ludovic Amruthalingam²[0000-0001-5980-5469], Philippe Gottfrois^{1,3}[0000-0001-8023-3207], Fabian Gröger^{1,2}[0000-0002-9699-688X], Marc Pouly^{2*}[0000-0002-9520-4799], and Alexander A. Navarini^{1,3*}[0000-0001-7059-632X]

¹ University of Basel, Switzerland

² Lucerne University of Applied Sciences and Arts, Switzerland

³ University Hospital of Basel, Switzerland

alvaro.gonzalezjimenez@usb.ch

Abstract. Medical anomaly detection has emerged as a promising solution to challenges in data availability and labeling constraints. Traditional methods extract features from different layers of pre-trained networks in Euclidean space; however, Euclidean representations fail to effectively capture the hierarchical relationships within these features, leading to suboptimal anomaly detection performance. We propose a novel yet simple approach that projects feature representations into hyperbolic space, aggregates them based on confidence levels, and classifies samples as healthy or anomalous. Our experiments demonstrate that hyperbolic space consistently outperforms Euclidean-based frameworks, achieving higher AUROC scores at both image and pixel levels across multiple medical benchmark datasets. Additionally, we show that hyperbolic space exhibits resilience to parameter variations and excels in few-shot scenarios, where healthy images are scarce. These findings underscore the potential of hyperbolic space as a powerful alternative for medical anomaly detection. The project website can be found at <https://hyperbolic-anomalies.github.io>

Keywords: Hyperbolic Learning · Anomaly Detection · Hierarchical

1 Introduction

Anomaly detection and localization plays a critical role in various domains, particularly in medical imaging, where distinguishing and localizing between normal and anomalous samples is crucial. A widely adopted approach involves training models exclusively on healthy images, identifying any deviation from this learned distribution as anomalous [32]. This strategy mitigates challenges associated with the scarcity of annotated lesion images while reducing annotation costs and biases inherent in training Artificial Intelligence (AI) models.

* These authors are joint last authors.

Among the most effective anomaly detection techniques are projection-based methods, which leverage pre-trained networks to map data into abstract representations, thereby enhancing the separation between normal and anomalous samples. One-class classification [29,23] defines a compact, closed distribution for normal samples, treating any deviations as anomalies. The teacher-student framework [35,9] employs a student network to learn normal sample representations from a teacher, using their representation discrepancy to identify anomalies. Memory Bank methods [8,28,22] store normal sample prototypes and apply statistical modeling or distance metrics to detect anomalies.

A common feature across these methods is the extraction of representations from specific layers of a pre-trained network. Each layer encodes hierarchical attributes, but conventional approaches rely on Euclidean space, which may not be the best option to capture hierarchical relationships [27,30]. This can lead to suboptimal feature representations and reduced anomaly identification performance.

Hyperbolic space, the geometry of constant negative curvature, is well suited to represent hierarchical structures due to its exponential expansion properties [6]. Recent advances have demonstrated the effectiveness of hyperbolic embeddings in domains such as few-shot learning [18], representation learning [11,12,21], and Out-of-Distribution (OOD) detection [33,14]. Given the hierarchical nature of medical image structures, including disease organization, progression, and anatomical relationships, we hypothesize that hyperbolic space can effectively embed these spatial relationships to enhance anomaly localization. This work aims to answer the following research question: **Can hyperbolic space effectively represent hierarchical features and improve anomaly localization performance?**

To address this question, we propose a novel framework that generate synthetic anomalies, extracts multi-layer features from a pre-trained network and projects them into hyperbolic space. These hyperbolic embeddings are aggregated by weighting features based on their confidence, specifically considering their distance from the origin [18,13], which encodes hierarchical depth. Finally, we construct a hyperplane in hyperbolic space to distinguish between normal and anomalous samples.

We validate our framework on multiple medical benchmark datasets including different imaging modalities such as Magnetic Resonance Imaging (MRI), Computer Tomography (CT), Optical Coherence Tomography (OCT) and X-Ray. Our results demonstrate that hyperbolic space consistently outperforms Euclidean space for anomaly detection and localization. Additionally, we find that hyperbolic space exhibits robustness to parameter tuning by adaptively learning the optimal curvature, further improving performance. Notably, our approach achieves state-of-the-art results in few-shot settings, where healthy images are scarce or unavailable.

The paper is organized as follows: In section 2, we introduce the motivation behind our framework and provide its mathematical formulation. Section 3 details the datasets, implementation, and training specifics, along with evalu-

ation metrics. Section 4 presents experimental findings, comparing hyperbolic and Euclidean-based methods and analyzing performance under few-shot conditions. Finally, section 5 summarizes our contributions and discusses broader implications.




Fig. 1. Overview of the anomaly localization methodology in the hyperbolic space, from medical anomaly synthesis to classification.

2 Methodology

This section details our framework to anomaly localization in hyperbolic space, illustrated in fig. 1. In section 2.1 we describe our method for synthesizing medical anomalies. Section 2.2 describes how we obtain patchified features from a pre-trained network. Section 2.3 outlines the mapping of patchified Euclidean features to hyperbolic space, followed by the hierarchical aggregation process. Finally, in section 2.4, we present our hyperbolic classifier, which leverages these aggregated features for classification.

2.1 Synthesis Anomalies

Given a training set of normal images $x_i \in \mathbb{R}^{H_0 \times W_0 \times C_0}$, we generate images x'_i with synthetic anomalies. These include

- **CutPaste** [23], random patches extracted and blended in another location with Poisson image editing [31,3];
- **Gaussian Intensity** [37], intensity variations introduced via Gaussian filtering to simulate anomalies such as tumors or cysts;
- **Source Deformation** [3], geometric deformations applied by shifting points within a mask, controlled by a scaling parameter.

2.2 Feature Extraction

Features are extracted from the anomaly image x'_i using a pre-trained network, typically a ResNet-like backbone. We select a subset of feature levels L , corre-

sponding to different layers in the network hierarchy. At each level, local features are computed by aggregating patch-wise neighborhoods using adaptive average pooling. The resulting feature maps are then upsampled to the highest spatial resolution among them for later aggregation to give a feature map $f_{i,l} \in \mathbb{R}^C$.

2.3 Euclidean to Hyperbolic Features

Hyperbolic geometry, characterized by constant negative curvature, effectively models hierarchical structures [27,30]. We employ the Lorentz model due to its simple expression for geodesics [21] and numerical robustness [26]. Minkowski space is the space of vector $\mathbf{z} = (z_0, \mathbf{z}) \in \mathbb{R} \times \mathbb{R}^n$ equipped with the Lorentz inner product $\langle \mathbf{z}, \mathbf{z}' \rangle_{\mathbb{L}} = \mathbf{z} \cdot \mathbf{z}' - z_0 z'_0$.

The Lorentz hyperboloid model \mathbb{L}_c^n of n -dimensional hyperbolic space with curvature c is the manifold that satisfies $\langle \mathbf{z}, \mathbf{z} \rangle_{\mathbb{L}} = -1/c$ with $z_0 > 0$. As the feature vectors lie on the Euclidean space we use the the exponential map to projects them onto the hyperboloid

$$\mathbf{x} = \exp_{\mathbf{z}}(\mathbf{v}) = \cosh(\sqrt{c} \|\mathbf{v}\|_{\mathbb{L}}) \mathbf{z} + \frac{\sinh(\sqrt{c} \|\mathbf{v}\|_{\mathbb{L}})}{\sqrt{c} \|\mathbf{v}\|_{\mathbb{L}}} \mathbf{v}, \quad (1)$$

so $\mathbf{f}_{i,l} = \exp_{\mathbf{O}}(f_{i,l})$, where $\mathbf{O} = (1/\sqrt{c}, \mathbf{0})$ is the hyperboloid origin.

We project the hyperbolic features to a lower-dimensional hyperbolic space, and adapt the features to the target domain with a hyperbolic linear layer [4], as the network is biased from the pre-training dataset *i.e.*, *ImageNet* [10] which is suboptimal in medical contexts. We aggregate features from different hierarchical levels to a single point in hyperbolic space \mathbf{z}_i using a weighted Lorentzian centroid [21]:

$$\mathbf{z}_i = \sqrt{c} \frac{\mathbf{z}'_i}{\|\mathbf{z}'_i\|_{\mathbb{L}}} \quad \text{with} \quad \mathbf{z}'_i = \sum_{l \in L} w_{i,l} \mathbf{f}_{i,l}. \quad (2)$$

The weights $w_{i,l}$ are the Euclidean L_2 norms of features $\mathbf{f}_{i,l}$ after transformation to the Poincaré ball, which are connected to model confidence [18,13,14].

2.4 Hyperbolic Classifier

We classify anomalous features using distances to hyperplanes in the Lorentz model. The hyperplane in \mathbb{L}_c^n perpendicular to \mathbf{w} is given by

$$H_{\mathbf{w}} = \{\mathbf{y} \in \mathbb{L}_c^n \mid \langle \mathbf{w}, \mathbf{y} \rangle_{\mathbb{L}} = 0\}, \quad (3)$$

and the distance of a point \mathbf{z} from the hyperplane reads

$$d_{\mathbb{L}}(\mathbf{z}, H_{\mathbf{w}}) = \frac{1}{\sqrt{c}} \left| \sinh^{-1} \left(\sqrt{c} \frac{\langle \mathbf{w}, \mathbf{z} \rangle_{\mathbb{L}}}{\|\mathbf{w}\|_{\mathbb{L}}} \right) \right|. \quad (4)$$

The patch-wise logit and probability for an image x_i with representation \mathbf{z}_i to be anomalous are then given by

$$\ell_{\mathbf{w}}(\mathbf{z}_i) = \text{sign}(\langle \mathbf{w}, \mathbf{z}_i \rangle_{\mathbb{L}}) \|\mathbf{w}\|_{\mathbb{L}} d_{\mathbb{L}}(\mathbf{z}_i, H_{\mathbf{w}}), \quad p_{\mathbf{w}}(\mathbf{z}_i) = [1 + \exp(\ell_{\mathbf{w}}(\mathbf{z}_i))]^{-1}. \quad (5)$$

The model constructs a hyperplane for robust class discrimination by optimizing it through binary cross-entropy

$$\mathcal{L} = -\mathbb{E}_{x_i \sim \mathcal{A}}[\log(p_{\mathbf{w}}(\mathbf{z}_i))] - \mathbb{E}_{x_i \sim \mathcal{N}}[\log(1 - p_{\mathbf{w}}(\mathbf{z}_i))] \quad (6)$$

where \mathcal{A} and \mathcal{N} are sets of anomalous and normal pixel centroids computed via eq. (2), respectively.

3 Experiments

This section describes how hyperbolic space is evaluated for anomaly detection and localization.

3.1 Datasets

We follow BMAD [2], a recent benchmark for medical anomaly detection and localization spanning different imaging modalities. It features defined dataset splits to facilitate reproducibility and prevent leakage. We only deviate by excluding the pathology dataset Camelyon16 due to known difficulties with memory, and by resizing all images to 224×224 pixels. The five datasets used in this work are summarized in table 1.

BraTS2021 [1] is a widely used dataset for brain tumor segmentation and classification in MRI, BMAD considers the FLAIR sequences for anomaly detection.

BTCV [20] and **Liver Tumor Segmentation (LiTs)** [5] focus on liver CT imaging. BMAD uses the anomaly-free BTCV set for training and LiTs for evaluation.

The **Retinal Edema Segmentation Challenge (RESC)** [16] provides OCT images for retinal pathology analysis.

OCT2017 [17] is a large-scale OCT dataset for retinal disease classification, comprising one normal category and three medical conditions. The latter are treated as a single abnormal class.

RSNA [34] contains chest X-rays labeled with one normal category and eight conditions, all of which are treated as a single abnormal class.

Table 1. Count of normal and anomalous samples across BMAD dataset splits.

Dataset	BraTS2021		BTCV+LiTs		RESC		OCT2017		RSNA	
Split ↓	norm.	anom.	norm.	anom.	norm.	anom.	norm.	anom.	norm.	anom.
Train	7,500	0	1,542	0	4,297	0	26,315	0	8,000	0
Valid	39	44	93	73	45	70	8	24	70	1,420
Test	640	3,075	833	660	1,041	764	242	726	781	16,413

3.2 Experimental Setup

We use a pre-trained WideResNet50 [36] as feature extractor in all experiments. To ensure a fair comparison with baseline methods, we refrain from using data augmentation, applying only ImageNet-based normalization [10]. The pre-trained network is frozen, and only the hyperbolic components are trained. We extract features from `layer_2` and `layer_3`, with a dimensionality of 1024, which are subsequently patchified using a patch size of 3.

The curvature parameter is trainable and initialized to $c = 1$. Training is conducted for 50 epochs across all datasets using the Adam optimizer with a learning rate of 10^{-3} and a batch size of 32. All experiments are performed on a single NVIDIA Tesla V100 GPU with 32 GB of memory.

We evaluate both image-level (detection) and pixel-level (localization) performance using Image-AUROC (I_{AUROC}) and Pixel-AUROC (P_{AUROC}) in percentage, respectively. In addition to our hyperbolic approach, we benchmark against several state-of-the-art Euclidean anomaly detection and localization models, including RD4AD [9], STTFPM [35], PaDiM [8], PatchCore [28], and CFA [22].

Finally, we test for statistical significance using the Mann-Whitney U test to compare the AUROC distributions between two models. We assume statistical significance for $p < 0.05$ and denote this with **bold**.

4 Results

Table 2. Comparison of anomaly detection and localization performance across medical datasets. The values represent the mean, the minimum (subscript), and maximum (superscript) over 5 different random seeds.

Methods	BraTS2021		BTCV + LiTs		RESC		OCT2017	RSNA
	I_{AUROC}	P_{AUROC}	I_{AUROC}	P_{AUROC}	I_{AUROC}	P_{AUROC}	I_{AUROC}	I_{AUROC}
RD4AD	89.52 ^{90.19} _{88.85}	96.36 ^{96.48} _{96.24}	59.14 ^{64.45} _{53.83}	91.40 ^{91.50} _{91.30}	88.25 ^{90.25} _{86.25}	96.18 ^{96.38} _{95.98}	94.88 ^{97.58} _{92.17}	67.63 ^{68.73} _{66.53}
STTFPM	84.25 ^{86.63} _{81.87}	96.03 ^{96.43} _{95.63}	61.48 ^{63.15} _{59.81}	96.26 ^{96.40} _{96.12}	87.26 ^{87.49} _{87.03}	94.90 ^{95.02} _{94.90}	91.88 ^{93.21} _{90.55}	69.31 ^{70.4} _{68.22}
PaDiM	79.62 ^{80.96} _{78.28}	94.22 ^{94.45} _{93.99}	50.91 ^{51.24} _{50.58}	90.48 ^{90.63} _{90.33}	75.15 ^{76.57} _{73.73}	91.22 ^{91.59} _{90.85}	90.17 ^{90.78} _{89.56}	74.48 ^{74.74} _{74.22}
PatchCore	92.02 ^{92.13} _{91.91}	95.53 ^{95.58} _{95.48}	59.33 ^{59.47} _{59.19}	95.00 ^{95.01} _{94.99}	90.54 ^{90.64} _{90.44}	95.87 ^{95.91} _{95.83}	97.45 ^{98.10} _{96.80}	75.67 ^{75.87} _{75.47}
CFA	84.99 ^{85.15} _{84.83}	96.61 ^{96.65} _{96.57}	53.89 ^{58.13} _{49.65}	97.40 ^{97.46} _{97.34}	72.47 ^{74.74} _{70.20}	92.49 ^{93.57} _{91.41}	79.10 ^{79.66} _{78.54}	66.65 ^{66.80} _{66.50}
Ours	92.49 ^{93.02} _{91.96}	95.56 ^{95.63} _{95.49}	65.94 ^{67.99} _{63.89}	96.49 ^{99.11} _{93.87}	90.71 ^{91.28} _{90.14}	95.32 ^{95.56} _{95.08}	97.85 ^{98.12} _{97.58}	79.46 ^{80.20} _{78.72}

Table 2 presents the experimental results for anomaly detection and localization. We observe acceptable agreement of the methods based on Euclidean geometry with the results reported by the BMAD in table 2 [2]. Among these, PatchCore achieves the most consistent performance across datasets, although the somewhat lower performance for BTCV+LiTs could be interpreted as sensitivity to a distribution shift. In contrast, the proposed hyperbolic framework shows the best performance on whole images across all datasets, even if this is not

always statistically significant, and it remains robust even for BTCV+LiTs. At the pixel level, the hyperbolic approach remains competitive with other methods, even though different Euclidean baselines outperform it in specific cases. However, in medical practice, misdiagnosing an entire image is generally more problematic than minor pixel-wise mismatches.

4.1 Ablation Study on Model Parameters

We conduct ablation studies on curvature, patch size, and dimensionality of the hyperbolic space using the BraTS dataset. Figure 2 presents the impact of these variations on performance.

We first investigate the role of curvature by fixing it to $c = \{0.01, 0.1, 1, 10, 100\}$. The first plot indicates that constraining the curvature leads to a decline in performance, with better results observed at lower curvature values. This underscores the advantage of a learnable curvature, which allows the model to adaptively optimize the geometry of the representation space for anomaly identification.

Next, we analyze the effect of the patch size $\{1, 2, 3, 4, 5, 6\}$ when aggregating local features. Increasing the patch size negatively impacts both I_{AUROC} and P_{AUROC} . This suggests that fine-grained feature extraction is preferable for capturing subtle anomalies, whereas overly large patches may dilute local information critical for accurate anomaly localization.

Lastly, hyperbolic space has been shown to efficiently encode representations in lower-dimensional embeddings, making it advantageous for memory-constrained scenarios [19,14]. To evaluate this, we reduce the feature dimensionality to $\{512, 128, 16, 8, 2\}$. The last plot reveals that while I_{AUROC} is more sensitive to extreme dimensionality reduction, P_{AUROC} remains relatively stable.




Fig. 2. Ablation study on key components of our hyperbolic framework: fixed curvature, patch size variations, and hyperbolic layer dimensionality.

4.2 Few-Shot Anomaly Detection and Localization

We further evaluate the robustness of our framework in a few-shot setting, where only a limited number of normal images are available for training. We experi-

ment with $\{1, 3, 5, 10, 25\}$ normal images and compare our performance against PaDiM [8] and PatchCore [28]. The results in Figure 3 demonstrate that our hyperbolic model significantly outperforms both baselines, particularly in extreme data scarcity scenarios.

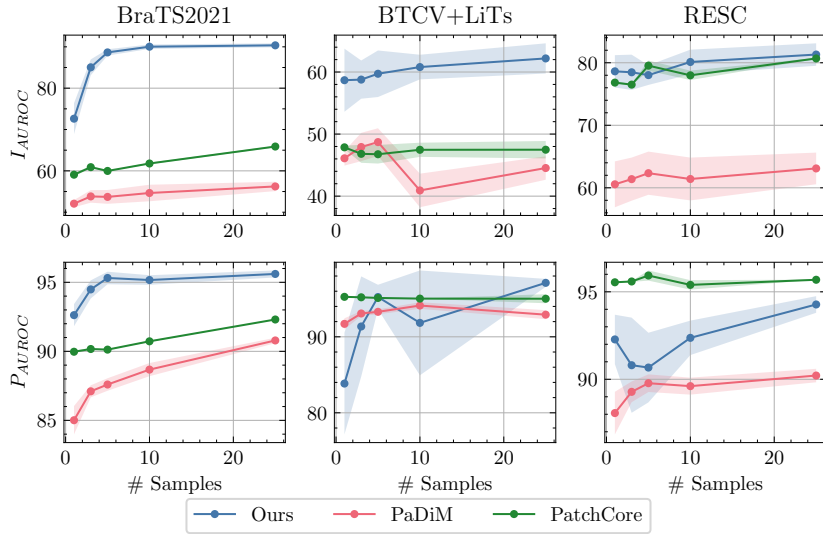


Fig. 3. Few-shot evaluation with varying normal image counts $\{1, 3, 5, 10, 25\}$. Our hyperbolic model outperforms PaDiM and PatchCore in scarce data scenarios. Error bands are obtained with five different random seeds, without changing the training set.

5 Conclusions

In this work, we introduced a hyperbolic anomaly detection and localization framework that leverages the unique geometric properties of hyperbolic space to enhance both classification and localization of medical anomalies.

Our evaluation across multiple medical imaging datasets demonstrates that our method consistently outperforms state-of-the-art anomaly identification approaches in terms of I_{AUROC} , and is competitive with the best ones for localization as demonstrated by P_{AUROC} . Additionally, we show that hyperbolic embeddings retain strong performance in low dimensions enabling efficient deployment in resource-constrained environments, and consistently outperforms Euclidean baselines in few-shot data regimes.

One key area for future investigation is the incorporation of features from earlier layers of the model, which could help leverage the hierarchical information embedded throughout the network. Although fully hyperbolic networks have

been shown to outperform hybrid architectures [4,7,33], this remains a developing research area, and challenges related to stability and reproducibility persist. Furthermore, while our study focused on feature-based anomaly baselines, expanding the comparison to reconstruction-based [25,24] or gradient-based [15] with other performance metrics could provide a more comprehensive assessment of the model’s effectiveness.

Additionally, integrating multi-modal data, such as radiology reports or genomic information, could provide a richer anomaly characterization, improving interpretability and clinical utility.

Our work contributes to the advancement of AI-driven medical anomaly detection and localization, with a particular emphasis on improving the accuracy and localization of anomalies, especially in few-shot settings. These findings have the potential to significantly enhance medical image quality assessment and facilitate the quantitative analysis of rare diseases, ultimately leading to more precise and data-efficient diagnostic models for clinical applications.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The RSNA-ASNR-MICCAI BraTS 2021 Benchmark on Brain Tumor Segmentation and Radiogenomic Classification (Sep 2021)
2. Bao, J., Sun, H., Deng, H., He, Y., Zhang, Z., Li, X.: BMAD: Benchmarks for Medical Anomaly Detection. In: CVPRW (Apr 2024)
3. Baugh, M., Tan, J., Müller, J.P., Dombrowski, M., Batten, J., Kainz, B.: Many tasks make light work: Learning to localise medical anomalies from multiple synthetic tasks. In: MICCAI (Jul 2023)
4. Bdeir, A., Schwethelm, K., Landwehr, N.: Fully Hyperbolic Convolutional Neural Networks for Computer Vision. In: ICLR (Oct 2023)
5. Bilic, P., Christ, P., Li, H.B., Vorontsov, E., Ben-Cohen, A., Kaissis, G., Szeskin, A., Jacobs, C., Mamani, G.E.H., Chartrand, G., et al.: The Liver Tumor Segmentation Benchmark (LiTS). Medical Image Analysis (Feb 2023)
6. Bridson, M.R., Haeffliger, A.: Metric Spaces of Non-Positive Curvature (1999)
7. Chen, W., Han, X., Lin, Y., Zhao, H., Liu, Z., Li, P., Sun, M., Zhou, J.: Fully hyperbolic neural networks. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (2022)
8. Defard, T., Setkov, A., Loesch, A., Audigier, R.: PaDiM: A Patch Distribution Modeling Framework for Anomaly Detection and Localization. In: ICPRW (2021)
9. Deng, H., Li, X.: Anomaly Detection via Reverse Distillation from One-Class Embedding. In: CVPR (Mar 2022)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (Jun 2009)
11. Desai, K., Nickel, M., Rajpurohit, T., Johnson, J., Vedantam, S.R.: Hyperbolic Image-text Representations. In: ICML (2023)

12. Ganea, O.E., Bécigneul, G., Hofmann, T.: Hyperbolic Entailment Cones for Learning Hierarchical Embeddings. In: ICML (Jun 2018)
13. GhadimiAtigh, M., Schoep, J., Acar, E., van Noord, N., Mettes, P.: Hyperbolic Image Segmentation. In: CVPR (Mar 2022)
14. Gonzalez-Jimenez, A., Lionetti, S., Bazazian, D., Gottfrois, P., Gröger, F., Pouly, M., Navarini, A.: Hyperbolic Metric Learning for Visual Outlier Detection. In: ECCVW (Sep 2024)
15. Gonzalez-Jimenez, A., Lionetti, S., Pouly, M., Navarini, A.A.: Sano: Score-based diffusion model for anomaly localization in dermatology. In: CVPRW. pp. 2988–2994 (2023)
16. Hu, J., Chen, Y., Yi, Z.: Automated segmentation of macular edema in OCT using deep neural networks. Medical Image Analysis (Jul 2019)
17. Kermany, D.S., Goldbaum, M., Cai, W., Valentim, C.C.S., Liang, H., Baxter, S.L., McKeown, A., Yang, G., Wu, X., Yan, F., et al.: Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. Cell (Feb 2018)
18. Khulkov, V., Mirvakhabova, L., Ustinova, E., Oseledets, I., Lempitsky, V.: Hyperbolic Image Embeddings. In: CVPR (Mar 2020)
19. Kusupati, A., Bhatt, G., Rege, A., Wallingford, M., Sinha, A., Ramanujan, V., Howard-Snyder, W., Chen, K., Kakade, S., Jain, P., et al.: Matryoshka Representation Learning. In: NeurIPS (Feb 2024)
20. Landman, B., Xu, Z., Igelsias, J., Styner, M., Langerak, T., Klein, A.: Multi-Atlas Labeling Beyond the Cranial Vault - Workshop and Challenge. In: MICCAI (2015)
21. Law, M., Liao, R., Snell, J., Zemel, R.: Lorentzian Distance Learning for Hyperbolic Representations. In: ICML (2019)
22. Lee, S., Lee, S., Song, B.C.: CFA: Coupled-hypersphere-based Feature Adaptation for Target-Oriented Anomaly Localization. IEEE Access (2022)
23. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: CutPaste: Self-Supervised Learning for Anomaly Detection and Localization. In: CVPR (2021)
24. Liu, J., Ma, Z., Wang, Z., Zou, C., Ren, J., Wang, Z., Song, L., Hu, B., Liu, Y., Leung, V.C.M.: A survey on diffusion models for anomaly detection. arXiv preprint arXiv: 2501.11430 (2025)
25. Livernoche, V., Jain, V., Hezaveh, Y., Ravanbakhsh, S.: On diffusion modeling for anomaly detection. ICLR (2023)
26. Mishne, G., Wan, Z., Wang, Y., Yang, S.: The Numerical Stability of Hyperbolic Representation Learning. In: ICML (2023)
27. Nickel, M., Kiela, D.: Poincaré Embeddings for Learning Hierarchical Representations. In: NeurIPS (2017)
28. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards Total Recall in Industrial Anomaly Detection. In: CVPR (May 2022)
29. Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S.A., Binder, A., Müller, E., Kloft, M.: Deep One-Class Classification. In: ICML (2018)
30. Sala, F., Sa, C.D., Gu, A., Re, C.: Representation Tradeoffs for Hyperbolic Embeddings. In: ICML (2018)
31. Schlüter, H.M., Tan, J., Hou, B., Kainz, B.: Natural Synthetic Anomalies for Self-Supervised Anomaly Detection and Localization. In: ECCV (Jul 2022)
32. Tschuchnig, M.E., Gadermayr, M.: Anomaly Detection in Medical Imaging - A Mini Review. In: Data Science – Analytics and Applications (2022)
33. van Spengler, M., Berkhout, E., Mettes, P.: Poincare ResNet. In: ICCV (2023)
34. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. In: CVPR (2017)

35. Yamada, S., Hotta, K.: Reconstruction Student with Attention for Student-Teacher Pyramid Matching (Mar 2022)
36. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016)
37. Zhang, X., Xu, M., Qiu, D., Yan, R., Lang, N., Zhou, X.: MediCLIP: Adapting CLIP for Few-shot Medical Image Anomaly Detection. In: MICCAI (2024)