

HypOProto: Hyperbolic Ordinal Prototypes for Left Ventricular Filling Pressure Classification

Victoria Wu¹, Nima Hashemi¹, Hooman Vaseli¹, Christina Luong², Purang Abolmaesumi¹, and Teresa S. M. Tsang²

¹ Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada

² Vancouver General Hospital, Vancouver, BC, Canada
purang@ece.ubc.ca **

Abstract. Echocardiography (echo) is a widely used imaging modality for assessing cardiac function, with Left Ventricular Filling Pressure (LVFP) serving as a critical physiological marker for conditions such as heart failure. Standard LVFP classification into normal *vs.* elevated categories relies on the Doppler-derived E/e' ratio, which is operator-dependent and often unavailable in resource-limited settings, motivating methods that infer LVFP directly from B-mode echo. Existing deep learning approaches achieve high performance but remain largely black-box, limiting clinical interpretability. We propose HypOProto, a hyperbolic, ordinal prototype-based framework for interpretable LVFP classification using a frozen, explainable foundation model backbone. HypOProto arranges prototypes along the physiological E/e' scale, placing borderline cases near the hyperboloid root where small angular differences separate similar cases, while normal and elevated cases occupy outward positions reflecting increasing diagnostic certainty. This hyperbolic geometry encodes clinically meaningful ordinal relationships and improves interpretability. We also introduce a novel Hyperbolic Prototype Angular Separation (HyperPAS) loss, enforcing inter-class prototype separation in hyperbolic space. HypOProto achieves SOTA performance while maintaining transparency, and highlights clinically relevant regions in visualizations. This work represents the first prototype-based framework for LVFP classification in echo. Our code can be found at <https://github.com/DeepRCL/HypOProto>.

Keywords: Prototypical Neural Networks · Hyperbolic embedding space · Ordinal learning · Echocardiography.

1 Introduction

Echocardiography (echo) is the most widely used imaging modality for evaluating cardiac function due to its non-invasive nature, portability, and ability to visualize the heart in real time. Beyond anatomical assessment, echo provides functional information that is critical for diagnosing diastolic dysfunction.

** T. S.M. Tsang and P. Abolmaesumi are joint senior authors.

Among these functional markers, Left Ventricular Filling Pressure (LVFP) is a primary indicator across various cardiac pathologies including heart failure [2]. Accurate estimation of LVFP is therefore essential in both clinical and research settings, as it directly informs diagnosis, prognosis, and treatment planning. In clinical practice, LVFP is typically estimated using Doppler echocardiography, particularly through the ratio of mitral inflow velocity (E) to the early diastolic mitral annular velocity (e'). This E/e' ratio is then typically used to classify patients as having normal or elevated LVFP [2]. While effective, the measurement is highly operator-dependent, sensitive to image quality and acquisition angle, and often unavailable in resource-limited or point-of-care settings [13]. These limitations motivate the exploration of inferring LVFP directly from B-mode echo images, which are more consistent and widely accessible.

Despite a growing number of works on using deep learning (DL) for B-mode echo analysis, there are currently no methods explicitly targeting LVFP estimation. The feasibility of this task has been indirectly demonstrated by Akerman et al. [1], who successfully diagnosed heart failure with preserved ejection fraction (HFpEF) from apical four-chamber (A4C) videos alone. Given that LVFP is the primary physiological signal for HFpEF, this result implies that B-mode echo encodes sufficient information for learning LVFP.

Most prior DL approaches in echocardiography are black-box models that achieve high performance but limited transparency, which hinders clinical adoption. Prototype-based networks offer an explainable AI (XAI) alternative by providing case-based reasoning in cardiac imaging [5,15]. However, these interpretable methods have not been applied to LVFP estimation.

A further challenge in scaling DL to large echo datasets is the reliance on fully supervised, end-to-end training, which proves impractical given limited labeled data. EchoPrime [18] addressed this by providing robust foundation model representations through self-supervised pretraining on large-scale echo data. However, its pooling architecture remains black-box and limits explainability.

Recent works have demonstrated that hyperbolic geometry improves performance in medical image classification by embedding naturally hierarchical relationships [6,10,16]. In hyperbolic space, there is a characteristic training tendency in which more ambiguous samples remain closer to the manifold origin (root), while more distinctive and class-specific patterns move outward along curved branches toward the boundary, preserving semantic hierarchical structure without Euclidean distortions [9,11]. This radial encoding of ambiguity has been explored in prior hyperbolic contrastive and metric learning methods [4,7]. This behavior is well suited to LVFP classification. Although typically framed as a binary task (normal *vs.* elevated), LVFP is inherently ordinal, governed by the continuous E/e' ratio with a clinical threshold at 14. Cases near this boundary are often more ambiguous, and many elevated studies cluster just above the cut-off, making them particularly difficult to distinguish. Placing these cases closer to the hyperbolic root allows the model to allocate more angular capacity for separating subtle directional differences. This design differs from existing hyperbolic prototype approaches: Busemann-style classification [4] positions prototypes at




Fig. 1. a) Method overview: Videos are encoded with DINOv3 to obtain patch features, then passed to an E/e' prediction head, which determines the hyperbolic radius for exponential mapping. Resulting embeddings are compared with hyperbolic prototypes for classification. b) Ordinal structure of the hyperbolic space: cases near the diagnostic threshold lie close to the root, while more confidently normal and elevated cases are positioned farther away.

the ideal boundary with implicitly learned radii, while other works [16,17] learn prototype positions without external supervision. In contrast, we explicitly anchor radii to E/e' , enforcing alignment with the clinical threshold and encoding ordinal severity by design.

We propose **HypOPProto**, a hyperbolic, ordinal prototype-based framework for interpretable LVFP classification from B-mode echo using a frozen, explainable patch-based foundation model backbone. HypOPProto arranges prototypes ordinally along the physiological E/e' scale, placing borderline and difficult cases where $E/e' \approx 14$ near the hyperboloid root, where small angular differences capture subtle distinctions, while clearly normal cases and cases with very high E/e' move outward with increasing diagnostic certainty (Figure 1b). This geometry encodes clinically meaningful ordinal structure and improves interpretability. Our key contributions are: (1) the first prototype-based echocardiography framework with a frozen foundation-model backbone and the first interpretable model for LVFP classification, (2) a hyperbolic ordinality mechanism that organizes prototypes according to the E/e' scale to better handle ambiguous cases, and (3) a novel Hyperbolic Prototype Angular Separation (HyperPAS) loss that promotes discriminative prototype directions in hyperbolic space.

2 Method

Our goal is to classify LVFP by comparing visual features with a fixed set of learnable prototypes (Figure 1a). This is an inherently ordinal task and to capture this structure, both features and prototypes are projected into hyperbolic space, where radial distance from the origin reflects proximity to the clinical cut-off. Each prototype has a class label and a learnable E/e' value for hyperbolic projection, and all prototypes are jointly optimized with the feature embeddings.

Feature Extraction. Unlike prior prototypical approaches that train the backbone end-to-end [5,15,16], we use a frozen pretrained DINOv3 encoder [14] trained on 10,000 multi-view echo studies to generate clip-level embeddings, without any task-specific exposure to LVFP during pretraining. For a video $x \in \mathbb{R}^{T_0 \times H_0 \times W_0 \times 3}$, the extractor outputs $F \in \mathbb{R}^{T \times H \times W \times D}$, computed once and fixed during training. This design eliminates the computational cost of backbone optimization. We then apply spatiotemporal attention over the embeddings to identify k salient regions of interest (ROIs), A , that guide prototype assignment. The attended representations are obtained via the element-wise product $Z = A \odot F$. Additional details on the feature extraction pipeline can be found in [15].

Hyperbolic Ordinality Projection. To capture the inherent ordinality of LVFP, we project features into a Lorentzian hyperbolic space where the radial distance from the origin reflects proximity to the clinical E/e' cutoff [12]. Borderline cases lie near the root of the hyperboloid, while clearly normal or elevated cases are mapped farther away, enabling the model to encode severity information directly in the feature geometry.

We use a lightweight prediction head to estimate a scalar $\hat{E}$ corresponding to the E/e' value. This estimate is converted to a hyperbolic radius r using a smooth distance function centered at the clinical cutoff:

$$r = r_{\min} + (r_{\max} - r_{\min}) \tanh \left(\frac{\alpha \sqrt{(\hat{E} - 14)^2 + \epsilon^2}}{r_{\max}} \right), \quad (1)$$

where $r_{\min} > 0$ ensures numerical stability, $r_{\max}$ bounds the radius, α controls growth away from the decision boundary, and ϵ is a smoothing factor that preserves gradients near the cutoff.

Following [3], we represent each Euclidean representation Z in the Lorentz hyperboloid as:

$$H(Z) = [Z_{\text{space}}, Z_{\text{time}}]. \quad (2)$$

Features are lifted to the hyperboloid using an exponential map at the origin with the learned radius r that encodes ordinal information based on the predicted E/e' value (Eq. 1). Let $\mathbf{u} = \frac{Z}{\|Z\|_2}$ denote the unit vector in the direction of Z . The hyperbolic space components are then

$$Z_{\text{space}} = \frac{\sinh(\sqrt{c}r)}{\sqrt{c}} \mathbf{u}, \quad Z_{\text{time}} = \sqrt{1/c + \|Z_{\text{space}}\|_2^2}, \quad (3)$$

where c is the hyperboloid curvature. This preserves angular information from the Euclidean features while encoding ordinality through the radial coordinate.

Prototype Classification. Along with the feature representations Z , each prototype P is projected to the Lorentz hyperboloid using the same procedure,

with the radius determined from the prototype’s E/e' value. For a given sample, we compute pairwise geodesic distances between its hyperbolic embeddings $H(Z_k)$ and all prototype embeddings $\{H(P_k)\}_{k=1}^K$ based on the Lorentzian inner product:

$$d(H(Z_k), H(P_k)) = \frac{1}{\sqrt{c}} \operatorname{arccosh} \left(-c \langle H(Z_k), H(P_k) \rangle_{\mathcal{L}} \right). \quad (4)$$

These distances are then converted to similarities, which allows the model to perform classification by directly comparing sample embeddings to prototypes while respecting the ordinal structure encoded in hyperbolic space.

Prototypes are periodically projected onto the nearest training features in hyperbolic space. Specifically, each prototype embedding is replaced by the hyperbolic representation of the closest sample belonging to the same class and with an E/e' value within a tolerance Δ_l of the prototype’s current value. This enables direct visualization while preserving the hyperbolic structure.

Hyperbolic Optimization Objectives. Our training objective combines complementary losses to promote accurate classification, structured prototypes, and coherent attention. We use cross-entropy for classification, hyperbolic prototype cluster and separation losses, and standard ℓ_1 and transformation-consistency regularization (see [17] for details). We also apply a mean absolute error (MAE) loss to the predicted E/e' values.

In addition, we introduce a novel Hyperbolic Prototype Angular Separation (HyperPAS) loss, adapted from [5]. This loss encourages prototypes of the same class to align in similar directions while pushing prototypes of different classes apart. Specifically, prototypes P_i are log-mapped to the tangent space at the origin and normalized. For each prototype, an InfoNCE-style angular objective separates same-class from different-class prototypes:

$$\mathcal{L}_i = -\log \frac{\sum_{j \in \mathcal{P}_i^+} \exp \left(\frac{\operatorname{sim}(\hat{p}_i, \hat{p}_j)}{\tau} \right)}{\sum_{j \in \mathcal{P}_i^+} \exp \left(\frac{\operatorname{sim}(\hat{p}_i, \hat{p}_j)}{\tau} \right) + \sum_{k \in \mathcal{P}_i^-} \exp \left(\frac{\operatorname{sim}(\hat{p}_i, \hat{p}_k)}{\tau} \right)}, \quad (5)$$

where τ is a temperature hyperparameter, $\mathcal{P}_i^+$ denotes same-class prototypes, and $\mathcal{P}_i^-$ denotes different-class prototypes.

The overall loss is averaged over all prototypes:

$$\mathcal{L}_{\text{HyperPAS}} = \frac{1}{K} \sum_{i=1}^K \mathcal{L}_i. \quad (6)$$

This formulation promotes inter-class separation while preserving class-specific directions in hyperbolic space.

The final training objective is a weighted combination of all terms:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda_1 \mathcal{L}_{\ell_1} + \lambda_2 \mathcal{L}_{\text{trans}} + \lambda_3 \mathcal{L}_{\text{MAE}} + \lambda_4 \mathcal{L}_{\text{cls/sepn}} + \lambda_5 \mathcal{L}_{\text{HyperPAS}}, \quad (7)$$

where $\lambda_1, \dots, \lambda_5$ are hyperparameters controlling the weight of each term.

3 Experiments and Results

3.1 Dataset and Implementation

We evaluate our method on AP4 videos from a private echo dataset collected with institutional approval from a tertiary care center. In total, the dataset contains 141,086 unique cines from 32,818 studies. LVFP labels were derived from clinical text reports, yielding a class distribution of 85% normal and 15% elevated. Data were split at the study level into training, validation, and test sets using a 70:15:15 ratio to prevent patient-level leakage.

Experiments use fixed-length 32-frame video clips ($\sim$ one cardiac cycle) with precomputed, frozen DINOv3 features of size $16 \times 14 \times 14 \times 512$ [14]. We employ 10 prototypes per class, with initial E/e' values uniformly distributed between 1 and 25. Optimization runs for 31 epochs with batch size 32, using a learning rate of 1×10^{-3} for the E/e' head and 1×10^{-4} for all other parameters. Training uses class-balanced sampling, and loss weights largely follow ProtoASNet [15], with additional weights of 0.2 for $\mathcal{L}_{\text{MAE}}$ and 0.1 for $\mathcal{L}_{\text{HyperPAS}}$.

3.2 Evaluation

Quantitative Results. Tables 1 and 2 summarize quantitative results at cine- and study-levels. As the dataset exhibits significant class imbalance, we report balanced accuracy rather than standard accuracy to ensure that performance on minority classes is not obscured by majority-class dominance. We compare HypOProto against four baselines: EchoPrime + MLP [18], a linear classifier on frozen EchoPrime embeddings; DINO + MLP [14], a linear classifier on frozen DINO embeddings; Proto + DINO [14,15], a Euclidean prototypical network using the same DINO embeddings; and our implementation of the fully trained closed-source Akerman et al. [1] model (denoted “E2E” in the tables, as it is trained end-to-end from scratch rather than fine-tuned from a foundation model, requiring substantially greater training time).

At the cine level, HypOProto achieves the best performance across all baselines. The most notable improvement is in the elevated LVFP class (F1 0.63), indicating that the hyperbolic ordinal formulation better handles elevated cases, which tend to be closer to the diagnostic boundary.

At the study level, computed as the mean of the softmax probabilities from cine-level logits, HypOProto attains the best mean and class-wise F1 for normal and elevated LVFP. While EchoPrime achieves the highest balanced accuracy, it lacks intrinsic interpretability. It also surpasses the Akerman et al.[1] model, while requiring substantially less training effort due to frozen features.

Qualitative Results. Qualitative visualizations (Figure 2) support that the model learns clinically meaningful and consistent representations. Prototype activation maps (a) show concentrated attention around the mitral valve inflow region, demonstrating anatomically aligned and physiologically grounded explanations. In the CoSNE projection (center-top) [8], most elevated prototypes lie

Table 1. Cine-level ($N = 21,120$) balanced accuracy and F1 scores reported on the test set. Results are reported as mean (stdev) over three runs, with best results in bold.

Method	Cine-level				E2E	XAI
	bACC $\uparrow$	Avg F1 $\uparrow$	Normal F1 $\uparrow$	Elevated F1 $\uparrow$		
EchoPrime+MLP [18]	80.02(0.19)	0.70(0.01)	0.87(0.01)	0.54(0.01)	$\times$	$\times$
DINO+MLP [14]	76.13(0.20)	0.67(0.01)	0.84(0.01)	0.50(0.01)	$\times$	$\times$
Proto+DINO [14,15]	79.45(0.46)	0.72(0.00)	0.88(0.01)	0.56(0.01)	$\times$	$\checkmark$
Akerman et al. [1]	79.65(0.53)	0.71(0.01)	0.87(0.01)	0.56(0.01)	$\checkmark$	$\times$
HypOProto (Ours)	82.57(0.43)	0.76(0.01)	0.91(0.01)	0.63(0.02)	$\times$	$\checkmark$

Table 2. Study-level ($N = 4,894$) balanced accuracy and F1 scores reported on the test set. Results are reported as mean (stdev) over three runs, with best results in bold.

Method	Study-level				E2E	XAI
	bACC $\uparrow$	Avg F1 $\uparrow$	Normal F1 $\uparrow$	Elevated F1 $\uparrow$		
EchoPrime+MLP [18]	84.33(0.40)	0.75(0.01)	0.88(0.01)	0.61(0.01)	$\times$	$\times$
DINO+MLP [14]	80.70(0.24)	0.73(0.02)	0.88(0.02)	0.58(0.02)	$\times$	$\times$
Proto+DINO [14,15]	81.49(0.02)	0.76(0.00)	0.90(0.00)	0.61(0.00)	$\times$	$\checkmark$
Akerman et al. [1]	82.60(0.58)	0.75(0.01)	0.89(0.01)	0.61(0.01)	$\checkmark$	$\times$
HypOProto (Ours)	83.97(0.15)	0.79(0.01)	0.91(0.01)	0.66(0.01)	$\times$	$\checkmark$

closer to the hyperbolic root while normal prototypes extend farther along the manifold. This reflects the dataset distribution where elevated studies cluster near the diagnostic threshold with a mean E/e' of 17, while normal studies lie farther below it with a mean of 8, making the elevated group harder to separate. The normal prototypes closer to the root correspond to borderline normal cases, while a few elevated prototypes farther out represent more extreme elevated pathology, demonstrating that the geometry captures the full ordinal spectrum rather than enforcing a rigid binary split. In contrast, the CoSNE embedding without E/e' -based ordinality (center-bottom) shows substantial mixing between normal and elevated prototypes, resulting in a less interpretable and separable representation. The hyperboloid plot (right) shows that both videos and prototypes are positioned along the manifold according to the predicted E/e' value that determines hyperbolic distance from the root. The color map reflects the ground-truth E/e' , and its smooth alignment with radial position indicates that the prediction head preserves the correct ordinal ordering. Notably, the dataset contains substantially more normal than elevated studies, with most normal cases far from the threshold and elevated cases clustered near it. The hyperbolic formulation accommodates this imbalance by allocating greater angular separation near the root, where borderline and mildly elevated cases concentrate.

Ablation Study. The ablation results show that incorporating hyperbolic ordinality and the HyperPAS loss improves prototype quality while maintaining competitive performance. The full model achieves the best average F1 and second-best balanced accuracy at the study level, indicating a strong balance be-




Fig. 2. a) Visualizations of learned prototypes. b) CoSNE plot of HypOPProto (top) *vs.* without radially enforced ordinality (bottom). c) Hyperboloid projections of videos and prototypes, colored by ground-truth E/e' to illustrate preserved ordinal structure.

Table 3. Study-level ablation on the validation set ($N = 4,922$). Results are reported over three runs, with the best performance in **bold** and the second best underlined. Each experiment removes key hyperbolic and loss components of the proposed model.

Method	Study-level		Sparsity ↓	Diversity ↑
	bACC ↑	Avg F1 ↑		
Proto+DINO	82.05(0.77)	0.75(0.01)	0.68(0.02)	0.97(0.03)
w/o Ord	83.31(1.53)	0.78(0.02)	0.53(0.09)	0.52(0.16)
w/o $\mathcal{L}_{\text{HyperPAS}}$	84.04(0.52)	0.77(0.01)	0.63(0.01)	0.63(0.03)
All parts (ours)	83.89(0.19)	0.78(0.01)	0.61(0.01)	0.70(0.09)

tween sensitivity and overall classification performance. Removing the ordinality component eliminates the predicted E/e' regression head; instead, features are projected to the hyperboloid using the norm of Z to determine the radius, following the strategy used in [16]. This slightly decreases balanced accuracy and prototype diversity, and removes the clinically grounded interpretation of the radial dimension, decoupling hyperbolic geometry from severity information. As a result, the embedding space becomes less structured, as shown in Fig. 2 (center). Omitting $\mathcal{L}_{\text{HyperPAS}}$ yields slightly increased balanced accuracy but produces less sparse and more redundant prototypes. HyperPAS primarily acts as a prototype organization regularizer rather than a pure performance enhancer: it promotes directional diversity in the tangent space, yielding sparser and more interpretable prototypes even though classification gains are modest. Compared with the Euclidean Proto+DINO baseline, all hyperbolic variants improve classification metrics. Overall, the proposed formulation attains top-tier predictive performance while consistently promoting sparser and more diverse prototypes.

3.3 Discussion

Failure Analysis. We stratified a subset of 18,677 test videos by whether the LVFP label agreed with the ground-truth E/e' threshold of 14. Concordant cases, where both signals aligned, achieved high accuracy (87.9% for normal with $E/e' < 14$, $n=16,079$; 83.3% for elevated with $E/e' > 14$, $n=2,062$). Discordant cases performed substantially worse (59.7% for elevated with $E/e' < 14$, $n=362$; 31.6% for normal with $E/e' > 14$, $n=174$). These cases likely reflect label noise or patients whose physiology deviates from standard guidelines.

Conclusion. To address the challenges of interpretable LVFP classification from echo, we proposed HypOProto, a hyperbolic, ordinal prototype-based framework built on a frozen, explainable patch-based foundation model. By arranging prototypes along the physiological E/e' scale, the model encodes clinically meaningful ordinal structure, placing borderline cases near the hyperboloid root and clearly normal or elevated cases outward to reflect diagnostic certainty. Additionally, visualizations highlight anatomically relevant regions. Future work will focus on these cases and explore a third, indeterminate class—not present in this dataset—to capture patients who are neither normal nor elevated due to diverse physiological factors, and will further evaluate the framework on public echo datasets as they become more readily available.

Acknowledgments. This research was supported in part by the Canadian Institutes of Health Research (CIHR) and the Natural Sciences and Engineering Research Council of Canada (NSERC), and through computational resources and services provided by Advanced Research Computing at the University of British Columbia.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Akerman, A.P., Porumb, M., Scott, C.G., Beqiri, A., Chatsias, A., Ryu, A.J., Hawkes, W., Huntley, G.D., Arystan, A.Z., Kane, G.C., et al.: Automated echocardiographic detection of heart failure with preserved ejection fraction using artificial intelligence. *JACC: Advances* **2**(6), 100452 (2023)
2. American Society of Echocardiography: Recommendations for the evaluation of left ventricular diastolic function by echocardiography. <https://www.asecho.org/wp-content/uploads/2025/07/Left-Ventricular-Diastolic-Function.pdf> (July 2025), aSE guideline document
3. Desai, K., Nickel, M., Rajpurohit, T., Johnson, J., Vedantam, S.R.: Hyperbolic image-text representations. In: International Conference on Machine Learning. pp. 7694–7731. PMLR (2023)
4. Ghadimi Atigh, M., Keller-Ressel, M., Mettes, P.: Hyperbolic busermann learning with ideal prototypes. *Advances in neural information processing systems* **34**, 103–115 (2021)

5. Ghamary, Y., Wu, V., Vaseli, H., Luong, C., Tsang, T., Bigdeli, S.A., Abolmaesumi, P.: Protoefnet: Dynamic prototype learning for inherently interpretable ejection fraction estimation in echocardiography. In: International Workshop on Interpretability of Machine Intelligence in Medical Image Computing. pp. 149–159. Springer (2025)
6. Gonzalez-Jimenez, A., Lionetti, S., Amruthalingam, L., Gottfrois, P., Gröger, F., Pouly, M., Navarini, A.A.: Is hyperbolic space all you need for medical anomaly detection? In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 312–322. Springer (2025)
7. Gonzalez-Jimenez, A., Lionetti, S., Bazazian, D., Gottfrois, P., Gröger, F., Navarini, A., Pouly, M.: Hyperbolic metric learning for visual outlier detection. In: European Conference on Computer Vision. pp. 327–344. Springer (2024)
8. Guo, Y., Guo, H., Yu, S.X.: Co-sne: Dimensionality reduction and visualization for hyperbolic data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21–30 (2022)
9. Hu, Y., Chen, Y., Xing, X., Zhang, J., Yezhanuly, B.M., Matkerim, B., Xia, Y.: Hyperbolic geometry-driven robustness enhancement for rare skin disease diagnosis. *IEEE Journal of Biomedical and Health Informatics* **29**(3), 2161–2171 (2024)
10. Huang, P., Huang, Y., Zhao, W., He, J., Yu, L.: Hyperpath: Knowledge-guided hyperbolic semantic hierarchy modeling for wsi analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 262–272. Springer (2025)
11. Khrulkov, V., Mirvakhabova, L., Ustinova, E., Oseledets, I., Lempitsky, V.: Hyperbolic image embeddings. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6418–6428 (2020)
12. Law, M., Liao, R., Snell, J., Zemel, R.: Lorentzian distance learning for hyperbolic representations. In: International conference on machine learning. pp. 3672–3681. PMLR (2019)
13. Minners, J., Allgeier, M., Gohlke-Baerwolf, C., Kienzle, R.P., Neumann, F.J., Jander, N.: Inconsistent grading of aortic valve stenosis by current guidelines. *Heart* **96**(18), 1463–1468 (2010)
14. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. *arXiv preprint arXiv:2508.10104* (2025)
15. Vaseli, H., Gu, A.N., Ahmadi Amiri, S.N., Tsang, M.Y., Fung, A., Kondori, N., Saadat, A., Abolmaesumi, P., Tsang, T.S.: Protoasnet: Dynamic prototypes for inherently interpretable and uncertainty-aware aortic stenosis classification in echocardiography. In: International conference on medical image computing and computer-assisted intervention. pp. 368–378. Springer (2023)
16. Vaseli, H., Wu, V., Kim, D., Tsang, M.Y., Gu, A.N., Luong, C., Abolmaesumi, P., Tsang, T.S.: Hiprotonet: Hyperbolic hierarchy-aware part prototypes for aortic stenosis severity classification. In: International Workshop on Advances in Simplifying Medical Ultrasound. pp. 197–207. Springer (2025)
17. Vaseli, H., Wu, V., Kondori, N., To, N.N.M., Fung, A., Gu, A.N., Abolmaesumi, P.: Happi: Hyperbolic hierarchical part prototypes for image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 685–694 (2025)
18. Vukadinovic, M., Chiu, I.M., Tang, X., Yuan, N., Chen, T.Y., Cheng, P., Li, D., Cheng, S., He, B., Ouyang, D.: Comprehensive echocardiogram evaluation with view primed vision language ai. *Nature* pp. 1–3 (2025)