Topology-Aware Learning of Tubular Manifolds via SE(3)-Equivariant Network on Ball B-Spline Curve

¹School of Artificial Intelligence, Beijing Normal University, China ²LILY Research Centre, Nanyang Technological University, Singapore

Abstract

Tubular-like system shape analysis is quite difficult in geometry and topology, while it is widely used in plants and organs analysis in practice. However, traditional discrete representations such as voxels and point clouds often require substantial storage and may lead to the loss of fine-grained geometric and topological details. To address these challenges, we propose SE(3)-BBSCformerGCN, a novel framework for learning shape-aware representations from continuous tubular topological manifolds with equivariance to rotations and translations. Our approach leverages Ball B-Spline Curve (BBSC) to define tubular manifolds and its functional space. We provide a formal mathematical definition and analysis of the resulting manifolds and the BBSC functional space, and incorporate an equivariant mapping that preserves geometric and topological stability. Compared to the point cloud and voxel based representations, our manifold-based formulation significantly reduces data complexity while preserving geometric attributes together with topological features. We validate our method on the branch classification task for Circle of Willis (CoW) on the TopCoW 2024 dataset and the clinical dataset. Our method consistently outperforms voxel and point cloud based baselines in terms of classification performance, generalization ability, convergence speed, and robustness to overfitting.

1 Introduction

Tubular structures play a vital role in a wide range of tasks across computer vision, computer graphics, and interdisciplinary domains, including vascular segmentation [1], classification [2], and traffic network planning [3, 4]. However, the analysis of tubular-like models is quite difficult in both geometry and topology. From a geometric perspective, the branch morphology within tubular-like models varies significantly, with marked differences in tortuosity among individuals for branches of the same designation. From a topological perspective, the morphology and connectivity of the branches are not unique. Certain branches are prone to stenosis, absence, or duplication, and tubular-like model structures frequently exhibit fusion or crossover. Meanwhile, both the geometric structure and topological configuration of tubular-like models, such as vascular networks, may vary across different subjects or even between different scans of the same subject.

The B-Spline model [5, 6, 7], along with its extensions such as Ball B-Spline Curves (BBSC) [8], serves as an industry-standard tool widely adopted in computer graphics and computer-aided design (CAD). In this work, we propose a compact and differentiable representation for tubular structures based on BBSC, using a small set of control points, radii, and knot vectors. Compared to existing voxel-based methods [9] and point cloud-based methods [10], our model not only reduces storage

^{*}First author. Email: 202431081034@mail.bnu.edu.cn

[†]Corresponding author. Email: wangxingce@bnu.edu.cn

overhead, but also enables accurate and interpretable geometric encoding. To enhance interpretability, we further define a smooth manifold structure over BBSC. To mitigate the sensitivity of geometric feature extraction to rigid transformations, we adopt a group-equivariant approach [11] and propose SE(3)-BBSCformer, a model that operates directly on the BBSC manifold and is equivariant under rotations and translations. By incorporating SE(3)-equivariance[12], our method ensures consistent feature extraction under arbitrary rotations and translations, leading to improved stability, accuracy, and generalization in tubular structure analysis.

To model the inter-branch topological connectivity, we use Graph Convolutional Networks (GCNs) [13] as the predominant approach for learning topological features. Building on this foundation, we innovatively propose a tubular manifold representation based on BBSC and construct a topology-aware SE(3)-equivariant network SE(3)-BBSCformerGCN on this manifold. Our main contributions are as follows: 1) We define a compact and interpretable tubular manifold based on BBSC, construct its associated functional space \mathcal{M} , and propose a metric on \mathcal{M} . An SE(3)-equivariant mapping, which enhances geometric consistency and symmetry preservation, is constructed from \mathcal{M} to a higher-dimensional functional space \mathcal{S} , forming the architecture of SE(3)-BBSCformer. 2) We define a topologically glued structure composed of multiple BBSC manifolds and design a corresponding GCN to model the topological relationships between manifolds, which plays a critical role in handling structures with significant topological heterogeneity. 3) We conduct experiments on a multi-class classification task of the Circle of Willis (CoW) [14, 15]. Our proposed SE(3)-BBSCformerGCN demonstrates superior performance against voxel and point cloud state-of-the-art (SOTA) baselines—in terms of accuracy, computational efficiency, and generalization ability—on both a public dataset and a clinical dataset collected in collaboration with a medical institution.

2 Related Work

Tubular structure extraction. Deep learning has been widely applied to tubular structure analysis [16] DSCNet [17] captures local vessel tortuosity via dynamic snake convolution, while DUNet [18] integrates deformable convolutions into UNet [19] for retinal vessel segmentation. DDT [20] employs a distance transform to enable geometry-aware voxel segmentation. Topological constraints have also been proven essential. GraphMorph [21] and TopoLab [22] explicitly model topological connections to improve performance on complex tubular structures. Meanwhile, TaG-Net [23] combines PointNet [24] with GNN [25] to extract both geometric and topological features for more accurate classification. The shape graph [26] models curves as edges and their intersections as nodes, enabling mathematical analysis of graph shape variations. Moreover, it introduces a multi-scale representation to simplify the expression of complex graphs. However, all these methods rely on discrete representations—pixels, voxels, or point clouds. In contrast, we propose a differentiable and compact representation using BBSC, enabling tubular structures to be modeled as manifolds with faithful geometric and topological expressiveness. Specifically, we represent tubular branches as BBSC-based nodes, while the connections between branches are treated as edges. This design avoids the complex process of defining edge weights and instead focuses on capturing the geometric characteristics of individual branches.

Group equivariant network. Group-equivariant models [11, 12] have become central in geometric deep learning [27] due to their robustness and interpretability. Group-equivariant Convolutional Networks (G-CNNs) [28] first introduced group symmetry into CNNs, showing strong performance on symmetry-rich imaging tasks. This idea has since been extended to 3D shape feature extraction, where models like Tensor Field Networks (TFN) [29], SE(3)-Transformer [30] incorporate SE(3)-equivariance into point cloud processing, significantly boosting both performance and stability. EquiTrack [31] achieves equivariance via steerable CNNs to process temporal sequences. SpaER [32] propose the construction of equivariant spatial means using steerable CNN filters and introduced an innovative use of self-attention mechanisms to process temporal sequences. MPerformer [33] employs an SE(3)-Transformer for molecular perception. However, the potential of applying group-equivariant models for geometric feature extraction in tubular structures has not been widely explored. To the best of our knowledge, we are the first to construct a topology-aware and group-equivariant network on manifolds and apply it to the shape analysis and feature extraction of tubular structures.

3 Preliminaries

Ball B-Spline Curves (BBSCs). BBSC is an extension of B-Spline curves to spheres, with each sphere defined by a control radius. This formulation provides a smooth representation for tubular objects with varying thickness. Given a set of control points $\{P_1, P_2, \ldots, P_n\}$, control radii $\{r_1, r_2, \ldots, r_n\}$, and a set of basis functions N(t) which are computed by a knot vector $\{t_1, t_2, \ldots, t_{n+p+1}\}$, a BBSC of degree p is defined as follows:

$$BBSC(t) = \sum_{i=0}^{n} N_{i,p}(t)C_i = \sum_{i=0}^{n} N_{i,p}(t) (P_i ; r_i) = \left(\sum_{i=0}^{n} N_{i,p}(t)P_i ; \sum_{i=0}^{n} N_{i,p}(t)r_i\right).$$
(1)

The basis function of degree 0 is initialized as $N_{i,0} = \begin{cases} 1, & \text{if } t_i \leq t < t_{i+1}, \\ 0, & \text{otherwise.} \end{cases}$, and the higher degree

basis function $N_{i,p}(t)$ can be calculated by iteratively interpolating the knot vector t_i, \ldots, t_{i+p+1} and $N_{i,p-1}(t), N_{i+1,p-1}(t)$ as follows:

$$N_{i,p}(t) = \frac{t - t_i}{t_{i+p} - t_i} N_{i,p-1}(t) + \frac{t_{i+p+1} - t}{t_{i+p+1} - t_{i+1}} N_{i+1,p-1}(t).$$
 (2)

A BBSC consists of two parts: the center B-Spline curve $\gamma(t) = \sum_{i=0}^n N_{i,p}(t) P_i$ and the radius function $\sigma(t) = \sum_{i=0}^n N_{i,p}(t) r_i$. Because $\gamma(t)$ and $\sigma(t)$ are calculated using the same basis function (Equation (1) and Figure 1), they inherently inherit the mathematical properties of B-Spline curves, including differentiability, local control, the convex hull property, and favorable topological characteristics. More importantly, compared to voxels and point clouds, the BBSC representation can continuously and differentiably represent a tubular object with a compact set of parameters, which not only reduces storage and computational complexity, but also enhances theoretical interpretability and enables accurate geometric and analytical modeling.

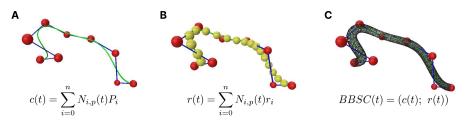


Figure 1: An example of BBSC. **A.** The 3D B-Spline curve is represented as a linear combination of control points (red balls) and their associated basis functions. **B.** The tubular surface is formed by sweeping balls (yellow balls) parameterized by control radii along the B-Spline curve. **C.** The BBSC providing a smooth and flexible curve representation in 3D Euclidean space.

Group Representation. Group [34] is a concept from abstract algebra [35] used to describe symmetries and related properties. A group G consists of a set equipped with a binary operation \circ and satisfies the following properties: closure, associativity and the existence of an identity element and an inverse element. The representation of a group G is a homomorphic mapping [36] ρ : $G \mapsto GL(\mathcal{V}), \mathcal{V} \subseteq \mathbb{R}^D$, which satisfies $\forall g_1, g_2 \in G, \rho(g_1 \circ g_2) = \rho(g_1) \ \rho(g_2). \ GL(\mathcal{V})$ denotes the set of all invertible linear transformations on \mathcal{V} and $\rho(g)$ corresponds to the matrix representation of group element g. The abstract group operation can be realized as the corresponding matrix multiplication.

Group Equivariance. A function $\varphi : \mathcal{X} \to \mathcal{Y}$ is called G-equivariant if $\forall g \in G$, given a set of transformation $\rho_{\mathcal{X}}(g) : \mathcal{X} \to \mathcal{X}$, \exists a transformation $\rho_{\mathcal{Y}}(g) : \mathcal{Y} \to \mathcal{Y}$, such that

$$\varphi(\rho_{\mathcal{X}}(q)x) = \rho_{\mathcal{V}}(q)\varphi(x), \quad x \in \mathcal{X}. \tag{3}$$

Equation (3) can be simplified as $\varphi(\rho_{\mathcal{X}}(g)x) = \varphi(x)$ when the matrix representation of $\rho_{\mathcal{Y}}(g)$ is an identity matrix I. In this case, φ is called invariant. If the function φ is defined in a 3D Euclidean space \mathbb{R}^3 and it is equivariant under both rotation and translation transformations in this space, the function is referred to as 3D roto-translation equivariant (SE(3)-equivariant), the group is thus denoted as SE(3).

SE(3)-Transformer. The SE(3)-Transformer [30] is a deep learning-based method that applies group-equivariance to auxiliary geometric features on 3D point clouds. It constructs a set of SE(3)-equivariant basis kernels $W^{c_oc_i} \in \mathbb{R}^{(2c_o+1)\times(2c_i+1)}$ using spherical harmonics [37], corresponding Clebsch-Gordan coefficients [38], and a learnable radial function. These are then integrated with an attention mechanism [39] to form a group-equivariant network (see details in Appendix B.2). Given an input $f_{in}^{c_i}$ with c_i channels, the output $f_{out}^{c_o}$ with c_o channels is obtained as follows:

$$f_{out,i}^{c_o} = W_V^{c_o c_o} f_{in,i}^{c_o} + \sum_k \sum_{j \in \mathcal{N}_i} \alpha_{i,j} W_V^{c_o c_i} (x_i - x_j) f_{in,j}^{c_i}, \tag{4}$$

where the attention matrix is computed as follows:

$$\vec{q_i} = \bigoplus_{c_o} \sum_{c_i} W_Q^{c_o c_i} f_{in,i}^{c_i}, \vec{k}_{i,j} = \bigoplus_{c_o} \sum_{c_i} W_K^{c_o c_i} (x_i - x_j) f_{in,j}^{c_i}, \alpha_{ij} = \frac{\exp(\vec{q_i}^T \vec{k}_{i,j})}{\sum_{j' \in \mathcal{N}_i} \exp(\vec{q_i}^T \vec{k}_{i,j})}.$$
(5)

The direct sum \bigoplus represents the concatenation of vectors. The attention mechanism is proven to be SE(3)-equivariant [30], as it relies on the inner product of α_{ij} and $W_V^{c_oc_i}(x_i-x_j)$, which is an SE(3)-equivariant operation.

4 Method and Analysis

Our method consists of two main components. First, we construct smooth BBSCs from discrete 3D data and, from a manifold perspective, define and analyze both their properties and the space they form. Second, we extend the SE(3)-Transformer to this manifold, yielding the SE(3)-BBSCformer for learning on continuous structures. We further employ a GCN to capture topological relations and design a pipeline (Figure 2) for branch classification in the CoW, a tubular arterial network with specific topological connections.

4.1 BBSC Functional Space and SE(3) Equivariant Mapping

Definition 1. We define the p-degree BBSC function space \mathcal{M}_p with given P, r, τ from their respective vector spaces as follows:

$$\mathcal{M}_p = \{ \beta(P, r, \tau) \mid (P, r, \tau) \in \Theta \}, \quad \Theta = \{ P \in \mathbb{R}^{3 \times n}, \, r \in \mathbb{R}_{>0}, \, \tau \in \mathbb{T} \}.$$
 (6)

where $\mathbb{T}=\{t_i\in[0,1]|t_1\leq t_2\leq\cdots\leq t_{n+p+1}\}\in[0,1]^{n+p+1},$ and $\beta(P,r,\tau)$ is the total structure of BBSC (see Equation (1)), which is a two-dimensional manifold. In order to further stabilize BBSC, we impose certain constraints on $\mathbb{T}=\{t_i\in[0,1]|t_{p+1}< t_{p+2}<\cdots< t_n,t_1=\cdots=t_{p+1}=0,t_n=\cdots=t_{n+p+1}=1\}.$ The first and last p+1 entries of the knot vector are set to 0 and 1, respectively, to ensure that the two endpoints of the curve coincide with the first and last control points. The parameter t in Equation (1) specifies positions along the curve but does not affect the overall BBSC shape, serving only as an internal parameter independent of $\beta\in\mathcal{M}_p$. Because $\beta(P,r,\tau)$ is a linear combination of P and r, it has infinitely many continuous partial derivatives. By excluding the first and last p+1 entries of τ , i.e., in the open subset of \mathcal{M}_p that omits two boundary points, $N_{i,p}(t)$ becomes a recursive fractional linear function of τ , and is differentiable with respect to τ of arbitrary order. Moreover, under the constraint that the middle entries of τ are strictly monotonically increasing, $\beta(P,r,\tau)$ possesses infinitely many continuous partial derivatives. According to Equation (1), P,r, and τ are in a separable product form. Consequently, their higher-order mixed partial derivatives can be expressed as products of higher-order partial derivatives, which remain continuous. Therefore, \mathcal{M}_p can be regarded as $C^\infty(\Theta,\mathcal{M}_p)$. The complete proof is provided in Appendix C.

The metric on \mathcal{M}_p can be simply defined as a weighted sum of the L2 norms of the parameters as follows:

$$g = \|P - \tilde{P}\|_2 + \alpha \|r - \tilde{r}\|_2 + \eta \|\tau - \tilde{\tau}\|_2. \tag{7}$$

Although the metric based on L2 norms provides a simple way to measure parameter variations, it ignores the perturbations on \mathcal{M}_p . Moreover, $\gamma(t)$ and $\sigma(t)$, which determine the BBSC shape, are independent (hence P and r are independent), whereas τ simultaneously influences both. Thus, the coupling between τ and P, r should be considered when defining the metric. To address this, we

propose a metric that accounts for perturbations on \mathcal{M}_p and parameter coupling. Concretely, we first characterize perturbations in the tangent space of \mathcal{M}_p as $\delta \gamma$ and $\delta \sigma$ as follows:

$$\delta \gamma = \sum_{i=1}^{n} N_{i,p}(t) \delta P_i + \sum_{i=1}^{n} P_i \frac{\partial N_{i,p}(t)}{\partial \tau} \delta \tau, \ \delta \sigma = \sum_{i=1}^{n} N_{i,p}(t) \delta r_i + \sum_{i=1}^{n} r_i \frac{\partial N_{i,p}(t)}{\partial \tau} \delta \tau.$$
 (8)

Given geometric perturbations $\delta\beta=(\delta\gamma,\delta\sigma)$ and $\tilde{\delta}\beta=(\tilde{\delta}\gamma,\tilde{\delta}\sigma)$ in the tangent space of \mathcal{M}_p , we define the following inner product as follows:

$$\langle \delta \beta, \tilde{\delta} \beta \rangle_{\mathcal{M}_p} = \int_0^1 \left(\langle \delta \gamma, \tilde{\delta} \gamma \rangle + \eta_{\gamma} \langle \frac{\partial \delta \gamma}{\partial t}, \frac{\partial \tilde{\delta} \gamma}{\partial t} \rangle + \alpha \delta \sigma \tilde{\delta} \sigma + \alpha \eta_{\sigma} \frac{\partial \delta \sigma}{\partial t} \frac{\partial \tilde{\delta} \sigma}{\partial t} \right) \omega(t) dt, \tag{9}$$

where $\alpha \geq 0$, η_{γ} , $\eta_{\sigma} \geq 0$, and $\omega(t)$ denotes a weighting function. One may simply set $\omega(t)=1$ to perform unweighted integration. However, we recommend using $\omega(t)=\|\gamma'(t)\|$, which corresponds to arc-length weighting, thereby mitigating the effect of reparameterization and enhancing geometric interpretability. The proposed inner product is inspired by the Sobolev H^1 inner product [40, 41], which incorporates not only the Euclidean inner product of $(\delta\gamma,\delta\sigma)$ and $(\tilde{\delta}\gamma,\tilde{\delta}\sigma)$ themselves, but also their first-order derivatives. Because $\langle\cdot,\cdot\rangle$ denotes the Euclidean inner product, it inherently possesses rotation and translation invariance. Consequently, $\langle\cdot,\cdot\rangle_{\mathcal{M}_p}$ inherits these desirable geometric properties. Moreover, this inner product preserves the separability of parameters P and P0, while simultaneously coupling them with P1 through Equation (8). Let P1 and P3 and P3 and P4 are P5. Denote the Jacobian of P6 by P6 by P7 benote the pullback metric can be finally written as follows:

$$g = \langle \delta\theta, \tilde{\delta}\theta \rangle_{\Theta} = \langle \delta\beta, \tilde{\delta}\beta \rangle_{\mathcal{M}_p} = \langle D\beta(\theta)\delta\theta, D\beta(\theta)\tilde{\delta}\theta \rangle_{\mathcal{M}_p}$$
 (10)

Similarly, we can define the entire BBSC functional space \mathcal{M} as follows.

Definition 2. \mathcal{M} is called a BBSC functional space if it contains all existing \mathcal{M}_p , as denoted as $\mathcal{M} = \bigcup_{n \geq 0} \mathcal{M}_p$.

Definition 3. A mapping $\varphi : \mathcal{M}_p \to \mathcal{S}$ is deemed SE(3)-equivariant on \mathcal{M}_p , if $\forall \beta(P, r, \tau) \in \mathcal{M}_p$ and \forall transformations $g \in SE(3)$, the following equation holds:

$$\varphi(\beta(\rho_{\mathcal{M}_p}(g)P, r, \tau)) = \rho_{\mathcal{S}}(g)\varphi(\beta(P, r, \tau)). \tag{11}$$

If φ is infinitely differentiable with respect to P, r, and τ , φ is a C^{∞} mapping on \mathcal{M}_p , meaning it is smooth as well.

Definition 4. \mathcal{GM} is a manifold with topological relations constructed from the direct sum of disjoint BBSC manifolds $\beta_1, \beta_2, \dots \beta_k \in \mathcal{M}$, which is defined as follows:

$$\mathcal{GM} = \left(\bigsqcup_{i \in I} \beta_i\right) / \sim, \sim \subseteq \bigsqcup_{i \in I} \beta_i \times \bigsqcup_{i \in I} \beta_i. \tag{12}$$

Specifically, \mathcal{GM} is a topological manifold composed of multiple \mathcal{M} with certain topological relationships between them. \mathcal{GM} can be viewed as a manifold with a gluing structure [42, 43, 44].

4.2 BBSC Manifold Construction

We follow [45] and adopt a similar BBSC fitting algorithm. Given a set of points $V \in \mathbb{R}^{N \times 3}$ from tubular objects, we apply DiffusionNet [46] (details in Appendix B.1) to embed them and output a BBSC, represented by control points $P = \{P_1, P_2, \dots, P_n\}$, control radii $r = \{r_1, r_2, \dots, r_n\}$, and a knot vector $\tau = \{t_1, t_2, \dots, t_n\}$. To ensure $t_{p+2} < t_{p+3} < \dots < t_{n-1}$ and $t_1 = \dots = t_{p+1} = 0$, $t_n = \dots = t_{n+p+1} = 1$, we predict the differences of the inner knot vector $\nabla \tau$ using a softmax function, and then obtain the full knot vector by applying a cumulative sum and concatenating the boundary zeros and ones. Considering all possible control parameter sets $P \in \mathbb{R}^{3 \times n}$ and $r \in \mathbb{R}^n$, together with constrained knot vectors τ , we construct the functional space \mathcal{M}_p endowed with a smooth mapping $\beta \in C^\infty(\Theta, \mathcal{M}_p)$.

We introduce an additional template spline $\overline{\beta}$ to constrain the shape of the BBSC centerline. We minimize a loss function composed of two terms, the Hausdorff distance D_H [47] between sampled

points on the constructed BBSC β and the raw surface points, and the metric between β and $\overline{\beta}$. Note that the Hausdorff distance can be replaced by the Chamfer distance [48], and a regularization term $\int_0^1 \kappa^2(t) \, dt$, which can be used to enforce the global smoothness of the BBSC centerline and can also be incorporated into the loss function. Based on these, we further incorporate topological relationships among the β to construct the topological manifold \mathcal{GM} .

4.3 SE(3)-BBSCformer and Graph Convolutional Networks

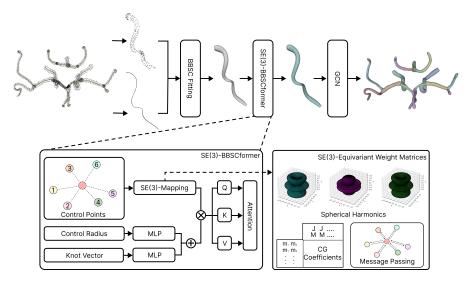


Figure 2: The SE(3)-BBSCformerGCN takes vascular centerlines and surface point clouds as inputs. A fitting module is first employed to construct BBSC. These spline-based manifold representations are then mapped into a high-dimensional manifold space via the SE(3)-BBSCformer. Finally, a GCN is utilized to process the topological relationships among the manifolds.

SE(3)-Transformer can be viewed as a mapping operating in a discrete space represented by a point cloud. In contrast, the BBSC model is a smooth and continuous representation. To integrate the advantages of both, we reconstruct the model and term it SE(3)-BBSCformer, which can be interpreted as an SE(3) mapping defined on \mathcal{M} . First, we compute the centroid of the control points $\overline{P} = \frac{\sum_{i=1}^n P_i}{n}$ to serve as the virtual anchoring point, acting as a central hub for transmitting information between the control points and generating the unit vector $\frac{P_i - \overline{P}}{\|P_i - \overline{P}\|}$ needed by SE(3)-BBSCformer. We initialize the c_{i_0} channel feature $f_{in}^{c_{i_0}}$ as a combination of vectors generated from r and τ as follows:

$$f_{in}^{c_{i_0}} = (MLP(r), MLP(\tilde{\tau})), \ \tilde{\tau} = ((t_1, \dots, t_{p+1}), \dots, (t_n, \dots t_{n+p+1}) \in \mathbb{R}^{n \times (p+1)}.$$

Then we view each control point P_i as being connected with \overline{P} and replace the x_i and x_j with P_i and \overline{P} in Equation (4) and (5). The resulting SE(3)-BBSCformer is defined as follows:

$$f_{out,i}^{c_o} = W_V^{c_o c_o} f_{in,i}^{c_o} + \sum_k \alpha_{i,j} W_V^{c_o c_i} (P_i - \overline{P}) f_{in,j}^{c_i}, \tag{13}$$

where the attention matrix is computed as follows:

$$\vec{q}_i = \bigoplus_{c_o} \sum_{c_i} W_Q^{c_o c_i} f_{in,i}^{c_i}, \ \vec{k}_j = \bigoplus_{c_o} \sum_{c_i} W_K^{c_o c_i} (P_j - \overline{P}) f_{in,j}^{c_i}, \ \alpha_{ij} = \frac{\exp(\vec{q}_i^T \vec{k}_j)}{\sum_{j=1}^n \exp(\vec{q}_i^T \vec{k}_j)}.$$
(14)

For convenience, we abbreviate SE(3)-BBSCformer as the mapping $\varphi:\mathcal{M}_p\to\mathcal{S}$, where \mathcal{S} denotes the high-dimensional functional space. Variable $\varphi(\beta)$ can be interpreted as the high-dimensional representation of BBSC in \mathcal{S} . If an object is a manifold with topological relationships, composed of multiple β components as described in **Definition 4**, then GCN can be applied to capture the topological relationships. Each BBSC manifold $\beta\in\mathcal{GM}$ is treated as a node in the graph G and

 $\varphi(\beta)$ serves as the initial node feature v. For simplicity, we initialize the edge attribute e as the difference between features of the connected node features. Then, the GCN is formulated as follows:

$$e_{i,j} = MLP([v_i, v_j, e_{i,j}]), \ v_i = MLP\left(\left[v_i, \frac{1}{\|j\|} \sum_{j \in \mathcal{N}_j} MLP([v_j, e_{i,j}])\right]\right). \tag{15}$$

Taking the CoW as an input example, we propose the SE(3)-BBSCformerGCN. The pipeline of our model is shown in Figure 2.

4.4 Theoretical Analysis and Discussion of Advantages

The BBSC mainfold representation β is highly compact, allowing complex tubular shapes to be described using only a small number of control parameters and knot vector, thereby substantially reducing storage and computational costs. Moreover, the BBSC manifold exhibits strong controllability because its shape can be easily adjusted by modifying the control parameters and the knot vector. Due to the smoothness of β , geometric and differential quantities such as normals, curvature, and torsion can be accurately calculated at any position in the object, without relying on numerical approximations that are prone to errors when estimated from discrete data [49, 50]. This property endows the representation with stronger expressive power and improved theoretical interpretability. Moreover, β can be shown to admit continuous higher-order partial and mixed derivatives for all parameters, i.e., $\beta \in C^{\infty}$. This ensures that $\mathcal M$ forms a continuous and infinitely differentiable functional space. Benefiting from these properties, the BBSC manifold representation is particularly well-suited for modeling complex geometric tubular structures.

The SE(3)-equivariant mapping $\varphi:\mathcal{M}\to\mathcal{S}$ can be shown to be a C^∞ mapping. Spherical harmonics are also C^∞ [51], so the differentiability of φ primarily depends on that of the MLP. In turn, the MLP's differentiability is determined by the choice of activation function. In this work, all MLPs are constructed using GELU activations and linear layers, both of which are C^∞ mappings. Consequently, φ is a C^∞ mapping from $\mathcal M$ to $\mathcal S$ that satisfies SE(3) equivariance. The group-equivariant φ preserves the intrinsic symmetries of $\mathcal M$ when mapping to $\mathcal S$, ensuring the stability of its shape features. Moreover, the GCN designed to capture topological relationships maintains the overall topological stability of the manifold space, thereby enhancing the model's generalization and robustness while remaining lightweight.

5 Experiments and Analysis

We evaluate the performance of our proposed model on the branch classification task for the CoW, a critical arterial structure located at the base of the human brain. Experiments are conducted on two datasets: the TopCoW dataset from the MICCAI 2024 Challenge [52], which contains 125 samples across 13 anatomical classes, and a clinical dataset collected from a collaborating medical institution, comprising 1,182 samples across 22 classes. We focus on the CoW branch classification task because it is a complex tubular structure characterized by both geometric and topological properties, and necessitate efficient computational approaches. Moreover, CoW exhibits a highly bilateral symmetry while simultaneously presenting complex topological heterogeneity, making group equivariant networks and GCNs particularly well-suited for modeling such structures. (see detailed descriptions of the CoW datasets, preprocess and training protocols in D)

5.1 Experiments on TopCoW 2024 MICCAI Challenge

We evaluate our model on the publicly available TopCoW dataset [52]. We include several SOTA voxel-based and point cloud-based classification methods as baselines. In addition, we conduct ablation studies to investigate the effects of incorporating SE(3)-equivariant BBSC mapping and the use of GCN for modeling topological relationships. Given the small sample size of the TopCoW 2024 dataset, we employ 5-fold cross-validation to enhance training robustness. We calculate the mean and standard deviation, highlighting the highest mean and lowest standard deviation in bold. Our model achieves the highest mean and lowest variance across AUC-ROC, Recall, and F1 score, while also demonstrating competitive performance in Accuracy and Precision (see results in Table 1).

We illustrates the class-wise prediction accuracy on the CoW dataset (see Figure 3). Across the majority of labels, our model consistently outperforms baseline methods. Due to the highly heterogeneous

Table 1: Performance comparison of different methods on the TopCoW 2024 dataset

Method	Accuracy%	AUC-ROC%	Precision %	Recall%	F1 Score%
3D DenseNet [53]	97.02 ± 0.57	99.63 ± 1.94	95.86 ± 0.51	95.42 ± 0.41	97.47 ± 0.48
3D ResNet [54]	96.54 ± 0.54	99.61 ± 0.11	96.79 ± 0.76	96.44 ± 0.83	96.52 ± 0.80
PointNet [24]	94.10 ± 0.88	99.46 ± 0.20	94.81 ± 1.36	94.69 ± 0.92	94.57 ± 1.10
PointNet++ [55]	96.61 ± 0.53	99.61 ± 0.09	$96.72 \pm \textbf{0.49}$	96.52 ± 0.55	96.45 ± 0.60
CurveNet [56]	97.14 ± 1.16	99.72 ± 0.20	97.31 ± 1.16	97.07 ± 1.19	97.09 ± 1.22
RepSurf-U [57]	97.41 ± 0.68	99.86 ± 0.07	96.83 ± 0.52	96.61 ± 0.54	96.50 ± 0.55
PointMLP [58]	94.10 ± 2.27	99.03 ± 0.56	94.41 ± 2.35	94.05 ± 2.31	93.94 ± 2.40
SE(3)-Transformer [30]	61.35 ± 5.80	90.93 ± 1.31	62.36 ± 5.28	60.77 ± 5.70	60.58 ± 5.69
SE(3)-BBSCformer	87.91 ± 1.35	98.24 ± 0.31	88.97 ± 1.66	88.01 ± 1.32	87.98 ± 1.26
BBSCformerGCN	97.95 ± 1.22	99.89 ± 0.02	97.25 ± 0.79	98.34 ± 0.60	97.78 ± 0.70
SE(3)-BBSCformerGCN	97.53 ± 0.83	$\textbf{99.99} \pm \textbf{0.01}$	97.36 ± 0.53	$\textbf{98.44} \pm \textbf{0.31}$	$\textbf{97.87} \pm \textbf{0.41}$

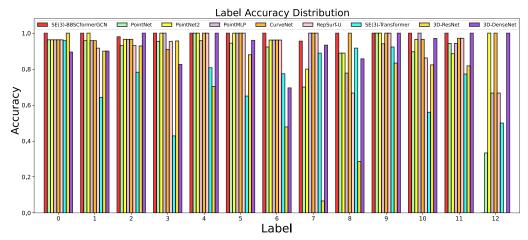


Figure 3: Per-class prediction accuracy of various models evaluated on TopCoW 2024 dataset.

anatomical structure of the CoW [59, 60, 61], the distribution of different branches label is notably imbalanced. For example, in the TopCoW 2024dataset, only 5% of the 125 samples contain the segment corresponding to Label 12 (see details in Appendix D.1). This extreme scarcity introduces a high degree of randomness during testing. However, we observe that this issue is effectively mitigated by increasing the total number of training samples, as further demonstrated in Section 5.2.

5.2 Experiments on Real-World Clinical Data

Table 2: Performance comparison of different methods on the clinical dataset

Method	Accuracy%	AUC-ROC%	Precision%	Recall%	F1 Score%
PointNet [24]	69.20 ± 0.98	96.50 ± 0.24	70.02 ± 0.77	69.17 ± 0.96	68.90 ± 1.02
PointNet++ [55]	70.27 ± 0.16	96.60 ± 0.17	70.82 ± 0.43	70.28 ± 0.20	70.05 ± 0.14
CurveNet [56]	81.91 ± 0.14	93.80 ± 1.13	62.92 ± 3.14	60.13 ± 4.25	58.51 ± 4.58
RepSurf-U [57]	78.22 ± 0.17	98.40 ± 0.11	78.34 ± 0.28	77.99 ± 0.29	78.00 ± 0.25
PointMLP [58]	82.70 ± 0.10	98.66 ± 0.06	82.86 ± 0.14	82.70 ± 0.10	82.67 ± 0.14
SE(3)-Transformer [30]	61.96 ± 1.45	92.32 ± 0.96	61.05 ± 2.07	61.94 ± 1.47	60.77 ± 2.57
SE(3)-BBSCformer	71.06 ± 1.39	93.25 ± 0.60	70.25 ± 1.16	71.04 ± 1.39	70.45 ± 1.21
BBSCformerGCN	95.45 ± 1.09	99.34 ± 0.21	95.10 ± 1.48	95.38 ± 1.13	95.23 ± 1.29
SE(3)-BBSCformerGCN	96.11 \pm 1.01	99.38 \pm 0.19	96.11 \pm 0.93	96.08 ± 0.97	96.02 ± 1.05

We further validate the generalization ability of our model on clinical data collected from the collaborating medical institution. This dataset contains approximately ten times more number of samples than the TopCoW 2024 MICCAI Challenge dataset. Although the inherent heterogeneity

of CoW still leads to imbalanced class distributions, the substantially larger sample size helps alleviate the randomness in predictions for underrepresented classes. The results presented in Table 2 demonstrate that our model exhibits stronger generalization ability compared to voxel-based and point-based approaches. While other methods experience a performance drop of 10–20%, our model maintains a consistently high level of evaluation metrics. Notably, SE(3)-BBSCformerGCN outperforms BBSCformerGCN in both mean and standard deviation across all evaluation metrics, further validating the stability and generalization advantages introduced by the SE(3)-equivariant mapping. Additionally, the comparison between SE(3)-BBSCformerGCN and SE(3)-BBSCformer highlights the critical role of topological information. Collectively, these findings suggest that both geometric and topological stability are the key factors underpinning the generalization capability of our model. Similarly, we provide the per-class prediction accuracy across all label categories (Figure 4). As shown, our model consistently outperforms all baselines across all 22 classes.

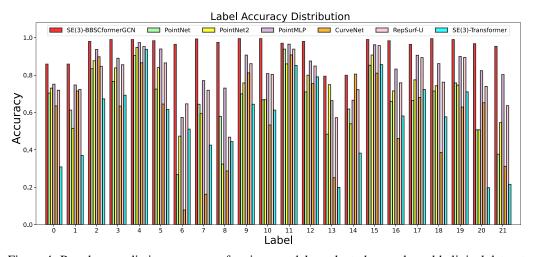


Figure 4: Per-class prediction accuracy of various models evaluated on real-world clinical dataset.

5.3 Training Efficiency, Stability, and Computational Cost Analysis

We illustrates the training and testing accuracy curves on the clinical datasets (see Figure 5). Our method consistently achieves faster convergence and greater stability across all methods. These advantages are largely attributed to the compact input representation, which preserves fine-grained geometric details without the need for aggressive downsampling. More critically, while other models exhibit severe overfitting, both SE(3)-BBSCformer and the SE(3)-Transformer demonstrate strong resistance to overfitting, highlighting the robustness and generalization power of equivariant architectures. Moreover, leveraging BBSCs to capture complex geometric features and topological structures of tubular systems further enhances our model's performance over existing baselines.

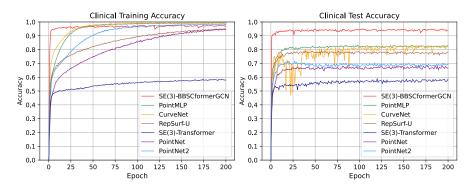


Figure 5: Epoch-wise training and testing accuracy evaluated on the clinical dataset.

To validate the efficiency of BBSC for tubular object representation, we compared our model against several baselines on evaluation time, FLOPs, and parameter count (see Table 3). SE(3)-

BBSCformerGCN achieves superior generalization and classification performance while significantly reducing inference time and memory usage. These results underscore the practical benefits of our design, particularly for fast and cost-effective deployment in clinical settings.

Table 3: Parameters amount and computing cost comparison of different methods

Method	Evaluation Time(ms)	FLOPs(M)	Parameters(M)
3D DenseNet [53]	8.03	9743.76	18.56
3D ResNet [54]	10.95	11463.08	85.23
PointNet [24]	0.88	450.38	3.46
PointNet++ [55]	1.54	4067.53	1.74
CurveNet [56]	81.95	269.56	2.13
RepSurf-U [57]	59.84	911.32	1.48
PointMLP [58]	38.39	15733.95	13.23
SE(3)-Transformer [30]	19.03	456.33	0.12
SE(3)-BBSCformer	1.92	1.72	0.35
BBSCformerGCN	1.14	88.98	3.60
SE(3)-BBSCformerGCN	3.02	59.01	2.70

6 Discussion and Future Work

We leverage an extension of B-spline curves to the spherical domain, which constructs the BBSC, to model tubular structures. We formulate a mathematically well-defined smooth manifold based on this representation, construct a functional space for tubular manifold, and discuss the associated metric on this space. Furthermore, we propose a rotation and translation equivariant mapping, which is applied to the anatomical classification of the CoW. While our experiments focus on this clinically significant cerebrovascular structure, the proposed framework is readily applicable to a wide range of other tubular structures. Additionally, the BBSC structure is naturally extensible. The control vectors can be lifted into an n-dimensional space, where the B-spline basis functions can be reused to construct smooth and high-dimensional manifolds. Notably, SplineCNN [62] has previously utilized this insight to define smooth convolutional filters via B-Spline kernels. Lastly, our model maps $\mathcal M$ defined by BBSC into a high-dimensional space $\mathcal S$, preserving SE(3)-equivariance of the 3D parameters P. However, a limitation of our current model is that the graph convolution operations used to model topological relations on $\mathcal S$ are not equivariant, which could compromise the overall stability and full equivariance of the model. It is worth investigating into the design of equivariant mappings in high-dimensional topological manifold spaces as a line of future work.

7 Conclusion

We propose SE(3)-BBSCformerGCN, a novel deep learning architecture that achieves group equivariance on tubular manifolds. Our model takes as input a BBSC manifold, parameterized by a small set of 4D control parameters and knot vectors. This representation not only enables more effective extraction of complex geometric and topological features from tubular structures, but also reduces computational complexity. This design offers strong mathematical interpretability and improved computational efficiency. We validate our model on a clinically important yet computationally underexplored task: classifying anatomical configurations of the CoW, and demonstrate its superior performance in classification, generalization, and resistance to overfitting. Our code will be available on https://github.com/niuyixuan/SE-3-BBSCformerGCN.

8 Acknowledgments

This research is partially supported by the National Natural Science Foundation of China (Grant No. 62072045), the Beijing Municipal Science and Technology Commission and the Zhongguancun Science Park Management Committee (Grant No. Z221100002722020), the Natural Science Foundation of Beijing (Grant No. 7242167), and the Teaching Reform Project of Beijing Normal University. We would also like to express our sincere gratitude to Peking Union Medical College Hospital (PUMCH) for kindly providing the clinical CoW scan data used in this study.

References

- [1] Joes Staal, Michael D Abràmoff, Meindert Niemeijer, Max A Viergever, and Bram Van Ginneken. Ridge-based vessel segmentation in color images of the retina. *IEEE transactions on medical imaging*, 23(4):501–509, 2004.
- [2] Volodymyr Mnih. Machine learning for aerial image labeling. University of Toronto (Canada), 2013.
- [3] Volodymyr Mnih and Geoffrey E Hinton. Learning to detect roads in high-resolution aerial images. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part VI 11*, pages 210–223. Springer, 2010.
- [4] Luisa Costa Sousa, Catarina F Castro, Carlos Conceição António, and Rui Chaves. Blood flow simulation and vascular reconstruction. *Journal of biomechanics*, 45(15):2549–2555, 2012.
- [5] Carl De Boor. A practical guide to splines, volume 27. springer New York, 1978.
- [6] Carl De Boor. On calculating with b-splines. *Journal of Approximation theory*, 6(1):50–62, 1972.
- [7] Michael Unser, Akram Aldroubi, and Murray Eden. B-spline signal processing. i. theory. *IEEE transactions on signal processing*, 41(2):821–833, 1993.
- [8] Zhongke Wu, Hock Soon Seah, and Mingquan Zhou. Skeleton based parametric solid models: Ball b-spline curves. In 2007 10th IEEE International Conference on computer-aided design and computer graphics, pages 421–424. IEEE, 2007.
- [9] Rafic Nader, Romain Bourcier, and Florent Autrusseau. Using deep learning for an automatic detection and classification of the vascular bifurcations along the circle of willis. *Medical image analysis*, 89:102919, 2023.
- [10] Dong Wang, Zhao Zhang, Ziwei Zhao, Yuhang Liu, Yihong Chen, and Liwei Wang. Pointscatter: Point set representation for tubular structure extraction. In *European conference on computer vision*, pages 366–383. Springer, 2022.
- [11] Jan E Gerken, Jimmy Aronsson, Oscar Carlsson, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and Daniel Persson. Geometric deep learning and equivariant neural networks. *Artificial Intelligence Review*, 56(12):14605–14662, 2023.
- [12] Carlos Esteves. Theoretical aspects of group equivariant neural networks. *arXiv preprint arXiv:2004.05154*, 2020
- [13] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [14] Bernard J Alpers, Richard G Berry, and Richard M Paddison. Anatomical studies of the circle of willis in normal brain. AMA Archives of Neurology & Psychiatry, 81(4):409–418, 1959.
- [15] Sumaiya Iqbal. A comprehensive study of the anatomical variations of the circle of willis in adult human brains. *Journal of clinical and diagnostic research: JCDR*, 7(11):2423, 2013.
- [16] Jie Mei, Rou-Jing Li, Wang Gao, and Ming-Ming Cheng. Coanet: Connectivity attention network for road extraction from satellite imagery. *IEEE Transactions on Image Processing*, 30:8540–8552, 2021.
- [17] Yaolei Qi, Yuting He, Xiaoming Qi, Yuan Zhang, and Guanyu Yang. Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6070–6079, 2023.
- [18] Qiangguo Jin, Zhaopeng Meng, Tuan D Pham, Qi Chen, Leyi Wei, and Ran Su. Dunet: A deformable network for retinal vessel segmentation. *Knowledge-Based Systems*, 178:149–162, 2019.
- [19] 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.
- [20] Yan Wang, Xu Wei, Fengze Liu, Jieneng Chen, Yuyin Zhou, Wei Shen, Elliot K Fishman, and Alan L Yuille. Deep distance transform for tubular structure segmentation in ct scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3833–3842, 2020.

- [21] Zhao Zhang, Ziwei Zhao, Dong Wang, and Liwei Wang. Graphmorph: Tubular structure extraction by morphing predicted graphs. Advances in Neural Information Processing Systems, 37:68472–68499, 2024.
- [22] Zhixing Zhang, Ziwei Zhao, Dong Wang, Shishuang Zhao, Yuhang Liu, Jia Liu, and Liwei Wang. Topology-preserving automatic labeling of coronary arteries via anatomy-aware connection classifier. In *International conference on medical image computing and computer-assisted intervention*, pages 759–769. Springer, 2023.
- [23] Linlin Yao, Feng Shi, Sheng Wang, Xiao Zhang, Zhong Xue, Xiaohuan Cao, Yiqiang Zhan, Lizhou Chen, Yuntian Chen, Bin Song, et al. Tag-net: topology-aware graph network for centerline-based vessel labeling. IEEE transactions on medical imaging, 42(11):3155–3166, 2023.
- [24] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [25] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
- [26] Aditi Basu Bal, Xiaoyang Guo, Tom Needham, and Anuj Srivastava. Statistical analysis of complex shape graphs. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(12):8788–8805, 2024.
- [27] Jiaqi Han, Jiacheng Cen, Liming Wu, Zongzhao Li, Xiangzhe Kong, Rui Jiao, Ziyang Yu, Tingyang Xu, Fandi Wu, Zihe Wang, et al. A survey of geometric graph neural networks: Data structures, models and applications. *arXiv preprint arXiv:2403.00485*, 2024.
- [28] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.
- [29] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. arXiv preprint arXiv:1802.08219, 2018.
- [30] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. Advances in neural information processing systems, 33:1970–1981, 2020.
- [31] Benjamin Billot, Neel Dey, Daniel Moyer, Malte Hoffmann, Esra Abaci Turk, Borjan Gagoski, P Ellen Grant, and Polina Golland. Se (3)-equivariant and noise-invariant 3d rigid motion tracking in brain mri. *IEEE transactions on medical imaging*, 43(11):4029–4040, 2024.
- [32] Jian Wang, Razieh Faghihpirayesh, Polina Golland, and Ali Gholipour. Spaer: Learning spatio-temporal equivariant representations for fetal brain motion tracking. In *International Workshop on Preterm, Perinatal* and Paediatric Image Analysis, pages 3–13. Springer, 2024.
- [33] Fanmeng Wang, Hongteng Xu, Xi Chen, Shuqi Lu, Yuqing Deng, and Wenbing Huang. Mperformer: An se (3) transformer-based molecular perceptron. In *Proceedings of the 32nd ACM International Conference* on Information and Knowledge Management, pages 2512–2522, 2023.
- [34] William Raymond Scott. Group theory. Courier Corporation, 2012.
- [35] Joseph Gallian. Contemporary abstract algebra. Chapman and Hall/CRC, 2021.
- [36] Ryszard Raczka and Asim Orhan Barut. Theory of group representations and applications. World Scientific Publishing Company, 1986.
- [37] Claus Müller. Spherical harmonics, volume 17. Springer, 2006.
- [38] Johan J de Swart. The octet model and its clebsch-gordan coefficients. In *The Eightfold Way*, pages 120–143. CRC Press, 2018.
- [39] 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.
- [40] Robert A Adams and John JF Fournier. Sobolev spaces, volume 140. Elsevier, 2003.
- [41] Haim Brezis and Haim Brézis. Functional analysis, Sobolev spaces and partial differential equations, volume 2. Springer, 2011.

- [42] Victor Guillemin and Alan Pollack. Differential topology, volume 370. American Mathematical Soc., 2010.
- [43] James R Munkres. Elements of algebraic topology. CRC press, 2018.
- [44] Allen Hatcher. Algebraic topology., 2005.
- [45] Zhongke Wu, Xingce Wang, Yan Fu, Junchen Shen, Qianqian Jiang, Yuanshuai Zhu, and Mingquan Zhou. Fitting scattered data points with ball b-spline curves using particle swarm optimization. *Computers & graphics*, 72:1–11, 2018.
- [46] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)*, 41(3):1–16, 2022.
- [47] Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. IEEE Transactions on pattern analysis and machine intelligence, 15(9):850–863, 1993.
- [48] Harry G Barrow, Jay M Tenenbaum, Robert C Bolles, and Helen C Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. Technical report, 1977.
- [49] Weihuan Chen, Shiing-shen Chern, and Kai S Lam. Lectures on differential geometry, volume 1. World Scientific Publishing Company, 1999.
- [50] Richard L Bishop and Richard J Crittenden. *Geometry of manifolds: Geometry of Manifolds*, volume 15. Academic press, 2011.
- [51] George B. Arfken, Hans J. Weber, and Frank E. Harris. Mathematical Methods for Physicists: A Comprehensive Guide. Academic Press, 2011.
- [52] Kaiyuan Yang, Fabio Musio, Yihui Ma, Norman Juchler, Johannes C Paetzold, Rami Al-Maskari, Luciano Höher, Hongwei Bran Li, Ibrahim Ethem Hamamci, Anjany Sekuboyina, et al. Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for cta and mra. *ArXiv*, pages arXiv–2312, 2024.
- [53] Sen Liang, Rongguo Zhang, Dayang Liang, Tianci Song, Tao Ai, Chen Xia, Liming Xia, and Yan Wang. Multimodal 3d densenet for idh genotype prediction in gliomas. *Genes*, 9(8):382, 2018.
- [54] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In proceedings of the IEEE International Conference on Computer Vision, pages 5533–5541, 2017.
- [55] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [56] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 915–924, 2021.
- [57] Haoxi Ran, Jun Liu, and Chengjie Wang. Surface representation for point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18942–18952, 2022.
- [58] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. arXiv preprint arXiv:2202.07123, 2022.
- [59] James R Ayre, Peter J Bazira, Mohammed Abumattar, Haran N Makwana, and Katherine A Sanders. A new classification system for the anatomical variations of the human circle of willis: A systematic review. *Journal of Anatomy*, 240(6):1187–1204, 2022.
- [60] Joshua D Jones, Pedro Castanho, Peter Bazira, and Katherine Sanders. Anatomical variations of the circle of willis and their prevalence, with a focus on the posterior communicating artery: a literature review and meta-analysis. *Clinical Anatomy*, 34(7):978–990, 2021.
- [61] Lars B Hindenes, Tor Ingebrigtsen, Jørgen G Isaksen, Asta K Håberg, Liv-Hege Johnsen, Marit Herder, Ellisiv B Mathiesen, and Torgil R Vangberg. Anatomical variations in the circle of willis are associated with increased odds of intracranial aneurysms: The tromsø study. *Journal of the Neurological Sciences*, 452:120740, 2023.
- [62] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 869–877, 2018.

- [63] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.
- [64] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International conference on learning representations*, 2020.
- [65] Oscar Méndez-Lucio, Mazen Ahmad, Ehecatl Antonio del Rio-Chanona, and Jörg Kurt Wegner. A geometric deep learning approach to predict binding conformations of bioactive molecules. *Nature Machine Intelligence*, 3(12):1033–1039, 2021.
- [66] Jingxuan Wang, Zongzhao Qiu, Xuan Zhang, Zhenghe Yang, Wei Zhao, and Xuefeng Cui. Boosting deep learning-based docking with cross-attention and centrality embedding. In 2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), pages 360–365. IEEE, 2022.
- [67] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. arXiv preprint arXiv:2104.13478, 2021.
- [68] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124, 2017.
- [69] Farhad Arbab, Ivan Herman, and Pål Spilling. An overview of manifold and its implementation. *Concurrency: practice and experience*, 5(1):23–70, 1993.
- [70] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895, 2020.
- [71] Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv* preprint *arXiv*:1910.03193, 2019.
- [72] K Kapoor, VK Kak, and B Singh. Morphology and comparative anatomy of circulus arteriosus cerebri in mammals. *Anatomia, histologia, embryologia*, 32(6):347–355, 2003.
- [73] Zvonimir Vrselja, Hrvoje Brkic, Stefan Mrdenovic, Radivoje Radic, and Goran Curic. Function of circle of willis. *Journal of Cerebral Blood Flow & Metabolism*, 34(4):578–584, 2014.
- [74] Julie Rosner, Vamsi Reddy, and Forshing Lui. Neuroanatomy, circle of willis. 2018.
- [75] Minghui Zhang, Xin You, Hanxiao Zhang, and Yun Gu. Topology-aware exploration of circle of willis for cta and mra: Segmentation, detection, and classification. *arXiv* preprint arXiv:2410.15614, 2024.
- [76] Li Chen, Thomas Hatsukami, Jenq-Neng Hwang, and Chun Yuan. Automated intracranial artery labeling using a graph neural network and hierarchical refinement. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, pages 76–85. Springer, 2020.
- [77] Eugene Wigner. *Group theory: and its application to the quantum mechanics of atomic spectra*, volume 5. Elsevier, 2012.
- [78] Lars B Hindenes, Asta K Håberg, Liv Hege Johnsen, Ellisiv B Mathiesen, David Robben, and Torgil R Vangberg. Variations in the circle of willis in a large population sample using 3d tof angiography: The tromsø study. *PLoS One*, 15(11):e0241373, 2020.
- [79] Monique J Krabbe-Hartkamp, Jeroen Van der Grond, Frank Erik De Leeuw, Jan Cees De Groot, Ale Algra, Berend Hillen, MM Breteler, and WP Mali. Circle of willis: morphologic variation on three-dimensional time-of-flight mr angiograms. *Radiology*, 207(1):103–111, 1998.
- [80] Mika Okahara, Hiro Kiyosue, Hiromu Mori, Shuichi Tanoue, Michihumi Sainou, and Hirohumi Nagatomi. Anatomic variations of the cerebral arteries and their embryology: a pictorial review. *European radiology*, 12:2548–2561, 2002.
- [81] Ana Maria Dumitrescu, Claudia Florida Costea, Cristina Furnică, Mihaela Dana Turliuc, Andrei Ionuț Cucu, Camelia Margareta Bogdănici, Şerban Turliuc, Gabriela Florența Dumitrescu, and Anca Sava. Morphological aspects of the vasculogenesis and angiogenesis during prenatal edification of the circle of willis: a review. *Romanian Journal of Morphology and Embryology*, 62(3):679, 2022.

- [82] JJ Van Overbeeke, B Hillen, and CA Tulleken. A comparative study of the circle of willis in fetal and adult life. the configuration of the posterior bifurcation of the posterior communicating artery. *Journal of Anatomy*, 176:45, 1991.
- [83] Zhenhong Liu, Xingce Wang, Zhongke Wu, Yi-Cheng Zhu, and Alejandro F Frangi. Mri joint superresolution and denoising based on conditional stochastic normalizing flow. *IEEE Transactions on Artificial Intelligence*, 2024.
- [84] Luca Antiga, Bogdan Ene-Iordache, and Andrea Remuzzi. Computational geometry for patient-specific reconstruction and meshing of blood vessels from mr and ct angiography. *IEEE transactions on medical* imaging, 22(5):674–684, 2003.
- [85] Scott MacDonald Black, Craig Maclean, Pauline Hall Barrientos, Konstantinos Ritos, and Asimina Kazakidi. Reconstruction and validation of arterial geometries for computational fluid dynamics using multiple temporal frames of 4d flow-mri magnitude images. Cardiovascular Engineering and Technology, 14(5):655–676, 2023.
- [86] Richard Izzo, David Steinman, Simone Manini, and Luca Antiga. The vascular modeling toolkit: a python library for the analysis of tubular structures in medical images. *Journal of Open Source Software*, 3(25):745, 2018.
- [87] Arunachalam Narayanaswamy, Saritha Dwarakapuram, Christopher S Bjornsson, Barbara M Cutler, William Shain, and Badrinath Roysam. Robust adaptive 3-d segmentation of vessel laminae from fluorescence confocal microscope images and parallel gpu implementation. *IEEE transactions on medical imaging*, 29(3):583–597, 2009.
- [88] Michal Chlebiej, Anna Zurada, Jerzy Gielecki, Mikolaj A Pawlak, and Maciej Szkulmowski. Customizable tubular model for n-furcating blood vessels and its application to 3d reconstruction of the cerebrovascular system. *Medical & Biological Engineering & Computing*, 61(6):1343–1361, 2023.
- [89] Dieuwertje Alblas, Iris N Vos, Micha M Lipplaa, Christoph Brune, Irene C van der Schaaf, Mireille RE Velthuis, Birgitta K Velthuis, Hugo J Kuijf, Ynte M Ruigrok, and Jelmer M Wolterink. Deep-learning-based extraction of circle of willis topology with anatomical priors. Scientific reports, 14(1):31630, 2024.
- [90] Md Shakib Shahariar Junayed, Kazi Shahriar Sanjid, Md Tanzim Hossain, M Monir Uddin, and Sheikh Anisul Haque. Topology-aware anatomical segmentation of the circle of willis: Hunet unveils the vascular network. *IET Image Processing*, 18(10):2745–2753, 2024.
- [91] Chih-Yang Hsu, Mahsa Ghaffari, Ali Alaraj, Michael Flannery, Xiaohong Joe Zhou, and Andreas Linninger. Gap-free segmentation of vascular networks with automatic image processing pipeline. Computers in biology and medicine, 82:29–39, 2017.
- [92] Jacob R Bumgarner and Randy J Nelson. Open-source analysis and visualization of segmented vasculature datasets with vesselvio. *Cell reports methods*, 2(4), 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and Section 1 accurately reflect the paper's contribution and scope. Section 1 summarizes the main contributions of this paper as well.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations of the work are discussed in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The proofs are in Appendix C.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The architecture of proposed model is described in Section 4.2 and 4.3, and the pipeline is showed in Figure 2. The implementation details of the model are provided in Appendix D.2. The code and TopCoW 2024 dataset with its BBSC format will be released after the review period, but we cannot release the clinical dataset because we have signed a confidentiality agreement with hospital to protect patient privacy.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The Non-privacy dataset and our implementation(model, train and test, visualization) will be released publicly soon.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The train and test details are provided in Appendix D.2 and D.3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: In the experiments presented in Table 1 and Table 2 of Section 5.1 and 5.2, we conduct 5-fold cross-validation and calculate the mean and standard deviation. The data distributions for each class are summarized in Appendix D.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: This part is provided in Section 5.3 and Appendix D.3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Research conducted in this paper conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We highlight the clinical value of our model in Appendix A.2 and Section 5, including its potential to aid CoW diagnosis and reduce medical workload. Abstract and Section 1 shows its benefits for tubular structure analysis and manifold learning tasks.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We will not release the clinical dataset involving privacy issues, so no part of this paper is with high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The creators or original owners of assets, used in the paper, are all properly credited. The license and terms of use are explicitly mentioned and properly respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The non-privacy dataset and the code in our work will be released after reviewing.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve research with human subjects.

Guidelines:

 The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: Usage of LLMs do not impact the core methodology, scientific rigorousness, or originality of the research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

A Background

A.1 Geometric Deep Learning and Continuous Function Learning

Geometric deep learning [27] is a branch of deep learning that leverages deep neural networks to model and extract intrinsic geometric features of structured data. To date, geometric deep learning has yielded impressive results in a range of fields, notably physics, biology, and computer graphics. In physics, GNS [63] and MESHGRAPHNET [64] investigate the application of geometric deep learning in the simulation of fluid and cloth, aiming to capture their intricate physical dynamics. In biology, DeepDock [65] and its enhanced version, caDeepDock [66], predict the binding conformation by learning the distributions based on Euclidean distance. In computer graphics, PointNet [24] and PointNet++ [55] represent major breakthroughs in point cloud processing by effectively capturing geometric features directly from raw point clouds. Most geometric deep learning methods operate on discrete data in Euclidean space, such as point clouds, by constructing geometric graphs[67, 68] or defining local point neighborhoods. These approaches typically assign geometric attributes (e.g., spatial coordinates) to the nodes of the graph and the points, enabling the extraction of meaningful geometric features.

A topological manifold [69] is a special type of geometric space, formally defined as a topological space that is locally Euclidean. Compared to modeling discrete data such as point cloud in 3D Euclidean space, modeling directly on continuous manifold domains can avoid the errors introduced by numerical estimation or implicit representations learned via neural networks. This approach enables models to achieve greater expressiveness, interpretability, and generalization capability. Function learning models [70, 71], with the goal of constructing high-dimensional mappings from functions to functions, preserve the analytical properties of the functions themselves, such as smoothness, and has found increasing application [70, 71]. Inspired by the aforementioned approaches, we construct a 3D rotation and translation equivariant mapping on smooth surface manifolds represented by Ball B-Spline Curve(BBSC) [8]. Owing to its compactness, local modifiability, smoothness, convex hull property, the BBSC provides an efficient and accurate representation of tubular geometric structures.

A.2 Circle of Willis

CoW is a ring-shaped arterial structure located at the base of the brain in birds, reptiles, and mammals [72]. It supplies blood to the brain and surrounding structures by connecting the bilateral internal carotid artery (ICA) systems with the vertebrobasilar system. In cases of stenosis or occlusion of a major artery (e.g., the ICA), the CoW provides collateral circulation to maintain perfusion in ischemic regions, serving a compensatory function. Additionally, the CoW plays a critical role in balancing and regulating intracranial arterial pressure and in protecting against acute cerebrovascular events [73]. Notably, the CoW is also a frequent site for cerebral aneurysms, which, when ruptured, can lead to stroke and potentially fatal outcomes. In its idealized form, the standard CoW consists of several key arteries: the anterior cerebral arteries (left and right), the anterior communicating artery, the distal termini of the internal carotid arteries (left and right), the P1 segments of the posterior cerebral arteries (left and right), and the posterior communicating arteries (left and right) [74]. Depending on clinical needs, more detailed sub-regional definitions may be adopted. However, due to factors such as genetic variation and developmental anomalies, the CoW often exhibits significant topological heterogeneity in its anatomical structure. For example, incomplete CoWs may result from the absence of one or more arterial segments; developmental variants may manifest as arteries that are duplicated, abnormally narrow, or overly dilated. From a topological perspective, the CoW is generally classified into three categories: complete circles, partially closed, and open configurations. Among them, open CoWs are often associated with poor collateral capacity and are more prone to widespread ischemia or infarction in the presence of arterial narrowing, occlusion, or aneurysm rupture. This high degree of topological variability poses considerable challenges for clinical diagnosis and increases the importance of automated and robust computational methods.

The segmentation and classification of the CoW are both critical and well-suited to computational approaches. Accurate and efficient methods not only reduce the workload of medical professionals but also enhance diagnostic precision, playing an essential role in the detection and early monitoring of CoW-related cerebrovascular diseases. To the best of our knowledge, there are currently few computational models specifically designed for CoW segmentation and classification. Although some methods developed for general vascular analysis can be adapted, their performance on CoW remains

unstable due to its inherent topological heterogeneity [59, 60, 61]. Therefore, designing effective models for CoW-related tasks requires a joint analysis of both geometric and topological features. Such models [75, 9, 76] are crucial for capturing the structural variations and delivering reliable diagnostic support in real-world clinical scenarios.

B Preliminary Model Architecture

B.1 DiffusionNet

DiffusionNet [46] is an effective method to learn point cloud features using the principle of heat diffusion. It mainly consists of two modules: diffusion range learning and gradient feature learning.

Learning diffusion. The diffusion learning can be interpreted as a filter that allows each point to propagate information radially and the diffusion range $H_t(u_0)$ with initial state u_0 , mainly depends on a learnable diffusion time t, is described by formula $H_t(u_0) = e^{t\Delta}u_0$. DiffusionNet effectively discretizes and simulates the diffusion process using the implicit Euler time step method,

$$h_t(u) = (M + tL)^{-1} M u$$
 (16)

where M and L represent mass matrix and weak Laplace matrix. To avoid the inconvenience of solving a large linear system separately for each channel, DiffusionNet offers spectral acceleration which leveages the precomputed basis of low frequency Laplacian eigenfunctions to represent the diffusion approximately. Let $\Phi = [\phi_i] \in \mathbb{R}^{V \times k}$ represent the concatenated matrix of the eigenvectors corresponding to the first k smallest eigenvalues λ_i where $[\phi_i]$ are solutions to $L\phi_i = \lambda_i M\phi_i$. The coefficients c of spectral basis can be obtained by $c = \Phi^T Mu$ and u can be recovered by $u = \Phi c$, and the coefficients c_i^t after t time diffusion can be calculated by formula $c_i^t = e^{-\lambda_i} c_i^0$. Then diffusion can be simplified as follows,

$$h_t(u) = \Phi c^t = \Phi(e^{-\lambda_0 t}, e^{-\lambda_1 t}, \dots, e^{-\lambda_k t})^T \cdot (\Phi^T M u)$$
(17)

Learning gradient feature. In additional, DiffusionNet constructs a learning scaling gradient feature for each point to enhance the feature learning ability of the filter. First, the gradient operator which is assembled into a sparse matrix G. The matrix is independent of the vertex feature and can be precomputed for each shape by computing the least-squares approximation of the function values of each point and its neighbors on the tangent plane. Note that in order to conveniently represent the tangent vector in the calculation of the gradient matrix, DiffusionNet uses a complex expression. Multiply each channel feature u of vertex v by the gradient matrix G and stack them up to get the gradient feature $w_v = [Gu]$. Then a learned scaling gradient feature g_v shown as follow:

$$g_v = tanh(Real(\langle \overline{w}_v, Aw_v \rangle)) = tanh\left(Real\left(\sum_{i=1}^D \sum_{j=1}^D \overline{w}_v(i)A_{i,j}w_v(j)\right)\right), \ A \in \mathbb{R}^{D \times D}$$
 (18)

where A is the learned real matrix which describe the scaling and D is the number of the feature channel. Whether to choose a real or complex matrix A mainly depends on whether the point has a normal. Because multiplying a complex scalar do both the scaling and rotation transformations. The real part of the scalar represents scaling, while the imaginary part represents rotation. Generally, the direction of rotation deponds on the outward normal and orientation, so for points without normal, a real matrix A is sufficient.

It is worth noting that even though only low-frequency basis is used in the diffusion spectral acceleration part, DiffusionNet can still learn high-frequency features through gradient learning and MLPs (Multi-Layer Perceptrons).

B.2 SE(3)-Transformer

The SE(3)-Transformer [30] integrates the principles of 3D rotation and translation equivariance by embedding Tensor Field Network (TFN) [29] into the computation of the key, query, and value vectors within the Transformer architecture. This design endows the Transformer with SE(3) equivariance in the 3D Euclidean space, ensuring that its outputs are consistent under rigid transformations. First, the SE(3)-Transformer leverages spherical harmonics [37] to map 3D point coordinates into a higher-dimensional representation space. Spherical harmonics form a set of orthogonal Fourier bases defined

on the unit sphere S^2 . Specifically, an l degree spherical harmonics Y^l can map any 3D unit vector to a 2l+1-dimensional vector in \mathbb{R}^{2l+1} , expressed as $Y^l(\vec{x}) = \{Y^l_{-l}(\vec{x}), Y^l_{-l+1}(\vec{x}), \dots, Y^l_{l-1}(\vec{x}), Y^l_{l}(\vec{x})\}$. Then, the Wigner-D matrices [77], which are unitary and irreducible, are introduced to achieve equivariance of high-dimensional vectors generated by spherical harmonics under rotations in SO(3):

$$Y^{l}(R_{q}\vec{x}) = W^{l}(g)Y^{l}(\vec{x}), \quad g \in SO(3), R \in \mathbb{R}^{3\times3}, W^{l}(g) \in \mathbb{R}^{(2l+1)\times(2l+1)}.$$
 (19)

Furthermore, Clebsch-Gordan tensor product [38] provides a mechanism to aggregate vectors from spaces of different dimensions while preserving equivariance. For $\vec{\mathbf{V}}^{l_1} \in \mathbb{R}^{(2l_1+1) \times C_1}$ and $\vec{\mathbf{V}}^{l_2} \in$ $\mathbb{R}^{(2l_2+1)\times C_2}, \vec{\mathbf{V}}^l = \vec{\mathbf{V}}^{l_1} \otimes_{ca}^l \vec{\mathbf{V}}^{l_2}$ can be expanded as:

$$V_{m,c}^{l} = \sum_{c_{1}=1,c_{2}=1}^{C_{1},C_{2}} w_{c_{1},c_{2},c} \sum_{m_{1}=-l_{1}}^{l_{1}} \sum_{m_{2}=-l_{2}}^{l_{2}} Q_{(l_{1},m_{1})(l_{2},m_{2})}^{(l,m)} v_{m_{1}}^{(l_{1},c_{1})} v_{m_{2}}^{(l_{2},c_{2})},$$
(20)

where $Q^{(l,m)}_{(l_1,m_1)(l_2,m_2)}$ are the Clebsch-Gordan coefficients. SE(3)-Transformer constructs a set of learnable radial functions $\psi:\mathbb{R}_{\geq 0}\to\mathbb{R}$ and leveages the Clebsch-Gordan coefficients to form the basis kernel $W^{c_oc_i}\in\mathbb{R}^{(2c_o+1)\times(2c_i+1)}$ which can transform $(2c_i + 1)$ channels input feature to $(2c_o + 1)$ channels output feature:

$$W^{c_o c_i}(\vec{x}) = \sum_{l=|c_i - c_o|}^{c_i + c_o} \psi_l^{c_o c_i}(\|\vec{x}\|) \sum_{m=-l}^l Y_m^l(\frac{\vec{x}}{\|\vec{x}\|}) Q_{lm}^{c_o c_i}, \tag{21}$$

where Y_m^l is the mth element in Y^l , $Q_{lm}^{c_oc_i} \in \mathbb{R}^{(2c_o+1)\times(2c_i+1)}$. With the basis kernel $W^{c_oc_i}$, c_i channels input feature $f_{in}^{c_i}$ of points x, SE(3)-Transformer can generate a c_o channels feature $f_{out}^{c_o}$, as shown in Equation (4) and (5).

Proof of Infinite Differentiability in BBSC Mapping

Proposition 1. Let $U = \{P \in \mathbb{R}^{3 \times n}, r \in \mathbb{R}_{>0}, \tau \in \mathbb{T}\}$, where $\mathbb{T} = \{t_i \in [0,1] | t_{p+1} < t_{p+2} < \cdots < t_n, t_1 = \cdots = t_{p+1} = 0, t_n = \cdots = t_{n+p+1} = 1\}$. Then, the mapping $\beta : U \to \mathcal{M}$, defined according to Equation (1) and Equation (2), is a C^{∞} mapping.

Proof. The p-degree BBSC manifold can be written as:

$$\beta(P, r, \tau) = \sum_{i=0}^{n} N_{i,p}(t) (P_i ; r_i) = \left(\sum_{i=0}^{n} N_{i,p}(t) P_i ; \sum_{i=0}^{n} N_{i,p}(t) r_i\right) = (\sigma(t), \sigma(t)), \quad (22)$$

and the basis function is as follows:

$$N_{i,p}(t) = \frac{t - t_i}{t_{i+p} - t_i} N_{i,p-1}(t) + \frac{t_{i+p+1} - t}{t_{i+p+1} - t_{i+1}} N_{i+1,p-1}(t), N_{i,0} = \begin{cases} 1, & \text{if } t_i \le t < t_{i+1}, \\ 0, & \text{otherwise.} \end{cases}, (23)$$

where $P_i \in \mathbb{R}^3$, $r_i \in R$, $t_{p+1} < t_{p+2} < \dots < t_n$, $t_1 = \dots = t_{p+1} = 0$, $t_n = \dots = t_{n+p+1} = 1$. We prove the smoothness of the manifold \mathcal{M} by showing that β admits infinitely many continuous partial derivatives with respect to P, r, and τ .

 β is a linear combination of the control points P_i , and thus admits infinitely many continuous partial derivatives with respect to each P_i , i.e., $\forall P_i \in \mathbb{R}^3$,

$$\frac{\partial \beta}{\partial P_i} = N_{i,p}(t), \quad \frac{\partial^2 \beta}{\partial P_i^2} = \dots = \frac{\partial^\infty \beta}{\partial P_i^\infty} = 0. \tag{24}$$

Similarly, since β admits infinitely many continuous partial derivatives with respect to each control

For τ , to avoid the denominator being 0, we do not consider the endpoints of the BBSC which means we only consider the differentiability of strictly increasing sequence $t_{p+1}, t_{p+2}, \ldots, t_n$. Now we need to prove that the corresponding basis function $N_{i,p}$ which contains t_i has continuous infinite partial derivatives. As observed in Equation (23), the basis function $N_{i,p}$ is a recursively defined composite function, and its differentiation can be handled via the chain rule. Since each knot value t_i appears in the basis function only through two types of expressions, simply denoted as $f_1 = \frac{t - t_i}{a - t_i}$

and $f_2 = \frac{t-b}{t_i-b}$ (maybe the former needs to be added with a minus sign, but this does not affect the derivative calculation), it is sufficient to verify whether both types possess infinitely many continuous partial derivatives. Thus β admits infinitely many continuous partial derivatives with respect to knot vector τ and can be calculated as follows:

$$\frac{d^m f_1}{dt_i^m} = -\frac{m!}{(a-t_i)^{m+1}}, \quad \frac{d^m f_2}{dt_i^m} = (-1)^m \frac{m! \cdot (t-b)}{(t_i-b)^{m+1}}, \quad a \neq t_i, \quad b \neq t_i.$$
 (25)

By the chain rule, any higher-order derivative of $N_{i,p}$ with respect to t_i can be expressed as a combination of products and sums of the derivatives of the f_1 and f_2 .

For P and τ , the function β is separable with respect to these variables. Specifically, it can be written as the product of G(P) = P and $N(\tau) = N_p(t)$:

$$\beta = N(\tau)G(P)^T. \tag{26}$$

Its mixed partial derivatives are then expressed as

$$\frac{\partial^{l} \beta}{\partial P^{m} \partial \tau^{n}} = \frac{\partial^{m} \beta}{\partial P^{m}} \frac{\partial^{n} \beta}{\partial \tau^{n}} = \frac{d^{m} G}{d P^{m}} \frac{d^{n} N}{d \tau^{n}}, \quad \text{where } m + n = l. \tag{27}$$

Since the partial derivatives of G(P) and $N(\tau)$ exist for all orders and are continuous, the mixed partial derivatives of β with respect to P and τ are also continuous for all orders.

Similarly, arbitrary-order continuous mixed partial derivatives with respect to r and τ can be obtained. For P and r, we have

$$c(P) = N_p(t)P^T = N_p(t)P^T \cdot 1 = c(P)H(r), \text{ where } H(r) = 1,$$
 (28)

which ensures that β also admits arbitrary-order continuous mixed partial derivatives with respect to P and r. The same reasoning applies to r itself. Thus, it can be proven that $\forall \beta \in \mathcal{M}, \beta$ is a C^{∞} mapping.

D Details of the CoW Dataset and Training Procedure

D.1 Per-Class Distribution of Dataset

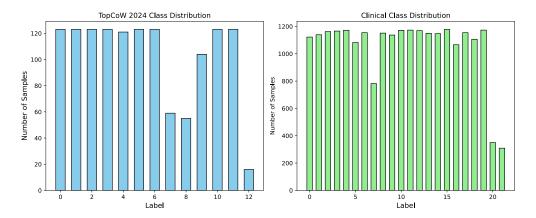


Figure 6: Per-class distribution of TopCoW 2024 and clinical dataset.

The CoW exhibits a highly heterogeneous anatomical structure, which arises not only from inter-ethnic and inter-population variability[78] but also from genetic factors[79], embryonic development[80], environmental influences[81], and hemodynamics-driven vascular remodeling[82]. This inherent heterogeneity leads to significant class imbalance in data representations of the CoW. As illustrated in Figure 6 we present an overview of the class distribution in both the TopCoW 2024 MICCAI Challenge dataset and our internal clinical dataset. Notably, in the TopCoW dataset, category 12 is severely underrepresented, with vessels labeled as class 12 comprising only 5% of all samples. Even

when training the model with a uniformly stratified split between training and testing sets, the model's learning of class-12 vessels remains highly stochastic—e.g., in extreme cases, the test set might contain only a single sample with class-12 branches. In contrast, as shown in Figure 4, our method demonstrates robust performance on the clinical dataset, even for rare classes such as 20 and 21, achieving accurate segmentation despite their limited presence in the data. While data augmentation is a common strategy to address such imbalance, we currently do not employ it. This is primarily due to the use of SE(3)-equivariant architectures, which are inherently invariant to translations and rotations, rendering many conventional augmentations redundant in our context.

D.2 Preprocess of Clinical Data

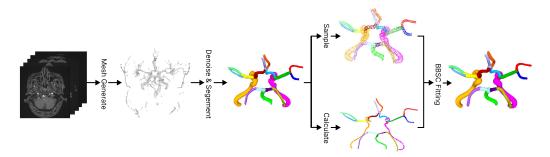


Figure 7: Clinical data preprocessing pipeline. A mesh is constructed from each MRI scan. The denoise and clean operations are performed to remove noise. Then the branches of the CoW are segmented. Surface point clouds and centerlines are computed from the segmented mesh, followed by BBSC fitting.

Compared with the TopCoW2024 dataset, clinical CoW scans are more challenging due to imaging artifacts and spurious small vessel components. Consequently, a series of preprocessing steps is required before BBSC fitting and shown in Figure 7. First, when severe MRI artifacts such as motion blur or metal-induced distortions degrade mesh quality, denoising methods can be applied [83]. Vascular meshes are then reconstructed from MRI either using geometric methods [84, 85] or the VMTK package [86]. The reconstructed meshes often contain noise from small vessel components and exhibit disconnected fragments. These are removed through cleaning [87] and repaired using mesh fix algorithms [88]. Branch segmentation [89, 90] and ground truth labeling follow, with manual verification and correction by domain experts to ensure accuracy. The majority of preprocessing operations can be implemented with open-source packages such as VMTK or medical CAD tools [91, 92]. Finally, surface point clouds and centerlines are extracted from the cleaned meshes for BBSC fitting. It is worth noting that, in our pipeline, the mesh is first reconstructed and then subjected to cleaning, denoising, and segmentation. Alternatively, these preprocessing operations can also be applied directly to the MRI images prior to mesh reconstruction.

D.3 Training Detail

BBSC construction. For each branch of the CoW, we individually construct a BBSC representation. These are then processed using DiffusionNet to obtain the three key components of the BBSC: control points, control radius, and the knot vector. For notational convenience, we denote the concatenation of control points and radius as the control parameters, represented as 4D vectors. In general, the length of the control parameters is either proportional to the branch's arc length or determined by a threshold-based allocation strategy. However, in our experiments, we observed that assigning a fixed control parameter length of 13 for all branches significantly simplified and accelerated the BBSC construction process, while having negligible impact on downstream classification performance. When constructing BBSCs, we consistently set the spline degree to 3. For the TopCoW 2024 dataset, we assign all branches a control parameter length of 13. For the clinical dataset from hospital, we adopt a length allocation scheme based on arc length thresholds, assigning control parameter lengths of 13, 10, and 5, respectively. Additionally, during the BBSC fitting process, we adopt the computationally efficient and easy-to-implement L2 norm metric (Equation (7)). The term $||P - \tilde{P}||_2$ in this metric can also be replaced with $\langle P, \tilde{P} \rangle$, yielding a comparable fitting performance. The

weighting parameters α and η are set to 1. During training, all control parameters are padded to a fixed length of 13 to ensure consistency across samples. It is important to note that we pad with zeros. To distinguish between padded zeros and the zero entries in the knot vector, we apply a value shift of +1 to all elements in the knot vector.

SE(3)-BBSCformerGCN. Due to the local support property of the p-degree B-Spline, where the influence of the i-th control parameter is restricted to $[t_i,\ldots,t_{i+p+1}]$, we first construct an $n\times(p+1)$ matrix $\tilde{\tau}$, where the i-th row corresponds to $[t_i,\ldots,t_{i+p+1}]$ (p=3). We then apply several MLP jointly to $\tilde{\tau}$ and the control radius to produce a (2l+1)-dimensional vector, where $l=0,1,\ldots,c_{i_0}$ $(c_{i_0}=3)$. The direct sum of c_{i_0} vectors constitutes the rotation and translation equivariant input $f_{in}^{c_{i_0}}$ of SE(3)-BBSCformer defined in Section 4.3. The degree of spherical harmonics is also c_{i_0} . The architecture of our graph convolutional network (GCN) follows a hybrid of the GCN and Res-GCN design used in DeepDock, as illustrated in Figure 8 (we set $H_1=3$ and $H_2=10$).

$$\longrightarrow H_1 \times \bigcirc GCN \longrightarrow H_2 \times \bigcirc Residual GCN \longrightarrow$$

Figure 8: Architecture of GCN in SE(3)-BBSCformerGCN.

Train and test detail. We train and evaluate all models presented in this paper on a single NVIDIA 3090Ti GPU. For the SE(3)-BBSCformerGCN, we set the number of SE(3)-BBSCformer layers to 1. During training, we used the Adam optimizer with a fixed learning rate of 0.001 and a weight decay of 0.01. The number or workers is 1. The loss function is set to cross-entropy for all models, and the number of training epochs is fixed at 200. For the TopCoW 2024 dataset, lightweight models such as SE(3)-BBSCformer, BBSCformerGCN, and SE(3)-BBSCformerGCN are trained and tested using a 5-fold cross-validation protocol, and their mean and standard deviation is computed and showed in Figure 3. During training, we split the TopCoW 2024 dataset into 100 samples for training and 20 for testing. Similarly, the clinical dataset is partitioned into training and test sets with a 4:1 split.

E Per-Class AUC-ROC

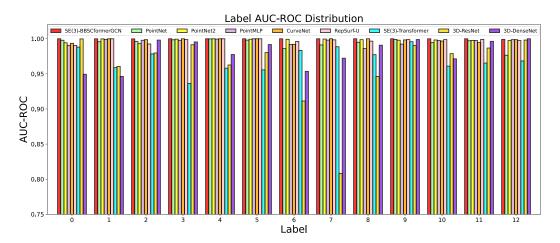


Figure 9: Per-class AUC-ROC of various models evaluated on the TopCoW 2024 Circle of Willis dataset.

We further present the AUC-ROC scores of different models across various label categories on both the TopCoW 2024 dataset and a clinical dataset, aiming to assess the per-class classification performance in more detail. As shown in Figure 10 and Figure 4, the y-axis is restricted to the range of 0.75 to 1.0 for improved visual clarity, allowing for more effective comparison between models in high-performance regimes. SE(3)-BBSCformerGCN consistently achieves the highest AUC-ROC scores across all anatomical label categories, clearly outperforming baseline methods such as voxel-based and point-based networks. This consistent superiority is observed not only on the relatively small-scale TopCoW 2024 dataset, but also on the large-scale real-world clinical dataset,

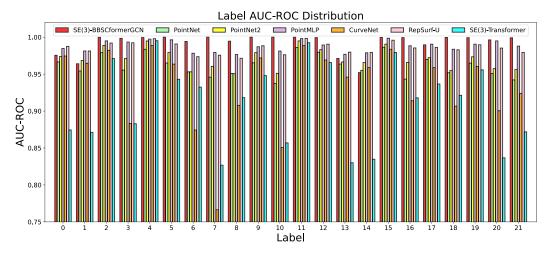


Figure 10: Per-class AUC-ROC of various models evaluated on the Circle of Willis dataset collected from clinical imaging at hospital.

highlighting the model's strong generalization ability. The results further demonstrate the advantages of our BBSC representation and the SE(3)-equivariant architecture, which together enable precise modeling of complex geometric and topological structures while maintaining robustness across domains. Importantly, even for challenging categories with high anatomical variability, our model maintains high AUC-ROC scores, indicating its effectiveness in capturing fine-grained structural differences and mitigating overfitting.

F Case Study of Clinical Data

We conduct a case study on Circle of Willis (CoW) clinical data collected from a collaborating medical institution, focusing on samples where SE(3)-BBSCformerGCN makes prediction errors and which exhibit certain topological heterogeneity. As shown in Figure 11, the first row presents a complete CoW structure, a configuration that only accounts for about 20%–25% of real-world cases. In the remaining samples with topological variations, most commonly observed is the absence of the bilateral posterior communicating arteries (PcoA-L and PcoA-R), which connect the upper and lower parts of the CoW. In the fourth row, a more severe case of unilateral absence is observed.

Beyond topology, geometric heterogeneity within the same type of branch also poses challenges. For instance, the MCA-L2 branch highlighted in green in the third row is noticeably longer than its counterparts in other samples, while the ICA-R1 branch in the fourth row, marked in ochre, is significantly thicker than in the other cases. Such variations in both topology and geometry contribute to the clinical difficulty of accurately classifying CoW branches.

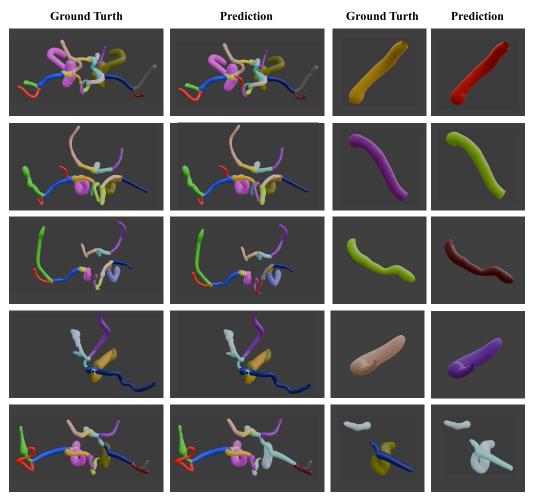


Figure 11: Clinical CoW samples with topological heterogeneity and the corresponding mispredicted structures by SE(3)-BBSCformerGCN. The first two columns show the complete CoW structures, while the last two columns highlight the misclassified branches.