

PrIntMesh: Precise Intersection Surfaces for 3D Organ Mesh Reconstruction

Deniz Sayin Mercadier¹, Hieu Le², Yihong Chen¹, Jiancheng Yang³, Udaranga Wickramasinghe⁴, Pascal Fua¹

¹EPFL ²UNC Charlotte ³ELLIS Institute Finland ⁴Adis SA, Lausanne, Switzerland

Abstract

Human organs are composed of interconnected substructures whose geometry and spatial relationships constrain one another. Yet, most deep-learning approaches treat these parts independently, producing anatomically implausible reconstructions. We introduce *PrIntMesh*, a template-based, topology-preserving framework that reconstructs organs as unified systems. Starting from a connected template, *PrIntMesh* jointly deforms all substructures to match patient-specific anatomy, while explicitly preserving internal boundaries and enforcing smooth, artifact-free surfaces. We demonstrate its effectiveness on the heart, hippocampus, and lungs, achieving high geometric accuracy, correct topology, and robust performance even with limited or noisy training data. Compared to voxel- and surface-based methods, *PrIntMesh* better reconstructs shared interfaces, maintains structural consistency, and provides a data-efficient solution suitable for clinical use.

1. Introduction

Human organs are not single solids but interconnected systems. The heart, for example, consists of four chambers—2 atria and 2 ventricles—separated by thin walls, linked by valves, and threaded with vessels. These shapes are tightly coupled; the deformations of one impact the others. Modeling the heart thus requires capturing these relationships, not just segmenting individual chambers. Even small misalignments can compromise anatomical plausibility.

Most current deep-learning approaches ignore this. They break organs apart, predicting substructures such as the left atrium or septal wall independently and ignoring their spatial or functional interdependence. This disrupts continuity and limits use in downstream tasks such as biomechanics [28, 31], electrophysiology [19, 27], and surgical planning [26]. Voxel-based models like nnU-Net [13] assign labels independently across the volume, often leaving gaps and disconnected fragments, as in Fig. 1a. Surface-based methods [4, 16, 17, 34, 35] reconstruct each chamber separately, with no guarantee of proper alignment at shared




Figure 1. **Voxel-based vs Mesh-based approaches.** (a) Voxel-based approaches such as nnU-Net [13] currently are the dominant ones for organ segmentation. However, they can yield gaps, floaters, and jagged boundaries, requiring post-processing to produce usable meshes. (b) Existing mesh-based methods generate smooth surfaces but model each class separately, leading to interpenetrations and misaligned boundaries between components. (c) *PrIntMesh* deforms a joint template for all components and precisely reconstructs shared surfaces, producing a smooth, water-tight and topologically correct mesh, without post-processing.

interfaces, as in Fig. 1b. The resulting meshes are often anatomically implausible and clinically unreliable.

To remedy this, we propose *PrIntMesh*, a method that reconstructs organs as unified, connected systems rather than collections of parts. Instead of attempting to reconcile independently predicted substructures as most earlier approaches do, *PrIntMesh* starts from a template that encodes full anatomical connectivity. It acts as a topological scaffold, ensuring structural integrity throughout deformation. Guided by volumetric image features, the network deforms this scaffold to match individual anatomy while preserving its internal organization.

PrIntMesh comprises two modules, a feature extractor

that processes volumetric imagery and a mesh deformation network that deforms the whole connected template. Unlike earlier models that predict each component independently [17, 34], it deforms all components together, preserving spatial relationships by construction. Anatomical realism is maintained by explicitly supervising pairwise interfaces between substructures to preserve coherent internal walls, and by applying geometric regularization terms that encourage smooth, artifact-free surfaces.

On high-resolution cardiac data, *PrIntMesh* achieves both high geometric accuracy and guaranteed topological consistency—a combination unmatched by competing methods. While voxel models may yield similar voxel-wise scores, they fail to maintain connectivity or realistic internal boundaries. We further tested *PrIntMesh* on hippocampus and lung datasets, demonstrating generality across anatomies with distinct topologies—from the folded hippocampus to disconnected lungs with imperfect annotations. In all cases, *PrIntMesh* produces structurally valid, smooth reconstructions even with limited training data.

Crucially, the structural prior embedded in the template enables data-efficient learning. For lungs, *PrIntMesh* yields a 2–3× lower Chamfer error than nnU-Net [13] while guaranteeing topological correctness and requiring fewer than 100 samples—highlighting its practicality for clinical deployment, with scarce annotated data and rapid adaptation required.

In short, our contributions are as follows:

- We propose a template-based, topology-preserving framework that reconstructs organs as unified, connected systems rather than independent parts.
- We use our templates to strictly enforce anatomical interfaces and geometric regularity, ensuring smooth, artifact-free surfaces and correct internal boundaries.
- We demonstrate high-accuracy, data-efficient reconstruction across diverse anatomies—including heart, hippocampus, and lungs—even given limited or imperfect training data.

The code will be made publicly available.

2. Related Work

Deep learning has become the de facto standard for medical image segmentation, including segmenting the liver [39], the whole heart [35], and brain tumors [29, 38]. Most methods output voxel-wise labels using volumetric networks such as nnU-Net [13] or atlas-based registration approaches [2, 12, 37]. These representations assign a single label per voxel, making it impossible to directly model shared surfaces or multi-part structures. While this may not severely affect standard metrics like Dice, it often leads to topological errors requiring post-processing.

Topology, however, is crucial for downstream modeling. Hemodynamic simulations demand watertight, con-

nected geometries for stable CFD [31]; biomechanics modeling relies on anatomically continuous walls to predict stresses [28]; electrophysiological models require consistent atrial and ventricular meshes for accurate excitation propagation [19, 27]; and surgical planning or disease progression studies depend on coherent 3D anatomy [21, 26]. As large-scale digital twins emerge [22, 23], manual mesh repair usually involving voxel-to-surface conversion followed by manual correction [5, 10], becomes infeasible.

Implicit representations based on Signed Distance Fields (SDFs) [18, 32, 36] offer an alternative by learning latent shape priors refined to match voxel segmentations. However, these methods require topologically correct training data or geometric losses, lack hard topology guarantees, and remain computationally expensive since the refinement occurs at inference time.

Another direction is to directly predict triangulated meshes [4, 16, 34, 35] when internal structures are irrelevant, which yields accurate boundaries. When they are important as in the heart, separate templates are often used per chamber [4, 16, 17, 34, 35], sometimes with per-template scaling or translation. Independent deformation of these templates, however, causes interpenetrations and gaps that are difficult to fix post hoc.

In contrast, our approach deforms a single unified template encoding internal structures, enabling topologically consistent multi-component reconstructions in a single forward pass without post-processing.

3. Method

Our method reconstructs anatomically consistent multi-part organ surfaces from volumetric medical images using a single, unified mesh template. This template encodes the correct topology of the organ and is deformed using learned image features to fit a specific patient’s anatomy. The key components of our approach are:

1. A **topologically-correct mesh template** that models all substructures and their interfaces jointly.
2. A **feature-guided mesh deformation network** that aligns and adapts this template to medical imaging data.
3. A **training strategy with combined loss functions** that encourages both overall shape accuracy and precise inter-substructure surface modeling.

We describe each one below, using heart reconstruction for illustration purposes. The method applies similarly to hippocampus and lung reconstruction by simply changing the template we deform.

3.1. Building a Topologically Correct Template

Our algorithm starts with a template that captures the topological structure of the target organ. For the heart, this includes the four chambers—left/right ventricles and atria—connected by biologically plausible shared walls,

such as the myocardium and septa. Our approach to designing the template is to start from a basic primitive geometry that naturally supports defining interconnected components.

To this end, we use a rhombicuboctahedron [15], illustrated in Fig. 2(a). It has 26 faces, consisting of 8 equilateral triangles and 18 squares, which makes it a good approximation of a sphere while also being easy to connect to other rhombicuboctahedra by having them share their square faces. We exploit this property by positioning four rhombicuboctahedra that share faces as shown in Fig. 2(b). We then subdivide all faces into smaller facets labeled with a class label denoting one or more substructures. Finally we adjust the position of individual vertices to make the four chambers initially spherical. This yields the initial template of Fig. 2(c), which reflects the anatomical layout, that is, adjoining chambers, within a single, connected mesh.




Figure 2. **Building the heart template.** (a) Initial rhombicuboctahedron. (b) Four rhombicuboctahedra glued together. (c) Their vertices are uniformly displaced to create four spheres.

For the hippocampus and its two parts, we built the template of Fig. 3(a), using only two spheres sharing one intersection instead of four. For the lung template of Fig. 3(b), we used two spheres for the two lobes of the right lung and three connected ones for the three lobes of the left lung.




Figure 3. Templates for (a) the hippocampus and (b) the lungs.

3.2. Deforming the Template to Match the Data

A correct, unified template is our starting point. A core challenge then lies in deforming it to accurately reconstruct the anatomy of a specific patient from volumetric medical images, while preserving both topology and mesh regularity as described below. To this end, we use the two-stream network architecture of Fig. 4 to compute translation vectors for each template vertex by combining voxel-based semantic understanding with mesh-based surface generation, as in

Voxel2Mesh [34].

3.2.1. Network Architecture

The first stream is a 3D U-Net [25] that outputs a dense voxel-wise segmentation and multi-scale feature maps. These maps are fed to the second stream to guide the mesh deformation and to compute an affine transformation that coarsely aligns the template to the target anatomy. Because our template is not radially symmetric, this alignment is critical to avoid unnecessary deformation. We estimate the transform from the predicted class centroids of the segmentation; the computation is fast, differentiable, and performed at every training iteration.

The second stream is a mesh decoder that deforms the aligned template in a coarse-to-fine fashion. At each stage, features are sampled from the voxel stream and used to update vertex positions. Instead of learning curvature-aware unpooling as in [34], we adopt a simpler subdivision scheme with specific layers quadrupling the number of triangles by splitting each face uniformly. This allows the mesh to gradually increase resolution while preserving class labels on vertices and faces, including those on shared interfaces.

To implement both streams, we started from the publicly available Voxel2Mesh code [34] and modernized it by replacing the hand-designed adjacency matrix based graph convolutions by newer, more efficient *PyTorch 3D* versions [24]. This greatly reduced memory usage and running time, allowing us to work with larger input image resolutions and higher output vertex counts.

3.2.2. Preserving Topology

Unconstrained deformation would yield anatomically invalid results, especially for shared surface such as the septal walls between heart chambers. This is because each component is encouraged to match its own target shape based on local image evidence. Without constraints enforcing agreement across interfaces, these shared surfaces can be pulled in conflicting directions, leading to drift, misalignment, or collapse. To prevent this, we explicitly model and supervise the surfaces shared between anatomical structures, such as the wall between the left and right ventricles, and not just the individual structures. Each shared surface is represented as a distinct region in the unified template mesh and is assigned a dedicated label. This shifts the learning objective: The network is not simply encouraged to place adjacent structures near each other, but to reconstruct their exact shared interface. The result is a single, continuous surface that inherently preserves the topology of both the individual parts and their points of contact.

These regions are supervised directly during training, just like the primary anatomical components by minimiz-




Figure 4. **Network Architecture.** The voxel encoder and decoder follow a U-Net architecture including skip-connections, and are trained with cross-entropy loss. The initial heart template goes through a simple alignment procedure based on the segmentation, and is then deformed and subdivided in multiple steps by the mesh decoder. The mesh decoder samples features from the voxel decoder to apply correct deformations, and the shape loss is applied to each intermediate mesh.

ing a Chamfer-based distance shape matching loss

$$\mathcal{L}_{\text{match}}^l = \sum_{c=1}^C \mathcal{L}_{\text{Chamfer}}^{l,c}, \quad (1)$$

where the index l denotes the fact that the losses are evaluated after each intermediate stage of the shape decode and averaged, instead of only being computed on the final mesh.

Minimizing the Chamfer term drives geometric alignment between the predicted mesh and the ground truth. Since the mesh represents all substructures and their interfaces, we compute the Chamfer loss separately for each labeled region c , including both anatomical parts and intersection surfaces. At training time, we randomly sample points from both predicted and ground truth meshes and compute a symmetric Chamfer distance per region. This helps the model learn accurate geometry for both individual structures and shared surfaces. Each class is weighted equally to assign sufficient importance to the shared surfaces, whose contribution would otherwise be dwarfed due to being significantly smaller than the primary components.

3.2.3. Preserving Mesh Regularity

In addition to preserving topology, geometric integrity must also be maintained. Without it, one is likely to observe

vertex collapse, irregular triangle areas, and tangled regions, especially near points of contact between substructures. This would not only degrade reconstruction quality but also render the mesh unsuitable for downstream tasks such as simulation or quantitative analysis. Critically, once such degradations occur, they are difficult or impossible to repair in post-processing. To eliminate such artifacts, we introduce a geometric regularization loss $\mathcal{L}_{\text{reg}}$ that, when minimized, forces the preservation of mesh quality by ensuring overall smoothness, avoiding vertex collapse, and promoting a uniform triangle structure. This eliminates the need for post hoc alignment or mesh repair steps, which are typically error-prone and inconsistent. We take this loss to be

$$\mathcal{L}_{\text{reg}}^l = \lambda_1 \mathcal{L}_{\text{Edge}}^l + \lambda_2 \mathcal{L}_{\text{EdgeUnif}}^l + \lambda_4 \mathcal{L}_{\text{Norm}}^l + \lambda_5 \mathcal{L}_{\text{Lapl}}^l, \quad (2)$$

where l again denotes the stage of the mesh decoder, the λ coefficients are scalar weights, and $\mathcal{L}_{\text{Edge}}^l$, $\mathcal{L}_{\text{Norm}}^l$, and $\mathcal{L}_{\text{Lapl}}^l$ are standard regularization terms that we define in the supplementary. $\mathcal{L}_{\text{EdgeUnif}}^l$ is specific to our approach and penalizes edge lengths deviations to keep the triangles uniform.

$$\mathcal{L}_{\text{EdgeUnif}}^l = \sqrt{\frac{1}{E-1} \sum_{i=1}^E (|e_i| - \mathcal{L}_{\text{Edge}}^l)^2},$$

Because our component submeshes are open surfaces and do not lend themselves well to smoothing by regularizers designed for closed surfaces, we compute these losses on the full mesh, rather than for each class separately, as we do for $\mathcal{L}_{\text{Chamfer}}^{l,c}$ in Section 3.2.2. Shared vertices receive contributions from neighbors in different sub-meshes, with appropriate weights.

3.2.4. Global Loss

In the previous sections we introduced $\mathcal{L}_{\text{reg}}^l$ and $\mathcal{L}_{\text{match}}^l$ whose minimization promotes mesh regularity and topology preservation. To encourage an accurate fit to the training annotations, we also define $\mathcal{L}_{CE}$, the cross-entropy loss, and $\mathcal{L}_{\text{Dice}}$ is the Dice loss [9] with respect to the ground truth, both computed voxel-wise. When training the network, we minimize the global loss

$$\mathcal{L} = \mathcal{L}_{CE} + \mathcal{L}_{\text{Dice}} + \frac{1}{L} \sum_{l=1}^L (\lambda_{\text{match}} \mathcal{L}_{\text{match}}^l + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}^l). \quad (3)$$

4. Experiments

We use three datasets to validate our approach on heart, lung, and hippocampus reconstruction, all of which incorporate internal structures that need to be modeled accurately. The template we use for heart reconstruction has been introduced in Section 3.1 and is depicted by Fig. 2. The hippocampus and lung templates are shown in Fig. 3. Further preprocessing and architecture details are provided in the supplementary.

4.1. Datasets

MM-WHS Heart Dataset [40]. It contains 20 whole-heart annotated CT scans for training and 40 CT scans, without publicly available annotations. Seven classes are defined: the left ventricle blood cavity (LV-), the right ventricle blood cavity (RV), the left atrium blood cavity (LA), the right atrium blood cavity (RA), the myocardium of the left ventricle (Myo), the ascending aorta (Ao), and the pulmonary artery (PA). In our experiments, we only model the four heart chambers. Thus, we combine left ventricle blood cavity (LV-) and its myocardium (Myo) into a single left ventricle class (LV), and ignore the Ao and PA classes. We also define three *intersection classes* that correspond to the surfaces shared by adjacent substructures. Specifically we consider $\text{LV} \cap \text{RV}$, $\text{LV} \cap \text{LA}$, and $\text{RV} \cap \text{RA}$ but not $\text{LA} \cap \text{RA}$ because the atrial septum is not consistently labeled in the dataset due to the annotations being intended for blood

cavities rather than the complete atrial structure. This yields a total of seven classes that we explicitly model.

Without test set annotations, we cannot directly compute metrics on it. Instead, we perform 5-fold cross-validation on the training set, splitting into 5 sets of 16 training and 4 validation samples. We train separate models for each set and report the mean and standard deviation for each metric.

MSD Hippocampus Dataset [1]. It comes from the Medical Segmentation Decathlon and comprises 260 small voxel volumes from MRI scans, with the hippocampus split into its anterior and posterior components. We define Anterior $\cap$ Posterior as the only intersection class, which yields a total of three distinct classes. We split the dataset into 200 training samples and 60 test ones.

TS Lung Dataset [33]. For lung reconstruction, we use images and annotations from the TotalSegmentator data. It comprises 1228 CT scans annotated and refined using an active learning setup with a human-in-the-loop. We extract those having viable lung annotations, yielding 348 annotated lung scans. Five classes are defined, the lower, middle and upper right lobe, which we designate as LR, MR and UR, and the lower and upper left lobe, designated as LL and UL. We also consider the four intersection classes $\text{LR} \cap \text{MR}$, $\text{LR} \cap \text{UR}$, $\text{MR} \cap \text{UR}$ and $\text{LL} \cap \text{UL}$, yielding nine separate classes.

4.2. Baselines

We compare against three state-of-the-art baselines that epitomize the three main approaches used to produce 3D segmentations, that is, relying on volumetric segmentation, triangulated surface meshes, or tetrahedral meshes.

nnU-Net [13]. It is a state-of-the-art voxel segmentation model, recently updated with deeper residual models [14]. As such, it cannot generate self-intersections, but is still subject to inconsistencies such as "floaters", that is, connected components separated from the main model, leading to potential gaps and holes in intersecting regions.

MeshDeformNet [17]. It is a multi-class improvement over the original Voxel2Mesh [34] that focuses on the heart. Instead of using the same spherical template for each class, it uses a per-class translation and scaling factor in advance, which makes the template slightly different for each class. As a mesh-based method, it can generate erroneous intersections between classes, but not nnU-Net-like floaters.

MedTet [6]. It is a recent method that relies on deep-marching tetrahedra [30] instead of 3D meshes to model organs. This tends to yield better reconstruction accuracy thanks to an hybrid vertex grid and signed-distance approach. However, as the outputs are still independent tetrahedral meshes, it can still make the same topological errors as the mesh-based approaches.

4.3. Evaluation Metrics

The primary goal of our approach is to produce anatomically faithful 3D reconstructions with correct topology—maintaining the intended connectivity between structures, avoiding interpenetration, and preserving shared surfaces. These structural properties are essential for both clinical relevance and downstream applications such as simulation. However, they are not adequately captured by standard segmentation metrics, which typically focus on voxel-wise or surface-level similarity. To fully assess reconstruction quality, we report both **topology-aware metrics**, which directly evaluate structural correctness, and **standard metrics**, which provide complementary assessments of overall volume alignment. This dual evaluation highlights the limitations of conventional metrics in avoiding anatomical errors and underscores the strengths of our method in preserving topological integrity. We describe both kinds below.

4.3.1. Topology Metrics

Topological correctness is difficult to assess using standard metrics. Structural errors such as unintended overlaps or missing connections often have little effect on global scores like Dice or Chamfer, which average over entire volumes or surfaces. Localized shifts at boundaries may only impact a small number of voxels, but these small inconsistencies can lead to artifacts that compromise anatomical validity and limit the utility of the reconstruction in downstream tasks.

Such topological errors arise in different ways depending on the representation. Mesh-based methods may generate interpenetrating surfaces between structures that should only be adjacent, while voxel-based methods—particularly those that segment each structure independently—can leave small gaps or disconnected regions where continuity is required. Because enforcing topological correctness has not been a primary focus in the field of organ reconstruction, there are no standard metrics to quantify topological mistakes. Thus, we designed two such metrics to quantify structural inconsistencies at interfaces: one measuring the volume of mesh intersections, and another capturing the presence of undesired gaps between components.

Volume Intersections. For mesh-based methods that generate erroneous intersections, we use mesh boolean operations to extract the intersection between meshes corresponding to different classes and compute its volume relative to the total mesh volume.

Unwanted Gaps. Voxel-based methods can produce unintended empty regions inside the heart due to disconnected or misaligned surfaces. We quantify these artifacts by counting the number of isolated background (empty) connected components that are entirely enclosed within the structure—excluding the outer background. In this context, each such enclosed component corresponds to a “hole” or gap.

4.3.2. Standard Metrics

In addition to topology metrics, we report standard evaluation metrics commonly used in segmentation benchmarks. While these do not directly assess structural integrity, they offer a baseline for geometric similarity. They come in two main flavors, volumetric and surface-based. Thus we consider the following two.

Chamfer Distance. It is calculated between two point clouds. For the mesh-based methods, we obtain them by sampling 1024K points from the output surfaces. For the volumetric ground-truth data and the output of nnU-Net, we get them by running the Marching Cube algorithm [20]. We calculate it for every class, including interfaces between substructures. With intersection surfaces not explicitly modeled by either MeshDeformNet or MedTet, we use boolean operations to extract approximate ones from the meshes to compute Chamfer distances. All calculations are done with vertex coordinates normalized into the $[-1, 1]$ range, and we report values multiplied by 10^{-3} .

Normal Consistency [11]. It is calculated in a very similar way to Chamfer distance by comparing the orientation of surface normals at matching points rather than their distance. Based on the cosine distance, its best value is 1 when the normals match perfectly, and 0 when they are completely orthogonal. Normals at random points on the mesh surface are obtained by interpolating vertex normals.

Dice Coefficient [8]. It measures volumetric overlap between ground truth and prediction, while being more sensitive to class imbalance than the Jaccard index [3]. We compute it for each segmentation class. For mesh based methods such as ours, we first convert the mesh representation into a voxel-based one to compute this metric.

4.4. Comparative Results

Heart Reconstruction. In Tab. 1, we report per-class Chamfer Distance (CD) and Normal Consistency (NC) values for our method and the baselines on the MM-WHS-4 dataset both intersection classes—regions where anatomical structures meet—and other classes—the actual heart chambers. Lower CD and higher NC indicate better surface alignment and smoother geometry.

Our method achieves the best overall performance across the majority of intersection classes, yielding the lowest average Chamfer Distance and highest or near-highest Normal Consistency. Notably, it ranks first in 5 out of 6 CD columns and achieves excellent NC performance. This indicates that it not only captures accurate surfaces but also ensures smooth transitions in geometrically complex areas, which is critical when modeling realistic cardiac interfaces. In terms of volume reconstruction accuracy as reported in Tab. 2, our method performs comparably to nnU-Net, which is unsurprising because the differences between reconstructions occur mostly at the boundaries between substructures,

Table 1. Per-class Chamfer distance ($\times 10^{-3}$) and normal consistency on MM-WHS-4. *Blue denotes the best method and red the second best. Our method is consistently one of the two and, when it is second-best, it is only by a small margin.*

Method	Intersection Classes								Other Classes	
	LV $\cap$ LA		LV $\cap$ RV		RV $\cap$ RA		Average		Average	
	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$
Ours	0.7 $\pm$ 0.3	0.95 $\pm$ 0.01	1.4 $\pm$ 0.7	0.95 $\pm$ 0.01	1.9 $\pm$ 1.2	0.94 $\pm$ 0.02	1.3 $\pm$ 0.6	0.95 $\pm$ 0.01	1.0 $\pm$ 0.3	0.93 $\pm$ 0.01
MedTet	3.4 $\pm$ 2.3	0.96 $\pm$ 0.01	5.6 $\pm$ 3.4	0.95 $\pm$ 0.01	2.7 $\pm$ 2.0	0.95 $\pm$ 0.01	3.9 $\pm$ 2.4	0.95 $\pm$ 0.01	1.1 $\pm$ 0.3	0.94 $\pm$ 0.01
MeshDeformNet	2.4 $\pm$ 1.5	0.91 $\pm$ 0.02	6.0 $\pm$ 4.1	0.90 $\pm$ 0.05	7.5 $\pm$ 9.5	0.89 $\pm$ 0.04	5.3 $\pm$ 4.8	0.90 $\pm$ 0.03	2.5 $\pm$ 1.4	0.88 $\pm$ 0.02
nnU-Net	0.9 $\pm$ 0.4	0.90 $\pm$ 0.00	1.6 $\pm$ 0.6	0.88 $\pm$ 0.01	2.4 $\pm$ 1.4	0.89 $\pm$ 0.01	1.6 $\pm$ 0.7	0.89 $\pm$ 0.01	1.1 $\pm$ 0.3	0.87 $\pm$ 0.01

Table 2. Comparing Dice coefficients. *Our meshes are watertight. After voxelization, which effectively reduces precision, their volumetric accuracy is similar to that of nnU-Net.*

Method	MM-WHS					MSD-Hippocampus		
	LV	RV	LA	RA	Mean	Anterior	Posterior	Mean
Ours (Mesh)	0.95	0.90	0.92	0.88	0.91	0.87	0.85	0.86
nnU-Net	0.95	0.90	0.92	0.89	0.92	0.88	0.86	0.87

which only represent a small fraction of the total number of voxels. Furthermore, its voxel-based nature results in visible artifacts when turning the reconstructed volume into 3D surfaces. As can be seen in the third column of Fig. 5, nnU-Net outputs exhibit rough and uneven walls, especially in the "LV + Walls" and "LV (No Walls)" views. This translates into the comparatively lower NC score of Tab. 2. Such artifacts do not exist in real human anatomy. In reality, organs like the heart have naturally smooth and continuous surfaces, shaped by biological processes and soft tissue mechanics. This problem is not specific to nnU-Net; it is inherent to fixed-resolution voxel-based representations.

Conversely, MedTet performs significantly worse than us in terms of CD score, especially in challenging regions such as LV-RV and RV-RA, but outperforms us slightly in terms of NC scores in the intersection classes. This suggests that while MedTet produces smooth normals, it is less good at precisely aligning surfaces.

These observations are further supported by the topological metrics presented in Tab. 3, which quantify structural correctness. Our method achieves zero intersection volume and no unwanted gaps, indicating clean, watertight surface reconstructions. In contrast, nnU-Net, while free from surface self-intersections, introduces an average of nearly one unwanted surface gap per case—further highlighting the challenges of voxel-based reconstruction in topologically complex regions. MedTet and MeshDeformNet introduce overlapping volumes between structures, and MeshDeformNet frequently outputs degenerate meshes, making a correct volume calculation difficult. The number of gaps is ill-defined for these because they are made of several independent connected components.

In short, by jointly working with surface and volume representations, our method produces smooth, coherent, and




Figure 5. **Qualitative comparison of 4-chamber heart outputs.** Left column describes the anatomical view. Each row compares methods visually. Our method produces smoother and more consistent geometry, especially in edge-sensitive regions.

Table 3. Topology metrics on MM-WHS-4.

Method	Intersection Volume	# of Unwanted Gaps
Ours	0	0
nnU-Net	0	0.85 $\pm$ 0.02
MedTet	0.18% $\pm$ 0.02%	undefined
MeshDeformNet	degenerate meshes	undefined

anatomically plausible geometry, while closely matching what we expect from actual biological structures. This is especially important in sensitive areas like the LV-RV wall, where precise surface detail matters both visually and functionally.

Hippocampus and Lung Reconstruction. We report our hippocampus reconstruction results in Tab. 4 with a focus


Table 4. Per-class surface metrics on MSD-Hippocampus.

	Ant. $\cap$ Post.		Anterior		Posterior	
	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$
Ours	3.3	0.98	1.3	0.90	1.3	0.90
nnU-Net	3.5	0.98	1.3	0.80	1.2	0.80

Table 5. Per-class surface metrics on TS-Lung.


	Intersection Classes								Other Classes	
	LR $\cap$ MR	LR $\cap$ UR	MR $\cap$ UR	LL $\cap$ UL	Average	Average	Average	Average	Average	Average
Method	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$	CD $\downarrow$	NC $\uparrow$
Ours	3.79	0.96	1.03	0.98	1.04	0.96	2.88	0.94	2.18	0.96
nnU-Net	3.57	0.90	0.96	0.89	0.93	0.90	2.26	0.90	1.93	0.90

on the junction between anterior and posterior segments, along with quantitative results in Fig. 6. While all methods yield similar CD numbers, ours delivers significantly higher NC at this junction, indicating a smoother and more realistic transition. Our method’s ability to model this boundary as a clean, topologically consistent surface—without jagged artifacts or mesh discontinuities—demonstrates its strength in capturing anatomically realistic intersections, even in small, structurally delicate regions. We observe similar behavior for lung reconstruction results, as can be seen in Tab. 5, with superior intersection reconstruction compared to mesh-based methods, and superior normal consistency compared to voxel segmentation.


Figure 6. **Qualitative comparison of hippocampus outputs.** The ground truth remains irregular even after smoothing due to the low data resolution.

4.5. Effect of Training Set Size

In medical applications, large training sets are not always available and being able to train a network using a relatively small dataset is desirable. With this in mind, we ran the hippocampus and lung reconstruction experiments,

Figure 7. **Lung reconstructions.** *PrIntMesh* yields a smooth reconstruction free of residual and topological irregularities and with well-defined intersections. In contrast, the nnU-Net output is much rougher.

using only a fraction of the training database to train our approach and nnU-Net. We report the results in Fig. 8. In the low-data regime, our approach clearly outperforms nnU-net not only in NC terms but also in CD terms. This is due to the strong prior our approach imposes on the organ shape with the built-in template. While nnU-Net produces artifacts and disconnected regions in the low-data regime and greatly benefits from more data, our approach remains consistent even with few training samples.



(a) Lung reconstruction.

(b) Hippocampus reconstruction.

Figure 8. **Performance in the low-data regime.** With few training samples available, our approach yields more accurate surface reconstructions than nnU-Net. We also show the reconstruction performance of using a data template for the lung, which is very similar to the performance of using a hand-made template.

5. Conclusion

We introduced *PrIntMesh*, a unified framework for jointly reconstructing multiple organ components. Unlike existing surface-based methods that model each component using independent, potentially overlapping meshes, *PrIntMesh* deforms a single multi-part template that preserves topology while accurately modeling shared interfaces. This enables state-of-the-art reconstruction accuracy with guaranteed topological consistency.

Our approach moves toward fully automatic generation of anatomically correct models suitable for downstream simulations without manual post-processing. Future work will focus on automating template creation to handle more complex anatomies, such as incorporating the aorta and pulmonary artery and capturing their temporal deformations.

References

- [1] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. The Medical Segmentation Decathlon. *Nature communications*, 13(1):4128, 2022. 5
- [2] Wenjia Bai, Wenzhe Shi, Declan P O’regan, Tong Tong, Haiyan Wang, Shahnaz Jamil-Copley, Nicholas S Peters, and Daniel Rueckert. A Probabilistic Patch-Based Label Fusion Model for Multi-Atlas Segmentation with Registration Refinement: Application to Cardiac MRI Images. *IEEE Transactions on Medical Imaging*, 32(7):1302–1315, 2013. 2
- [3] J. Bertels, T. Eelbode, M. Berman, D. Vandermeulen, F. Maes, R. Bisschops, and M. Blaschko. Optimizing the Dice Score and Jaccard Index for Medical Image Segmentation: Theory and Practice. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2019. 6
- [4] F. Bongratz, A.-M. Rickmann, S. Pölsterl, and C. Wachinger. Vox2cortex: Fast Explicit Reconstruction of Cortical Surfaces from 3D Mri Scans with Geometric Deep Neural Networks. In *Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2
- [5] Jerome Charton, Stephen Baek, and Youngjun Kim. Mesh Repairing using Topology Graphs. *Journal of Computational Design and Engineering*, 2021. 2
- [6] Y. Chen, J. Yang, D. S. Mercadier, H. Le, and P. Fua. Medtet: An Online Motion Model for 4D Heart Reconstruction. In *arXiv Preprint*, 2024. 5
- [7] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit Fairing of Irregular Meshes Using Diffusion and Curvature Flow. In *ACM SIGGRAPH*, 1999. 1
- [8] R.L. Dice. Measures of the Amount of Ecologic Association Between Species. *Ecology*, 1945. 6
- [9] T. Falk, D. Mai, R. Bensch, Özgün Çiçek, A. Abdulkadir, Y. Marrakchi, A. Böhm, J. Deubner, Z. Jäckel, K. Seiwald, A. Dovzhenko, O. Tietz, C. D. Bosco, S. Walsh, D. Saltukoglu, T. L. Tay, M. Prinz, K. Palme, M. Simons, I. Diester, T. Brox, and O. Ronneberger. U-Net: Deep Learning for Cell Counting, Detection, and Morphometry. *Nature Methods*, 16:67–70, 2018. 5
- [10] B. Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012. 2
- [11] G. Gkioxari, J. Malik, and J. Johnson. Mesh R-CNN. In *International Conference on Computer Vision*, 2019. 6
- [12] Juan Eugenio Iglesias and Mert R Sabuncu. Multi-Atlas Segmentation of Biomedical Images: a Survey. *Medical Image Analysis*, 24(1):205–219, 2015. 2
- [13] F. Isensee, P.F. Jaeger, S.A.A. Kohl, J. Petersen, and K.H. Maier-Hein. Nnu-Net: A Self-Configuring Method for Deep Learning-Based Biomedical Image Segmentation. *Nature Methods*, 18(2):203–211, 2021. 1, 2, 5
- [14] F. Isensee, T. Wald, C. Ulrich, M. Baumgartner, S. Roy, K. Maier-Hein, and P. Jaeger. Nnu-Net Revisited: A Call for Rigorous Validation in 3D Medical Image Segmentation. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2024. 5
- [15] J. Kepler. *Harmonice Mundi*. Linz, 1619. 3
- [16] F. Kong and S. C. Shadden. Whole Heart Mesh Generation for Image-Based Computational Simulations by Learning Free-From Deformations. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 550–559, 2021. 1, 2
- [17] F. Kong, N. Wilson, and S. C. Shadden. A Deep-Learning Approach for Direct Whole-Heart Mesh Reconstruction. *Medical Image Analysis*, 74, 2021. 1, 2, 5
- [18] Hieu Le, , Nicolas Talabot, Jiancheng Yang, and Pascal Fua. Enforcing Topological Interaction Between Implicit Surfaces via Uniform Sampling. In *arXiv Preprint*, 2023. 2
- [19] Lei Li, Julia Camps, Blanca Rodriguez, and Vicente Grau. Solving the inverse problem of electrocardiography for cardiac digital twins: A survey. *IEEE Reviews in Biomedical Engineering*, 2024. 1, 2
- [20] W.E. Lorensen and H.E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In *ACM SIGGRAPH*, pages 163–169, 1987. 6
- [21] Daniel H Pak, Minliang Liu, Theodore Kim, Caglar Ozturk, Raymond McKay, Ellen T Roche, Rudolph Gleason, and James S Duncan. Robust automated calcification meshing for personalized cardiovascular biomechanics. *NPJ Digital Medicine*, 7(1):213, 2024. 2
- [22] Shuang Qian, Devran Ugurlu, Elliot Fairweather, Laura Dal Toso, Yu Deng, Marina Strocchi, Ludovica Cicci, Richard E Jones, Hassan Zaidi, Sanjay Prasad, et al. Developing cardiac digital twin populations powered by machine learning provides electrophysiological insights in conduction and repolarization. *Nature Cardiovascular Research*, 4(5):624–636, 2025. 2
- [23] Mengyun Qiao, Kathryn A McGurk, Shuo Wang, Paul M Matthews, Declan P O’Regan, and Wenjia Bai. A personalized time-resolved 3d mesh generative model for unveiling normal heart dynamics. *Nature Machine Intelligence*, pages 1–12, 2025. 2
- [24] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv Preprint*, 2020. 3
- [25] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 234–241, 2015. 3
- [26] Kensuke Sakata, Ryan P Bradley, Adityo Prakosa, Carolyn AP Yamamoto, Syed Yusuf Ali, Shane Loeffler, Brock M Tice, Patrick M Boyle, Eugene G Kholmovski, Ritu Yadav, et al. Assessing the arrhythmogenic propensity of fibrotic substrate using digital twins to inform a mechanisms-based atrial fibrillation ablation strategy. *Nature cardiovascular research*, 3(7):857–868, 2024. 1, 2
- [27] Kensuke Sakata, Carolyn AP Yamamoto, Adityo Prakosa, Brock M Tice, Syed Yusuf Ali, Shane Loeffler, Eugene G Kholmovski, Sunil Kumar Sinha, Joseph E Marine, Hugh Calkins, et al. Digital twins enable stratification of persistent atrial fibrillation patients for ablation diminishing unnecessary heart damage. *npj Digital Medicine*, 8(1):256, 2025. 1, 2

- [28] Matteo Salvador, Marina Strocchi, Francesco Regazzoni, Christoph M Augustin, Luca Dede', Steven A Niederer, and Alfio Quarteroni. Whole-heart electromechanical simulations using latent neural ordinary differential equations. *NPJ Digital Medicine*, 7(1):90, 2024. [1](#), [2](#)
- [29] D. She, Y. Zhang, Z. Zhang, H. Li, Z. Yan, and X. Sun. Eoformer: Edge-Oriented Transformer for Brain Tumor Segmentation. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2023. [2](#)
- [30] T. Shen, J. Gao, K. Yin, M.-Y. Liu, and S. Fidler. Deep Marching Tetrahedra: A Hybrid Representation for High-Resolution 3D Shape Synthesis. In *Advances in Neural Information Processing Systems*, 2021. [5](#)
- [31] Adam Updegrove, Nathan M Wilson, Jameson Merkow, Hongzhi Lan, Alison L Marsden, and Shawn C Shadden. Simvascular: an open source pipeline for cardiovascular simulation. *Annals of biomedical engineering*, 45(3):525–541, 2017. [1](#), [2](#)
- [32] J. Verhulsdonk, T. Grandits, F. S. Costabal, T. Pinetz, R. Krause, A. Auricchio, G. Haase, S. Pezzuto, and A. Effland. Shape of My Heart: Cardiac Models through Learned Signed Distance Functions. In *Medical Imaging with Deep Learning*, 2024. [2](#)
- [33] Jakob Wasserthal, Hanns-Christian Breit, Manfred T Meyer, Maurice Pradella, Daniel Hinck, Alexander W Sauter, Tobias Heye, Daniel T Boll, Joshy Cyriac, Shan Yang, et al. Totalsegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. *Radiology: Artificial Intelligence*, 2023. [5](#)
- [34] U. Wickramasinghe, E. Remelli, G. Knott, and P. Fua. Voxel2mesh: 3D Mesh Model Generation from Volumetric Data. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2020. [1](#), [2](#), [3](#), [5](#)
- [35] H. Yang, R. Tam, and X. Tang. Whole-Heart Reconstruction with Explicit Topology Integrated Learning. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 106–115, 2023. [1](#), [2](#)
- [36] J. Yang, E. Sedykh, J. Adhinarta, H. Le, and P. Fua. Generating Anatomically Accurate Heart Structures via Neural Implicit Fields. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2024. [2](#)
- [37] S. Yang, Z. Pan, T. Amert, K. Wang, L. Yu, T. Berg, and M. Lin. Physics-Inspired Garment Recovery from a Single-View Image. *ACM Transactions on Graphics*, 37(5):1–14, 2018. [2](#)
- [38] X. Zeng, P. Zeng, C. Tang, P. Wang, B. Yan, and Y. Wang. Dbtrans: A Dual-Branch Vision Transformer for Multimodal Brain Tumor Segmentation. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 502–512, 2023. [2](#)
- [39] Xukun Zhang, Yang Liu, Sharib Ali, Xiao Zhao, Mingyang Sun, Minghao Han, Tao Liu, Peng Zhai, Zhiming Cui, Peixuan Zhang, et al. Anatomical-Aware Point-Voxel Network for Couinaud Segmentation in Liver CT. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 465–474, 2023. [2](#)
- [40] X. Zhuang. Multivariate Mixture Model for Myocardial Segmentation Combining Multi-Source Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(12):2933–2946, 2018. [5](#)

PrIntMesh: Precise Intersection Surfaces for 3D Organ Mesh Reconstruction

Supplementary Material

A. Geometric Regularization Loss Terms

Edge Losses $\mathcal{L}_{\text{Edge}}$ and $\mathcal{L}_{\text{EdgeUnif}}$. We take them to be

$$\mathcal{L}_{\text{Edge}}^l = \frac{1}{E} \sum_{i=1}^E |e_i|, \quad (4)$$

$$\mathcal{L}_{\text{EdgeUnif}}^l = \sqrt{\frac{1}{E-1} \sum_{i=1}^E (|e_i| - \mathcal{L}_{\text{Edge}}^l)^2},$$

where E is the total number of edges in the mesh and e_i are its edges. Minimizing them discourages the elongation of edges by penalizing both their average length and deviations from this average length, thereby promoting the regularity of the mesh facets.

Normal Loss $\mathcal{L}_{\text{Norm}}$. We write it as

$$\mathcal{L}_{\text{Norm}} = \sum_{i=1}^F \sum_{j=1}^F \text{Adj}(f_i, f_j) (1 - \cos(n_{f_j}, n_{f_i})), \quad (5)$$

$$\text{Adj}(f_i, f_j) = \begin{cases} 1, & \text{if } f_i, f_j \text{ are adjacent faces} \\ 0, & \text{otherwise} \end{cases},$$

where F is the number of faces in the mesh and f_i, f_j are faces. Minimizing it further regularizes pairs of adjacent triangle faces by harmonizing their normals.

Laplacian Loss $\mathcal{L}_{\text{Lapl}}$. We take it to be

$$\mathcal{L}_{\text{Lapl}} = \frac{1}{V} \sum_{i=1}^V \left\| v_i - \frac{1}{|N_i|} \sum_{j \in N_i} v_j \right\|,$$

where v_i are the vertices, V the number of vertices in the mesh and N_i the set of indices of the vertices neighboring v_i . Minimizing it promotes smoothness by pushing each vertex towards the mean of its neighbors [7].

B. Architecture and Preprocessing

For our segmentation network, we use nnU-Net in our experiments on the hippocampus and lung, with 5 levels containing [32, 64, 128, 256, 320] channels for the lung and 4 levels for the hippocampus. For our heart experiments, we use a smaller UNet with fewer channels which we find faster to train for comparable results. For the GCN, we use the same structure as in [34] with a hidden dimension of 32 and as many layers as the segmentation network decoders, and using learned neighborhood sampling [34]. We use the Adam optimizer with a learning rate of 10^{-4} , $\beta_1 = 0.9$ and

Table S1. regular template compared to a data template on MMWHS-4. Results are similar in terms of CD ($\times 10^{-3}$) despite visual differences.

Template	Intersection Classes				Other Classes
	LV $\cap$ LA	LV $\cap$ RV	RV $\cap$ RA	Average	Average
Sphere-Based	0.68 $\pm$ 0.44	1.85 $\pm$ 1.12	2.30 $\pm$ 1.24	1.61 $\pm$ 0.68	1.18 $\pm$ 0.30
Data-Based	0.65 $\pm$ 0.30	1.65 $\pm$ 0.77	2.60 $\pm$ 1.30	1.63 $\pm$ 0.55	1.24 $\pm$ 0.22

$\beta_2 = 0.999$. We set loss multipliers on a per-dataset basis by picking the best from a random search. All experiments are performed using a single V100 GPU.

For the MM-WHS-4 dataset, We preprocess each CT scan by cropping each heart around its bounding box with a 10% extra margin added around the cube for padding. The resulting volume is resized to $128 \times 128 \times 128$ and values are normalized per instance to have zero-mean and unit standard deviation.

For the MSD hippocampus we resize each volume to $64 \times 64 \times 64$ and normalize gray-level values using the mean and standard deviation of the whole training dataset.

For the TS Lung dataset, we resize all volumes to $128 \times 128 \times 128$, and use a random 260/28/60 train/validation/test split for our experiments, with values being normalized per-instance.

C. Ablation Studies

C.1. Data-Based Template Generation

We use the hand-designed templates of Figs. 2 and 3, which we have to re-define for each new organ. Instead, we could derive them directly from labeled ground truth data. To demonstrate this, we triangulated a representative lung segmentation from MM-WHS-4, and decimated the result to generate approximately as many vertices as in our hand-designed template. We then trained our model to use this new template as we did before and evaluate on 4 out of 5 splits, to avoid evaluating on the split the segmentation comes from. In terms of reconstruction loss, using the data template underperforms only very slightly as shown in Tab. S1, but the reconstructed meshes are significantly less regular as shown in Fig. S1.

C.2. Shared-Surface Supervision

While we advocate training by separately supervising the shared surfaces, we could instead use a simpler setup where we keep the unified template, but supervise only the main classes. In this case, the shared surface components are supervised additively by both main classes they belong to. We report the resulting reconstruction performance on a single




Figure S1. **Regular vs Data-Based Template.** (a) Reconstruction using the hand-crafted template of Fig. 3(b). (b) Reconstruction using a data-based template. While it yields similar numerical reconstruction performance is similar, the triangles in the data-based version are much less regular and the resulting meshes less suitable for downstream applications, such as simulation.

Table S2. The effect of not supervising shared-surfaces in MMWHS-4. *Performance is greatly reduced in terms of Chamfer distance ($\times 10^{-3}$)*

Template	Intersection Classes				Other Classes
	LV $\cap$ LA	LV $\cap$ RV	RV $\cap$ RA	Average	Average
Shared-surface	0.31	2.59	0.75	1.22	1.05
Main-surface	0.57	22.07	2.06	8.23	1.09

split of MM-WHS-4 in Tab. S2. There is a significant drop in the performance of the shared-surfaces. This degradation arises from the ambiguity and loss of identity of the shared surfaces during training. As they are not explicitly supervised, parts of the main classes can "leak" to cover parts of the ground truth shared surface, leaving the shared surface partially reconstructed in the deformed template.