

Deep Medial Voxels: Learned Medial Axis Approximations for Anatomical Shape Modeling

Antonio Pepe, Richard Schussnig, Jianning Li, Christina Gsaxner, Dieter Schmalstieg, *Fellow, IEEE*, and Jan Egger

Abstract—Shape reconstruction from imaging volumes is a recurring need in medical image analysis. Common workflows start with a segmentation step, followed by careful post-processing and, finally, ad hoc meshing algorithms. As this sequence can be time-consuming, neural networks are trained to reconstruct shapes through template deformation. These networks deliver state-of-the-art results without manual intervention, but, so far, they have primarily been evaluated on anatomical shapes with little topological variety between individuals. In contrast, other works favor learning implicit shape models, which have multiple benefits for meshing and visualization. Our work follows this direction by introducing deep medial voxels, a semi-implicit representation that faithfully approximates the topological skeleton from imaging volumes and eventually leads to shape reconstruction via convolution surfaces. Our reconstruction technique shows potential for both visualization and computer simulations¹.

Index Terms—convolutional neural networks, convolution surface, distance function, centerline, medial axis transform, implicit model

I. INTRODUCTION

THE reconstruction of anatomical regions is a prerequisite for numerous tasks in medical image analysis, ranging from visualization [1], [2] to implant design [3] and physics-based simulations of pathological processes [4], [5]. Traditionally, this is performed after the initial binary segmentation of an imaging volume [1], [3]. With an emphasis on blood vessels, different studies and challenges analyzed the need for a smooth, organic-looking and artifact-free representation in clinical visualizations [1], [2] and patient-specific computer simulations [1], [6]–[8]. The recent SEG.A. challenge² showed that the quantitatively best-performing

algorithms were not necessarily the solution of choice for clinicians and biomedical simulations. Less attention has been given to methods that do not only guarantee a correct shape reconstruction, but also address the additional requirements needed for visual assessment and computer simulations. Some studies have successfully analyzed the role of *skeleton-based* implicit representations and how they can play a key role in both domains [1], [2], [8]. These representations satisfy different needs. Implicit representations describe shapes analytically in the form of level sets. The analytical formulation allows to enforce the surface continuity, a requirement for most computer simulations [7], [8], while the skeleton allows one to quantify vascular growth [9] but also to parametrically³ model and mesh the geometry [1], [8] for use cases like in-silico trials [10]. In these studies, the skeletons were simply represented as centerlines and first required binary segmentation. While centerlines provide a sufficient approximation for tubular shapes, like blood vessels and nerves, a more generalized skeletal structure is needed for non-tubular shapes. A possible solution is given by the medial axis transform (MAT), which represents a super-set of the classic centerline..

Additionally, recent work showed that discretized distance functions, a type of implicit representation, can be learned by neural networks at least as efficiently as binary representations for different organs as they tend to outperform other methods that only predict a binary output [11], [12]. Less attention has been paid to learning implicit representations based on skeletal structures, such as the MAT-based representations [12]–[14]. In this paper, we introduce *deep medial voxels* (DMV), the first effort to create parametric meshes directly from imaging volumes, with several advantages:

- 1) No explicit binary segmentation step is required; instead, a segmentation of the underlying geometry is intrinsically provided as a level-set representation.
- 2) We generate a skeleton mesh and apply surface convolution [15] to obtain a surface mesh. Due to the parameterization provided by the skeleton, the surface is smooth, watertight and unaffected by voxel discretization.
- 3) A differentiable formulation of a shape skeleton, which allows one to analyze topology information during training and back-propagate a topological loss.

Manuscript received XXX; revised XXX; accepted XXX. Date of publication XXX; date of current version XXX. This work received funding from the TU Graz LEAD Project *Mechanics, Modeling and Simulation of Aortic Dissection*. (Corresponding author: Antonio Pepe.)

A. Pepe and C. Gsaxner are with the Graz University of Technology, Institute of Computer Graphics and Vision, Inffeldgasse 16a, 8010 Graz, Austria. (e-mail: pepe@student.tugraz.at, gsaxner@tugraz.at).

R. Schussnig is with the Ruhr University Bochum, Faculty of Mathematics, 44780, Bochum, Germany. (e-mail: richard.schussnig@rub.de).

D. Schmalstieg is with University of Stuttgart, VISUS, Allmandring 19, 70569 Stuttgart, Germany (e-mail: dieter.schmalstieg@visus.uni-stuttgart.de).

J. Li and J. Egger are with the Essen University Hospital, Institute for Artificial Intelligence in Medicine, Girardetstraße 2, 45131 Essen, Germany. (e-mail: {jianning.li, jan.egger}@uk-essen.de).

¹Code available at <https://github.com/apepe91/dmv>

²SEG.A. 2023 <https://multicenteraorta.grand-challenge.org/>

³A parametric representation is here considered as a lower-dimensional description of a surface, e.g., skeleton-based parametrization.

Moreover, DMV allows to

- 4) Co-segment the reconstructed shape by clustering the predicted skeleton
- 5) Identify in- and outflow interfaces for computer simulation of vasculature by analyzing the skeleton connectivity
- 6) Produce meshes that can be readily applied in numerical simulations

II. BACKGROUND

Conventionally, polygonal surface meshes of segmented volumes are generated with algorithms such as marching cubes [3], [16] and often suffer from discretization artifacts (Figure 1) [7], [17] that limit their use for visualization [2], [3] and simulation [7]. Various meshing and editing methods have been designed to overcome such limitations. Some of them were integrated into software such as SimVascular and VMTK [4], [18]. A recent trend is to model anatomical shapes by learning differentiable prior deformations [19], [20], bypassing binary segmentation steps without losing accuracy. Yet, such learned mesh deformations operate at a fixed resolution and require a shape prior. As an alternative, we rely on implicit representations, which have been shown to provide optimal quality in both visualization and simulation applications [1], [6]. Here, we provide an overview of implicit shape modeling, introduce the medial axis transform as a shape description tool, and discuss how our work is related to the current state of the art in learned shape modeling.

A. Implicit shape modeling

Implicit representations describe shapes as a level set of a function $f(P) : \mathbb{R}^3 \rightarrow \mathbb{R}$ [17] and a parameter $C \in \mathbb{R}$:

$$f(P) - C = 0. \quad (1)$$

Examples are distance transforms [11], metaballs [21], and convolution surfaces [2]. Their analytical formulation allows the extraction of high-resolution meshes without discretization artifacts [2], [22]. They can easily model branching structures and ensure smooth surfaces by relying on convolution operations. Given a skeletal structure $\Gamma = \bigcup_i \Gamma_i$, described as the union of a set of primitives (e.g., polylines) with thickness R_i , a convolution surface is defined as

$$f(P) = \sum_i \int_{\Gamma_i} e^{-k \frac{\|P-s\|_2^2}{R_i^2}} ds, \quad (2)$$

where the convolution kernel (e.g., e^{-kx}) is \mathcal{C}^0 and strictly decreasing [2], [15]. A more detailed description of convolution surfaces was recently provided by Suárez et al. [15].

Here, we use skeletons represented as a *simplicial complex* (Figure 2), a topological representation that is gaining popularity in different fields [23]. A simplicial complex is a non-manifold mesh formed by a set of vertices, lines, and triangles. Each subset of $n + 1$ interconnected vertices forms an n -dimensional simplex and is the basic unit of the complex [24].

B. Medial axis transform

Centerlines are common for the topological analysis of tubular shapes [1], [2], [11]. However, centerlines represented as splines or polylines do not hold enough descriptive power to enclose all the geometric features of generic shapes with no tubular prior. An alternative framework is provided by the *medial axis transform* (MAT). The MAT can be thought of as a generalization of centerlines to generic shapes. The MAT of a closed surface S is defined as the set of all maximally inscribed spheres such that each sphere is tangent to S in at least two points [25]. For the case of an ideally tubular shape, the MAT corresponds to its centerline. However, the ideal MAT may capture too many details for more complex shapes. For example, in a 3D reconstruction even small amounts of noise can cause large variations in the MAT [26]. Various studies have investigated more stable and simpler approximations of the MAT [25]–[28]. An early approximation, λ -MAT, considers only spheres with radius $r \geq \lambda$. Its drawback is that topological changes are introduced if regions with a maximum medial radius $r_M \leq \lambda$ are present [25].

The *scale axis transform* considers exactly those spheres which are centered on the topological skeleton [25]. Another robust representation is *LS-MAT* [29], which estimates a set of MAT spheres through a least-squares optimization. *Q-MAT* approximates the MAT through a *simplicial complex* [26] with a user-defined number of vertices. Each simplex, line or triangle models a truncated cone or a medial slab that approximates the volume of the removed medial spheres that were centered along the simplex.

Skeletal structures based on simplicial complexes not only allow to describe shapes with a limited amount of parameters, but they have also proven to be useful in shape co-segmentation [30] and landmark detection [14]. *Voxel cores* [28] define a MAT over a binary voxel grid by its distance transform. The distance value d_i of a voxel v_i represents the radius of the medial sphere centered in v_i . A threshold is used to prune all spheres with radius $r \leq \lambda$ that do not lead to a topological change. Here, we extend these notions and develop a differentiable approximation of the MAT for a topology-aware training of neural networks.

C. Learning shape representations

Implicit neural representations were introduced to encode a given geometric shape within the weights of a neural network [31], [32] or to infer a parametric description of a given mesh [33]. Point2Skeleton [14] and P2MAT-Net [34] infer a MAT approximation from a sparse surface point cloud. Their output is represented as a simplicial complex, which allows reconstructing surfaces as the union of medial primitives [26]. For shape *completion*, such methods show advantages over the common Poisson reconstruction [35]. Learning strategies have also been suggested for the reconstruction of shapes from volume grids [36]. Volume grids are used extensively in medical imaging. For shapes such as the hippocampus [19], learned template deformations generate explicit meshes from imaging volumes [19], [20], [37].

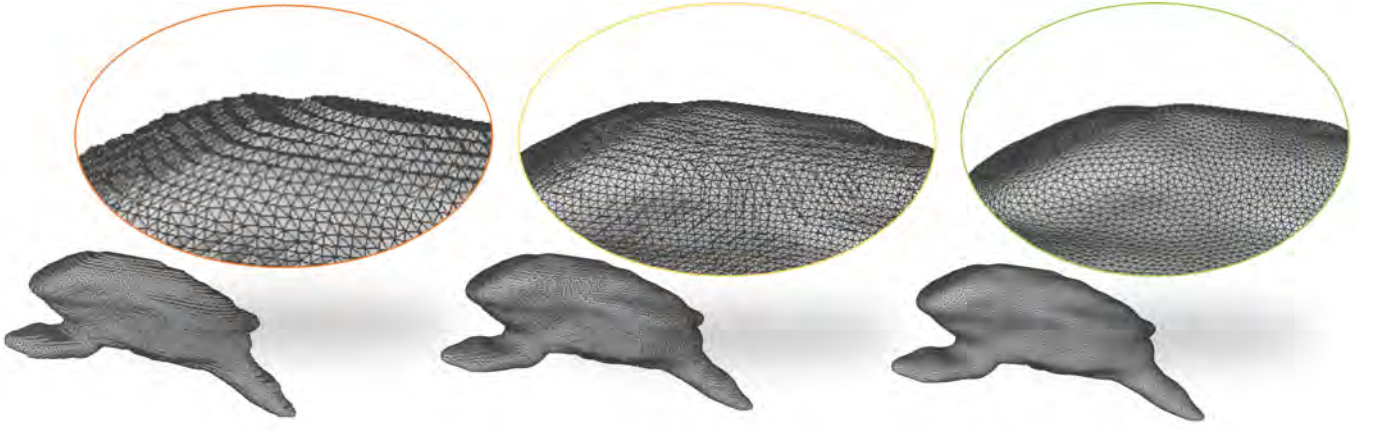


Fig. 1. A surface mesh is obtained (left, orange) directly from a binary segmentation using marching cubes, (middle, yellow) from a discrete distance field generated from the binary segmentation. In both cases, voxel artifacts can be noticed in different amounts. (right, green) The convolution surface of the medial voxels obtained with our differentiable DMV approach.

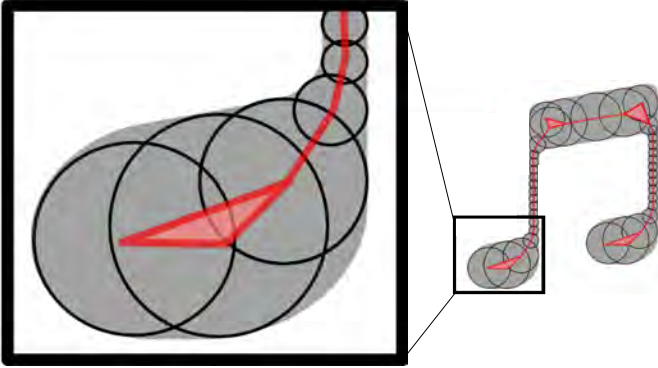


Fig. 2. A simplicial complex (red) over a 2D shape (gray). The black circles are a representative subset of the medial axis transform. If two circles overlap, their centers are connected by a line. If three circles overlaps, their centers are connected by a triangle. The same concept holds for 3D shapes where circles are replaced by spheres.

Learning to infer implicit representations from such volumes has received less attention, although some literature suggests a high potential [38]. An example is *Deep Distance Transform* (DDT) [11], which infers the distance transform of *tubular* shapes, like the aorta. Li et al. [12] learn a normalized distance function from MRI data to reconstruct the left cardiac atrium. They show that adversarial learning of distance functions can improve the generalization of the network. More recently, Lin et al. [39] suggested a skeleton-aware distance transform for 2D instance segmentation by learning a weighted energy functional derived from the distance transform. Their work showed the importance of considering topological information when learning implicit representations. Zhang et al. [40] showed the importance of preserving topology when learning to segment tubular structures such as coronary arteries. However, as remarked by Lin et al. [39], none of these methods infers a parametric 3D representation, and all are limited to a specific organ as they were each designed for a specific class of shapes. We fill this gap by learning differentiable MAT approximations from volumes and reconstruct the original shape by means of

convolution surfaces.

III. COMPUTING MAT FROM IMAGING VOLUME

In the first stage, our method accepts a raw voxel grid as input and automatically computes an approximate MAT of the target shape represented in the input volume. The algorithm approximates the MAT with configurable fidelity, supporting the extraction of meshes with different levels of detail.

A. CNN backbone and UDF regularization

In a supervised setting, a U-Net-like CNN backbone is trained to jointly infer binary segmentation and UDF from a raw imaging volume. UDF levels are binned into K classes, with class 0 representing the background.

The CNN, fed with a 3D patch, is trained to predict the discretized UDF. In our implementation, given a volume V delimited by the surface S and a point $P = (x, y, z) \in \mathbb{R}^3$, we define the UDF as

$$u(P) = \begin{cases} \lfloor \inf_{k \in S} \|P - k\|_2 + 0.5 \rfloor, & \text{if } P \in V \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

The UDF is then multiplied by the predicted binary segmentation to remove spurious noise. Now, for the definition of UDF, each voxel v with distance $u(v) > 0$ must have at least one neighbor $w \in \mathcal{N}(v)$ with distance $u(w) = u(v) - 1$. To guarantee spatial coherence in the prediction, we recompute the output as

$$u_q(v) = \min_{w \in \mathcal{N}(v)} (u(w)) + 1. \quad (4)$$

B. Differentiable relaxation of MAT

Unlike previous work [11], [12], [19], [38], we promote the preservation of topology during training by introducing a differentiable approximation of MAT. The definition of MAT implies that the UDF will present a local maximum, or *ridge*, at the center of each sphere, where the local UDF value is the radius of the sphere [41]. However, depending on the

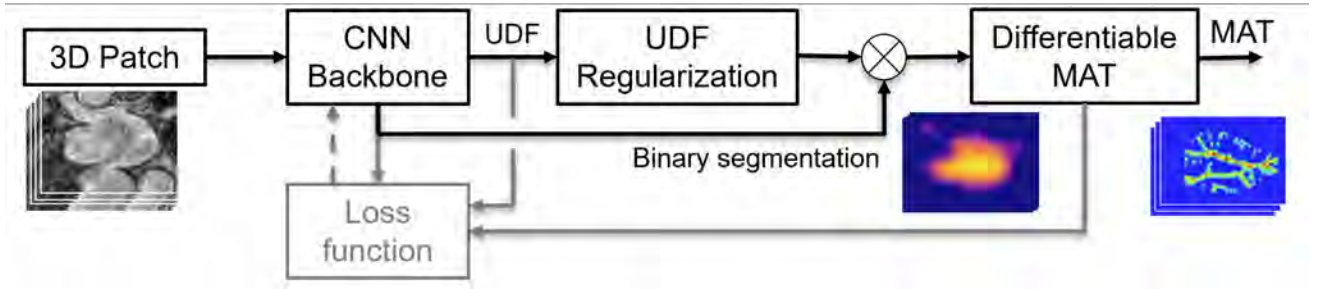


Fig. 3. Illustration of our pipeline. A CNN backbone is used to predict the binary segmentation and the UDF of the target organ. The UDF is regularized to ensure local continuity and a differentiable MAT is extracted, which provides topology information in the loss function. The gray steps are only executed during training.

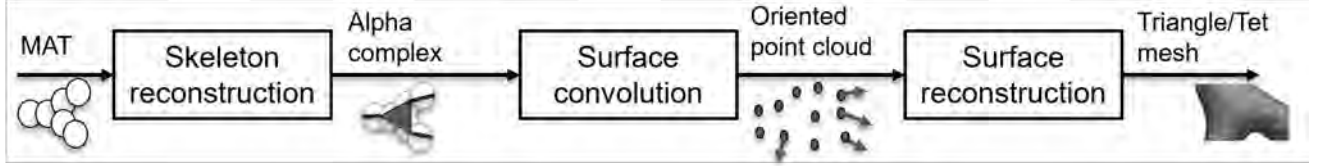


Fig. 4. Following the prediction of the MAT as shown in Figure 3, we convert the MAT to a weighted alpha complex. The convolution surface of the complex provides an oriented point cloud of the surface shape. This point cloud allows reconstructing both triangular and tetrahedral mesh representations.

resolution of the grid, local maxima may not be present for smaller structures or may be neglected by the CNN. Therefore, we introduce a relaxation of this definition and retrieve all MAT *candidate points* that are locally dominant

$$C_{rel} = \{(P, r) : u(P) > \mu(\mathcal{N}_u(P)) \wedge r = u(P)\}, \quad (5)$$

where μ is the average UDF value in a neighborhood \mathcal{N}_u of P . This definition includes all local maxima $C \subseteq C_{rel}$. Any non-maximally inscribed sphere will not negatively impact the shape reconstruction. To allow differentiation during neural network training, we approximate the inequality in Equation 5 as

$$C_d = \{(P, r) : \sigma(u(P) - \mu(\mathcal{N}_u(P))) > 0 \wedge r = u(P)\}, \quad (6)$$

where $\sigma(x) = (1 + e^{-kx})^{-1}$ is the scaled Sigmoid function. For $k = 50$, $\sigma(x)$ provides a good differentiable approximation of the step function [42].

C. Topology and ridge-preserving loss

Previous work focused only on distance values [11], [12] or, for tubular structures, on a soft centerline to preserve topology [43]. For more general shapes, we must preserve the entire MAT.

$$\mathcal{L}(\hat{y}_s, \hat{y}_u, y_s, y_u) = l_s(\hat{y}_s, y_s) + l_u(\hat{y}_u, y_u) + l_L(\hat{y}_u, y_u) + l_m(\hat{y}_u, y_u) \quad (7)$$

where \hat{y}_s is the predicted segmentation, \hat{y}_u is the UDF, and y_s and y_u are the respective ground truths. The function $l_s(\cdot, \cdot)$ is the Dice score similarity loss. For the UDF, the function $l_u(\cdot, \cdot)$ is a stochastic distance loss

$$l_u(\hat{y}, y) = - \sum_{i=1}^N \sum_{k=1}^K \left[\mathbf{1}(\hat{y}_{s,i} = k) \cdot \mathcal{P}(\hat{y}_{s,i} = k | y_{s,i}) \right], \quad (8)$$

where $\mathbf{1}(\cdot, \cdot)$ is an indicator function, and $\mathcal{P}(\hat{y}_{s,i} = k | y_{s,i})$ is the conditional probability of $\hat{y}_{s,i}$, given $y_{s,i}$ and a normal distribution with variance $\sigma = 1$. To preserve UDF continuity, we minimize the L_2 distance

$$l_L(\hat{y}_u, y_u) = L_2(\nabla^2 \hat{y}_u, \nabla^2 y_u) \quad (9)$$

of the Laplacian

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}. \quad (10)$$

Finally, we evaluate the MAT candidate points to preserve the topology as

$$l_m(\hat{y}_u, y_u) = l_u(C_d(\hat{y}_u), C_d(y_u)), \quad (11)$$

which increases the UDF loss for the MAT points $C_d(\cdot)$, the topological skeleton.

IV. COMPUTATION OF MESH FROM MAT

In the second part of our method, we use the predicted MAT to generate a skeleton and, from this skeleton, reconstruct the entire shape.

A. Skeleton generation

The MAT is a cloud of 4D points (x, y, z, r) , i.e., the ridges of the discrete UDF (Figure 4). To generate a skeleton, different methodologies on how to link medial spheres have been suggested [14]. The Nerve theorem guarantees that a cloud of spheres can be dually represented as a Čech complex, a simplicial complex where two overlapping spheres are connected by an edge, and, three, by a triangle [24]. As Čech complexes are computationally expensive, research focused on how to approximate them [14], [27], [34]. We approximate the Čech complex by building a *Laguerre-Voronoi diagram* from

the MAT points [27]. If the cells of three points are adjacent in pairs in the diagram, they form a triangle; two form a segment. This form leads to a simplicial complex known as a weighted alpha complex [24]. As the complete MAT can present a large number of elements, we rely on Q-MAT [26] to simplify the simplicial complex if necessary.

B. Shape reconstruction as convolution surface

The MAT and the skeleton provide a sparse and lightweight representation of the underlying shape. For shape reconstruction, we extend the notion of convolution surfaces to simplicial complexes. Convolution guarantees shape continuity of C^0 and higher orders; this property improves visual quality [2] and is a frequent prerequisite for finite element analysis [8], [22]. Convolution over lines and triangles requires different formulations of the partition integrals [44] given in Equation 2.

Segment convolution: A line segment \mathbf{AB} with vertex radii $r_A, r_B \in \mathbb{R}^+$ can be described in the form $s(x) = A + lx\mathbf{u}$, where l is its length and \mathbf{u} its versor. The integration over the segments for a point $P \in \mathbb{R}^3$ is therefore

$$f_{line}(P) = \int_0^1 K(\|P - s(x)\|_2, r_s) dx, \quad (12)$$

where K is the convolution kernel, and r_s is the interpolation of the radii at $s(x)$.

Triangle convolution: A general triangle T is split into two right-angled triangles T_1 and T_2 . These triangles are computed independently as

$$f_{triang}(P) = \int_0^1 \int_1^{1-y} K(\|P - P_{xy}\|_2, r_{xy}) dx dy, \quad (13)$$

where $P_{xy} \in \mathbb{R}^3$ is a point in the triangle with normalized projections x and y on the two catheti, and $r_{xy} \in \mathbb{R}$ is a bilinear interpolation of the radii at P_{xy} . Both cases rely on the kernel function $K(d, r) = r^2/(d^2 + r^2)$. The desired level set has distance $d = r$ and, therefore, following Equation 1, $C = K(r, r) = 0.5$.

Mesh generation: The convolution surface definition allows us to derive explicit surface and volumetric mesh representations. We do this in two steps:

1) We compute the convolution field on a volume grid. We extract the *oriented point cloud* from the voxels of level set C , where the point normal is opposite to the convolution surface gradient. To account for grid discretization, we optimize the point location along its normal according to the scheme of Algorithm 1, in which the hyperbolic tangent both modulates the step size and inverts the orientation if the surface level has been exceeded.

2) We rely on local Poisson reconstruction for watertight surface triangulation of the oriented point cloud [45], optionally followed by volume tetrahedralization [4], [46].

C. Differentiability

Our pipeline is trained to compute the MAT from an imaging volume. The CNN infers the SDF of an organ, and an approximation of the MAT is extracted by means of a differentiable filter. For validation and loss computation,

Algorithm 1 Iterative point cloud refinement.

ω : maximum tolerated error, C : surface level, k : maximum step size.

for each: cloud point P

0: $\mathbf{N} = -\frac{\partial \hat{S}(P)}{\partial P}$

0: **do:**

0: $P = P + k \cdot \tanh(\hat{S}(P) - C) \cdot \frac{\mathbf{N}}{\|\mathbf{N}\|}$

0: **while** $\|\hat{S}(P) - C\| \geq \omega = 0$

the SDF and MAT of the ground truth segmentation masks were precomputed. The differentiable MAT filter allowed us to focus the loss on topologically-relevant voxels despite possible topology changes of the MAT. The second part of the pipeline, which converts the MAT into a tetrahedral mesh, is designed as a post-processing step. As long as a differentiable convolution kernel is used, the post-processing is also differentiable. However, differentiation of the simplified mesh skeleton would increase memory requirements, as differentiation on both the voxel and mesh structures is required. We leave this optimization to future developments.

V. EXPERIMENTS

Following previous work [11], we carried out different experiments on four anatomical shapes and compared the results with four methods [11], [12], [19], [47] with a five-fold cross-validation. All experiments were developed using PyTorch v.1.8.0 on a desktop PC (CPU: Intel i7 3.20 GHz, 64 GB RAM, GPU: NVIDIA Titan RTX).

A. Datasets

We tested our method on labeled public datasets of 154 left cardiac atria (MRI, *2018 Left Atrium Segmentation Challenge* [48]), 270 hippocampi (MRI, *Medical Segmentation Decathlon* [49]), 40 aortae (CT, *SegThor Challenge* [50]), and 100 aortic dissections (CT, *Yao et al.* [51]). All datasets allow usage for research purposes. Volumes were resampled to an isotropic resolution of 1 mm and randomly augmented during training with rotation and cropping.

B. Ablation study and quantitative analysis

We performed an ablation study of the components of the loss function (Table I) on the left atria dataset. We used five-fold cross-validation, learning rate $l_r = 10^{-2}$, patch size $112 \times 112 \times 80$, batch size 8 and stochastic gradient descent. V-Net [47] was used as a backbone, as it previously showed good performance in learning implicit representations [12]. Ablation showed a clear contribution of both Laplacian and MAT information to the loss function. Regularization of the UDF further increased the prediction quality. This result shows how joint topology and gradient information provide a more robust training than explicit shape matching only. With this final configuration, we tested DMV against the original backbone (V-Net) and the two most relevant methods based on implicit models (DDT [11], SASSNet [12]) in different modalities (Table II). Furthermore, we compared DMV with

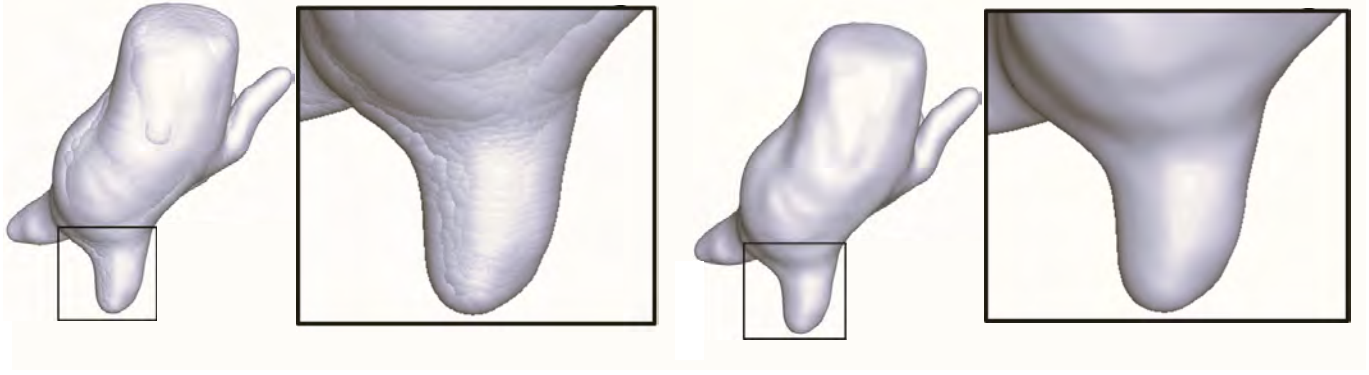


Fig. 5. Shape reconstruction from the MAT is usually computed as the union of medial primitives [14], [26]. (left) The union operation does not guarantee surface continuity and smoothness, leading to a series of “bumps”. (right) Instead, we combine the MAT with a convolution surface. The convolution guarantees a smooth and watertight reconstruction.

0: TABLE I

ABLATION STUDY. THE BASELINE CONSISTS OF A V-NET [47] BACKBONE TRAINED USING THE DICE SCORE FOR THE BINARY OUTPUT AND THE MEAN SQUARED ERROR FOR THE UDF OUTPUT. ALL TESTS ARE EVALUATED ON THE LEFT ATRIUM DATASET USING A FIVE-FOLD CROSS VALIDATION.

Method	Dice Score \uparrow [%]	Hausdorff Distance \downarrow [mm]
Baseline (BL)	85.53 \pm 9.86	12.24 \pm 6.01
BL + Laplacian Loss (LAPL)	88.23 \pm 8.67	8.01 \pm 5.47
BL + MAT Loss (MATL)	89.80 \pm 7.44	7.38 \pm 5.06
BL + MATL + LAPL	91.25 \pm 4.29	5.47 \pm 4.03
BL + MATL + LAPL + UDF Reg. (DMV)	96.02 \pm 3.20	3.64 \pm 2.93

TABLE II

COMPARISON AGAINST STATE-OF-THE-ART METHODS FOR DIFFERENT ORGANS AND MODALITIES. AO: AORTA. AD: AORTIC DISSECTION. LA: CARDIAC LEFT ATRIUM. HI: HIPPOCAMPUS.

Method	Dice Score \uparrow [%]			Hausdorff Distance \downarrow [mm]		
	AO, CT	AD, CT	LA, MRI	AO, CT	AD, CT	LA, MRI
V-Net [47]	78.53 \pm 8.01	85.33 \pm 6.54	89.01 \pm 12.23	19.16 \pm 11.51	32.60 \pm 20.15	11.34 \pm 7.10
SASSNet [12]	82.84 \pm 6.12	89.08 \pm 5.49	89.50 \pm 9.00	18.21 \pm 10.60	21.60 \pm 10.82	8.25 \pm 5.22
DDT [11]	90.83 \pm 4.14	88.93 \pm 4.19	88.53 \pm 9.11	9.02 \pm 3.76	8.12 \pm 4.02	9.27 \pm 6.01
DMV (ours)	91.60 \pm 3.99	90.36 \pm 3.87	96.02 \pm 3.20	4.35 \pm 2.82	3.84 \pm 2.68	3.64 \pm 2.93
Method	Dice Score \uparrow [%]		Chamfer weighted \downarrow		All predictions watertight	
	HI, MRI		HI, MRI			
Voxel2Mesh [19]	85.30 \pm 3.20		(1.3 \pm 0.3) $\times 10^{-3}$		No	
DMV (ours)	89.23 \pm 3.08		(1.0 \pm 0.1) $\times 10^{-3}$		Yes	

Voxel2Mesh [19], which directly generates surface meshes for the hippocampus.

All methods were trained under the same conditions with regards to pre-processing and hyper-parameter settings, unless a different setting was suggested in the original paper. In particular, SASSNet and Voxel2Mesh were fully trained using their original settings. As some methods generate a volumetric grid and some generate a surface mesh, all outputs were first converted to a volumetric grid before calculating the similarity metrics. This also considered the fact that all datasets provide a ground truth in the form of a binary volumetric grid. While all methods performed sufficiently well, DMV quantitatively outperformed all methods in these scenarios. Additionally, the previous methods were conceived for only one of the evaluated target shapes, like the left atrium (SASSNet [12]) or the aorta (DDT [11]). The results show how DMV is more ductile and performs well in different scenarios. Furthermore, DMV guarantees the generation of smooth and watertight surfaces thanks to the properties of convolution surfaces. Additional tests were performed on the VascuSynth dataset [52] evaluate the use of

DMV on vascular shapes with high topological changes. The input volumes were highly corrupted with uniform noise in the range [0, 1] and the original VascuSynth volumes were expected as output. Although DDT [11] performs fairly well, as this was conceived for tubular shapes, DMV shows that further improvements are possible.

C. Qualitative analysis

For visual understanding, Figure 6 provides a qualitative overview of the DMV output as skeleton and the reconstructed tetrahedral meshes. Example outputs are compared with the corresponding ground truth. It can be seen that the ground truth presents a higher number of discretization artifacts. Furthermore, a Likert scale questionnaire administered to senior biomedical research scientists showed great acceptance of the shapes generated for simulation, such as computational fluid dynamics (CFD) and visualization. The summary is provided in Table III. Seven of the 12 scientists (58.3%) had at least five years of experience and six (50.0%) had a senior or lead

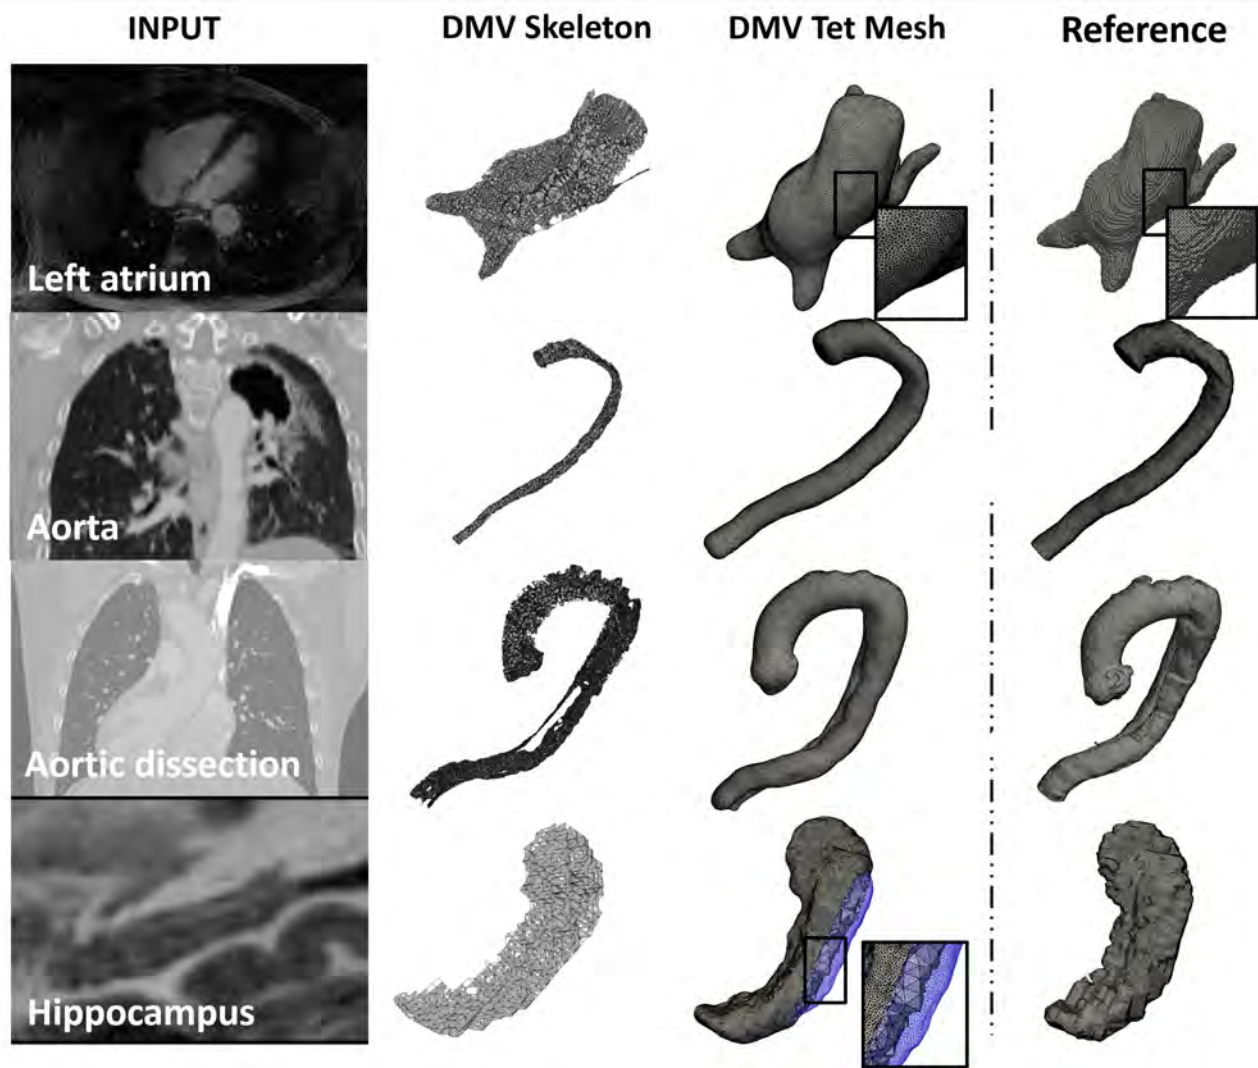


Fig. 6. Qualitative results of test cases. From left to right: input section, predicted skeleton as a simplicial complex, tetrahedral mesh from the convolution surface of the skeleton, post-processed marching cubes triangulation as a test reference. The blue section cut of the tetrahedral mesh shows the inner volume.

TABLE III

QUALITATIVE EVALUATION OF RECONSTRUCTION QUALITY. THE STUDY INVOLVED 12 FIELD EXPERTS WITH MULTIPLE YEARS OF EXPERIENCE IN CARDIOVASCULAR MODELING AND SIMULATION. LEGEND: 1 - COMPLETELY DISAGREE; 2 - SOMEWHAT DISAGREE; 3 - NEUTRAL; 4 - SOMEWHAT AGREE; 5 - COMPLETELY AGREE.

Question	Answer					Median	Inter-quartile range
	1	2	3	4	5		
Small structures are well preserved.	0	0	2	6	4	4	1
Artifacts are absent or negligible.	0	1	1	4	6	4,5	1
Topology is well preserved.	0	0	0	5	6	5	1
The mesh could easily be used in CFD simulations.	0	0	1	4	7	5	1
The cluster-based parametrization improves visualization.	0	0	0	5	7	5	1
The cluster-based parametrization is helpful in CFD.	0	0	1	4	7	5	1
Participants	Years					Median	Inter-quartile range
	1-2	3-4	5-6	7-8	9+		
Years of work experience in the field.	2	2	6	0	1	5	1
	Job level						
	Lead scientist		Senior research scientist		Research scientist		
Current job level.	1		5		5		Research intern 1

TABLE IV

QUANTITATIVE EVALUATION OF THREE METHODS ON THE SYNTHETIC DATASET VASCU SYNTH [52]. THE IMAGES WERE HIGHLY CORRUPTED WITH UNIFORMLY-DISTRIBUTED NOISE DURING TRAINING AND THE ORIGINAL BINARY VESSEL TREES WERE EXPECTED AS OUTPUT. THE VESSEL TREES IN THE DATASET PRESENT A HIGH DEGREE OF TOPOLOGICAL VARIATION.

Method	Dice Score \uparrow [%]	Hausdorff Distance \downarrow [mm]
VNet [47]	91.62 \pm 0.04	3.21 \pm 3.09
DDT [11]	95.13 \pm 0.01	2.40 \pm 2.07
DMV (ours)	97.01 \pm 0.02	2.39 \pm 2.59

research position in industry or academia. All participants provided a neutral to extremely positive rank for the mesh outputs with the exception of one participant, who could still notice minor artifacts in some cases.

D. Co-segmentation and parameterization

A particularity of the simplicial complex is that it contains topological information and can therefore be used for shape co-segmentation [30]. In Figure 8, we provide examples of how our simplified skeletons can be combined with SEG-MAT [30] for unsupervised shape co-segmentation [30], without the need for any additional post-processing. We also provide an example of cluster-based shape parameterization (Figure 7), where the co-segmented cluster of the aortic root (blue) has been deformed using a linear parametrization along the cluster. The resulting effect simulates the presence of an aneurysmatic dilatation. Such parametrizations are often non-trivial [53], but relevant in sensitivity analysis of biomechanics simulations [54]. The dilation effect was obtained by multiplying the radii of the MAT points by a dilation coefficient modulated by a Gaussian distribution centered on the junction between the blue (root) and the light blue (ascending aorta) clusters. The modulation of the MAT radii allows one to recompute a new convolution surface and generate an updated mesh that is still smooth and watertight.

E. Definition of inflow and outflow interfaces

For computational fluid dynamics (CFD), the generated meshes require the definition of inflow and outflow interfaces. Manual definition of such interfaces can be time-consuming. A common alternative is to place them at the extremities (leaves) of the shape centerline and orient them orthogonally to the centerline itself. The predicted DMV skeletons are generally too complex to quickly determine such locations. However, the interface points are mostly tubular, and oversimplifying the MAT leads to a simplicial complex where the leaf elements are depicted as simple segments, allowing quick determination of the leaf elements and their orientations (Figure 9).

F. Numerical simulation

Patient-specific simulations of the cardiovascular system and its pathologies are of increasing importance in clinical practice, as they allow virtual surgery, facilitate enhanced planning or can be employed in medical device design. In this

context, the construction of suitable (volumetric) meshes is a major challenge. Mesh quality significantly affects the time to find a solution and thereby influences the clinical relevance of computational methods in medicine. For example, the triangular surface representations derived via the proposed framework can be readily employed for biomedical applications such as fluid and solid mechanics or fluid–structure interaction (see, e.g., [4], [55]–[57]).

For the tetrahedralization step, we follow a standard procedure and generate tetrahedral grids via Gmsh [46] based on the surface triangulation. Note here that the mesh quality of the surface triangulation immediately transfers to the volumetric mesh generated. Hence, improving surface triangulation quality is crucial for practical engineering problems, as mesh quality impacts solution accuracy, linear solver performance, and numerical stability [58], [59]. Fig. 10 depicts some standard quality metrics [60] of the resulting surface triangulations generated by the DMV and alternative methods. The DMV leads to improved quality metrics, on top of the quantitatively better shape approximation (higher dice score and lower Hausdorff distance, see Tab. II).

Typically, surface mesh construction, volumetric meshing and numerical simulation are completely separated tasks within a computational engineering workflow. Our pipeline yields data in industry standard stereolithography format (STL, 3D Systems Inc.) describing the exterior triangulated surface of the domain of interest. Hence, standard interfaces are available, rendering the obtained results directly available from an application point of view. Volumetric meshing can then be carried out choosing from this vast variety of tools. Finite element packages such as Ansys (Ansys Inc.), Abaqus (Dassault Systèmes Simulia Inc.), COMSOL Multiphysics (Comsol Inc.) or open-source alternatives such as SimVascular [61], FEBio [62], or others provide such functionality either by custom implementations or interfacing libraries such as Gmsh [46], coreform Cubit (Coreform LLC) or ICEM CFD (Ansys Inc.). Depending on the application, further preprocessing steps are required to fully determine the boundary value problem to be solved, which lies beyond the scope herein.

Lastly, let us demonstrate suitability of the grids for CFD, specifically blood flow in a stationary left atrium. In a simplified setting, identical volumetric flow rates are prescribed in the left and right inferior and superior pulmonary veins (PV) together with a zero reference pressure at the mitral valve (MV), following Dueñas-Pamplona et al. [63].

Blood velocity \mathbf{u} and pressure p in the domain Ω are governed by the Navier–Stokes equations for incompressible flow of a generalized Newtonian fluid,

$$\rho \frac{\partial}{\partial t} \mathbf{u} + \rho \mathbf{u} \cdot \nabla \mathbf{u} - 2 \nabla \cdot (\mu \nabla^S \mathbf{u}) + \nabla p = \mathbf{0} \quad \text{in } \Omega, \quad (14)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega, \quad (15)$$

where $\rho = 1060 \text{ kg/m}^3$ is the density; μ , the viscosity, and ∇^S , the symmetric gradient. The viscosity is given by

$$\mu(\dot{\gamma}) = \mu_\infty + (\mu_0 - \mu_\infty) \left(1.0 + (\lambda \dot{\gamma})^2 \right)^{\frac{n-1}{2}}, \quad (16)$$

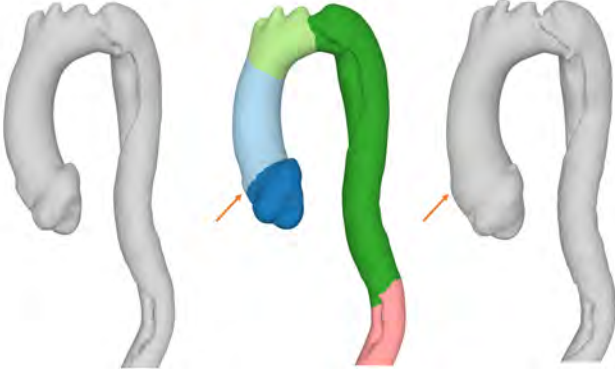


Fig. 7. The co-segmentation has been used to modify the radii associated with the MAT points of the aortic root and the ascending aorta. As a result, it is possible to simulate the onset of pathological processes like aneurysmatic enlargements.



Fig. 8. Examples of shape co-segmentation from the Q-MAT simplification [26] of the predicted medial skeletons for a left cardiac atrium (top) and a dissected aorta (bottom). From the left, the reconstructed surface, the medial vertices, the simplified medial skeleton and the co-segmented surface. Without prior information, the co-segmentation correctly splits the shape into meaningful regions. The dissected aorta is split into aortic root, ascending aorta, aortic arch, descending aorta, and abdominal aorta.

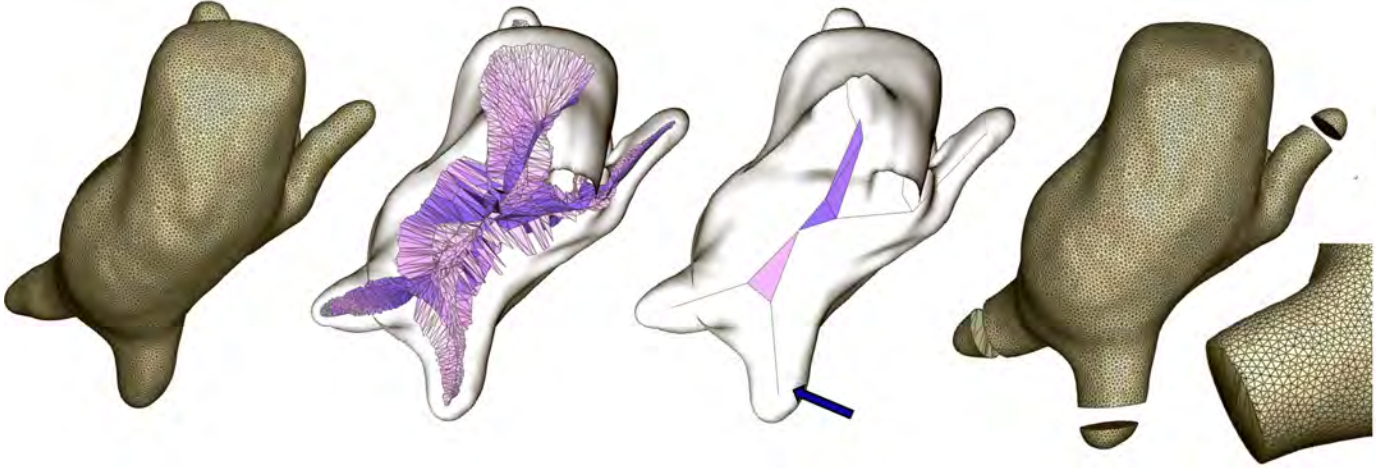


Fig. 9. From left to right: Predicted DMV mesh of a left atrium, MAT represented as simplicial complex, Q-MAT simplification where the extremities are described with segments, truncated mesh where inflow and outflow interfaces are modeled as flat cuts, rotated view showing additional interface.

with the shear rate $\dot{\gamma} = \sqrt{2\nabla^S \mathbf{u} : \nabla^S \mathbf{u}}$. Physiological parameters $\mu_\infty = 3.22$ mPas, $\mu_0 = 27.21$ mPas, $\lambda = 1.556$ 1/s and $n = 0.462$ are selected [64]. Simulations are carried out in ExaDG [65], a multi-purpose, higher-order discontinuous Galerkin flow solver based on the deal.II finite element library [66], employing Taylor–Hood finite elements and hybrid multigrid methods [67], [68].

Snapshots of the flow field and pressure distribution at three distinct points (diastole, peak systole and late systole) in the fourth cardiac cycle considered show mainly laminar flow from the PV to the MV. During diastole, the flow becomes disorganized, and the expected vortical structures are observed (Figure 12). Flow in the left atrium appendage is low, i.e., almost stagnant throughout the cardiac cycle. Due to the regular grid, iteration counts in the linear solvers stay low throughout the entire simulation, while the approximation quality seems adequate using elements with a target edge length of 0.8–1.8 mm and second-order velocity interpola-

tion. Further aspects such as a detailed analysis of the fluid dynamics, a grid convergence study and model verification are omitted for brevity. However, decent grid quality as required for CFD is achieved as demonstrated through this example.

VI. CONCLUSION

We presented an automatic strategy for adaptive and parametric modeling of anatomical shapes from raw 3D imaging data. We showed how an approximation of the medial axis transform can be generated through regularized voxel-to-voxel transformations. The presented strategy is the first to address the problem of topology preservation for adaptive geometry reconstruction from learned 3D discrete implicit fields of different organs and imaging modalities. Additionally, we showed how this representation could provide a versatile medium for various tasks ranging from volumetric meshing to shape co-segmentation. A qualitative expert evaluation based on a Likert scale questionnaire showed a positive perception of the

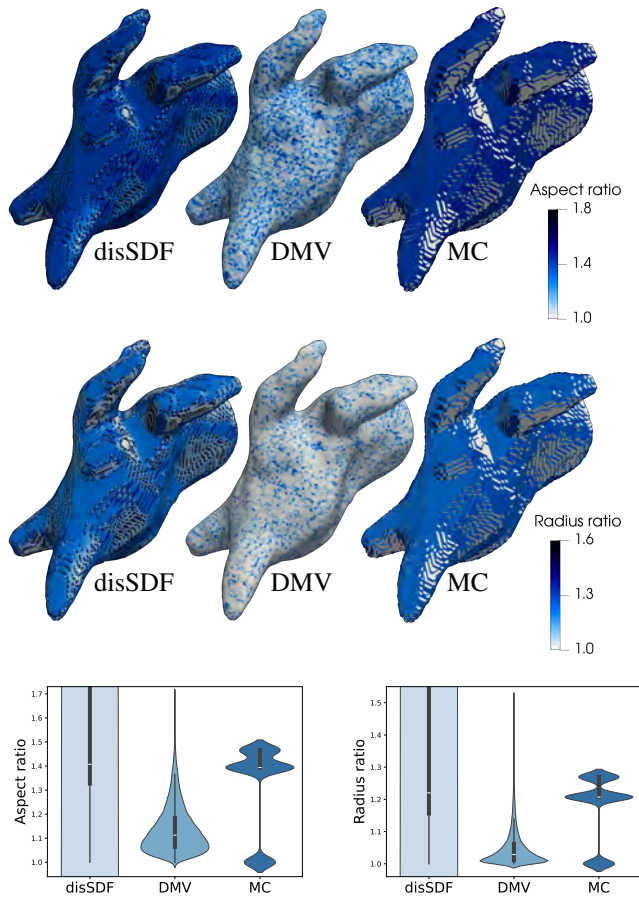


Fig. 10. Quality metrics aspect ratio and radius ratio obtained via disSDF, DMV and MC. Values of 1.0 are optimal and lower is better. DMV performs better than MC and disSDF, which have similar medians, but disSDF contains outliers emphasized in the violin plot. MC: surface mesh generated using marching cubes. disSDF: surface mesh generated from a discretized SDF with smooth embedding. DMV: our approach.

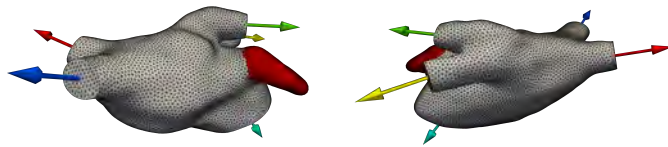


Fig. 11. Tetrahedral grid employed for fluid mechanics simulation: anterior (top) and posterior view (bottom) with the left atrium appendage highlighted in red and outward normals on the LSPV (green), LIPV (yellow), RSPV (blue), RIPV (red) and MV (cyan).

generated shapes. Additionally, a CFD simulation confirmed the direct applicability of the DMV output in FEM-based numerical simulations. In future work, we will further examine the capabilities of DMV for procedural and neural shape augmentation, with a focus on pathological vascular trees, and the possibility of relying on DMV for the generation of structured hexahedral meshes for different shapes [8], [69] and their application in in-silico trials [10].

This work was supported by the Alexander von Humboldt Foundation funded by the German Federal Ministry of Education and Research and the German Science Foundation *DFG* (grant EXC 2075).

REFERENCES

- [1] G. Mistelbauer, C. Rössl, K. Bäumler, B. Preim, and D. Fleischmann, "Implicit modeling of patient-specific aortic dissections with elliptic fourier descriptors," in *Comput Graph Forum*, vol. 40, pp. 423–34, 2021.
- [2] S. Oeltze and B. Preim, "Visualization of vasculature with convolution surfaces: method, validation and evaluation," *IEEE Trans Med Imaging*, vol. 24, no. 4, pp. 540–48, 2005.
- [3] J. Li *et al.*, "Autoimplant 2020-first miccai challenge on automatic cranial implant design," *IEEE Trans Med Imaging*, vol. 40, no. 9, pp. 2329–42, 2021.
- [4] K. Bäumler *et al.*, "Fluid–structure interaction simulations of patient-specific aortic dissection," *Biomechanics and Modeling in Mechanobiology*, vol. 19, no. 5, pp. 1607–28, 2020.
- [5] R. Shad *et al.*, "Patient-specific computational fluid dynamics reveal localized flow patterns predictive of post-left ventricular assist device aortic incompetence," *Circulation: Heart Failure*, vol. 14, no. 7, p. e008034, 2021.
- [6] L. Antiga, M. Piccinelli, L. Botti, B. Ene-Iordache, A. Remuzzi, and D. A. Steinman, "An image-based modeling framework for patient-specific computational hemodynamics," *Medical & biological engineering & computing*, vol. 46, pp. 1097–1112, 2008.
- [7] E. Kerrien, A. Yureidini, J. Dequidt, C. Duriez, R. Anxionnat, and S. Cotin, "Blood vessel modeling for interactive simulation of interventional neuroradiology procedures," *Med Image Anal*, vol. 35, pp. 685–98, 2017.
- [8] D. Bošnjak, A. Pepe, R. Schussnig, D. Schmalstieg, and T.-P. Fries, "Higher-order block-structured hex meshing of tubular structures," *Engineering with Computers*, pp. 1–21, 2023.
- [9] N. S. Burris *et al.*, "Vascular deformation mapping for CT surveillance of thoracic aortic aneurysm growth," *Radiology*, vol. 302, no. 1, pp. 218–25, 2022.
- [10] A. Sarraimi-Foroushani *et al.*, "In-silico trial of intracranial flow diverters replicates and expands insights from conventional clinical trials," *Nature Communications*, vol. 12, no. 1, pp. 1–12, 2021.
- [11] Y. Wang *et al.*, "Deep distance transform for tubular structure segmentation in CT scans," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3833–42, 2020.
- [12] S. Li, C. Zhang, and X. He, "Shape-aware semi-supervised 3d semantic segmentation for medical images," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 552–61, Springer, 2020.
- [13] J. Hu, B. Wang, L. Qian, Y. Pan, X. Guo, L. Liu, and W. Wang, "Mat-net: Medial axis transform network for 3d object recognition.," in *IJCAI*, pp. 774–781, 2019.
- [14] C. Lin, C. Li, Y. Liu, N. Chen, Y.-K. Choi, and W. Wang, "Point2skeleton: Learning skeletal representations from point clouds," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4277–86, 2021.
- [15] A. J. F. Suárez, E. Hubert, and C. Zanni, "Anisotropic convolution surfaces," *Computers & Graphics*, vol. 82, pp. 106–16, 2019.
- [16] T. S. Newman and H. Yi, "A survey of the marching cubes algorithm," *Computers & Graphics*, vol. 30, no. 5, pp. 854–79, 2006.
- [17] Q. Hong, Q. Li, B. Wang, K. Liu, and Q. Qi, "High precision implicit modeling for patient-specific coronary arteries," *IEEE Access*, vol. 7, pp. 72020–29, 2019.
- [18] M. Piccinelli, A. Veneziani, D. A. Steinman, A. Remuzzi, and L. Antiga, "A framework for geometric analysis of vascular structures: Application to cerebral aneurysms," *IEEE Transactions on Medical Imaging*, vol. 28, no. 8, pp. 1141–1155, 2009.
- [19] U. Wickramasinghe, E. Remelli, G. Knott, and P. Fua, "Voxel2mesh: 3d mesh model generation from volumetric data," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 299–308, 2020.
- [20] D. H. Pak *et al.*, "Distortion energy for deep learning-based volumetric finite element mesh generation for aortic valves," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 485–94, Springer, 2021.
- [21] Y. Kanamori, Z. Szego, and T. Nishita, "Gpu-based fast ray casting for a large number of metaballs," *Comput Graph Forum*, vol. 27, no. 2, pp. 351–60, 2008.
- [22] T. Fries, S. Omerović, D. Schöllhammer, and J. Steidl, "Higher-order meshing of implicit geometries, Part I: Integration and interpolation in cut elements," *Comput Methods Appl Mech Eng*, vol. 313, pp. 759–84, 2017.

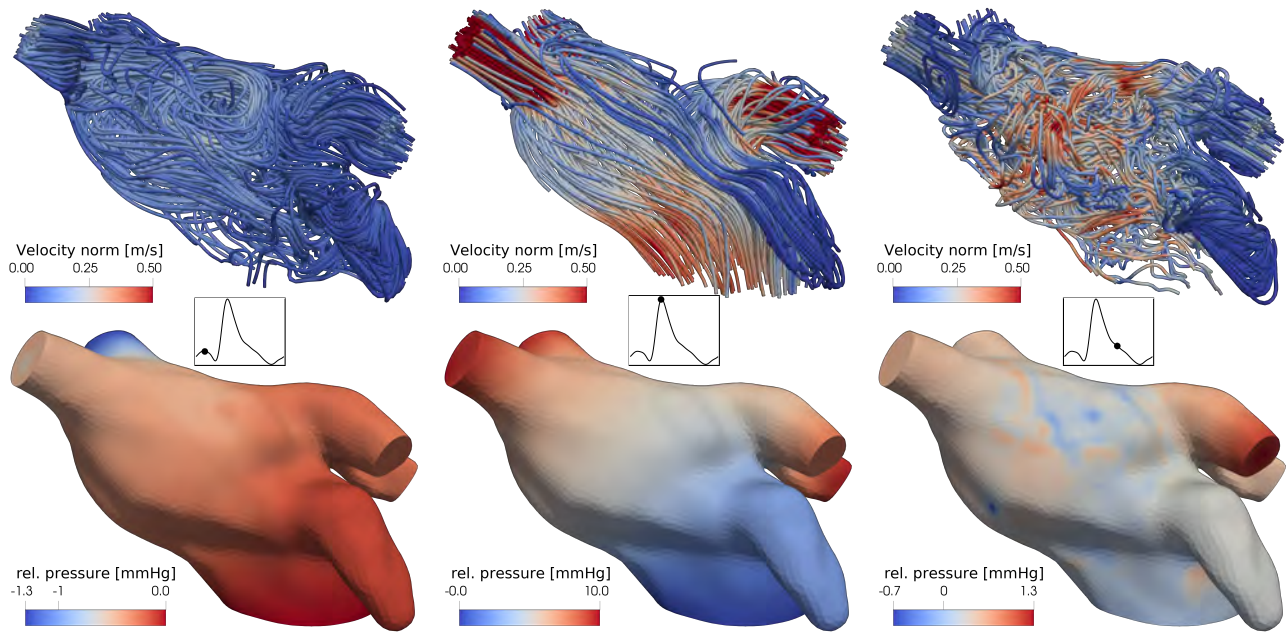


Fig. 12. Streamlines colored by velocity norm (top row) and relative pressure (bottom row) in diastole (left column), peak systole (middle column) and late systole (right column) as indicated by the inflow profile scale. Streamlines show a largely laminar flow, which is disorganized in diastole and stagnant in the left atrium appendage, while the pressure gradient from PV to MV reaches 10 mmHg in systole.

- [23] J. J. Torres and G. Bianconi, "Simplicial complexes: higher-order spectral dimension and dynamics," *J. phys. Complex*, vol. 1, no. 1, p. 015002, 2020.
- [24] H. Edelsbrunner and J. L. Harer, *Computational topology: an introduction*. American Mathematical Society, 2022.
- [25] J. Giesen, B. Miklos, M. Pauly, and C. Wormser, "The scale axis transform," in *Proc. Annual Symposium on Computational Geometry*, pp. 106–15, 2009.
- [26] P. Li, B. Wang, F. Sun, X. Guo, C. Zhang, and W. Wang, "Q-mat: Computing medial axis transform by quadratic error minimization," *ACM Trans. Graph.*, vol. 35, no. 1, pp. 1–16, 2015.
- [27] Z. Dou *et al.*, "Coverage axis: Inner point selection for 3d shape skeletonization," *Comput Graph Forum*, vol. 41, no. 2, pp. 2329–42, 2022.
- [28] Y. Yan, D. Letscher, and T. Ju, "Voxel cores: Efficient, robust, and provably good approximation of 3d medial axes," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 1–13, 2018.
- [29] D. Rebain *et al.*, "LSMAT least squares medial axis transform," in *Comput Graph Forum*, vol. 38, pp. 5–18, 2019.
- [30] C. Lin *et al.*, "Seg-mat: 3d shape segmentation using medial axis transform," *IEEE Trans Vis Comput Graph*, 2020.
- [31] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 165–74, 2019.
- [32] D. Rebain, K. Li, V. Sitzmann, S. Yazdani, K. M. Yi, and A. Tagliasacchi, "Deep medial fields," *arXiv preprint arXiv:2106.03804*, 2021.
- [33] K. Genova, F. Cole, A. Sud, A. Sarna, and T. Funkhouser, "Local deep implicit functions for 3d shape," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4857–66, 2020.
- [34] B. Yang *et al.*, "P2MAT-Net: Learning medial axis transform from sparse point clouds," *Comput Aided Geom Des*, vol. 80, p. 101874, 2020.
- [35] P. Erler, P. Guerrero, S. Ohrhallinger, N. J. Mitra, and M. Wimmer, "Points2surf: Learning implicit surfaces from point clouds," in *European Conference on Computer Vision*, pp. 108–124, 2020.
- [36] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, "Convolutional occupancy networks," in *European Conference on Computer Vision*, pp. 523–40, 2020.
- [37] J. Zhao, J. Zhao, S. Pang, and Q. Feng, "Segmentation of the true lumen of aorta dissection via morphology-constrained stepwise deep mesh regression," *IEEE Trans Med Imaging*, 2022.
- [38] J. Ma *et al.*, "How distance transform maps boost segmentation CNNs: an empirical study," in *Medical Imaging with Deep Learning*, pp. 479–92, PMLR, 2020.
- [39] Z. Lin, D. Wei, A. Gupta, X. Liu, D. Sun, and H. Pfister, "Structure-preserving instance segmentation via skeleton-aware distance transform," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 529–539, Springer, 2023.
- [40] X. Zhang, K. Sun, D. Wu, X. Xiong, J. Liu, L. Yao, S. Li, Y. Wang, J. Feng, and D. Shen, "An anatomy- and topology-preserving framework for coronary artery segmentation," *IEEE Transactions on Medical Imaging*, 2023.
- [41] H. Xia and P. G. Tucker, "Fast equal and biased distance fields for medial axis transform with meshing in mind," *Applied Mathematical Modelling*, vol. 35, no. 12, pp. 5804–19, 2011.
- [42] M. Liao, Z. Wan, C. Yao, K. Chen, and X. Bai, "Real-time scene text detection with differentiable binarization," in *Proc. AAAI conference on artificial intelligence*, vol. 34, pp. 11474–81, 2020.
- [43] S. Shit *et al.*, "cldice-a novel topology-preserving loss function for tubular structure segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16560–16569, 2021.
- [44] J. McCormack and A. Sherstyuk, "Creating and rendering convolution surfaces," in *Computer Graphics Forum*, vol. 17, pp. 113–20, 1998.
- [45] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Proc. EG Symposium on Geometry Processing*, vol. 7, 2006.
- [46] C. Geuzaine and J. Remacle, "Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities," *Int J Numer Methods Eng*, vol. 79, no. 11, pp. 1309–1331, 2009.
- [47] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *International conference on 3D vision (3DV)*, pp. 565–71, 2016.
- [48] Z. Xiong *et al.*, "A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging," *Medical Image Analysis*, vol. 67, p. 101832, 2021.
- [49] A. L. Simpson *et al.*, "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *arXiv:1902.09063*, 2019.
- [50] Z. Lambert, C. Petitjean, B. Dubray, and S. Kuan, "Segthor: Segmentation of thoracic organs at risk in ct images," in *Proc. IPTA*, pp. 1–6, IEEE, 2020.
- [51] Z. Yao *et al.*, "Imagetbad: A 3d computed tomography angiography image dataset for automatic segmentation of type-b aortic dissection," *Front. Physiol.*, vol. 12, 2021.
- [52] P. Jassi and G. Hamarneh, "Vascusynth: Vascular tree synthesis software," *Insight Journal*, vol. January-June, pp. 1–12, 2011.

- [53] C. Zhang, L. Yang, L. Xu, G. Wang, and W. Wang, "Real-time editing of man-made mesh models under geometric constraints," *Computers & Graphics*, vol. 82, pp. 174–182, 2019.
- [54] G. García-Isla *et al.*, "Sensitivity analysis of geometrical parameters to study haemodynamics and thrombus formation in the left atrial appendage," *Int J Num Method Biomed Eng*, vol. 34, no. 8, p. e3100, 2018.
- [55] R. Schussnig, D. Pacheco, M. Kaltenbacher, and T.-P. Fries, "Semi-implicit fluid–structure interaction in biomedical applications," *Comput Methods Appl Mech Eng*, vol. 400, p. 115489, 2022.
- [56] R. Schussnig, D. Pacheco, and T. Fries, "Efficient split-step schemes for fluid-structure interaction involving incompressible generalised newtonian flows," *Comput Struct*, vol. 260, p. 106718, 2022.
- [57] R. Schussnig, M. Rolf-Pissarczyk, K. Bäumler, T. Fries, G. Holzapfel, and M. Kronbichler, "On the role of tissue mechanics in fluid–structure interaction simulations of patient-specific aortic dissection," *Int J Numer Methods Eng*, 2024 (submitted).
- [58] J. R. Shewchuk, "What is a Good Linear Element? Interpolation, Conditioning, and Quality Measures," in *11th International Meshing Roundtable, {IMR} 2002*, (Ithaca, United States), 2002.
- [59] D. Braess, *Finite elements: Theory, fast solvers, and applications in solid mechanics*. Cambridge: Cambridge University Press, 3rd ed., 2007.
- [60] P. P. Pébay and T. J. Baker, "Analysis of triangle quality measures," *Math Comput*, vol. 72, no. 244, pp. 1817–1839, 2003.
- [61] A. Updegrove, N. M. Wilson, J. Merkow, H. Lan, A. L. Marsden, and S. C. Shadden, "Simvascular: an open source pipeline for cardiovascular simulation," *Ann Biomed Eng*, vol. 45, pp. 525–541, 2017.
- [62] S. A. Maas, B. J. Ellis, G. A. Ateshian, and J. A. Weiss, "FEBio: Finite Elements for Biomechanics," *J Biomech Eng*, vol. 134, no. 1, p. 011005, 2012.
- [63] J. Dueñas-Pamplona, J. G. García, J. Sierra-Pallares, C. Ferrera, R. Agujetas, and J. R. López-Mínguez, "A comprehensive comparison of various patient-specific CFD models of the left atrium for atrial fibrillation patients," *Computers Biol Med*, vol. 133, p. 104423, 2021.
- [64] S. Ranftl, T. Müller, U. Windberger, W. von der Linden, and G. Brenn, "A Bayesian approach to blood rheological uncertainties in aortic hemodynamics," *Int J Numer Method Biomed Eng*, 2022.
- [65] D. Arndt, N. Fehn, G. Kanschat, K. K. M. Kronbichler, P. Munch, W. Wall, and J. Witte, "ExaDG: High-order discontinuous Galerkin for the exascale," in *Software for Exascale Computing - SPPEXA 2016-2019*, (Cham), pp. 189–224, Springer International Publishing, 2020.
- [66] D. Arndt, W. Bangerth, M. Bergbauer, M. Feder, M. Fehling, J. Heinz, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J. Pelteret, B. Turcksin, D. Wells, and S. Zampini, "The deal.II Library, Version 9.5," *J Numer Math*, 2023.
- [67] N. Fehn, W. A. Wall, and M. Kronbichler, "Robust and efficient discontinuous galerkin methods for under-resolved turbulent incompressible flows," *J Comput Physics*, vol. 372, pp. 667–693, 2018.
- [68] N. Fehn, P. Munch, W. Wall, and M. Kronbichler, "Hybrid multigrid methods for high-order discontinuous Galerkin discretizations," *J Comput Phys*, vol. 415, p. 109538, 2020.
- [69] P. Viville, P. Kraemer, and D. Bechmann, "Meso-skeleton guided hexahedral mesh design," *Comput Graph Forum*, p. e14932, 2023.