

AortaDiff: Volume-Guided Conditional Diffusion Models for Multi-Branch Aortic Surface Generation

Delin An , Pan Du , Jian-Xun Wang , and Chaoli Wang 

Abstract— Accurate 3D aortic construction is crucial for clinical diagnosis, preoperative planning, and computational fluid dynamics (CFD) simulations, as it enables the estimation of critical hemodynamic parameters such as blood flow velocity, pressure distribution, and wall shear stress. Existing construction methods often rely on large annotated training datasets and extensive manual intervention. While the resulting meshes can serve for visualization purposes, they struggle to produce geometrically consistent, well-constructed surfaces suitable for downstream CFD analysis. To address these challenges, we introduce AortaDiff, a diffusion-based framework that generates smooth aortic surfaces directly from CT/MRI volumes. AortaDiff first employs a volume-guided conditional diffusion model (CDM) to iteratively generate aortic centerlines conditioned on volumetric medical images. Each centerline point is then automatically used as a prompt to extract the corresponding vessel contour, ensuring accurate boundary delineation. Finally, the extracted contours are fitted into a smooth 3D surface, yielding a continuous, CFD-compatible mesh representation. AortaDiff offers distinct advantages over existing methods, including an end-to-end workflow, minimal dependency on large labeled datasets, and the ability to generate CFD-compatible aorta meshes with high geometric fidelity. Experimental results demonstrate that AortaDiff performs effectively even with limited training data, successfully constructing both normal and pathologically altered aorta meshes, including cases with aneurysms or coarctation. This capability enables the generation of high-quality visualizations and positions AortaDiff as a practical solution for cardiovascular research.

Index Terms—Conditional diffusion model, volume-guided surface generation, multi-branch vessel modeling

1 INTRODUCTION

The aorta is the largest blood vessel in the human body, originating from the heart and branching out to deliver oxygenated blood throughout the body. Like other blood vessels, it is susceptible to various pathological conditions, such as coarctation and aneurysm [58]. Coarctation arises from congenital defects causing aortic narrowing, while aneurysm results from chronic stress or atherosclerosis, leading to localized dilation and wall weakening. Anatomically, the aorta exhibits a multi-branch topology with complex surface features, encompassing significant variations in shape and scale. Although aortic modeling has been extensively studied, the characterization and analysis of multiscale aortic structures—particularly those with branching (e.g., the supra-aortic branches at the arch)—remains a nontrivial challenge [26, 44]. In the context of human anatomy, particularly the aortic arch, the major supra-aortic branches include the *right subclavian artery (RSA)*, *left subclavian artery (LSA)*, *right common carotid artery (RCCA)*, and *left common carotid artery (LCCA)*. As illustrated in Figure 1, these vessels emerge from the aortic arch and supply blood to the head and upper extremities. Given the aorta’s vital role, hemodynamic analysis is essential for clinical diagnosis and treatment planning to assess blood flow dynamics and identify potential abnormalities.

Computational fluid dynamics (CFD) simulations are a powerful tool for understanding aortic disease development and rupture risk. By computing critical hemodynamic parameters such as blood flow velocity, pressure distribution, and *wall shear stress (WSS)*, CFD analysis provides valuable insights for medical professionals. However, the ac-

curacy and reliability of CFD simulations are highly dependent on the quality of the underlying 3D aorta meshes. These meshes are typically manually extracted from *computed tomography (CT)* or *magnetic resonance imaging (MRI)* volumes using tools like SimVascular [60]. This manual process is labor-intensive and time-consuming, requiring the annotation of contours across hundreds or even thousands of CT/MRI slices, making it a significant bottleneck in clinical workflows.




Fig. 1: Extraction of 3D aorta meshes from a CT/MRI volume.

To minimize manual intervention, researchers have developed various segmentation methods [17, 20, 71]. While these methods, typically trained on large annotated CT/MRI datasets, achieve high segmentation accuracy, they fall short in directly constructing 3D aorta meshes. Point clouds offer a more accurate approximation of 3D aortic structures than semantic segmentation. Advances in computer vision have led to point cloud-based methods [35, 72] that perform well in generating general 3D objects, such as chairs, aircraft, and cars. However, these methods rely on large-scale training datasets, which are significantly harder to obtain for aortic data. A more direct approach is to generate 3D surfaces. Studies [6, 69] have shown that using statistical shapes and deformable models can effectively construct specific organs, such as the heart, due to their relatively simple structure. For modeling complex vascular structures, recent studies [11, 30, 65] have combined generative models (e.g., VAEs, GANs, diffusion models) with statistical shape modeling to synthesize multi-branch vessel geometries. However, these methods primarily focus on small-scale vasculatures and branching topology, often lacking detailed surface representations. Moreover, they serve purely as geometry synthesizers without leveraging medical imaging data. Mistelbauer et al. [42] employed *elliptic Fourier descriptors (EFDs)* to generate surface meshes for aortic dissections with separated lumens. However, creating CFD-ready meshes for healthy

- D. An and C. Wang are with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556, USA. E-mail: {dan3, chaoli.wang}@nd.edu.
- P. Du is with the Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, IN 46556, USA. E-mail: pdu@nd.edu.
- J.-X. Wang is with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14850, USA. E-mail: jw2837@cornell.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

aortas and other pathologies—featuring multiple branches and varying diameters—remains a significant challenge, demanding more flexible and generalizable modeling techniques.

In summary, existing solutions face three major challenges in fully automating aorta mesh extraction from CT/MRI volumes:

- **Limited availability of aortic data:** Publicly accessible datasets primarily provide segmentation labels rather than detailed 3D mesh annotations, restricting the development of fully automated construction methods.
- **High reliance on manual intervention:** Generating high-quality aorta meshes still requires considerable manual effort, including mesh smoothing, hole filling, and geometric refinement, making the process labor-intensive and time-consuming.
- **Lack of an end-to-end solution:** Existing methods cannot directly generate meshes from medical images that are suitable for both visualization and hemodynamic simulation of the main aorta and supra-aortic branches, necessitating post-processing to obtain a complete and usable 3D representation.

To address these challenges, we propose AortaDiff, an automatic framework for aorta mesh construction conditioned on CT/MRI volumes. AortaDiff integrates both point cloud and segmentation-based approaches and consists of three key stages: (1) A volume-guided *conditional diffusion model* (CDM) generates the aortic centerline from the input CT/MRI volume. (2) A contour extraction technique uses the generated centerline points as priors to guide a semi-automated segmentation method, transforming it into a fully automated process. (3) A surface construction method constructs the final aorta mesh by fitting a 3D *non-uniform rational B-spline* (NURBS) surface to the extracted contours. These stages are designed to address the limitations of existing approaches. In the first stage, inspired by the success of *denoising diffusion probabilistic models* (DDPM) [25], we employ a CDM [57] to generate centerline points rather than the entire aortic point cloud. This approach provides a compact yet morphologically informative representation of the aorta, significantly reducing computational overhead and enabling effective training on limited aortic datasets. In the second stage, the generated centerline points serve as prompts for a prompt-dependent semi-automated segmentation model, eliminating manual intervention in the segmentation process. The third stage completes the construction by generating a CFD-ready aorta mesh, making AortaDiff a fully end-to-end solution for aorta mesh construction.

The primary contributions of AortaDiff are as follows:

- **Fully automatic aorta mesh extraction:** AortaDiff offers an end-to-end solution for generating 3D aorta meshes directly from CT/MRI volumes, covering normal, aneurysmal, and coarctation cases, and removing the need for extensive manual intervention.
- **Minimal training data requirements:** Unlike deep learning-based segmentation and point cloud generation methods, AortaDiff operates effectively without requiring large annotated training datasets, making it practical for real-world applications.
- **High-quality mesh generation:** Parameterized NURBS-based meshes produced by AortaDiff are directly compatible with CFD solvers without additional post-processing, enabling efficient hemodynamic analysis, including the estimation of flow-derived biomarkers critical for diagnosis and treatment planning.
- **Support for parametric anatomical modification:** The flexible geometry construction framework enables customizable anatomical adjustments, facilitating *in silico* experimentation for device optimization and surgical planning.

By addressing key limitations of existing methods, AortaDiff provides a structured, fully automated approach to aorta mesh construction that renders manual contour annotation and post-processing unnecessary, thereby streamlining the workflow and enhancing the efficiency and consistency of hemodynamic analysis in clinical applications.

2 BACKGROUND AND RELATED WORK

Background. Hemodynamic analysis is essential for understanding cardiovascular conditions and guiding diagnosis and treatment decisions. It enables the estimation of biomarkers and flow characteristics linked to disease progression, such as regions of elevated WSS associated with

thrombus formation and aneurysm rupture risk, and abnormal pressure gradients or flow acceleration that help assess aortic coarctation severity and plan interventions. These analyses, typically based on CFD simulations of cardiovascular geometries, provide quantitative insights into vessel morphology, flow velocity, pressure, and WSS distributions, offering biomarkers beyond vessel diameter to better predict complications and evaluate the impact of surgical or endovascular treatments. Parametric geometry models further support *in silico* experimentation for optimizing device designs and planning procedures.

Numerous studies have explored visualizing and analyzing hemodynamic data for clinical applications. Lawonn et al. [31] developed a tool for occlusion-free blood flow animation to enhance visualization clarity. Oeltze-Jafra et al. [48] introduced a clustering-based approach for visualizing vortical flows in CFD simulations. Tao et al. [59] designed a web-based interface to visualize 4D hemodynamic data. Combining patient-specific hemodynamics with wall deformation, Meuschke et al. [41] developed a visualization tool integrating 2.5D and 3D representations. To support the exploration of multi-field aneurysm data, they also proposed an interactive tool to identify regions of interest on vessel walls [38]. Other notable contributions include CAVOCLA [40], which facilitates blood flow classification for cerebral aneurysms, and GUCCI [39], which identifies cohort-specific characteristics of the aorta using glyph-based depictions.

Despite these advancements, a significant gap remains in obtaining high-quality vascular geometries, which are fundamental for CFD simulations [14, 15]. The aorta, the largest blood vessel originating from the heart, has a complex structure, including the main and supra-aortic branches. Accurate 3D representations of aortic geometries are critical for clinical applications. However, constructing these geometries poses several challenges, including the complexity of multi-branch structures, the limited availability of medical imaging data, and the demand for high geometric fidelity.

Surface construction. Recent advancements in computer vision and deep learning have demonstrated significant potential in medical image processing and 3D surface construction. Most approaches primarily concentrate on image or volume segmentation to extract 3D surfaces, while others focus on the 3D construction of specific organs, such as the heart and liver. More recently, diffusion models have been utilized to generate 3D point clouds [35, 67]. Existing methods for 3D surface construction can be broadly classified into three categories: *segmentation*, *point cloud*, and *geometry* techniques.

Segmentation methods are used to identify vascular regions in 3D medical images, generating segmented volumes of specific organs. Techniques such as marching cubes [34] or surface nets [18], are then employed to extract 3D surfaces. Advanced 3D segmentation models, like DenseUNet [73], utilize densely connected layers to improve feature propagation, particularly for complex structures. UNetR++ [55], which incorporates transformer architectures to capture long-range dependencies, achieves state-of-the-art performance. Numerous studies have applied segmentation to vascular analysis. For instance, Lin et al. [33] introduced a geometry-constrained deformable attention network to enhance aortic structure segmentation. Montalt-Tordera et al. [45] trained a U-Net model for thoracic aorta hemodynamics analysis, tackling patient-specific blood flow assessment challenges. Myronenko et al. [46] developed an automated method for aorta segmentation using 3D CT images. Vagenas et al. [61] combined traditional image processing and machine learning techniques for precise aorta segmentation and construction. While these methods effectively capture detailed anatomical features, they rely on extensive annotated data and are susceptible to error propagation, limiting their utility in aortic surface construction. Although region-growing [23] methods have advanced aortic segmentation, they require seed point specification and are highly sensitive to intensity thresholds, affecting accuracy and consistency. *Segment anything model* (SAM)-based methods [28] rely on precise input prompts to achieve accurate segmentation, still requiring substantial human intervention.

Point cloud methods construct 3D objects by approximating surfaces using sampled points. Generative models, particularly diffusion-based approaches, enhance the fidelity of these constructed geometries.




Fig. 2: The framework of AortaDiff: (a) A volume-guided CDM generates a centerline by extracting a feature volume using a ViT encoder, retrieving feature vectors at the centerline points, and concatenating them with the 1D centerline image. (b) ScribblePrompt extracts contours for centerline points. (c) A 3D NURBS mesh is constructed from the contours representing the geometry. (d) The aorta mesh is then used in CFD simulations.

For example, Erler et al. [16] developed Points2Surf, which learns implicit surface representations from point cloud patches for high-quality construction. Zhou et al. [72] introduced the *point-voxel diffusion* (PVD) model, combining denoising diffusion models with point-voxel representations to achieve high-fidelity construction. Luo and Hu [35] approached point cloud generation as a thermodynamic process, improving robustness against noise. Li et al. [32] proposed a disentangled refinement method for point cloud upsampling, resulting in higher resolution and accuracy. Wu et al. [67] introduced *point straight flow* (PSF), simplifying the diffusion process into a single step, significantly accelerating generation while preserving quality. Despite their strengths, these methods require extensive training data and face challenges to ensure continuity and smoothness in constructed surfaces, particularly for complex geometries like the main aorta and its supra-aortic branches.

Geometry methods bypass intermediate steps like segmentation or point cloud creation, directly generating 3D geometries. For example, Selle et al. [54] and Hahn et al. [21] proposed pipelines for vascular analysis and visualization that combine region-growing-based segmentation with skeletonization and graph-based structural analysis. These methods produce effective visualizations of vessel branching patterns and diameters; however, their performance is limited by factors such as low image resolution, partial volume effects, and discontinuities in vessel diameter estimation. Kretschmer et al. [29] generated and refined smooth vascular models using centerline-based descriptions. Recent advancements in deep learning have further enhanced the capabilities of geometry methods. Lyu et al. [36] utilized GAN-based models to generate geometries for the aorta and carotid arteries. Xiong et al. [68] leveraged cascaded GANs for high-fidelity aortic geometry construction. Black et al. [9] constructed arterial geometries using 4D flow-MRI data to validate hemodynamic simulations. Although these methods offer end-to-end solutions, they often require extensive datasets, rely on structural assumptions, and lack flexibility for adjustments.

To address the challenges of aortic surface construction, we propose AortaDiff, a volume-guided CDM that decomposes the aorta into centerlines and contours for precise modeling. Diffusion models have shown significant success in image generation, with DDPM [25] serving as a foundational framework. Nichol and Dhariwal [47] improved DDPM by modeling reverse diffusion variances, while Chen et al. [10] introduced dynamic resolution image generation using neural fields. Beyond images, diffusion models have been applied to text-to-image generation [70] and video synthesis [56]. In visualization, tools like GAN Lab [27] have been used to interpret deep generative models. While diffusion models have shown promise in image generation, their application to vessel modeling and visualization remains underexplored.

Traditional aorta modeling methods, such as voxel-based reconstruction or manual parameterization, often introduce surface defects or require intensive manual effort, hindering scalability. AortaDiff integrates a volume-guided CDM for centerline generation with SAM-based contour extraction. By combining diffusion modeling with automated parameterization, it generates smooth, simulation-ready meshes with minimal manual intervention, enabling scalable geometry construction for CFD and surgical planning.

3 AORTADIFF

Figure 2 illustrates the AortaDiff framework, which consists of four key stages: (a) extracting the aortic centerline from CT/MRI volumes using a CDM, (b) obtaining vessel contours from cross-sections through image segmentation and edge detection, (c) constructing the 3D aorta mesh using NURBS, and (d) conducting downstream CFD simulation tasks with the generated aorta meshes. The algorithm description is provided in Section 1 of the appendix.

3.1 Centerline Representation

Point cloud methods are commonly used for 3D object construction but typically require large training datasets, often numbering in the thousands. In contrast, aortic datasets are generally limited (e.g., only tens of samples in our experiment), making direct learning from aortic point clouds challenging. To address this, we adopt a simple yet effective representation: *centerlines* and their corresponding *contours*. The centerline represents the central trajectory of the aorta as a 1D approximation of its 3D structure. It consists of a sequence of ordered points, each corresponding to a specific location along the vessel. This representation captures the global shape and topology of the aorta while significantly reducing data complexity. Compared to point clouds, the centerline is more compact, easier to process, and requires fewer computational resources. Additionally, it integrates seamlessly with other geometric features, such as cross-sectional contours, offering a more efficient and informative vessel representation. Empirically, we select 16 points per centerline for both the main aorta and supra-aortic branches, striking a balance between accuracy and efficiency. For the main aorta, this configuration effectively captures its curvature and branching patterns while avoiding unnecessary complexity or overlapping contours that could arise from overly dense sampling. For the supra-aortic branches, although they are shorter in length, they exhibit greater relative radius variation compared to the main aorta. Using 16 points ensures that sufficient detail is preserved while preventing excessive redundancy, which is supported by the study shown in Section 4 of the appendix. Unlike point clouds, we encode the centerline as a 1D image—a matrix with three channels corresponding to the centerline

points' x , y , and z coordinates. This format aligns well with image processing techniques and deep learning models, such as 1D-UNet, a variant of the U-Net architecture [52], for centerline generation.

3.2 Volume-Guided CDM

To generate the centerline, we employ a volume-guided CDM that progressively denoises a sequence of noisy centerline images. The CDM is trained to estimate the noise added at each timestep of the diffusion process, enabling it to learn the underlying aortic structure. Diffusion models, such as DDPM [25], are generative frameworks that iteratively reverse a noise diffusion process. During training, noise is incrementally added to the image, and the model is optimized to predict and remove this noise at each timestep. This process allows the generation of new samples by starting from a simple prior distribution (e.g., Gaussian) and iteratively refining them through denoising.

However, clinical applications often require constructing a patient-specific 3D aortic surface. We incorporate information from the 3D CT/MRI volume into the diffusion process to ensure the generated centerline aligns with the underlying anatomical structure. As shown in Figure 2 (a), a *vision transformer* (ViT) [13] encoder extracts a feature volume from the input CT/MRI volume, encoding high-level, voxel-wise semantic information. This ViT, adapted from MONAI for 3D medical imaging, consists of 12 layers, 8 attention heads, a hidden size of 16, and a multi-layer perceptron dimension of 3072. We aim to integrate this feature information for the centerline points as conditional inputs to guide the diffusion process. A naïve approach would associate the entire feature volume with each centerline point. In contrast, for efficiency and effectiveness, we look up the corresponding feature vector for each centerline point from the feature volume and concatenate it with the noisy centerline image's latent vector.

CDM training. AortaDiff integrates volumetric features with centerline points at each timestep of the diffusion process:

- **Input representation:** The 1D *ground-truth* (GT) centerline image, $\mathbf{CI}_0$, is encoded into a latent vector, $\mathbf{L}_0$. Noise is progressively added to $\mathbf{L}_t$ at each timestep t , producing a noisy sequence $(\mathbf{L}_0, \dots, \mathbf{L}_t, \dots, \mathbf{L}_T)$. The model input consists of $\mathbf{L}_t$ concatenated with feature vectors, $\mathbf{f}$, extracted from the feature volume, $\mathbf{FV}$.
- **Noise prediction:** The model is trained to predict the noise, $\epsilon_\theta(\mathbf{CI}_t, t, \mathbf{f})$, added at t , facilitating the denoising process.
- **Optimization objective:** The training aims to minimize the difference between the predicted and actual noise using the loss

$$\mathcal{L}_{\text{CDM}} = \mathbb{E}_{\mathbf{CI}_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(\mathbf{CI}_t, t, \mathbf{f})\|^2], \quad (1)$$

where $\mathbf{CI}_0$ and $\mathbf{CI}_t$ denote the clean and noisy centerline images, ϵ is the added noise, and ϵ_θ is the noise predicted by the model.

Inference with CDM. During inference, the input consists of $\mathbf{L}_T$ (the fully noisy latent vector) and the feature volume, $\mathbf{FV}$. The process follows the reverse of the training paradigm: at each timestep t , the noisy latent vector, $\mathbf{L}_t$, determines voxel locations, and relevant feature vectors, $\mathbf{f}$, are retrieved as conditional inputs. The model then iteratively denoises the latent vector. Eventually, at timestep 0, the final output, $\mathbf{CI}_0$, is constructed via a decoder.

3.3 Contour Extraction

Extracting contours around centerline points enables the construction of a 3D surface that accurately captures the detailed geometry of the vessel lumen. We utilize ScribblePrompt [66], a segmentation model that requires minimal user input (e.g., scribbles, clicks, and bounding boxes) to achieve precise segmentation.

Orthogonal slice extraction. For each centerline point p , we compute its tangent vector $\mathbf{t}_p$ using adjacent points. To maintain computational consistency, the last centerline point shares the tangent vector of its preceding one. Based on $\mathbf{t}_p$, we define a slicing plane Π_p orthogonal to $\mathbf{t}_p$. The plane's orientation is determined by a *local coordinate frame* $(\mathbf{t}_p, \mathbf{n}_p, \mathbf{b}_p)$, where $\mathbf{n}_p$ and $\mathbf{b}_p$ are the normal and binormal vectors spanning Π_p . We construct the rotation matrix $\mathbf{R}_p = [\mathbf{b}_p \ \mathbf{n}_p \ \mathbf{t}_p]$, which aligns Π_p with the *global coordinate system*. To extract a cross-sectional image S_p containing the vessel, we resample the input volume $\mathbf{V}$ onto Π_p using trilinear interpolation

$$S_p = \mathbf{V}(\mathbf{R}_p \cdot \mathbf{l}_q + \mathbf{g}_p), \quad (2)$$

where $\mathbf{g}_p$ and $\mathbf{l}_q$ denote the coordinates of p in the global coordinate system and any point q on Π_p in the local coordinate frame.

Segmentation with ScribblePrompt. The extracted slice S_p is segmented using ScribblePrompt, a SAM-based model trained on large-scale medical datasets with simulated scribbles, clicks, and bounding boxes to mimic realistic user interactions. We chose ScribblePrompt for its high segmentation accuracy and robustness to slight variations in input prompts. While ScribblePrompt typically relies on scribble-based input, we automate this step by directly using the centerline points generated from AortaDiff as prompts. This results in precise lumen segmentation, effectively adapting to anatomical variations, including geometric transitions like bifurcation. Examples of the segmentation results are shown in Figure 3.




Fig. 3: Examples of segmentation results using ScribblePrompt. The red point in the second row denotes the centerline point used as the prompt.

Contour extraction and processing. After obtaining the segmentation mask, we extract strong gradients corresponding to the vessel boundary by Canny edge detection. The detected edges are then processed with OpenCV's `findContours` function [2], which retrieves the outermost contour $\mathcal{C}_p^{2D}$ in a topologically consistent manner. Since NURBS requires a consistent number of control points across cross-sections, we uniformly resample $\mathcal{C}_p^{2D}$ to 32 points. This conservative choice exceeds the typical range used in manual segmentation (5 to 20 points), ensuring sufficient geometric accuracy for reliable hemodynamic analysis. It also maintains smooth, watertight surfaces, even when handling low-resolution or imperfect inputs.

Mapping contour back to 3D. Each 2D contour $\mathcal{C}_p^{2D}$ extracted from Π_p must be mapped back to 3D. Since Π_p is defined by the local coordinate frame $(\mathbf{t}_p, \mathbf{n}_p, \mathbf{b}_p)$ and centered at p , we transform each 2D contour point $\mathcal{C}_{p_i}^{2D}$ back to $\mathcal{C}_{p_i}^{3D}$ in the original 3D space using the inverse of $\mathbf{R}_p$, i.e., $\mathbf{R}_p^{-1} \cdot \mathbf{l}_c + \mathbf{g}_p$, where $\mathbf{l}_c$ denotes the coordinate of any contour point c in the local coordinate frame. This process ensures that the constructed contours align accurately with the input volume $\mathbf{V}$, preserving the anatomical vessel structure while maintaining a consistent contour representation for NURBS-based construction.

3.4 Aortic Surface Construction

After obtaining the aortic centerline and corresponding contours, we construct the 3D aortic surface using NURBS, achieving a smooth and anatomically accurate representation of the vessel.

Contour alignment. To ensure consistency across adjacent contours, we establish point-to-point correspondence using the *iterative closest point* (ICP) algorithm [8], which preserves spatial coherence along the vessel and provides a reliable foundation for surface construction.

Surface fitting. We implement a curve-fitting algorithm [50] to fit a cubic spline curve from the aligned contour points $\{\mathcal{C}_{p_i}^{3D}\}$, which gives the control points $\mathbf{p}_{i,j} \in \mathbb{R}^{m \times n \times 3}$ and weights $w_{i,j} \in \mathbb{R}^{m \times n}$

$$w_{i,j}, \mathbf{p}_{i,j} = \text{CurveFitting}(\{\mathcal{C}_{p_i}^{3D}\}), \quad (3)$$

where m and n represent the resolutions along the longitudinal (stream-wise) and circumferential directions, respectively. i denotes the longitudinal position along the centerline, while j represents the circumferential position around each cross-section. The estimated control points and weights can be directly used to create a 3D NURBS surface

$$S(u, v) = \frac{\sum_{i=0}^n \sum_{j=0}^m N_{i,p}(u) N_{j,q}(v) w_{i,j} \mathbf{p}_{i,j}}{\sum_{i=0}^n \sum_{j=0}^m N_{i,p}(u) N_{j,q}(v) w_{i,j}}, \quad (4)$$

where u represents the parametric direction along the aortic centerline, capturing the longitudinal dimension, and v represents the circumferential direction around each cross-sectional contour. $N_{i,p}(u)$ and $N_{j,q}(v)$ are B-spline basis functions, while $w_{i,j}$ and $\mathbf{p}_{i,j}$ represent the corresponding weights and control points. p and q denote the spline orders on u and v directions. The surface fitting process interpolates control points along u (centerline) and v (contours), generating a continuous parametric surface. This ensures the constructed aorta maintains geometric accuracy, effectively capturing subtle variations in vessel curvature, diameter, and morphology.

Multi-branch integration. We separately model the main aorta and supra-aortic branches and seamlessly integrate them into a unified surface. Overlapping regions are refined to eliminate discontinuities. NURBS provides a robust framework for vascular modeling, capable of capturing complex branching structures and detailed geometric variations with high precision.

3.5 CFD Simulation

To validate the feasibility of the constructed aorta meshes for downstream CFD simulations, we employ OpenFOAM [3], a widely used solver for simulating fluid dynamics in complex geometries. The simulation process consists of the following steps:

- **Geometry and mesh generation:** The constructed aorta geometry is remeshed into a high-quality triangulated surface using PyACVD [62], followed by volumetric tetrahedral meshing via VMTK’s TetGen-based algorithm [5]. The mesh is compatible with OpenFOAM and dense enough to capture key flow characteristics. The mesh comprises polyhedral elements with appropriate boundary definitions for the inlet, outlets, and vessel walls.
- **Governing equations and numerical discretization:** Each simulation solves steady-state incompressible Navier-Stokes equations

$$\begin{aligned} \nabla \cdot \mathbf{u} &= 0, \\ (\mathbf{u} \cdot \nabla) \mathbf{u} &= -\nabla p + \nu \nabla^2 \mathbf{u}, \end{aligned} \quad (5)$$

where $\mathbf{u}$ is the velocity field, p is the pressure, ν is the kinematic viscosity. We solve those equations using SimpleFoam on multiple GPUs in parallel via OpenMPI [19].

- **Boundary conditions:** Physiological boundary conditions are applied to replicate realistic aortic flow dynamics, including *inlet* (a prescribed velocity profile simulating either pulsatile or steady inflow conditions), *outlets* (either zero-gradient conditions or specified pressure values to allow natural flow development), and *vessel walls* (a no-slip boundary condition enforcing fluid adherence to the vessel wall).
- **Solver configuration:** The simulation employs the *semi-implicit method for pressure-linked equations* (SIMPLE) algorithm [49] in a steady-state formulation. To enhance computational efficiency, parallel domain decomposition is applied, distributing the computational workload across multiple processors.
- **Post-processing and flow analysis:** Once the simulation converges, key hemodynamic parameters, such as velocity magnitude and pressure distribution, are extracted to quantify fluid behavior. WSS is particularly important for assessing the risk of aneurysm rupture, thrombus formation, and disease progression, making it essential for CFD analysis to ensure the clinical relevance of the computational model.

4 RESULTS AND DISCUSSION

We conducted experiments on aortic data to evaluate AortaDiff’s performance and validate the design of our framework. Our evaluation

compares AortaDiff against state-of-the-art methods for centerline generation and contour extraction. The resulting aorta meshes are assessed through both quantitative metrics and qualitative analysis. To further demonstrate the practical applicability of AortaDiff, we performed CFD simulations on the generated meshes and compared the results with those obtained from GT meshes.

4.1 Experimental Setup

To train a 3D mesh model of the aorta, the dataset must meet the following criteria: (1) It should contain a sufficient number of aorta models with diverse morphologies. (2) It should include a variety of aortic structures, such as the main aorta and supra-aortic branches. (3) It should provide mesh labels for the aorta or, at the very least, segmentation labels. However, datasets containing aorta models are limited. Most available datasets provide only segmentation labels for the main aorta, while mesh labels are rare.

To facilitate the training of our AortaDiff model, we used the *vascular model repository* (VMR) dataset [64], which includes 3D aorta meshes along with corresponding MRI volumes. This dataset features various blood vessels, such as the aorta, cerebral arteries, pulmonary arteries, and coronary arteries. In total, it contains 275 blood vessel cases, including 91 aortas from both animals and humans. Since the dataset includes non-human aortas, we specifically selected 22 human aorta cases for our study, with 18 used for training and four for inference.

To evaluate the generalizability of AortaDiff, we additionally incorporated several cases from the *aortic vessel tree* (AVT) dataset [51], which consists of CT volumes. This dataset exclusively contains human aorta models, with a total of 56 cases. We selected eight cases with diverse aortic morphologies to ensure a more comprehensive evaluation. Since the AVT dataset does not provide preexisting aorta meshes, we generated them by converting the segmentation labels into surface meshes using the marching cubes algorithm [34]. We then performed manual post-processing to ensure surface smoothness and to eliminate topological artifacts such as holes.

All CT/MRI volumes were resampled with an isotropic spacing of $0.8 \times 0.8 \times 0.3$ mm, resized to $128 \times 128 \times 128$, and normalized to an intensity range of $[0, 1]$ while retaining the full volume. Data augmentation was applied using MONAI’s *RandFlipd* and *RandRotate90d* [1], each with a probability of 0.1. Model training was performed on a single NVIDIA A40 GPU with 48 GB of memory. We used the Adam optimizer with hyperparameters $\beta_1 = 0.9$ and $\beta_2 = 0.99$. The learning rate was initialized at 1×10^{-3} with a batch size of 16. The training process was run for a total of 100,000 iterations.

In the paper, the tables give quantitative results for all twelve test cases (four from the VMR dataset and eight from the AVT dataset), while the figures showcase qualitative results for six selected cases. The remaining six cases are included in Section 3 of the appendix.

Table 1: Average CD, HD, and EMD for centerline generation on all test cases from the VMR and AVT datasets. For all tables in this paper, values are in millimeters (mm) as mean $\pm$ standard deviation, with the best results highlighted in bold.

metric	VMR				AVT			
	DPM	PVD	Dit-3D	AortaDiff	DPM	PVD	Dit-3D	AortaDiff
supra-aortic branches								
CD $\downarrow$	2.68 $\pm$ 0.05	2.46 $\pm$ 0.06	2.25 $\pm$ 0.09	0.37$\pm$0.02	3.05 $\pm$ 0.25	2.80 $\pm$ 0.13	2.33 $\pm$ 0.25	0.55$\pm$0.02
HD $\downarrow$	3.90 $\pm$ 0.20	3.62 $\pm$ 0.15	3.12 $\pm$ 0.10	0.35$\pm$0.01	3.93 $\pm$ 0.15	3.87 $\pm$ 0.25	4.05 $\pm$ 0.10	0.37$\pm$0.01
EMD $\downarrow$	1.82 $\pm$ 0.05	1.75 $\pm$ 0.05	1.54 $\pm$ 0.09	0.19$\pm$0.02	1.57 $\pm$ 0.12	1.63 $\pm$ 0.08	1.60 $\pm$ 0.15	0.29$\pm$0.00
main aorta								
CD $\downarrow$	2.09 $\pm$ 0.03	1.97 $\pm$ 0.05	1.71 $\pm$ 0.06	0.29$\pm$0.01	1.85 $\pm$ 0.20	1.88 $\pm$ 0.12	1.89 $\pm$ 0.25	0.44$\pm$0.02
HD $\downarrow$	3.49 $\pm$ 0.15	3.11 $\pm$ 0.15	2.60 $\pm$ 0.10	0.26$\pm$0.00	3.30 $\pm$ 0.10	3.39 $\pm$ 0.22	3.55 $\pm$ 0.10	0.32$\pm$0.01
EMD $\downarrow$	1.40 $\pm$ 0.03	1.34 $\pm$ 0.05	1.15 $\pm$ 0.05	0.14$\pm$0.01	1.30 $\pm$ 0.13	1.30 $\pm$ 0.05	1.30 $\pm$ 0.13	0.23$\pm$0.02
overall								
CD $\downarrow$	2.39 $\pm$ 0.05	2.22 $\pm$ 0.06	1.98 $\pm$ 0.08	0.33$\pm$0.02	2.45 $\pm$ 0.23	2.34 $\pm$ 0.13	2.11 $\pm$ 0.25	0.50$\pm$0.02
HD $\downarrow$	3.70 $\pm$ 0.18	3.37 $\pm$ 0.15	2.85 $\pm$ 0.10	0.31$\pm$0.01	3.62 $\pm$ 0.13	3.63 $\pm$ 0.24	3.75 $\pm$ 0.10	0.35$\pm$0.01
EMD $\downarrow$	1.60 $\pm$ 0.05	1.52 $\pm$ 0.05	1.32 $\pm$ 0.07	0.17$\pm$0.02	1.45 $\pm$ 0.13	1.46 $\pm$ 0.06	1.45 $\pm$ 0.14	0.26$\pm$0.03

4.2 Centerline Generation

AortaDiff decomposes the aorta into a centerline and its corresponding orthogonal contours, with the centerline serving as the foundation for mesh construction. Following Section 3.1, we use a 1D-UNet to process the 1D image representation of the centerline. Alternatively, a point cloud generation model could be employed since the centerline can be viewed as a 3D point cloud. To validate our design choice, we compare AortaDiff against three baseline models: *diffusion probabilistic model* (DPM) [35], *point-voxel diffusion* (PVD) [72], and *Dit-3D* [43].




Fig. 4: Comparison of centerline generation methods on the selected test cases from the VMR and AVT datasets.

Table 2: Average Dice, ASD, and HD for segmentation on all test cases from the VMR and AVT datasets.

metric	VMR					AVT				
	SAM	MedSAM	TotalSeg	LoGB-Net	AortaDiff	SAM	MedSAM	TotalSeg	LoGB-Net	AortaDiff
Dice $\uparrow$	0.575 $\pm$ 0.078	0.757 $\pm$ 0.024	0.800 $\pm$ 0.007	0.890 $\pm$ 0.006	0.967$\pm$0.020	0.553 $\pm$ 0.075	0.736 $\pm$ 0.019	0.782 $\pm$ 0.020	0.870 $\pm$ 0.018	0.947$\pm$0.012
ASD $\downarrow$	1.720 $\pm$ 0.350	0.960 $\pm$ 0.148	0.715 $\pm$ 0.160	0.705 $\pm$ 0.170	0.502$\pm$0.035	1.638 $\pm$ 0.315	0.875 $\pm$ 0.170	0.635 $\pm$ 0.150	0.620 $\pm$ 0.160	0.457$\pm$0.055
HD $\downarrow$	9.150 $\pm$ 2.880	8.790 $\pm$ 3.000	6.580 $\pm$ 0.330	6.470 $\pm$ 0.340	5.752$\pm$0.200	8.420 $\pm$ 2.860	8.060 $\pm$ 2.990	6.100 $\pm$ 0.360	6.000 $\pm$ 0.370	5.312$\pm$0.230
supra-aortic branches										
Dice $\uparrow$	0.605 $\pm$ 0.078	0.787 $\pm$ 0.024	0.910 $\pm$ 0.007	0.918 $\pm$ 0.006	0.972$\pm$0.020	0.565 $\pm$ 0.075	0.746 $\pm$ 0.019	0.882 $\pm$ 0.020	0.880 $\pm$ 0.018	0.952$\pm$0.012
ASD $\downarrow$	1.690 $\pm$ 0.350	0.938 $\pm$ 0.148	0.705 $\pm$ 0.160	0.693 $\pm$ 0.170	0.490$\pm$0.035	1.628 $\pm$ 0.315	0.868 $\pm$ 0.170	0.625 $\pm$ 0.150	0.608 $\pm$ 0.160	0.450$\pm$0.055
HD $\downarrow$	9.110 $\pm$ 2.880	8.740 $\pm$ 3.000	6.520 $\pm$ 0.330	6.410 $\pm$ 0.340	5.720$\pm$0.200	8.390 $\pm$ 2.860	8.030 $\pm$ 2.990	6.050 $\pm$ 0.360	5.950 $\pm$ 0.370	5.280$\pm$0.230
main aorta										
Dice $\uparrow$	0.590 $\pm$ 0.078	0.772 $\pm$ 0.022	0.885 $\pm$ 0.007	0.904 $\pm$ 0.010	0.970$\pm$0.016	0.559 $\pm$ 0.075	0.741 $\pm$ 0.019	0.860 $\pm$ 0.020	0.880 $\pm$ 0.018	0.950$\pm$0.012
ASD $\downarrow$	1.705 $\pm$ 0.340	0.949 $\pm$ 0.160	0.710 $\pm$ 0.160	0.699 $\pm$ 0.170	0.496$\pm$0.045	1.633 $\pm$ 0.315	0.872 $\pm$ 0.170	0.630 $\pm$ 0.150	0.614 $\pm$ 0.160	0.454$\pm$0.055
HD $\downarrow$	9.130 $\pm$ 2.880	8.765 $\pm$ 3.000	6.550 $\pm$ 0.330	6.440 $\pm$ 0.340	5.736$\pm$0.215	8.405 $\pm$ 2.860	8.045 $\pm$ 2.990	6.075 $\pm$ 0.360	5.975 $\pm$ 0.370	5.296$\pm$0.230
overall										
Dice $\uparrow$	0.590 $\pm$ 0.078	0.772 $\pm$ 0.022	0.885 $\pm$ 0.007	0.904 $\pm$ 0.010	0.970$\pm$0.016	0.559 $\pm$ 0.075	0.741 $\pm$ 0.019	0.860 $\pm$ 0.020	0.880 $\pm$ 0.018	0.950$\pm$0.012
ASD $\downarrow$	1.705 $\pm$ 0.340	0.949 $\pm$ 0.160	0.710 $\pm$ 0.160	0.699 $\pm$ 0.170	0.496$\pm$0.045	1.633 $\pm$ 0.315	0.872 $\pm$ 0.170	0.630 $\pm$ 0.150	0.614 $\pm$ 0.160	0.454$\pm$0.055
HD $\downarrow$	9.130 $\pm$ 2.880	8.765 $\pm$ 3.000	6.550 $\pm$ 0.330	6.440 $\pm$ 0.340	5.736$\pm$0.215	8.405 $\pm$ 2.860	8.045 $\pm$ 2.990	6.075 $\pm$ 0.360	5.975 $\pm$ 0.370	5.296$\pm$0.230

These state-of-the-art baselines are chosen as their implementations are publicly available. However, because these models are inherently generative, the shape of their generated centerlines varies stochastically. To ensure a fair comparison, we adapt them into conditional generative models by following AortaDiff’s conditioning strategy. Specifically, we use a ViT [13] encoder to extract features from the input volume, align them with the centerline points, and concatenate them with the point cloud input at each diffusion step. The qualitative evaluation results are presented in Figure 4. We can see that AortaDiff is the clear winner across the selected cases from the VMR and AVT datasets.

To assess the similarity between the generated and GT point clouds across different methods, we employ three metrics: *chamfer distance* (CD) [7], *Hausdorff distance* (HD) [22], and *Earth mover’s distance* (EMD) [53]. CD measures the overall quality of the generated point clouds by computing the average nearest-neighbor distance. HD captures the worst-case discrepancy by quantifying the greatest minimum distance between two sets of points, providing insight into the maximum bias in the generated centerlines. EMD evaluates the global structural alignment between the generated and reference point clouds.

The quantitative evaluation results across all test cases are presented in Table 1. AortaDiff consistently outperforms all baselines. On the VMR dataset, AortaDiff achieves a CD of 0.33 mm, an HD of 0.31 mm, and an EMD of 0.17 mm, demonstrating significantly improved accuracy over the baseline models. Even on the AVT dataset, despite differences in image acquisition (CT modality), AortaDiff maintains superior performance, highlighting its robustness and generalization capability. These results confirm that AortaDiff produces more accurate and structurally consistent centerlines than baseline models. The centerline diffusion process is presented in Section 2 of the appendix.

4.3 Contour Extraction

The centerline generated in Section 3.1 does not form a perfectly smooth curve, leading to intersections at adjacent cross-sectional planes defined by normal vectors from neighboring centerline points. These intersections negatively impact contour extraction, potentially compromising the subsequent NURBS mesh construction. We apply NURBS curve fitting to address this issue, using the original centerline points as control points for a B-spline curve. A suitable knot vector is defined, and the fitted curve is uniformly resampled to produce a smooth and continuous centerline representation.

Building on the refined centerlines, AortaDiff uses ScribblePrompt to extract aortic contours on each cross-sectional plane of the input volume. This step is crucial for constructing the aorta mesh using the NURBS method. We compare AortaDiff with state-of-the-art segmentation approaches, including SAM-based methods—SAM and MedSAM [37]—as well as the supervised learning-based TotalSegmen-

tator [63] (TotalSeg) and LoGB-Net [4]. Multiple prompt formats, such as bounding boxes and points, can guide segmentation for SAM-based models. Here, we use the generated centerline points as point prompts to facilitate segmentation, ensuring spatial alignment between the extracted contours and the centerline. In contrast to TotalSeg and LoGB-Net’s direct volume-to-segmentation approach, which processes the entire CT volume as input, AortaDiff operates on orthogonal cross-sections, leveraging localized structural cues to enhance segmentation accuracy. To quantitatively evaluate segmentation performance, we employ three metrics: the *Dice coefficient* (Dice) [12], *average surface distance* (ASD) [24], and HD. The results are summarized in Table 2.

The results show that compared to SAM, MedSAM and AortaDiff significantly enhance segmentation performance, thanks to their training using large-scale medical imaging data, which strengthens their ability to capture anatomical structures. However, despite utilizing centerline-based prompts, MedSAM still struggles to accurately delineate the aortic region, particularly in the supra-aortic branches. LoGB-Net, a fully supervised model specifically designed for aorta segmentation, outperforms both SAM and MedSAM across all metrics. TotalSeg also achieves strong performance on the main aorta but detects only fragmented segments of the supra-aortic branches, as it is a general-purpose model not optimized for these structures. Nevertheless, both LoGB-Net and TotalSeg remain inferior to AortaDiff, likely because they cannot incorporate structural information from centerline point prompts. In contrast, AortaDiff benefits from ScribblePrompt’s strong prior knowledge gained through large-scale self-supervised pretraining, enabling it to generalize better to unseen vessel structures.

To further evaluate segmentation performance, we visualize the results of different methods in Figure 5. The results indicate that SAM-based methods struggle to accurately segment the aorta, particularly in the supra-aortic branches. For example, in the third and fourth rows of Figure 5, the segmentation results of SAM and MedSAM show significant deviations from the GT in the AVT dataset. In contrast, for the VMR dataset, these methods occasionally produce segmentations that approximate the GT in the supra-aortic branches. This inconsistency suggests that SAM-based methods lack the robustness needed for aorta segmentation across different datasets, likely due to their reliance on prompts that fail to fully capture the complex anatomical variations. Meanwhile, TotalSeg, LoGB-Net, and AortaDiff demonstrate superior segmentation performance. However, TotalSeg and LoGB-Net still fall short of AortaDiff. In particular, TotalSeg sometimes detects only fragments of the supra-aortic branches, resulting in discontinuous or incomplete segmentation in these regions. Compared to the fully supervised methods that directly extract contours for mesh construction, AortaDiff benefits from ScribblePrompt’s pretraining on large-scale




Fig. 5: Comparison of segmentation methods on the selected test cases from the VMR and AVT datasets. Each case includes segmentation results for both the main aorta and supra-aortic branches. Cross-sections are resampled on planes orthogonal to the centerline, with each plane centered on a corresponding centerline point.




Fig. 6: Comparison of contours extracted from different segmentation methods on the selected test cases from the VMR and AVT datasets.

medical imaging data and integrates centerline-based prompts, enhancing its robustness and adaptability to anatomical variability.

In Figure 6, we visualize the extracted aortic contours on each cross-sectional plane of the input volume. The results indicate that segmentation inaccuracies often produce highly distorted vessel contours, particularly in the supra-aortic branches. Notably, when SAM and MedSAM fail to recognize the aorta, they frequently generate excessively large contours that span most of the image, leading to severe segmentation artifacts. Refer to Figure 5. In the VMR dataset, the cross-sections exhibit relatively low noise, allowing SAM and MedSAM to precisely segment. However, the AVT dataset contains noisier cross-sections with additional tissues and organs, which adversely affect the segmentation performance of SAM and MedSAM, resulting in suboptimal outcomes. This explains the irregular and fragmented contours observed in their segmentation outputs. TotalSeg exhibits difficulty in detecting the supra-aortic branches, leading to the loss of their corresponding contours. In contrast, AortaDiff extracts contours with significantly greater accuracy and consistency across both the main aorta and supra-aortic branches. To further validate this observation, we conduct a quantitative analysis of contour extraction across different segmentation methods. As shown in Table 3, AortaDiff consistently produces contours that more closely align with the GT.

4.4 Aorta Mesh Construction

We construct the aorta mesh from the extracted contours using the NURBS method. Other segmentation methods produce suboptimal contours that fail to meet the requirements for NURBS interpolation and perform worse than those generated by AortaDiff. Therefore, we only compare the NURBS-constructed meshes from AortaDiff with the marching cubes [34] and surface nets [18] generated meshes from LoGB-Net [4], a state-of-the-art model for aorta segmentation. We

also compared with marching cubes-generated meshes from the EFDs-based signed distance field [42]. For the VMR dataset, GT meshes are provided, allowing direct comparison. In contrast, for the AVT dataset, we manually constructed the GT meshes from the segmentation results using SimVascular to ensure anatomical consistency. The marching cubes-generated meshes are obtained by converting segmentation labels into mesh representations. To ensure a fair comparison, only basic smoothing techniques are applied to reduce noise and eliminate small holes while preserving overall geometry. Each supra-aortic branch is constructed separately and assembled into a complete aorta mesh via NURBS. The quantitative results in Table 4 demonstrate that AortaDiff outperforms other methods in terms of CD, HD, and EMD.

As shown in Figure 7, the NURBS-constructed meshes from AortaDiff exhibit significantly greater smoothness and structural consistency than the marching cubes-constructed meshes from other methods.

While LoGB-Net produces segmentation results that closely resemble the GT, its marching cubes-generated meshes suffer from severe artifacts, particularly in the supra-aortic branches. Meshes from LoGB-Net+surface nets show slight improvements in smoothness and structural continuity compared to those from LoGB-Net+marcching cubes, but still exhibit noticeable artifacts. These artifacts include missing sections, jagged surfaces, and structural discontinuities, especially in regions with complex vessel bifurcations. Furthermore, the marching cubes method applied to the LoGB-Net output often causes shrinkage or even complete disappearance of supra-aortic branches with small radii. As a result, its accuracy is highly dependent on the resolution of the input volume. In contrast, AortaDiff demonstrates robustness against such artifacts, generating meshes that are complete, smooth, and topologically consistent. Although EFDs leverage smoothed curves to generate a better implicit field, the resulting meshes are still derived from

Table 3: Average CD, HD, and EMD for contour point clouds on all test cases from the VMR and AVT datasets.

metric	VMR					AVT				
	SAM	MedSAM	TotalSeg	LoGB-Net	AortaDiff	SAM	MedSAM	TotalSeg	LoGB-Net	AortaDiff
supra-aortic branches										
CD ↓	3.40±0.55	2.85±0.50	3.00±0.52	1.95±0.35	1.30±0.18	3.65±0.58	3.00±0.52	3.10±0.55	2.10±0.38	1.40±0.20
HD ↓	9.80±3.10	9.10±3.15	9.20±3.20	6.80±0.55	6.00±0.35	10.20±3.25	9.40±3.20	9.50±3.30	7.10±0.58	6.20±0.38
EMD ↓	1.85±0.45	1.30±0.35	1.35±0.38	1.00±0.25	0.68±0.12	2.00±0.48	1.40±0.38	1.45±0.40	1.10±0.28	0.72±0.14
main aorta										
CD ↓	3.20±0.50	2.70±0.45	1.95±0.32	1.85±0.30	1.20±0.12	3.45±0.52	2.80±0.48	2.10±0.35	1.98±0.32	1.30±0.14
HD ↓	9.60±3.00	8.90±3.05	6.80±0.52	6.60±0.45	5.90±0.27	10.00±3.20	9.20±3.15	7.20±0.55	6.90±0.50	6.05±0.32
EMD ↓	1.82±0.40	1.28±0.32	1.00±0.22	0.92±0.19	0.61±0.09	1.95±0.42	1.35±0.35	1.10±0.24	1.02±0.22	0.66±0.11
overall										
CD ↓	3.28±0.52	2.78±0.47	2.30±0.40	1.90±0.33	1.25±0.15	3.55±0.55	2.90±0.50	2.45±0.42	2.00±0.35	1.35±0.17
HD ↓	9.70±3.05	9.00±3.10	7.50±0.55	6.70±0.48	5.95±0.30	10.10±3.23	9.30±3.18	7.40±0.58	7.00±0.52	6.10±0.34
EMD ↓	1.84±0.42	1.29±0.34	1.10±0.23	0.96±0.20	0.64±0.10	1.98±0.44	1.38±0.36	1.20±0.25	1.05±0.23	0.69±0.12




Fig. 7: Comparison of mesh construction methods on the selected test cases from the VMR and AVT datasets. For each case, we overlay the opaque constructed mesh with the semi-transparent GT mesh to facilitate visual comparison. For the aneurysm case from the VMR dataset, we show a zoomed-in comparison of the mesh constructed by EFDs+marching cubes, LoGB-Net+surface nets, and AortaDiff. We also provide a zoomed-in view of the cross-sections of the main aorta and a supra-aortic branch. In these cross-sections, the red, gray, pink, blue, and green curves represent the meshes from GT, LoGB-Net+marching cubes, LoGB-Net+surface nets, EFDs+marching cubes, and AortaDiff, respectively.

Table 4: Average CD, HD, and EMD for the constructed aorta meshes on all test cases from the VMR and AVT datasets.

metric	VMR					AVT				
	LoGB-Net+ marching cubes	LoGB-Net+ surface nets	EFDs+ marching cubes	AortaDiff	supra-aortic branches	LoGB-Net+ marching cubes	LoGB-Net+ surface nets	EFDs+ marching cubes	AortaDiff	supra-aortic branches
CD ↓	0.65±0.09	0.62±0.08	0.58±0.07	0.31±0.04	0.80±0.11	0.76±0.10	0.72±0.09	0.43±0.06	0.75±0.11	0.75±0.11
HD ↓	2.30±0.27	2.25±0.25	2.10±0.23	1.30±0.22	2.60±0.32	2.55±0.30	2.40±0.28	1.75±0.45	2.40±0.28	2.40±0.28
EMD ↓	1.25±0.28	1.20±0.26	1.12±0.24	0.80±0.20	1.35±0.32	1.30±0.30	1.22±0.28	0.92±0.30	1.22±0.28	1.22±0.28
main aorta										
CD ↓	0.55±0.07	0.52±0.06	0.48±0.06	0.26±0.03	0.70±0.09	0.67±0.08	0.63±0.08	0.36±0.05	0.63±0.08	0.63±0.08
HD ↓	2.10±0.23	2.05±0.22	1.95±0.20	1.18±0.18	2.40±0.28	2.35±0.27	2.25±0.25	1.58±0.42	2.25±0.25	2.25±0.25
EMD ↓	1.15±0.23	1.12±0.22	1.08±0.21	0.72±0.18	1.25±0.28	1.22±0.27	1.18±0.25	0.84±0.28	1.18±0.25	1.18±0.25
overall										
CD ↓	0.60±0.08	0.57±0.07	0.53±0.07	0.28±0.04	0.75±0.10	0.72±0.09	0.68±0.09	0.39±0.06	0.68±0.09	0.68±0.09
HD ↓	2.20±0.25	2.15±0.24	2.05±0.22	1.23±0.20	2.50±0.30	2.45±0.29	2.35±0.27	1.66±0.44	2.35±0.27	2.35±0.27
EMD ↓	1.20±0.25	1.17±0.24	1.12±0.22	0.76±0.19	1.30±0.30	1.27±0.29	1.22±0.27	0.88±0.29	1.22±0.27	1.22±0.27

marching cubes and exhibit artifacts and poor smoothness—though less pronounced than those from LoGB-Net—making them unsuitable for CFD simulations. Additionally, the evaluation includes three challenging cases—two aneurysm cases and one coarctation case—from the VMR and AVT datasets, further validating AortaDiff’s capability to construct aorta meshes with complex geometries. Notably, AortaDiff was trained solely on the VMR dataset, while the AVT dataset was used exclusively for testing. Despite this, AortaDiff successfully constructs aorta meshes from the AVT dataset, producing smooth, complete surfaces that align with the vascular structures observed in CT images. The constructed meshes preserve anatomical details while maintaining structural continuity, even for previously unseen datasets.

4.5 Hemodynamic Simulation

Using OpenFOAM, we performed hemodynamic simulations on aorta meshes constructed with AortaDiff and compared them to GT meshes.

CFD simulations require smooth, watertight surface meshes with anatomically accurate features and sufficiently high resolution to ensure reliable flow field computation, though they do not impose a strict lower bound on angular discretization. This flexibility is a key advantage of our NURBS-based parameterization, which produces simulation-ready geometries that remain smooth and watertight even from sparse or noisy cross-sectional inputs, enabling robust integration with CFD workflows. To ensure simulation reliability, we conducted a mesh independence study, confirming that our chosen resolution—approximately 1 million cells—yields mesh-converged results. Each simulation was run until velocity and pressure residuals fell below 1×10^{-6} , with continuity errors well controlled at approximately 1×10^{-5} , ensuring numerical accuracy. These simulations capture blood flow velocity, pressure fields, and WSS distributions. To ensure a fair comparison, identical boundary conditions were applied to both GT and AortaDiff meshes. Meshes constructed from LoGB-Net segmentations were excluded from this experiment due to insufficient geometric fidelity and structural discontinuities, making them unsuitable for CFD analysis.

As shown in Figure 8, the visualization results demonstrate a high degree of consistency in flow dynamics and stress patterns between AortaDiff and GT meshes. In particular, the spatial distribution of velocity, pressure, and WSS indicates that AortaDiff constructions accurately preserve key anatomical features critical for precise hemodynamic modeling. For both the VMR and AVT datasets, velocity streamlines exhibit similar flow paths and magnitudes, with characteristic acceleration near the ascending aorta and arch, followed by deceleration toward the descending segment. Flow direction and lumen confinement remain consistent, with no noticeable divergence or dispersion




Fig. 8: Comparison of hemodynamic simulation results on the selected test cases from the VMR and AVT datasets between aorta meshes constructed using AortaDiff and GT meshes. The visualization includes velocity fields, pressure distributions, and WSS magnitudes.

artifacts. Pressure distributions show smooth gradients from the inlet to the distal outlets, with no abrupt discontinuities or localized anomalies in the AortaDiff meshes, closely matching the patterns observed in GT meshes. WSS visualizations reveal similar spatial variations between AortaDiff and GT results. Elevated WSS values consistently appear on the inner curvature of the arch and branch origins, while lower WSS zones are observed on outer walls and downstream regions. Such spatial patterns are preserved across all test cases and datasets, including high-curvature and bifurcation areas. These findings suggest that the geometric properties influencing local hemodynamic conditions—such as surface smoothness, curvature continuity, and vessel tapering—are well preserved in AortaDiff constructions. As a result, the simulations produce comparable flow, pressure, and WSS characteristics across all anatomical regions evaluated.

5 AD-HOC EXPERT FEEDBACK

AortaDiff was tested and evaluated by a CFD and cardiovascular flow expert with over ten years of research experience, who provided the following feedback: “AortaDiff offers an efficient and accurate method for generating meshes for hemodynamic simulations, a critical component of image-based CFD in clinical diagnosis and treatment planning. Unlike traditional machine learning-based segmentation models, which typically produce binary voxel images, AortaDiff directly generates simulation-ready meshes that can be seamlessly integrated into any downstream fluid solver. To verify this capability, I used AortaDiff to generate three distinct aortic geometries from three CT images. The resulting geometries closely aligned with the raw images—an impressive outcome given the speed and full automation of the process.”

The expert further noted “To assess the feasibility of these meshes for CFD, I employed OpenFOAM to conduct fluid simulations. The simulation solved the steady-state incompressible Navier-Stokes equations using a parabolic velocity profile at the inlet and a fixed pressure at the outlet. The simulation ran successfully, demonstrating a smooth transition from image processing to hemodynamic analysis. The flow solution revealed a reasonable pressure drop along the streamwise direction, with the highest pressure observed in the ascending aorta and the aortic arch, and the lowest pressure near the descending aorta. As expected, the pressure remained consistently high at junctions between the main aorta and its branches due to vortex formation and flow separation in these geometrically complex regions. Similarly, higher WSS values were observed near the branching points of the aortic arch, where flow acceleration and bifurcations occur. This suggests a potential risk of wall deterioration, which may contribute to cardiovascular conditions such as aneurysms and coarctations.”

The expert continued “This experiment confirms that AortaDiff produces high-quality, accurate, and CFD-compatible meshes that integrate seamlessly into CFD workflows. These meshes enable advanced simulations with higher-fidelity settings, such as fluid-structure interactions and Windkessel boundary conditions. Additionally, since the generated surface is a NURBS surface, alternative CFD approaches—such as isogeometric analysis—can further refine flow hemodynamics, expanding AortaDiff’s applicability in computational hemodynamics research. The parameterized NURBS surfaces also facilitate statistical

shape analysis when sufficient patient-specific data is available. This enables systematic quantification of anatomical variations, identification of shape-based risk factors, and development of personalized diagnostic and prognostic models.

Finally, the expert concluded “In conclusion, I highly recommend AortaDiff to researchers and engineers working on image-based CFD tasks, as it provides accurate, reliable, and highly efficient meshes.”

6 CONCLUSIONS AND FUTURE WORK

We have introduced AortaDiff, a novel framework for generating 3D aorta meshes directly from volumetric CT/MRI data. Constructing aortic geometries presents several challenges, including the scarcity of annotated data, the heavy reliance on manual intervention, and the absence of a complete end-to-end solution. To address these issues, AortaDiff employs a volume-guided CDM to generate aortic centerlines—a particularly effective approach given the limited training data available for point cloud generation. The generated centerline points serve as prompts for semi-automatic segmentation, enabling precise contour extraction from cross-sectional images and significantly reducing manual effort. We use these contours to construct a smooth 3D mesh via the NURBS method, ensuring an automated pipeline.

AortaDiff is designed as a modular framework, where each component, such as point cloud generation, contour segmentation, and mesh construction, can be independently improved or replaced to further enhance the overall mesh quality as new methods emerge. Our experiments show that AortaDiff produces high-quality aorta meshes to support hemodynamic simulations comparable to those derived from GT meshes, outperforming conventional marching cubes approaches. Furthermore, image-based CFD results generated from our meshes provide clinically valuable insights, including biomarkers derived from WSS and pressure drop analyses in aortic coarctation, supporting both diagnosis and treatment planning. The parametric, NURBS-based model enables flexible hemodynamic recomputation under customized anatomical modifications, enabling in silico experimentation for surgical planning and device design. Unlike traditional voxel- or parameterization-based manual methods, AortaDiff automates segmentation, accelerates mesh generation, and delivers higher-resolution surfaces without increasing manual effort—capabilities that are essential for scalable clinical integration.

However, AortaDiff still has some remaining limitations. In certain cases, the CDM-generated centerlines are not perfectly centered, though they still reside within the aortic region. While the robustness of the segmentation module mitigates the major impact on performance, this highlights an opportunity for improvement. Generating CFD-ready meshes for more complex aortic anatomies, such as aortic dissection, remains an open challenge and will be investigated through structural reasoning and lumen-aware modeling.

ACKNOWLEDGMENTS

This research was supported in part by the University of Notre Dame’s Scientific Artificial Intelligence (SAI) Graduate Fellowship, the U.S. National Science Foundation through grants IIS-1955395, OAC-2047127, IIS-2101696, OAC-2104158, and IIS-2401144, and the

U.S. National Institutes of Health through grant 1R01HL177814. Delin An would like to thank Drs. Emily Johnson and Donald Brower for their advice on this work through the SAI program. The authors would like to thank the anonymous reviewers for their insightful comments.

REFERENCES

- [1] Medical Open Network for AI (MONAI). <https://monai.io/>. 5
- [2] Open Computer Vision Library (OpenCV). <https://opencv.org/>. 4
- [3] Open Field Operation and Manipulation (OpenFOAM). <http://www.openfoam.com/>. 5
- [4] D. An, P. Du, P. Gu, J.-X. Wang, and C. Wang. Hierarchical LoG Bayesian neural network for enhanced aorta segmentation. In *Proceedings of IEEE International Symposium on Biomedical Imaging*, pp. 1–5, 2025. doi: 10.1109/ISBI60581.2025.10980947 6, 7
- [5] 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*, 46(11):1097–1112, 2008. doi: 10.1007/S11517-008-0420-1 5
- [6] M. Auer and T. C. Gasser. Reconstruction and finite element mesh generation of abdominal aortic aneurysms from computerized tomography angiography data with minimal user interactions. *IEEE Transactions on Medical Imaging*, 29(4):1022–1028, 2010. doi: 10.1109/TMI.2009.2039579 1
- [7] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. In *Proceedings of International Joint Conference on Artificial Intelligence*, pp. 659–663, 1977. doi: 10.5555/1622943.1622971 6
- [8] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. doi: 10.1109/34.121791 4
- [9] S. M. Black, C. Maclean, P. H. Barrientos, K. Ritos, and A. Kazakidi. Reconstruction and validation of arterial geometries for computational fluid dynamics using multiple temporal frames of 4D flow-MRI magnitude images. *Cardiovascular Engineering and Technology*, 14(5):655–676, 2023. doi: 10.1007/s13239-023-00679-x 3
- [10] Y. Chen, O. Wang, R. Zhang, E. Shechtman, X. Wang, and M. Gharbi. Image neural field diffusion models. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8007–8017, 2024. doi: 10.1109/CVPR52733.2024.00765 3
- [11] M. Danu, C. I. Nita, A. Vizitiu, C. Suciu, and L. M. Itu. Deep learning based generation of synthetic blood vessel surfaces. In *Proceedings of International Conference on System Theory, Control and Computing*, pp. 662–667, 2019. doi: 10.1109/ICSTCC.2019.8885576 1
- [12] L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945. doi: 10.2307/1932409 6
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16×16 words: Transformers for image recognition at scale. In *Proceedings of International Conference on Learning Representations*, 2021. 4, 6
- [14] P. Du, D. An, C. Wang, and J.-X. Wang. AI-powered automated model construction for patient-specific CFD simulations of aortic flows. *arXiv preprint arXiv:2503.12515*, 2025. doi: 10.48550/arXiv.2503.12515 2
- [15] P. Du, M. Xu, X. Zhu, and J.-X. Wang. HUG-VAS: A hierarchical NURBS-based generative model for aortic geometry synthesis and controllable editing. *arXiv preprint arXiv:2507.11474*, 2025. doi: 10.48550/arXiv.11474 2
- [16] P. Erler, P. Guerrero, S. Ohrhallinger, M. Wimmer, and N. J. Mitra. Points2Surf: Learning implicit surfaces from point cloud patches. *arXiv preprint arXiv:2007.10453*, 2020. doi: 10.48550/arXiv.2007.10453 3
- [17] A. Fantazzini, M. Esposito, A. Finotello, F. Auricchio, B. Pane, C. Basso, G. Spinella, and M. Conti. 3D automatic segmentation of aortic computed tomography angiography combining multi-view 2D convolutional neural networks. *Cardiovascular Engineering and Technology*, 11:576–586, 2020. doi: 10.1007/s13239-020-00481-z 1
- [18] S. F. F. Gibson. Using distance maps for accurate surface representation in sampled volumes. In *Proceedings of IEEE Symposium on Volume Visualization*, pp. 23–30, 1998. doi: 10.1145/288126.288142 2, 7
- [19] R. L. Graham, T. S. Woodall, and J. M. Squyres. Open MPI: A flexible high performance MPI. In *Proceedings of International Conference on Parallel Processing and Applied Mathematics*, pp. 228–239, 2006. doi: 10.1007/11752578_29 5
- [20] M. Habijan, I. Galić, H. Leventić, K. Romić, and D. Babin. Abdominal aortic aneurysm segmentation from CT images using modified 3D U-Net with deep supervision. In *Proceedings of International Symposium on Electronics in Marine*, pp. 123–128, 2020. doi: 10.1109/ELMAR49956.2020.9219015 1
- [21] H. K. Hahn, B. Preim, D. Selle, and H.-O. Peitgen. Visualization and interaction techniques for the exploration of vascular structures. In *Proceedings of IEEE Visualization Conference*, pp. 395–402, 2001. doi: 10.1109/VISUAL.2001.964538 3
- [22] F. Hausdorff. *Grundzüge der Mengenlehre*. Veit & Comp, Leipzig, 1914. 6
- [23] W. He, Y. Cao, Y. Li, Y. Miao, W. Shi, F. He, F. Yan, Z. Jiang, and H. Zhang. A study on CT aorta segmentation using vessel enhancement diffusion filter and region growing. In *Proceedings of International Conference on Natural Computation*, pp. 939–943, 2015. doi: 10.1109/ICNC.2015.7378117 2
- [24] T. Heimann, B. van Ginneken, M. Styner, Y. Arzhaeva, V. Aurich, C. Bauer, A. Beck, C. Becker, R. Beichel, G. Bekes, et al. Comparison and evaluation of methods for liver segmentation from CT datasets. *IEEE Transactions on Medical Imaging*, 28(8):1251–1265, 2009. doi: 10.1109/TMI.2009.2013851 6
- [25] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Proceedings of Advances in Neural Information Processing Systems*, 2020. 2, 3, 4
- [26] Y. Jin, A. Pepe, J. Li, C. Gsaxner, F. hua Zhao, K. L. Pomykala, J. Kleesiek, A. F. Frangi, and J. Egger. AI-based aortic vessel tree segmentation for cardiovascular diseases treatment: Status quo. *arXiv preprint arXiv:2108.02998*, 2021. doi: 10.48550/arXiv:2108.02998 1
- [27] M. Kahng, N. Thorat, D. H. P. Chau, F. B. Viégas, and M. Wattenberg. GAN Lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):310–320, 2019. doi: 10.1109/TVCG.2018.2864500 3
- [28] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al. Segment anything. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026, 2023. doi: 10.1109/ICCV51070.2023.00371 2
- [29] J. Kretschmer, C. Godenschwager, B. Preim, and M. Stamminger. Interactive patient-specific vascular modeling with sweep surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2828–2837, 2013. doi: 10.1109/TVCG.2013.169 3
- [30] T. P. Kuipers, P. Konduri, H. A. Marquering, and E. J. Bekkers. Generating cerebral vessel trees of acute ischemic stroke patients using conditional set-diffusion. In *Proceedings of International Conference on Medical Imaging with Deep Learning*, vol. 250, pp. 782–792, 2024. 1
- [31] K. Lawonn, S. Glaßer, A. Vilanova, B. Preim, and T. Isenberg. Occlusion-free blood flow animation with wall thickness visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):728–737, 2016. doi: 10.1109/TVCG.2015.2467961 2
- [32] R. Li, X. Li, P.-A. Heng, and C.-W. Fu. Point cloud upsampling via disentangled refinement. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 344–353, 2021. doi: 10.1109/CVPR46437.2021.00041 3
- [33] W. Lin, H. Liu, L. Gu, and Z. Gao. A geometry-constrained deformable attention network for aortic segmentation. In *Proceedings of International Conference on Medical Image Computing and Computer Assisted Interventions*, pp. 287–296, 2022. doi: 10.1007/978-3-031-16443-9_28 2
- [34] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Proceedings of ACM SIGGRAPH Conference*, pp. 163–169, 1987. doi: 10.1145/37401.37422 2, 5, 7
- [35] S. Luo and W. Hu. Diffusion probabilistic models for 3D point cloud generation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2837–2845, 2021. doi: 10.1109/CVPR46437.2021.00286 1, 2, 3, 5
- [36] J. Lyu, Y. Fu, M. Yang, Y. Xiong, Q. Duan, C. Duan, X. Wang, X. Xing, D. Zhang, J. Lin, et al. Generative adversarial network-based noncontrast CT angiography for aorta and carotid arteries. *Radiology*, 309(2):e230681, 2023. doi: 10.1148/radiol.230681 3
- [37] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang. Segment anything in medical images. *Nature Communications*, 15(1):654:1–654:9, 2024. doi: 10.1038/s41467-024-44824-z 6
- [38] M. Meuschke, T. Günther, P. Berg, R. Wickenhöfer, B. Preim, and K. Lawonn. Visual analysis of aneurysm data using statistical graphics. *IEEE*

- Transactions on Visualization and Computer Graphics*, 25(1):997–1007, 2019. doi: [10.1109/TVCG.2018.2864509](https://doi.org/10.1109/TVCG.2018.2864509) 2
- [39] M. Meuschke, U. Niemann, B. Behrendt, M. Gutberlet, B. Preim, and K. Lawonn. GUCCI - guided cardiac cohort investigation of blood flow data. *IEEE Transactions on Visualization and Computer Graphics*, 29(3):1876–1892, 2023. doi: [10.1109/TVCG.2021.3134083](https://doi.org/10.1109/TVCG.2021.3134083) 2
- [40] M. Meuschke, S. Oeltze-Jafra, O. Beuing, B. Preim, and K. Lawonn. Classification of blood flow patterns in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics*, 25(7):2404–2418, 2019. doi: [10.1109/TVCG.2018.2834923](https://doi.org/10.1109/TVCG.2018.2834923) 2
- [41] M. Meuschke, S. Voß, O. Beuing, B. Preim, and K. Lawonn. Combined visualization of vessel deformation and hemodynamics in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):761–770, 2017. doi: [10.1109/TVCG.2016.2598795](https://doi.org/10.1109/TVCG.2016.2598795) 2
- [42] G. Mistelbauer, C. Rössl, K. Bäumler, B. Preim, and D. Fleischmann. Implicit modeling of patient-specific aortic dissections with elliptic Fourier descriptors. *Computer Graphics Forum*, 40(3):423–434, 2021. doi: [10.1111/CGF.14318](https://doi.org/10.1111/CGF.14318) 1, 7
- [43] S. Mo, E. Xie, R. Chu, L. Hong, M. Nießner, and Z. Li. DiT-3D: Exploring plain diffusion transformers for 3D shape generation. In *Proceedings of Advances in Neural Information Processing Systems*, 2023. 5
- [44] S. Moccia, E. D. Momi, S. E. Hadji, and L. S. Mattos. Blood vessel segmentation algorithms - review of methods, datasets and evaluation metrics. *Computer Methods and Programs in Biomedicine*, 158:71–91, 2018. doi: [10.1016/j.cmpb.2018.02.001](https://doi.org/10.1016/j.cmpb.2018.02.001) 1
- [45] J. Montalt-Tordera, E. Pajazit, R. Jones, E. Sauvage, R. Puranik, A. A. V. Singh, C. Capelli, J. Steeden, S. Schievano, and V. Muthurangu. Automatic segmentation of the great arteries for computational hemodynamic assessment. *Journal of Cardiovascular Magnetic Resonance*, 24(1):57, 2022. doi: [10.1186/s12968-022-00891-z](https://doi.org/10.1186/s12968-022-00891-z) 2
- [46] A. Myronenko, D. Yang, Y. He, and D. Xu. Aorta segmentation from 3D CT in MICCAI SEG.A. 2023 challenge. In *Proceedings of International Conference on Medical Image Computing and Computer Assisted Interventions*, pp. 13–18, 2023. doi: [10.1007/978-3-031-53241-2_2](https://doi.org/10.1007/978-3-031-53241-2_2) 2
- [47] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of International Conference on Machine Learning*, pp. 8162–8171, 2021. 3
- [48] S. Oeltze-Jafra, J. R. Cebral, G. Janiga, and B. Preim. Cluster analysis of vortical flow in simulations of cerebral aneurysm hemodynamics. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):757–766, 2016. doi: [10.1109/TVCG.2015.2467203](https://doi.org/10.1109/TVCG.2015.2467203) 2
- [49] S. V. Patankar and D. B. Spalding. A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. In *Numerical Prediction of Flow, Heat Transfer, Turbulence and Combustion*, chap. 5, pp. 54–73. Elsevier, 1983. doi: [10.1016/B978-0-08-030937-8.50013-1](https://doi.org/10.1016/B978-0-08-030937-8.50013-1) 5
- [50] L. Piegl and W. Tiller. *The NURBS Book (Monographs in Visual Communication)*. Springer, second ed., 1996. 4
- [51] L. Radl, Y. Jin, A. Pepe, J. Li, C. Gsaxner, F.-H. Zhao, and J. Egger. AVT: Multicenter aortic vessel tree CTA dataset collection with ground truth segmentation masks. *Data in Brief*, 40:107801, 2022. doi: [10.1016/j.dib.2022.107801](https://doi.org/10.1016/j.dib.2022.107801) 5
- [52] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241, 2015. doi: [10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28) 4
- [53] Y. Rubner, C. Tomasi, and L. J. Guibas. A metric for distributions with applications to image databases. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pp. 59–66, 1998. doi: [10.1109/iccv.1998.710701](https://doi.org/10.1109/iccv.1998.710701) 6
- [54] D. Selle, B. Preim, A. Schenk, and H. Peitgen. Analysis of vasculature for liver surgical planning. *IEEE Transactions on Medical Imaging*, 21(11):1344–1357, 2002. doi: [10.1109/TMI.2002.801166](https://doi.org/10.1109/TMI.2002.801166) 3
- [55] A. M. Shaker, M. Maaz, H. A. Rasheed, S. H. Khan, M. Yang, and F. S. Khan. UNETR++: Delving into efficient and accurate 3D medical image segmentation. *IEEE Transactions on Medical Imaging*, 43(9):3377–3390, 2024. doi: [10.1109/TMI.2024.3398728](https://doi.org/10.1109/TMI.2024.3398728) 2
- [56] F. Shi, J. Gu, H. Xu, S. Xu, W. Zhang, and L. Wang. BIVDiff: A training-free framework for general-purpose video synthesis via bridging image and video diffusion models. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7393–7402, 2024. doi: [10.1109/CVPR52733.2024.00706](https://doi.org/10.1109/CVPR52733.2024.00706) 3
- [57] V. Singh, S. Jandial, A. Chopra, S. Ramesh, B. Krishnamurthy, and V. N. Balasubramanian. On conditioning the input noise for controlled image generation with diffusion models. *arXiv preprint arXiv:2205.03859*, 2022. doi: [10.48550/arXiv.2205.03859](https://doi.org/10.48550/arXiv.2205.03859) 2
- [58] A. Swillens, L. Lanoye, J. D. Backer, N. Stergiopoulos, P. R. Verdonck, F. Vermassen, and P. Segers. Effect of an abdominal aortic aneurysm on wave reflection in the aorta. *IEEE Transactions on Biomedical Engineering*, 55(5):1602–1611, 2008. doi: [10.1109/TBME.2007.913994](https://doi.org/10.1109/TBME.2007.913994) 1
- [59] J. Tao, X. Huang, F. Qiu, C. Wang, J. Jiang, C. Shene, Y. Zhao, and D. Yu. VesselMap: A web interface to explore multivariate vascular data. *Computers & Graphics*, 59:79–92, 2016. doi: [10.1016/J.CAG.2016.05.024](https://doi.org/10.1016/J.CAG.2016.05.024) 2
- [60] A. Updegrove, N. M. Wilson, J. Merkow, H. Lan, A. L. Marsden, and S. C. Shadden. SimVascular: An open source pipeline for cardiovascular simulation. *Annals of Biomedical Engineering*, 45(3):525–541, 2017. 1
- [61] T. P. Vagenas, K. Georgas, and G. K. Matsopoulos. Deep learning-based segmentation and mesh reconstruction of the aortic vessel tree from CTA images. In *Proceedings of International Conference on Medical Image Computing and Computer Assisted Interventions*, pp. 80–94, 2023. doi: [10.1007/978-3-031-53241-2_7](https://doi.org/10.1007/978-3-031-53241-2_7) 2
- [62] S. Valette, J. Chassery, and R. Prost. Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):369–381, 2008. doi: [10.1109/TVCG.2007.70430](https://doi.org/10.1109/TVCG.2007.70430) 5
- [63] J. Wasserthal, H.-C. Breit, M. T. Meyer, M. Pradella, D. Hinck, A. W. Sauter, T. Heye, D. T. Boll, J. Cyriac, S. Yang, et al. TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images. *Radiology: Artificial Intelligence*, 5(5):e230024, 2023. doi: [10.1148/ryai.230024](https://doi.org/10.1148/ryai.230024) 6
- [64] N. M. Wilson, A. K. Ortiz, and A. B. Johnson. The vascular model repository: A public resource of medical imaging data and blood flow simulation results. *Journal of Medical Devices*, 7(4):0409231, 2013. doi: [10.1115/1.4025983](https://doi.org/10.1115/1.4025983) 5
- [65] J. M. Wolterink, T. Leiner, and I. Išgum. Blood vessel geometry synthesis using generative adversarial networks. *arXiv preprint arXiv:1804.04381*, 2018. doi: [10.48550/arXiv:1804.04381](https://doi.org/10.48550/arXiv:1804.04381) 1
- [66] H. E. Wong, M. Rakic, J. V. Guttag, and A. V. Dalca. ScribblePrompt: Fast and flexible interactive segmentation for any biomedical image. In *Proceedings of European Conference on Computer Vision*, pp. 207–229, 2024. doi: [10.1007/978-3-031-73661-2_12](https://doi.org/10.1007/978-3-031-73661-2_12) 4
- [67] L. Wu, D. Wang, C. Gong, X. Liu, Y. Xiong, R. Ranjan, R. Krishnamoorthi, V. Chandra, and Q. Liu. Fast point cloud generation with straight flows. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9445–9454, 2023. doi: [10.1109/CVPR52729.2023.00911](https://doi.org/10.1109/CVPR52729.2023.00911) 2, 3
- [68] X. Xiong, Y. Ding, C. Sun, Z. Zhang, X. Guan, T. Zhang, H. Chen, H. Liu, Z. Cheng, L. Zhao, et al. A cascaded multi-task generative framework for detecting aortic dissection on 3-D non-contrast-enhanced computed tomography. *IEEE Journal of Biomedical and Health Informatics*, 26(10):5177–5188, 2022. doi: [10.1109/JBHI.2022.3190293](https://doi.org/10.1109/JBHI.2022.3190293) 3
- [69] M. Ye, D. Yang, M. Kanski, L. Axel, and D. Metaxas. Neural deformable models for 3D bi-ventricular heart shape reconstruction and modeling from 2D sparse cardiac magnetic resonance imaging. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pp. 14247–14256, 2023. doi: [10.1109/ICCV51070.2023.01310](https://doi.org/10.1109/ICCV51070.2023.01310) 1
- [70] Y. Zeng, V. M. Patel, H. Wang, X. Huang, T.-C. Wang, M.-Y. Liu, and Y. Balaji. JeDi: Joint-image diffusion models for finetuning-free personalized text-to-image generation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6786–6795, 2024. doi: [10.1109/CVPR52733.2024.00648](https://doi.org/10.1109/CVPR52733.2024.00648) 3
- [71] B. Zhang, Z. Lai, S. Liu, X. Xie, X. Zhou, Z. Hou, X. Ma, B. Liu, K. Li, and M. Song. SLDU-Net: A similarity-based dynamic linking network for the automated segmentation of abdominal aorta aneurysms and branching vessels. *Biomedical Signal Processing and Control*, 100:106991, 2025. doi: [10.1016/j.bspc.2024.106991](https://doi.org/10.1016/j.bspc.2024.106991) 1
- [72] L. Zhou, Y. Du, and J. Wu. 3D shape generation and completion through point-voxel diffusion. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pp. 5806–5815, 2021. doi: [10.1109/ICCV48922.2021.00577](https://doi.org/10.1109/ICCV48922.2021.00577) 1, 3, 5
- [73] Y. Zhou, H. Chang, X. Lu, and Y. Lu. DenseUNet: Improved image classification method using standard convolution and dense transposed convolution. *Knowledge-Based Systems*, 254:109658, 2022. doi: [10.1016/j.knsys.2022.109658](https://doi.org/10.1016/j.knsys.2022.109658) 2