

1 SeqSeg: Learning Local Segments for Automatic  
2 Vascular Model Construction

3 Numi Sveinsson Cepero<sup>1</sup> and Shawn C. Shadden<sup>1</sup>

4 <sup>1</sup>Department of Mechanical Engineering, University of California,  
5 Berkeley, Berkeley, 94720, CA, USA.

6 Contributing authors: [numi@berkeley.edu](mailto:numi@berkeley.edu); [shadden@berkeley.edu](mailto:shadden@berkeley.edu);

7 **Abstract**

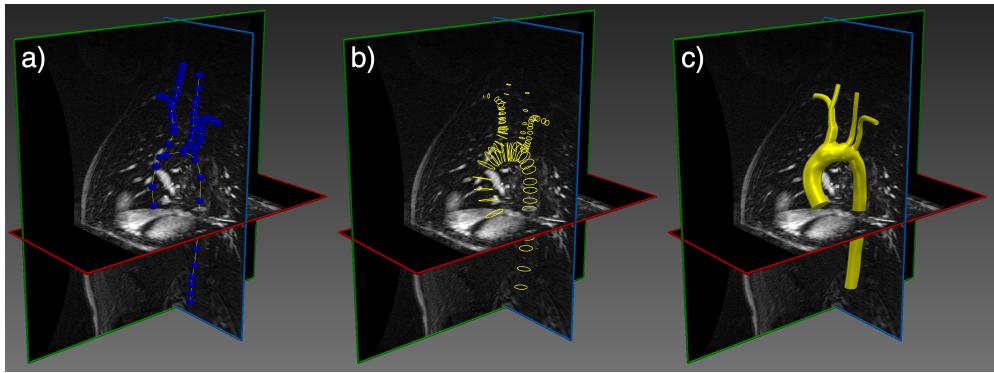
8 Computational modeling of cardiovascular function has become a critical part of  
9 diagnosing, treating and understanding cardiovascular disease. Most strategies  
10 involve constructing anatomically accurate computer models of cardiovascular  
11 structures, which is a multistep, time-consuming process. To improve the model  
12 generation process, we herein present SeqSeg (sequential segmentation): a novel  
13 deep learning based automatic tracing and segmentation algorithm for construct-  
14 ing image-based vascular models. SeqSeg leverages local U-Net-based inference to  
15 sequentially segment vascular structures from medical image volumes. We tested  
16 SeqSeg on CT and MR images of aortic and aortofemoral models and compared  
17 the predictions to those of benchmark 2D and 3D global nnU-Net models, which  
18 have previously shown excellent accuracy for medical image segmentation. We  
19 demonstrate that SeqSeg is able to segment more complete vasculature and is  
20 able to generalize to vascular structures not annotated in the training data.

21 **Keywords:** Vascular Model Construction, Medical Image Segmentation, Blood Vessel  
22 Tracking, Convolutional Neural Network, Deep Learning, Cardiovascular Simulation

23 **1 Introduction**

24 Image-based vascular modeling is used for a variety of purposes including diagnosis,  
25 personalized treatment planning and fundamental understanding of disease progres-  
26 sion [1–4]. Specialized software has been developed for such modeling, including  
27 SimVascular [5, 6], CRIMSON [7] and VMTK [8]. This modeling paradigm uses  
28 medical imaging, such as computed tomography (CT) or magnetic resonance (MR)  
29 angiography, to construct a patient-specific anatomical model of vessels of interest.

30 This geometric model is subsequently converted into a 3D computational mesh to  
 31 support detailed blood flow and/or tissue mechanics simulation and analysis. The con-  
 32 struction of an anatomical model from medical image data remains largely a manual  
 33 process [5]. Figure 1 shows a typical workflow for vascular model construction, start-  
 34 ing with the creation of centerlines along the vessels of interest, 2D segmentation of  
 35 the vessel lumen along the centerlines, and lofting of the 2D segmentations to gener-  
 36 ate a unified 3D model of the vascular geometry. Alternative segmentation approaches  
 37 exist, including region-growing or level-set methods [6]; however, these methods gener-  
 38 ally struggle in the segmentation of highly-branched structures such as blood vessels,  
 39 particularly in the context of limited image resolution, unclear boundaries and image  
 40 artifacts [9]. Additionally, when the model is constructed manually, substantial user  
 41 bias may result. Ultimately, despite the popularity and maturity of image-based cardio-  
 42 vascular modeling over the past 20 years, the process of deriving a simulation-suitable  
 43 anatomical model from medical image data has remained a primary bottleneck for  
 44 large-cohort studies or translational applications where timely results are needed.



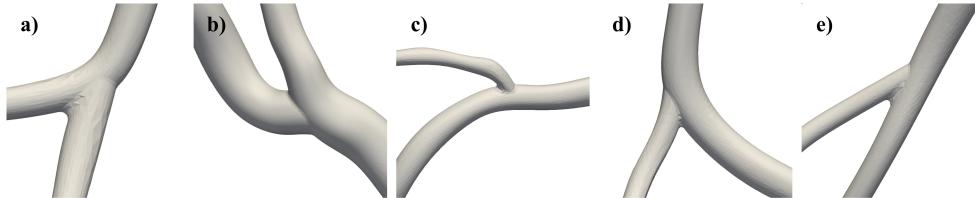
**Fig. 1** A typical vascular model construction workflow involves (a) creating vessel paths by manual selection of point (b) sequential segmentation of the vessel lumen boundary at discrete cross-sections along the paths and (c) lofting these segmentation rings into a unified model. This process is described in more detail in [5].

45 Recently, machine learning has been applied to automate and speed up image  
 46 segmentation. Note that while medical image segmentation is performed for a  
 47 variety of healthcare applications, we focus here on the purpose of generating a  
 48 *simulation-suitable* model that can be utilized to generate a computational domain for  
 49 physics-based simulation. Simulation suitable models have certain criteria that must  
 50 be met such as, being connected, sufficiently “smooth”, and able to be meshed (dis-  
 51 cretized) with quality elements. Most learning methods focus on pixel classification,  
 52 which often results in segmentations that are disconnected or have substantial arti-  
 53 facts that complicate, or prevent, generation of a mesh suitable to support numerical  
 54 simulation.

55 Most progress has been made when machine learning has been applied to isolated  
 56 anatomic vascular regions [10, 11] including for cardiac models [12, 13]. In the work  
 57 of Maher, et al. [14, 15] segmentation of branched vascular domains was achieved by

58 assuming the existence of vessel centerlines (cf. Fig. 1a). Under such assumptions,  
 59 these centerlines are traversed and local 2D cross-sectional segmentations of the lumen  
 60 boundary are generated using a trained network. This framework essentially auto-  
 61 mated step (b) shown in Fig. 1. However, for many vascular models, the generation  
 62 of vessel centerlines is the most labor intensive step. Moreover, with this approach,  
 63 segmentation is only performed at discrete 2D slices along the vessel, which provides  
 64 incomplete sampling and can be problematic when the cross-section is not connected  
 65 or the centerline is not sufficiently aligned with the vessel. And more importantly, dis-  
 66 crete cross-sectional segmentation performs poorly at vessel bifurcations, which are  
 67 present in almost all applications.

68 Herein, we present a novel method for segmenting branched vascular geometries  
 69 from medical image data utilizing local deep learning-based segmentation that does  
 70 not require *a priori* centerline information. This approach starts from a seed point and  
 71 generates a local 3D segmentation of the vessel(s) containing the seed point over a local  
 72 subvolume. Based on this local segmentation, we determine the orientation of the vessel  
 73 and any locally connected branches. We then step the subvolume along the determined  
 74 vessel direction (and new subvolumes along the identified local branch directions) to  
 75 generate a 3D segmentation of the neighboring segment(s). This approach is motivated  
 76 by the following considerations: when viewed locally by a subvolume that is centered on  
 77 a vessel and slightly larger than the vessel diameter, vessels of different sizes and from  
 78 different regions exhibit substantial geometric similarity (Figure 2), and consequently  
 79 learning to locally segment a portion of a vessel should be easier than learning to  
 80 segment an entire vascular network. While cropping of medical image volumes has  
 81 been performed previously, for example, for coronary tracking [16] [17], to the best of  
 82 our knowledge such approach has not been used to generate 3D segmentation or for  
 83 segmentation of general vascular geometries.



**Fig. 2** When viewed locally, vasculature of different sizes and anatomical regions exhibit substantial geometric similarity. a) the pulmonary artery ( $r = 1.5\text{mm}$ ), b) the brachiocephalic artery ( $r = 9\text{mm}$ ), c) the coronary artery ( $r = 1\text{mm}$ ), d) the cerebral artery ( $r = 2\text{mm}$ ) and e) the femoral artery ( $r = 3\text{mm}$ ) are presented

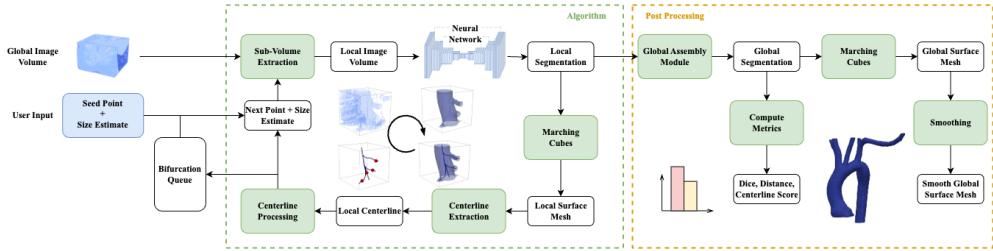
84 By processing local subvolumes, we simplify the deep learning task and introduce  
 85 beneficial inductive bias to the machine learning model, allowing it to generalize to  
 86 vasculature not present in training data. We test this method on a dataset of publicly-  
 87 accessible aortic and aortofemoral models, and the results are compared to benchmark  
 88 global 2D/3D nnU-Net neural network models that have previously shown excellent  
 89 results for medical image segmentation. The main contribution of this work is a new  
 90 method capable of:

- Tracing vasculature after initialization with a single point and vessel radius estimate.
- Segmenting vasculature while ensuring global connectivity to maintain physiologic topology.
- Detecting bifurcations, storing them and tracing them sequentially.
- Delivering a global surface mesh of segmented vasculature.
- Generalizing to segment parts of vasculature not annotated in training data.

## 2 Method

### 2.1 Algorithm

Figure 3 shows a schematic of the algorithm. Briefly, a “seed point”, (specifying a location and direction) and a rough diameter “size estimate” of the vessel containing the seed point are supplied by the user. A local subvolume surrounding the seed point is extracted from the global image volume. The vessel portion contained in the subvolume is segmented using a neural network. The segmentation is postprocessed and converted to a surface mesh, after which a centerline is extracted. The resulting centerline is subsequently used to choose the next subvolume location and size. These steps are explained in further detail below.



**Fig. 3** Overview of the tracing and segmentation algorithm with inputs of the global raw image and seed points for initialization. The algorithm takes steps, stores bifurcations in the queue during tracing, and outputs a global segmentation map for post processing

### 2.2 Segmentation

#### 2.2.1 Dataset, Sampling and Augmentations

To train the U-Net and test the algorithm, we utilized a dataset of 41 CT and 44 contrast enhanced MRI aortic and aortofemoral cases, which is commensurate with the amount of data typically provided in segmentation challenges. The breakdown of how many cases are used for training, validation and testing is specified in Table 1. The VMR datasets used for training are accessible from the open access Vascular Model Repository at <https://vascularmodel.com>. For further testing on CT images, we also use a subset of the AVT dataset [18], specifically the dataset obtained from Dongyang Hospital. Table 1 shows details on the datasets; modalities, purpose, training/test split,

**Table 1** The datasets of patients used for model training and method testing. Abbreviation are as follows: Datasets; VMR: Vascular Model Repository, AVT-D: Aortic Vessel Tree dataset, subset from Dongyang Hospital. Anatomy; AO:Aorta, AF:Aortofemoral. Disease; H:Healthy, AAA:Abdominal Aortic Aneurysm, MA: Marfan Syndrome, CA:Coarctation of Aorta, AOD:Aortoiliac Occlusive Disease, SVD:Single Ventricle Defect. Sex; M:Male, F:Female, U:Unknown. Sex and age information was not available for the AVT dataset.

Dataset	Modality	Purpose	Train/Test	Anatomy	Disease	Sex	Age(yr)
VMR	CT	Train/Test	33/8	25 AO, 16 AF	23 H, 15 AAA, 3 MA	23M 6F 12U	6 - 80 ave: 58
VMR	MR	Train/Test	37/7	38 AO, 6 AF	19 H, 14 CA, 5 SVD, 2 MA 4 AOD	30M 14F	0.6 - 67 ave: 17
AVT -D[18]	CT	Test	0/18	18 AO	18 H	-	-

117 anatomies, diseases (if present), sex ratio and age ranges. The datasets contain a 3D  
118 image volume and a respective “ground truth” vascular segmentation map (converted  
119 from segmentation surface models for the VMR data) and corresponding centerlines  
120 that served as ground truth labels for training and testing.

121 To generate training data for the local segmentation U-Net, the global 3D medical  
122 image volumes in the VMR training datasets were sampled along the centerlines and  
123 these subvolumes (i.e. “Patches”) were stored. Namely, two volumes were extracted  
124 at each Patch: 1) the original medical image data and 2) a binary segmentation of the  
125 subvolume based on the model representing the ground truth label.

126 To improve the learning process, we varied the samples in terms of centering and  
127 size. Briefly, some samples were centered along the centerline while others were shifted  
128 from the centerline, and the subvolume sizes varied from just capturing the lumen of  
129 the vessel to including more surrounding tissue. More specifically, each sample volume  
130  $s_i$  is a cube dependent on its side length and center, i.e.,  $s_i(L_i, \mathbf{c}_i)$  where  $L_i$  is its side  
131 length and  $\mathbf{c}_i$  is the center point of sample  $i$ . The side length and center are sampled  
132 as follows:

$$\begin{aligned} L_i &= R_i * \alpha_i \\ \mathbf{c}_i &= \mathbf{C}_i + \beta_i * R_i * \mathbf{w}_i \\ \alpha &\sim \mathcal{N}(\mu_r, \sigma_r^2) \\ \beta &\sim \mathcal{N}(\mu_s, \sigma_s^2) \end{aligned} \tag{1}$$

133 where  $R_i$  is the local radius of the vessel,  $\mathbf{C}_i$  is the point on the centerline,  $\mathbf{w}$  is a  
134 unit vector perpendicular to the centerline and  $\alpha, \beta$  represent the radius ratios used to  
135 enlarge or shift the sample.  $\mathbf{w}$  was chosen by sampling a random linear combination  
136 of orthogonal unit vectors  $\mathbf{u}, \mathbf{v}$  that defined a plane perpendicular to the centerline:

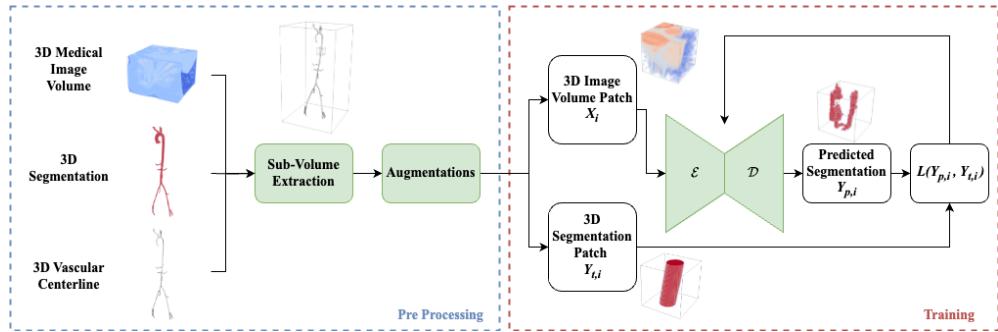
$$\mathbf{w}_i = \frac{a_i * \mathbf{u} + b_i * \mathbf{v}}{\|a_i * \mathbf{u} + b_i * \mathbf{v}\|}, \quad a, b \sim U[-1, 1] \tag{2}$$

137 where  $a, b$  are scalars sampled from a uniform distribution between  $[-1, 1]$ . We used  
 138  $(\mu_r, \sigma_r^2) = (5, 1)$  and  $(\mu_s, \sigma_s^2) = (0, 0.8)$  so that the mean sample was 5 times the size  
 139 of the radius and centered on the centerline. This stochasticity was purposefully added  
 140 to represent the variance that the tracing algorithm encounters during inference and  
 141 is intended to increase the robustness of the neural network. This process resulted in  
 142 a training dataset  $D$  consisting of  $N$  pairs of image subvolumes  $X_i$  and corresponding  
 143 blood vessel segmentations  $Y_{t,i}$ :

$$D = \{(X_1, Y_{t,1}), (X_2, Y_{t,2}), \dots, (X_N, Y_{t,N})\} \quad (3)$$

144 In total, we get  $D = 36289$  patches for CT data and  $D = 33603$  patches for MR data.  
 145 For the VMR dataset 15 patient datasets were excluded for final testing, i.e., 8 CT and  
 146 7 MR cases were not sampled for training or validation. The generation of subvolume  
 147 data for network training is shown schematically in Fig. 4.

148 Before training, MR image volumes were normalized via z-scoring, where each  
 149 voxel value,  $x$ , is subtracted from the image mean  $\mu$  and then divided by the image  
 150 standard deviation  $\sigma$ . CT volumes were clipped and z scored according to foreground  
 151 image values where  $\mu$  and  $\sigma$  are calculated only from voxels labelled as vessel in the  
 152 ground truth training data and held constant during inference [19]; see Table 2 for  
 153 details. The preferred image spacing was chosen as the median spacing across all cases  
 154 and all image volumes were resampled using a 3rd order spline. Segmentation maps  
 155 were resampled differently, using linear splines on one-hot encoded maps, and then  
 156 the argmax of the result.



**Fig. 4** Preprocessing involves extracting subvolumes along ground truth centerlines and data augmentation prior to neural network training. Thousands of samples are acquired from only a few dozen models. The neural network consists of an encoder  $\mathcal{E}$  followed by a decoder  $\mathcal{D}$ , which outputs the predicted segmentation map used to compute loss,  $L$ , during training

### 157 2.2.2 Neural Network Architecture and Training

158 The convolutional neural network (CNN) U-Net architecture was used for segmenta-  
 159 tion. The U-Net is tailored to the processing of medical images, going from the original  
 160 resolution down to a low-resolution, high-dimensional space and then back up to the

161 original resolution. Herein, a 3D version of U-Net was chosen for the SeqSeg segmen-  
 162 tation step. Note, we also compare our end results to those of two benchmark U-Net  
 163 models, i.e. a 2D U-Net and a 3D U-Net, trained on the global image volumes, see  
 164 Section 2.9 for further details. The U-Net learns features primarily through two differ-  
 165 ent mechanisms. First, by downsampling the original image data the model is forced to  
 166 retain only important global information when squeezed through a lower-dimensional  
 167 space. Second, by using skip connections across the neural network the model is able  
 168 to retain features related to finer details from the higher resolution image in its final  
 169 prediction. The skip connections are concatenations of blocks of the same resolution.  
 170 The final output is a pixelwise probability map indicating the likelihood that a pixel  
 171 corresponds to a target tissue.

172 Our U-Net was trained for binary classification: to predict whether voxel  $y_{i,j,k}$   
 173 inside segmentation mask  $Y_p \in \mathbf{R}^{W \times H \times D}$  belongs to a blood vessel:

$$Y_p = \{y_{i,j,k} \in [0, 1] \mid 0 \leq i < W; 0 \leq j < H; 0 \leq k < D\} \quad (4)$$

$$y_{i,j,k} = \begin{cases} 1 & \text{if belongs to vessel} \\ 0 & \text{otherwise} \end{cases}$$

174 where  $i, j, k$  refers to the index of a voxel in an image of width  $W$ , height  $H$  and depth  
 175  $D$ . In this section, lower case notation refers to individual nodes or voxels, e.g.  $y$ ,  
 176 whereas capital notation, e.g.  $Y$ , refers to a set of nodes or voxels such as composing  
 177 an image, segmentation mask or output from neural network layers.

178 In mathematical terms, the neural network is a parameterized function  $f_\theta$  that  
 179 transforms a raw image input  $X \in \mathbf{R}^{W \times H \times D}$  into a blood vessel segmentation map:

$$Y_p = f(X \mid \theta) \quad (5)$$

180 where  $\theta$  are the parameters of the neural network, which are optimized using training  
 181 data. The final output,  $Y_p$ , ranges between  $[0, 1]$  and can thus be interpreted as a  
 182 probability map of whether each voxel belongs to a blood vessel. This enables the  
 183 volume to be binarized by thresholding to a particular probability value.

184 We utilized the nnU-Net framework for hyperparameter specification and training  
 185 [19]. The framework automatically determines parameters such as image resampling  
 186 spacing, patch size and batch size based on training data and GPU memory size. The  
 187 underlying neural network architecture used is the U-Net, with additional constraints  
 188 on specific parameters. Table 2 lists the specifications of our implemented U-Net model  
 189 architectures and training parameters. Since the SeqSeg model is trained on smaller  
 190 volumes compared to the benchmarks, its required batch size can be larger, see Table  
 191 2. The nnU-Net framework utilizes stochastic gradient descent with Nesterov momen-  
 192 tum with an initial learning rate of 0.01 accompanied by a learning rate scheduler  
 193 of  $(1 - epoch/epoch_{max})^{0.9}$ , where  $epoch_{max} = 1000$  was chosen; see [19] for further  
 194 details. Training was performed using an NVIDIA Geforce RTX 2080ti GPU (11 GB  
 195 GPU memory) on the Savio High Performance Computing cluster at the University  
 196 of California, Berkeley.

**Table 2** The U-Net architecture and training specifications, for both the SeqSeg models and global benchmark models

Parameter	SegSeg CT	SeqSeg MR	3D CT	3D MR	2D CT	2D MR
Intensity Normalization	0.5/99.5% clip + foreground z-score	all image z-score	0.5/99.5% clip + foreground z-score	all image z-score	0.5/99.5% clip + foreground z-score	all image z-score
Image Target Spacing	0.200, 0.0547, 0.0547	0.0859, 0.0625, 0.0625	0.0800, 0.0488, 0.0488	0.0859, 0.0586, 0.0586	0.0488, 0.0488	0.0586, 0.0586
Patch Size	[20,80,80]	[40,48,48]	[96,160,160]	[56, 256, 160]	[512, 512]	[512,384]
Batch Size	33	57	2	2	12	16
Max Nr. Features	320	320	320	320	512	512
Nr. Stages Encoder	5	4	6	6	8	7
Nr. Stages Decoder	4	3	5	5	7	6
Nr. Layers per Stage	2	2	2	2	2	2
Nr. Pooling Ops. per Axis	[2,4,4]	[3,3,3]	[3,5,5]	[4,5,5]	[7,7]	[6,6]
Conv. Kernel Size	[3,3,3]	[3,3,3]	[3,3,3]	[3,3,3]	[3,3]	[3,3]

### 2.2.3 Loss Function

The loss function was a combination of Dice score ( $\mathcal{D}$ ) and binary cross-entropy ( $\mathcal{CE}$ ):

$$\mathcal{D}(Y_p, Y_t) = \frac{2 \cdot \|Y_p \cap Y_t\|}{\|Y_p\| + \|Y_t\|} \quad (6)$$

$$\mathcal{CE}(Y_p, Y_t) = \frac{1}{n} \sum_{y \in Y} (y_t \cdot \log y_p + (1 - y_t) \cdot \log (1 - y_p)) \quad (7)$$

where  $Y_p$  and  $Y_t$  are respective prediction and ground truth segmentation masks, respectively, and  $n$  is the total number of voxels.  $Y_t$  is defined similarly to  $Y_p$  in Eq. (4). Binary cross entropy is a common loss function for binary classification and we added Dice loss to regulate it for medical image segmentation. Namely, the Dice score helps counter the class imbalance that pixelwise classification problems face in medical image segmentation. This is critical when working with 3D images where the number of voxels belonging to a blood vessel is a small percentage of the total voxels

207 in the volume. It follows that our loss function is defined as

$$\mathcal{L} = \sum_i^{N_b} (1 - \mathcal{D}(Y_{p,i}, Y_{t,i}) - \mathcal{CE}(Y_{p,i}, Y_{t,i})) \quad (8)$$

208 for a batch size  $N_b$ , where each batch is a subset of the total dataset  $N_b < N$  described  
209 in Eq. (3). The data is batched to fit into GPU memory as described in Table 2. Each  
210 image in the batch is processed in parallel on a GPU and the loss is accumulated  
211 before taking a gradient step to update the model parameters.

## 212 2.3 Surface, Centerline Calculations and Step Taking

213 As mentioned above, the output of the U-Net is a binarized image subvolume. The  
214 marching cubes algorithm [20] can be applied to this binarized image subvolume to  
215 generate a local surface mesh of the vessel segment. The resulting surface was cut using  
216 the image subvolume boundary planes, which results in truncation boundaries for the  
217 vessel(s), i.e., “inlets” or “outlets”. One of these truncation boundaries is identified  
218 as the source (inlet) and others are identified as targets (outlets). This process was  
219 performed automatically using information from previous steps and from the direction  
220 of tracing. To do this, the centers of the truncation boundaries are calculated. The  
221 truncation boundary center closest to the previous stepping point is chosen as source  
222 and the rest as targets.

223 The surface mesh, with respective outlet labels, is used to automatically generate  
224 centerline(s) and radius estimates of the local vessel segment using a levelset based  
225 centerline extraction method. The method calculates centerline(s) as the path(s) that  
226 follow a wave propagation starting from a seed point [21]. The wave propagation is  
227 modeled by equation:

$$|\nabla T(x)|F(x) = 1 \quad (9)$$

228 where  $T(x)$ , the time it takes for wave to reach point  $x$ , is solved using a set “speed”  
229 function  $F(x)$ .  $F(x)$  is set to have values proportional to distance from vessel boundary,  
230 leading to higher value towards the center and lower closer to vessel walls. When  
231 Eq. 9 is solved with  $T(x_0) = 0$  at source point  $x_0$ , we obtain a solution with wave  
232 propagation faster in the center of vessels. Then, using that solution, we perform  
233 gradient descent starting from target point(s), where  $T(x)$  is high, until we reach the  
234 source point, where  $T(x)$  is low, and have therefore defined a centerline path(s). Since  
235 the “speed” function had higher values towards the center then so do the values of  
236  $\nabla T$  which forces the gradient descent towards the center of the vessel while tracing  
237 back, see [21] for details. Furthermore, we estimate the radius of the vessel at each  
238 point along the centerline by its distance to the surface.

239 The centerline extraction depends on well-defined outlet centers fed as seed points.  
240 Our method defines these outlet centers automatically, as described above. In the case  
241 of a bifurcation, a single outlet was labeled as the source based on the previous step  
242 and the direction of tracing. The point(s) to move to along the computed centerline(s)  
243 is chosen at 80% along each branch, see stepping point choice in Fig. 5. We found that  
244 this allowed for smaller step size, ensuring more accurate bifurcation detection.

245 **2.4 Choice of Subvolume Size and Chances**

246 The local centerlines calculated as described above contain 1D (lines) meshes in 3D  
 247 space connecting all outlets with radius information along them, see Figure 5. This  
 248 radius estimate is subsequently used to determine the size of the next subvolume to  
 249 extract. The length of the next subvolume is chosen as five times the radius estimated,  
 250 consistent with the size of the training samples as described in Section 2.2.1. Further-  
 251 more, to prevent underestimation of subvolume size, we let radius estimate carry on  
 252 from one step to the next. The subvolume sidelength  $L$  is calculated by an average  
 253 of the current radius estimate,  $r_i$ , and the estimate from the previous step,  $r_{i-1}$ :

$$L = 5 * (r_i + r_{i-1})/2 \quad (10)$$

254 Additionally, we use the segmentation prediction itself as an indicator of subvolume  
 255 size to vessel size ratio. If a high percentage of voxels within an image subvolume is  
 256 predicted as belonging to vessel, that may indicate a small subvolume to vessel ratio,  
 257 i.e., that the vessel occupies a large part of the volume. Utilizing this, we defined  
 258 a cutoff percentage,  $\gamma^*$ , for which if the percentage exceeds it then we enlarge the  
 259 subvolume size until it drops below, see Algorithm 1, where  $X$  is the subvolume, a  
 260 function of sidelength  $L$ .

---

**Algorithm 1** An algorithm to enlarge the subvolume based on the percentage of voxels predicted as vessel,  $\gamma$ .

---

```

 $R \leftarrow (r_i + r_{i-1})/2$ 
 $L_0 \leftarrow 5 * R$ 
 $Y_p \leftarrow f(X(L_0))$ 
 $\gamma \leftarrow \frac{1}{W*H*D} \sum_{y \in Y_p} y$ 
while  $\gamma \geq \gamma^*$  do
     $L \leftarrow 5 * R * 1.1$   $\triangleright$  Increase sidelength by 10%
     $Y_p \leftarrow f(X(L))$ 
     $\gamma \leftarrow \frac{1}{W*H*D} \sum_{y \in Y_p} y$ 
    if  $\frac{L}{L_0} > 1.3$  then
        break  $\triangleright$  Maximum increase of 30%
    end if
end while

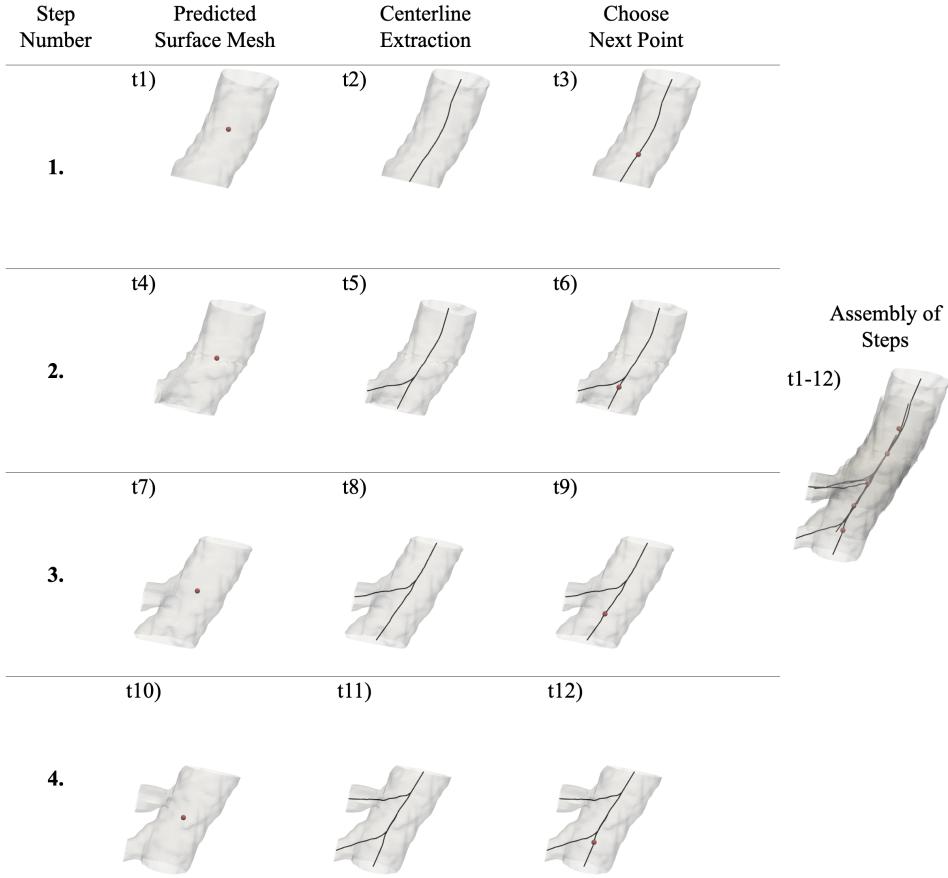
```

---

261 When SeqSeg encounters local subvolumes with image artifacts or unclear vessel  
 262 boundaries, the neural network model sometimes produces inaccurate segmentations.  
 263 However, in many cases, these inaccuracies are bound to those specific locations in  
 264 the image volume whereas the following downstream vasculature may be clearer and  
 265 easier to segment. To handle these situations, we implemented a “chances” feature  
 266 to SeqSeg. When SeqSeg fails to detect  $\geq 2$  outlets or fails to successfully compute  
 267 a centerline, we give the step another chance and move further in the same direction  
 268 and try again. Given a point  $\mathbf{p}_i$  with a corresponding vessel tangent  $\mathbf{t}_i$  and radius  $R_i$ ,  
 269 the next “chance” location  $\mathbf{p}_{i+1}$  is calculated as:

$$\mathbf{p}_{i+1} = \mathbf{p}_i + R * \mathbf{t}_i \quad (11)$$

270 We set a maximum number of chances to three. This allows SeqSeg to better move  
 271 past difficult regions of the image.

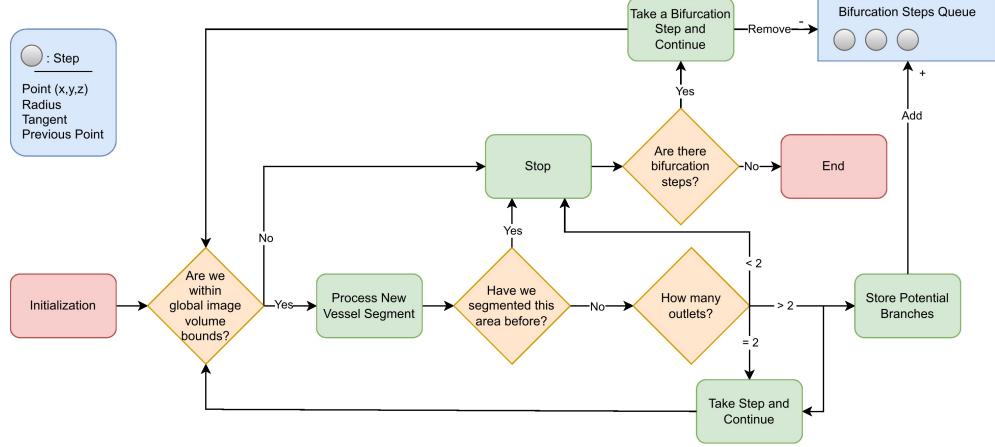


**Fig. 5** Automatic tracing using local surface mesh predictions for 3 steps, involving 12 calculation time steps. Centerlines are extracted and the next points are chosen to move to. These steps are subsequently assembled together to form the global vasculature model

## 2.5 Bifurcations and Retracing Prevention

272 Bifurcations are detected by counting the branches of the centerline successfully com-  
 273 puted. When bifurcations are detected, they are stored and returned to once other  
 274 branches have been traced. Namely, the largest radius outlet was chosen for continued  
 275 tracing while the others were saved as bifurcation points and were revisited once the  
 276 current vessel had been fully traced. These bifurcation points were periodically sorted  
 277

278 by radius to ensure prioritization of the largest vessels first, similar to how a human  
 279 would interrogate the vasculature; see Figure 6.



**Fig. 6** How the algorithm takes steps and handles bifurcations, starting from an initial seed point. The bifurcation points are stored in queue for subsequent tracing

280 Since the method detects outlets locally, it can sometimes detect the same bifurcations  
 281 multiple times. This occurs especially if a small step size is used to advance  
 282 the subvolume. To save computational time, we implemented a retracing prevention  
 283 technique that periodically checks the global segmentation to determine whether the  
 284 algorithm has segmented the current region before. We added buffers to the global  
 285 assembly module to ensure that these checks only applied to segmentations involving  
 286 past branches and not the current one.

## 287 2.6 Initialization

288 As mentioned above, the SegSeg method requires a seed point accompanied by a size  
 289 estimate and direction for initialization. For evaluation purposes, this seed point is  
 290 chosen at the ‘start’ of each vascular model, in the largest artery closest to the heart,  
 291 similar to how a user would define it.

## 292 2.7 Stop Criteria

293 Since SeqSeg is an automated tracing method, stop criteria are required to terminate  
 294 step-taking. The current framework has no explicit stop criteria; however, indirectly,  
 295 it stops when either of two scenarios occur:

- 296 1. The method reaches the global image boundary, thereby requesting image data  
 297 that does not exist.
- 298 2. The requested subvolume is of low resolution (e.g., resulting from tracing a small  
 299 vessel) or is of too low quality (e.g., blurry or has image artifacts) that results in

300 segmentation failure, or a centerline extraction failure from the resulting low-quality  
 301 segmentation.

302 In addition to the “forced” stop criteria mentioned above, we have implemented  
 303 *optional* criteria as well, that can be toggled if premature stopping is desired. The  
 304 optional stop criteria are as follows:

- 305 1. Define  $N_{max}$ , maximum number of steps taken and stop once number of steps taken  
 306 has exceeded  $N_{max}$ . We set  $N_{max} = 500$ .
- 307 2. Define  $R_{min}$ , minimum allowed radius, and stop tracing down a branch once radius  
 308 estimate is under  $R_{min}$ . We set  $R_{min} = 0.5\text{mm}$ .
- 309 3. Define  $NB_{max}$ , maximum number of branches to be traced down, and stop once  
 310 number of branches exceed  $NB_{max}$ .  $NB_{max}$  can, for example, be set as 15.

### 311 2.8 Global Assembly

312 Sampling subvolumes may overlap a given voxel multiple times depending on the step  
 313 size used to propagate the subvolume along the identified vessel (cf. superposition of  
 314 subvolumes on the right side of Figure 5), or because of a new subvolume introduced  
 315 to trace a bifurcation. The end result is that several predictions may exist for a given  
 316 voxel. Thus, all local segmentations are gathered globally by calculating a weighted  
 317 mean prediction for each voxel. During development, we noticed that segmentations  
 318 tend to be less accurate close to the subvolume boundary, so we added Gaussian  
 319 weighting giving more weight to voxel predictions closer to subvolume centers. This is  
 320 also consistent with the benchmark nnU-Net method [19]. For each subvolume predic-  
 321 tion of sidelength  $L$  and center point  $c$ , we define a weight map  $W$ , where each voxel  
 322 with indices  $i, j, k$  and point location  $p_{i,j,k}$  gets a weight  $w_{i,j,k}$  according to:

$$w_{i,j,k} = e^{-\frac{\|p_{i,j,k} - c\|_2}{2\sigma^2}}, \sigma = \frac{1}{4}L \quad (12)$$

323 Therefore, voxels close to  $c$  get a weight close to 1 while voxels close to the borders of  
 324 the subvolume get a weight value of approximately two standard deviations from the  
 325 mean, 0.02.

326 Given a set  $\mathcal{S}$ , of size  $N_{\mathcal{S}}$ , of segmentations in which local volumes included voxel  
 327  $y_{i,j,k}$ , the final global segmentation value was given by

$$y_{i,j,k} = \frac{1}{N_{\mathcal{S}} \sum_{s \in \mathcal{S}} w_{i,j,k}^s} \sum_{s \in \mathcal{S}} w_{i,j,k}^s \cdot s_{i,j,k} \quad (13)$$

328 where  $i, j, k$  refer to global voxel indices and  $w_{i,j,k}^s$  is the weight value for that voxel  
 329 associated with segmentation  $s$ . This was performed prior to thresholding so that the  
 330 resulting global segmentation retained voxel values ranging between [0, 1] depending  
 331 on the confidence. Finally, the segmentation was upsampled, and thresholded using a  
 332 value of  $t = 0.5$ :

$$y_{i,j,k} = \begin{cases} 1 & \text{if } y_{i,j,k} \geq t \\ 0 & \text{if } y_{i,j,k} < t \end{cases} \quad (14)$$

333 where the largest connected body is retained, converted to a surface mesh using  
 334 marching cubes and smoothed to remove voxel artifacts. For mesh smoothing, we  
 335 use a windowed sinc function interpolation kernel to move mesh vertices[22]. More  
 336 specifically, we perform ten iterations with a passband value of 0.01.

## 337 2.9 Experiments, Metrics and Statistical Analysis

338 We compared our results to those of two benchmark nnU-Net models, i.e. a 2D  
 339 nnU-Net and a 3D nnU-Net, trained on the global image volumes. The 3D nnU-  
 340 Net performs 3D convolutions whereas the 2D version performs 2D convolutions, and  
 341 outputs 2D predictions, solely on the image z-plane, the patient’s axial plane. A 3D  
 342 segmentation map output from the 2D nnU-Net is assembled by a z-stack of 2D seg-  
 343 mentations. The nnU-Net is arguably the most state-of-the-art method for medical  
 344 image segmentation and thus is chosen as a benchmark for comparison.

345 The metrics for comparing SeqSeg with the global nnU-Net benchmarks were as  
 346 follows:

$$347 \mathcal{D}(X, Y) == \frac{2 \cdot \|X \cap Y\|}{\|X\| + \|Y\|} \quad (15)$$

$$348 \mathcal{H}(X, Y) = \max\{d(X, Y), d(Y, X)\}, \quad (16)$$

where  $d(X, Y) = \sup_{x \in X} \inf_{y \in Y} d(x, y)$

$$\mathcal{CO}(Y, C_t) = \frac{\int_{C_t} Y dx}{\int_{C_t} dx} \quad (17)$$

349 where  $\mathcal{D}$  is the Dice score,  $\mathcal{H}$  is the Hausdorff distance, and  $\mathcal{CO}$  is the centerline over-  
 350 lap with  $C_t$  being the ground truth centerline and  $X, Y$  being segmentation maps.  
 351 Here  $X$  represents the ground truth segmentation map and  $Y$  the predicted segmen-  
 352 tation maps, either output from SeqSeg or the benchmark nnU-Net models. The Dice  
 353 score measures the overlap between two segmentations and ranges between  $[0, 1]$ . The  
 354 Dice score is common for medical image segmentation because of its ability to penal-  
 355 ize imbalanced datasets accurately. The Hausdorff distance measures the maximum  
 356 distance between two surfaces and has a minimum of 0 for identical surfaces. The  
 357 centerline overlap is a score ranging between  $[0, 1]$  and represents the percentage of  
 358 the ground truth centerline captured in the predicted segmentation. Both the  $\mathcal{H}$  and  
 359  $\mathcal{CO}$  metrics give additional insight into a method’s ability to capture bifurcations and  
 360 specifically small blood vessels compared to the Dice score. The Dice score compares  
 361 segmentations on a volumetric basis by comparing voxels, but since most voxels belong  
 362 to larger blood vessels, it results in indirect bias.

363 Not all blood vessels were annotated in the test datasets, or they were present but  
 364 truncated. We thus masked the outputs from all methods with the volume surrounding  
 365 the ground truth annotated vessels. We define this mask volume by labelling all pixels  
 366 within a six radius distance from the ground truth centerline.

367 For statistical analysis, we perform the Wilcoxon signed-rank test between the  
 368 resulting metrics scores of SeqSeg and the benchmark. This is a non-parametric test  
 369 similar to the paired t-test. But since the paired t-test has limitations when comparing  
 370 machine learning model performances, we opted for the Wilcoxon test [23]. Specifically,  
 371 we test the null hypothesis that the median of differences between the two sets of

372 sample results (metric scores) is zero. A p-value less than 0.05 was considered to  
373 reject the null hypothesis, and therefore indicate a statistically significant difference  
374 between the two sets. We both perform Wilcoxon tests between SeqSeg and 2D nnU-  
375 Net predictions and between SeqSeg and 3D nnU-Net predictions.

### 376 3 Results

377 We tested SeqSeg on 15 held-out VMR image volumes (test set), 8 of which were  
378 CT volumes and 7 of which were MR volumes, as well as additional 18 AVT CT  
379 image volumes. Namely, SeqSeg and the 2D and 3D nnU-Net benchmark methods  
380 were used to segment the vasculature from these image volumes, and those resulting  
381 segmentations were compared to manually-generated “ground-truth” segmentations  
382 from the open data Vascular Model Repository and the AVT dataset[18].

383 A quantitative evaluation of the Dice score ( $\mathcal{D}$ ), Hausdorff distance ( $\mathcal{H}$ ) and center-  
384 line overlap ( $\mathcal{CO}$ ) for segmentations generated from the VMR test set using SeqSeg and  
385 the nnU-Net benchmark methods is presented in Table 3. SeqSeg, on average, obtained  
386 higher Dice scores than the nnU-Net benchmarks in 11 test cases, lower Hausdorff dis-  
387 tance scores in 11 cases, and higher centerline overlap scores in 12 cases, all out of a  
388 total of 15 test cases. Specifically, SeqSeg on average obtained higher scores than the  
389 2D and 3D nnU-Net benchmarks in terms of Dice score by 0.017 and 0.029, respectively  
390 for CT, and 0.036 and 0.029, respectively for MR. For the Hausdorff distance, SeqSeg  
391 on average obtained lower distance than the 2D and 3D nnU-Net benchmarks by 0.59  
392 and 0.966 pixels, respectively for CT, and 0.565 and 0.872, respectively for MR data.  
393 For centerline overlap, SeqSeg obtained higher scores on average than the 2D and 3D  
394 nnU-Net benchmarks by capturing 3% and 10.9% more vessel segments, respectively  
395 for CT, and 9.4% and 10.4% more segments, respectively for MR. Improvements to  
396 metric averages that were statistically significant ( $p < 0.05$ ) are indicated by an aster-  
397 isk \* in Table 3. Beyond mean improvements, SeqSeg appeared far more robust. This  
398 can be observed from the the box plots in Figure A1, which demonstrates greater  
399 consistency in the performance of SeqSeg for all quantitative metrics.

400 Since the objective of segmentation is a unified, high-quality geometric model, dis-  
401 tilling the comparison down to any single metric is overly simplistic (if not deceptive).  
402 For broader perspective, Figure 7 provides a qualitative comparison, showing the CT  
403 and MR segmentations for which the 2D nnU-Net benchmark yielded the best, median  
404 and worst Dice scores, and includes comparison of these to the segmentations gener-  
405 ated from SeqSeg. (Note, since the 2D nnU-Net was superior to the 3D nnU-Net, the  
406 2D nnU-Net was considered the de facto benchmark.) This figure demonstrates that  
407 SeqSeg generally captures more of the vasculature, and particularly the connections  
408 to smaller branch arteries. The Appendix provides a visual comparison of the segmen-  
409 tations generated by SeqSeg and the nnU-Net benchmarks for all CT, and MR, VMR  
410 test cases in Figures A2, and A3, respectively.

411 For the AVT CT test set, the quantitative metric scores can be seen in Table 5  
412 between SeqSeg and the 2D nnU-Net benchmark. On average, SeqSeg obtained higher  
413 Dice scores by 0.065 and centerline overlap by 10.8%. On average, the benchmark

414 Hausdorff distance scores were lower by 0.401 pixels. In terms of statistical significance, both Dice and centerline overlap differences were found statistically significant  
415 whereas the difference in Hausdorff distance was not. Boxplots of the metric scores  
416 are shown in Figure 8, again showing a smaller spread (better robustness) for SeqSeg  
417 compared to the benchmark, particularly for Dice score and centerline overlap. For  
418 qualitative comparison, all resulting meshes are shown in Figure A4. As observed,  
419 SeqSeg captures more, and smaller, branches as compared to the benchmark, even  
420 branches not included in the ground truth (cf. cases 2, 4, 5, 6, 8, 9, 11, 15, 16, 17 and  
421 18).

422 Inference time also differed between SeqSeg and the benchmarks. If both are run  
423 on the same CPU, measured SeqSeg inference time ranged 20-80min, depending on  
424 the number of branches, whereas the nnU-Net benchmarks ranged 2-3hr.

### 426 3.1 Comparison with nnU-Net’s largest connected region

427 The ultimate goal in image-based modeling is to use a segmentation as the computational  
428 domain for numerical simulation. Simulations require domains to be unified  
429 and well defined. Since the nnU-Net segmentations are often disjoint, filtering and  
430 keeping only the largest connected body would be necessary to use the segmentation  
431 for simulation purposes. Thus, for a more practical comparison, in this subsection we  
432 compare SeqSeg segmentations with nnU-Net segmentations that have been filtered  
433 to retain the largest connected region.

434 We present the quantitative metric values for the largest connected region results  
435 in Table 4. From this table we can observe that SeqSeg on average obtained higher  
436 metric scores than the 2D/3D nnU-Net benchmark, respectively, as follows: the Dice  
437 coefficient improved by 0.062/0.032 for CT and 0.064/0.029 for MR; the Hausdorff  
438 distance improved by 1.812/2.002 for CT and 2.153/0.839 pixels for MR; and the  
439 global centerline overlap increased by 10.2/16.8% for CT and 18.7/13.8% for MR.  
440 Improvements to metric averages that were statistically significant ( $p < 0.05$ ) are  
441 indicated by an asterisk \* in Table 4. Differences in centerline overlap scores were  
442 found statistically significant between SeqSeg and both benchmark methods. The box  
443 plots of these metrics for all cases are shown in Figure A1 and again reveal far less  
444 spread in the metrics for SeqSeg compared to both benchmark models, indicating  
445 greater robustness in segmentation results for SeqSeg.

446 Figure 7 displays segmentation results for the best, median and worst case results  
447 for the 2D nnU-Net benchmark largest connected region, and compares to the segmen-  
448 tation predicted by SeqSeg. (Again, we excluded the 3D nnU-Net in this comparison  
449 as it generally performed worse than its 2D counterpart.) For all cases shown, SeqSeg  
450 generally captures a greater number of branches and greater extent of the vessels, even  
451 when compared to nnU-Net’s best results. This is further demonstrated in Figures A2  
452 and A3 in the Appendix for all test cases.

453 For the AVT CT test data, Table 5 and Figure A4 show results for the benchmarks  
454 after largest connected component filtering, quantitatively and qualitatively respect-  
455 fully. We obtain statistically significant difference between SeqSeg and the benchmark  
456 for all metrics; Dice, Hausdorff distance and centerline overlap. As shown in Fig. A4,  
457 SeqSeg produces better unified vascular trees in more instances than the benchmark.

**Table 3** Quantitative comparisons for the VMR test dataset between the two benchmark U-Net segmentation methods (2D, 3D) and SeqSeg using the Dice score ( $\mathcal{D}$ ), Hausdorff distance ( $\mathcal{H}$ ) and centerline overlap ( $\mathcal{CO}$ ). The case types were either aortofemoral (AF) or aortic (AO), and the number of branches segmented is also shown (Nr. Br.). \* indicates statistically significant difference ( $p < 0.05$ )

Mod.	Case	Type	Nr. Br.	$\mathcal{D} \uparrow$			$\mathcal{H} \downarrow$			$\mathcal{CO} \uparrow$			
				Seq- Seg	2D U-Net	3D U-Net	Seq- Seg	2D U-Net	3D U-Net	Seq- Seg	2D U-Net	3D U-Net	
CT	1	AF	9	<b>0.907</b>	0.885	0.846	<b>1.930</b>	2.208	3.526	<b>0.939</b>	0.906	0.657	
	2	AF	10	0.931	<b>0.941</b>	0.909	1.951	<b>1.406</b>	2.442	0.884	<b>0.928</b>	0.611	
	3	AF	10	<b>0.885</b>	0.860	0.855	<b>2.339</b>	4.452	3.973	<b>0.959</b>	0.864	0.791	
	4	AO	5	0.902	<b>0.923</b>	0.901	2.522	<b>1.281</b>	1.976	0.919	<b>0.951</b>	0.939	
	5	AO	5	<b>0.940</b>	0.845	0.865	<b>0.717</b>	1.544	1.804	<b>1.000</b>	0.916	0.992	
	6	AO	6	<b>0.951</b>	0.947	0.946	<b>0.867</b>	0.991	0.999	<b>0.980</b>	0.951	0.942	
	7	AO	5	<b>0.955</b>	0.951	0.938	<b>0.725</b>	3.237	3.216	<b>0.994</b>	0.911	0.864	
	8	AO	4	<b>0.954</b>	0.934	0.932	<b>0.708</b>	1.358	1.473	0.990	<b>1.000</b>	<b>1.000</b>	
	<b>Avg.</b>	-	-	<b>0.928</b>	0.911	0.899	<b>1.470</b>	2.060	2.426	<b>0.958</b>	0.928	0.849	
p-value				-	-	0.547	0.078	-	0.742	0.148	-	0.547	0.109
MR	1	AF	9	<b>0.877</b>	0.706	0.816	<b>1.429</b>	3.613	3.652	<b>0.977</b>	0.510	0.650	
	2	AO	5	<b>0.810</b>	0.759	0.766	<b>1.154</b>	1.515	2.337	<b>0.950</b>	0.796	0.719	
	3	AO	5	<b>0.836</b>	0.824	0.822	<b>0.657</b>	0.807	0.924	<b>0.844</b>	0.816	0.711	
	4	AO	5	0.923	<b>0.933</b>	0.898	1.132	<b>0.975</b>	2.487	<b>0.981</b>	0.929	0.855	
	5	AO	5	<b>0.909</b>	0.894	0.895	<b>1.371</b>	2.141	2.048	<b>1.000</b>	0.930	0.960	
	6	AO	5	0.923	<b>0.932</b>	0.920	0.704	<b>0.653</b>	0.986	<b>0.990</b>	0.966	0.927	
	7	AO	4	<b>0.945</b>	0.925	0.904	<b>0.780</b>	1.474	0.895	<b>1.000</b>	0.929	0.983	
	<b>Avg.</b>	-	-	<b>0.889</b>	0.853	0.860	<b>1.032</b>	1.597	1.904	<b>0.933</b>	0.839	0.829	
	p-value	-	-	-	0.078	0.016*	-	0.109	0.016*	-	0.016*	0.016*	

## 4 Discussion

U-Net learning models, and particularly the more recent nnU-Net, have shown excellent potential for automating image segmentation tasks. However, segmentation of branched vascular structures from medical image data is fraught with challenges since vessels typically compose relatively few pixels, vascular geometry varies considerably between patient and location, and maintaining connectivity of highly branched structures by pixel classification is tricky. We herein propose a sequential segmentation technique (SeqSeg) that leverages U-Net learning to locally build vascular models. We observed that SeqSeg generally outperformed current state-of-the-art global nnU-Net models when tested on typical vascular images used for image-based modeling, particularly when comparing overall extent of connected vasculature predicted. Namely, SeqSeg was superior in extending the segmentations into smaller branch arteries or distal segments when compared to the global nnU-Net benchmarks (or, in fact, “ground truth” segmentations). In addition, SeqSeg performance was generally more robust, as indicated by less variance in the quantitative results.

**Table 4** Quantitative comparison for the VMR test dataset *after* largest connected body filtering between the two benchmark U-Net segmentation methods (2D, 3D) and our method, SeqSeg, using the Dice score ( $\mathcal{D}$ ), Hausdorff distance ( $\mathcal{H}$ ) and centerline overlap ( $\mathcal{CO}$ ). The case types were either aortofemoral (AF) or aortic (AO), and the number of branches segmented is also shown (Nr. Br.). \* indicates statistically significant difference ( $p < 0.05$ )

Mod.	Case	Type	Nr. Br.	$\mathcal{D} \uparrow$			$\mathcal{H} \downarrow$			$\mathcal{CO} \uparrow$		
				Seq- Seg	2D U-Net	3D U-Net	Seq- Seg	2D U-Net	3D U-Net	Seq- Seg	2D U-Net	3D U-Net
CT	1	AF	9	<b>0.907</b>	0.879	0.830	<b>1.930</b>	2.989	6.723	<b>0.939</b>	0.824	0.521
	2	AF	10	0.931	<b>0.932</b>	0.893	1.951	<b>1.544</b>	5.966	<b>0.884</b>	0.800	0.439
	3	AF	10	<b>0.885</b>	0.858	0.846	<b>2.339</b>	4.810	5.060	<b>0.959</b>	0.831	0.733
	4	AO	5	<b>0.902</b>	0.887	0.916	<b>2.522</b>	3.596	2.878	<b>0.919</b>	0.869	0.865
	5	AO	5	<b>0.940</b>	0.704	0.865	<b>0.717</b>	8.163	1.804	<b>1.000</b>	0.750	0.992
	6	AO	6	<b>0.951</b>	0.945	0.946	<b>0.867</b>	1.335	0.999	<b>0.980</b>	0.865	0.942
	7	AO	5	0.955	<b>0.952</b>	0.939	<b>0.725</b>	2.464	2.875	<b>0.994</b>	0.907	0.831
	8	AO	4	<b>0.954</b>	0.934	0.932	<b>0.708</b>	1.358	1.473	0.990	<b>1.000</b>	<b>1.000</b>
	<b>Avg.</b>	-	-	<b>0.928</b>	0.886	0.896	<b>1.470</b>	3.282	3.472	<b>0.958</b>	0.856	0.790
p-value				-	-	-	-	0.109	0.039*	-	0.023*	0.039*
MR	1	AF	9	<b>0.877</b>	0.508	0.810	<b>1.429</b>	13.575	3.652	<b>0.977</b>	0.199	0.623
	2	AO	5	<b>0.810</b>	0.752	0.766	<b>1.154</b>	1.874	2.337	<b>0.950</b>	0.744	0.713
	3	AO	5	<b>0.836</b>	0.832	0.822	<b>0.657</b>	0.807	0.924	<b>0.844</b>	0.816	0.711
	4	AO	5	0.923	<b>0.931</b>	0.898	<b>1.132</b>	1.743	2.524	<b>0.981</b>	0.894	0.855
	5	AO	5	<b>0.909</b>	0.897	0.898	<b>1.371</b>	1.925	1.778	<b>1.000</b>	0.930	0.960
	6	AO	5	0.923	<b>0.930</b>	0.920	<b>0.704</b>	0.831	0.986	<b>0.990</b>	0.916	0.927
	7	AO	4	<b>0.945</b>	0.928	0.904	<b>0.780</b>	1.540	0.895	<b>1.000</b>	0.929	0.983
	<b>Avg.</b>	-	-	<b>0.889</b>	0.825	0.860	<b>1.032</b>	3.185	1.871	<b>0.963</b>	0.776	0.825
	p-value	-	-	-	0.156	0.016*	-	0.016*	0.016*	-	0.016*	0.016*

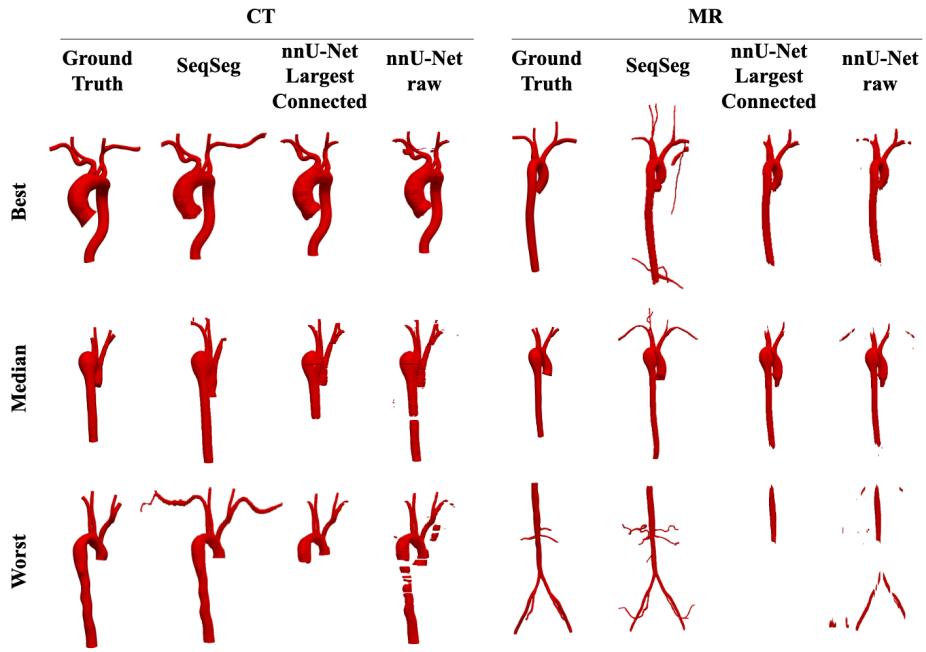
473 We note that when comparing the difference in the quantitative metrics across  
 474 the VMR test cases, the superiority of SeqSeg did not necessarily reach statistical  
 475 significance as measured by  $p < 0.05$ . This is perhaps due to the smaller number of  
 476 test cases we had access to. In addition, the VMR ground truth used for evaluating  
 477 these metrics were not as extensive as they could have been, which likely handicapped  
 478 the comparison since SeqSeg generally excelled at extending the segmentations further  
 479 down the vascular tree compared to the nnU-Net benchmarks. Namely, the ground  
 480 truth data tended to be limited to the aorta and proximal segments of branch arteries,  
 481 which are generally easier to segment, and both SeqSeg and the nnU-Net benchmarks  
 482 performed *on average* equally well in such “less-challenging” regions. Moreover, metrics  
 483 like Dice are inherently biased to larger vessels. While the ground truth segmentations  
 484 could have been manually altered to extend vessels, or include missing branches, this  
 485 can introduce potential bias; thus, we chose not to adulterate the ground truth from  
 486 the public repository. Lastly, we note that most prior publications in this field [24–  
 487 29] do not report whether their improvements to prior benchmarks were statistically  
 488 significant.

**Table 5** Quantitative comparison for the AVT dataset between the benchmark 2D U-Net segmentation method, raw output and after largest connected body filtering (LC), and SeqSeg using the Dice score ( $\mathcal{D}$ ), Hausdorff distance ( $\mathcal{H}$ ) and centerline overlap ( $\mathcal{CO}$ ). The number of branches segmented is shown (Nr. Br.) and \* indicates statistically significant difference ( $p < 0.05$ )

Mod.	Case	Nr. Br.	$\mathcal{D} \uparrow$			$\mathcal{H} \downarrow$			$\mathcal{CO} \uparrow$		
			Seq- Seg	2D U-Net	LC 2D U-Net	Seq- Seg	2D U-Net	LC 2D U-Net	Seq- Seg	2D U-Net	LC 2D U-Net
CT	1	15	<b>0.924</b>	0.87	0.843	67.9	<b>51.7</b>	172	<b>0.782</b>	0.506	0.301
	2	14	<b>0.951</b>	0.901	0.881	<b>28.2</b>	42.6	36.9	<b>1</b>	0.899	0.813
	3	14	<b>0.925</b>	0.567	0.299	60.4	<b>56.8</b>	132	<b>0.929</b>	0.568	0.338
	4	10	<b>0.912</b>	0.857	0.858	<b>17.4</b>	29.9	29.9	<b>1</b>	0.967	0.966
	5	10	<b>0.898</b>	0.823	0.758	<b>27.2</b>	30.4	96.5	<b>0.976</b>	0.834	0.581
	6	8	<b>0.87</b>	0.854	0.835	38.8	<b>29.7</b>	44.7	<b>0.973</b>	0.928	0.783
	7	13	<b>0.936</b>	0.787	0.778	<b>26.4</b>	39.2	42.7	<b>0.933</b>	0.887	0.812
	8	8	<b>0.911</b>	0.857	0.799	25	<b>17.7</b>	130	<b>0.981</b>	0.743	0.412
	9	8	0.875	0.884	<b>0.885</b>	26.3	<b>23.3</b>	22.3	0.988	<b>0.989</b>	<b>0.989</b>
	10	11	0.872	<b>0.874</b>	0.868	99.6	<b>52.4</b>	96.4	0.757	<b>0.874</b>	0.764
	11	11	<b>0.935</b>	0.86	0.795	19.9	<b>19.7</b>	120	<b>0.98</b>	0.705	0.459
	12	8	<b>0.892</b>	0.846	0.846	<b>22</b>	29.4	29.4	<b>0.992</b>	0.969	0.969
	13	9	<b>0.953</b>	0.892	0.867	<b>12.8</b>	23.9	72.2	<b>0.93</b>	0.815	0.502
	14	5	<b>0.918</b>	0.883	0.886	<b>23.4</b>	26.8	26.8	<b>1</b>	0.992	0.992
	15	7	<b>0.916</b>	0.902	0.899	30.1	22.5	<b>18.8</b>	<b>0.995</b>	0.978	0.894
	16	8	<b>0.914</b>	0.866	0.858	<b>17.3</b>	26.4	29.3	<b>0.995</b>	0.966	0.966
	17	11	<b>0.904</b>	0.863	0.792	<b>18.6</b>	25.6	69.2	<b>0.983</b>	0.796	0.572
	18	6	<b>0.941</b>	0.897	0.867	<b>12.9</b>	18.8	75.5	<b>0.998</b>	0.837	0.632
<b>Avg.</b>	-		<b>0.914</b>	0.849	0.812	31.899	<b>31.498</b>	69.120	<b>0.955</b>	0.847	0.708
p-value	-			5.3e-4*	3.8e-5*		0.609	4.2e-4*		3.8e-5*	2.3e-5*

489 The application of SeqSeg to the AVT CT test dataset offered an interesting application.  
490 Whereas SeqSeg and the benchmarks were trained on VRM data, with a subset  
491 of the VMR data held out for testing, the AVT data was a completely *de novo* data  
492 source unrelated to the training. For the application to this *de novo* data, SeqSeg more  
493 convincingly outperformed the nnU-Net benchmarks, achieving statistically significant  
494 higher accuracy in terms of *all* quantitative metrics. Moreover, the qualitative com-  
495 parisons shown in Fig. A4 demonstrate that SeqSeg was able to segment far more  
496 aortic branch arteries, and extend arteries further distally than the benchmark, or, in  
497 fact, even than the ground truth. It is unknown if this superior performance extends  
498 to similar *de novo* MR data since such data was not available.

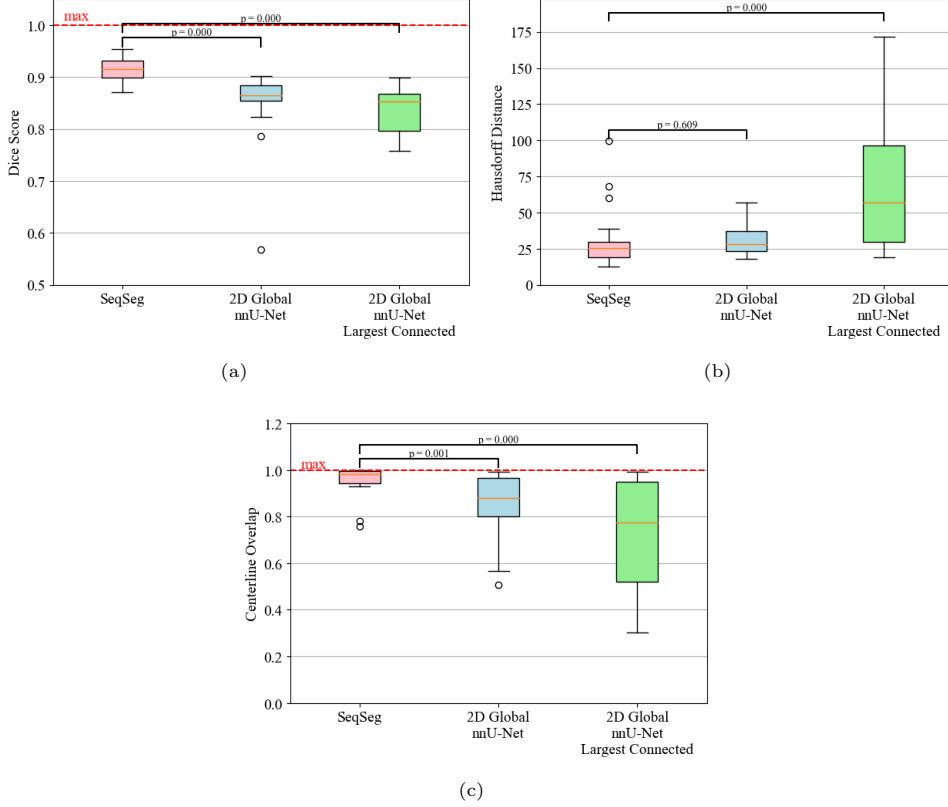
499 The advantage of SeqSeg is that it focuses the segmentation task locally around a  
500 vascular segment. Indeed, the same neural network architecture and training strate-  
501 gies were used for SeqSeg and the benchmarks. The centerline overlap metric, which  
502 give increased weight to smaller branches and bifurcations, showed generally better  
503 performance for SeqSeg. In fact, after largest connected region filtering, the difference  
504 in centerline overlap metric was statistically significant for all test datasets. This is



**Fig. 7** Qualitative comparison of the resulting meshes on the VRM test dataset, comparing the best, median and worst cases of the nnU-Net benchmark to those of SeqSeg. From Table 3 these are cases 7, 4 and 5 for CT and 4, 5, and 1 for MR data, respectively

505 likely because the nnU-Net benchmarks may sacrifice smaller branches in order to  
 506 accurately predict larger branches that carry more pixels. In contrast, SeqSeg deals  
 507 with all branches, if detected, more equitably, by focusing on one segment at a time.  
 508 Figures A2, A3, and A4 show another major advantage of assembling a vascular  
 509 network piece-wise: ensuring connectivity of the resulting model, which is crucial for  
 510 blood flow and tissue mechanics simulation purposes.

511 Since the overall objective is to produce models capable of physics based sim-  
 512 ulation, it is notable to mention SeqSegs superiority towards that goal. Firstly, as  
 513 mentioned above, SeqSeg surpasses the benchmark's ability to generate expansive *and*  
 514 single connected body models. Secondly, since SeqSeg traces the vasculature, it main-  
 515 tains information on branches and their connections relative to the global vascular  
 516 organization. This can be used to place boundary conditions (inlet and outlet condi-  
 517 tions), necessary for physics-based simulation setup. This information is not available  
 518 for typical CNN segmentation methods since all pixels are treated equally and vas-  
 519 culature organization is ignored. The authors note that this study does not directly  
 520 look at the effects of these methods on actual physics-based simulation results, which  
 521 is beyond the scope of the current study.



**Fig. 8** Quantitative metric scores for the AVT test dataset for SeqSeg and the 2D nnU-Net benchmark, both raw and after largest connected component filtering. (a) Dice (b) Hausdorff distance (in pixels) (c) Centerline Overlap. See Sec. 2.9 for definition of p-values.

522 One might assume an advantage of SegSeq is that because it uses local patches  
 523 of the image volume, the number of inputs for training is higher than for the global  
 524 nnUnets. However, during training nnUnet partitions the image volume into patches  
 525 and uses extensive augmentations, which greatly increases the “samples” for training.

526 The results from the benchmark 2D and 3D nnU-Nets show the limitations of 3D  
 527 convolutional neural networks for global vasculature segmentation—the problems of  
 528 class-imbalance and image size. Because global image volumes surpass GPU memory,  
 529 methods are forced to either downsample or split the image into patches to fit on a  
 530 GPU. Our method excels within the constraints of GPU memory because it processes  
 531 smaller subvolumes at each time, which rarely exceeds the GPU memory limit, see  
 532 Table 2 for larger possible batch sizes for example. Furthermore, in a global image  
 533 volume, the vascular pixels represent only a fraction of the total pixels, making training  
 534 difficult. Our method focuses on the pixels around the vasculature, which, by definition,  
 535 alleviates class imbalance.

536 On the other hand, the results also indicate that the source of better segmentation  
 537 is not simply locality. The benchmark models were trained on small patches that

538 do not undergo downsampling. By training locally *and* incorporating prior learned  
539 information, i.e. the location and size of the vessel, SeqSeg is generally able to segment  
540 with greater detail and accuracy, particularly in smaller vessels.

541 Another limitation that impacts global segmentation learning is that ground truth  
542 segmentation, being human-generated, in most cases did not contain segmentation of  
543 all branches or portions. This implies that some training data had certain arteries,  
544 e.g., the renal arteries, segmented while others did not, which can result in poor  
545 segmentation of test data. Since SeqSeg can utilize training patches around vessels,  
546 the training mostly encounters positive examples of arteries, e.g. the renal arteries  
547 only if they are present, and will not encounter negative (wrong) examples from less  
548 segmented images, e.g. where the renal arteries were not segmented. Thus, SeqSeg can  
549 be more efficient with training data, which is beneficial since annotated data collection  
550 is costly and time-consuming. This could also help explain the ability of SeqSeg to  
551 segment a greater number of smaller branches, even those not present in all training  
552 examples.

553 Additionally, SeqSeg may have been able to generalize to regions not present in  
554 the training data because vessels share similar image features when viewed locally. For  
555 example, Figures A2 and A3 show how SeqSeg managed to trace and segment small  
556 bifurcations not included in the ground truth as well as elongate already segmented  
557 vessels. Inspection confirmed that these vessels were present in the image data. In  
558 fact, the authors further confirmed this qualitatively by training a model solely on *one*  
559 branch per case (the aorta), and SeqSeg was able to generalize to other branches on  
560 test data. Furthermore, SeqSeg manages to detect and handle bifurcations, which has  
561 been a challenge for blood vessel tracking and tracing methods [16, 30–33]. Unlike  
562 other works, SeqSeg does not depend on explicit bifurcation detection, but instead  
563 handles them implicitly through surface representations expressing them. This makes  
564 handling complex junctions with multiple branches more tractable.

565 For further comparison to previous works, SeqSeg achieved better global Dice scores  
566 than Maher et al. who trained neural networks for 2D lumen segmentation on similar  
567 datasets [15]. Furthermore, the method of Maher et al. depended on previously user-  
568 generated centerlines, whereas our method automatically generates the centerlines  
569 while simultaneously segmenting the vasculature. This is significant since centerline  
570 generation is often the most time-consuming step of image-based model construction.

571 SeqSeg used a 3D U-Net neural network architecture for local segmentation predictions,  
572 however, other architectures, such as vision transformers [29], transfer learning  
573 models such as 3D MedNet [28] or V-Net [27], with residual connections, could  
574 possibly be used to perform this task. Similarly, future developments could include  
575 learning methods to determine step size or other parameters that are derived from  
576 deterministic procedures in our presented model. For example, deep learning can be  
577 applied to directly predict subsequent points [34], local centerline segments or surfaces  
578 using template-based approaches similar to what has been done for cardiac model  
579 construction [13]. Additionally, the SeqSeg method can be trained and tested for gen-  
580 eralization to other vascular anatomies such as coronary arteries, pulmonary arteries  
581 and cerebrovascular models. Since the training and testing occur locally, new data  
582 from different anatomies can be readily incorporated into the existing framework.

583 Limitations of the presented method include the dependence on voxel-based  
584 segmentation, the dependence on capturing bifurcation roots, and the possible com-  
585 putational cost. Voxel-based segmentation inevitably leads to staircase artifacts on  
586 the final surface, as shown in Figures A2 and A3. Since our stepwise approach relies  
587 on accurately capturing bifurcation roots, there is the possibility of missing whole  
588 branches if the root is difficult to segment, e.g. because of image artifacts. The method  
589 also requires neural network inference at each step, which has the potential to scale  
590 poorly for extensive vascular networks.

## 591 5 Conclusion

592 Despite its growing importance, reconstructing vascular models from medical image  
593 data in an accurate and rapid manner remains an open area of research. In this work  
594 we present SeqSeg; a novel image-based vascular model construction method based  
595 on building the vascular network in a stepwise manner to facilitate learning. SeqSeg  
596 is capable of automatically tracing and assembling a global segmentation and surface,  
597 depending only on a single seed point. We tested the method on CT and MR images  
598 of aortic and aortofemoral models and compared to state-of-the-art benchmark 2D  
599 and 3D U-Net segmentation methods, SeqSeg had similar or better accuracy in terms  
600 of Dice score, Hausdorff distance, and centerline overlap, but more notably was more  
601 robust and able to connect a greater extent of the vasculature. Our future work includes  
602 training and testing using other vascular anatomies as well as further optimizing local  
603 segmentation and bifurcation detection.

604 **Acknowledgments.** We thank Prof. Alison Marsden and Dr. Nathan Wilson for  
605 their input to this work, Dr. Fanwei Kong for advice, comments and support early on  
606 and we thank Dr. Martin Pfaller for valuable input and dataset curation and access.  
607 This research used the Savio computational cluster provided by the Berkeley Research  
608 Computing program at the University of California, Berkeley.

## 609 Declarations

### 610 *Funding*

611 Earlier stages of this work were supported by the National Institute of Health,  
612 Award No. 5R01LM013120. We also acknowledge support by the National Science  
613 Foundation, Award No. 1663747 and 2310910 for later stages of this work.

### 614 *Conflict of interest*

615 The authors have no competing interests to declare that are relevant to the content  
616 of this article.

### 617 *Availability of data and materials*

618 Data used for this study is available at [vascularmodel.com](http://vascularmodel.com)

619 **Code availability**

620 Code written and used for this manuscript is available at <https://github.com/numisveinsson/SeqSeg>

622 **Authors' contributions**

623 NSC and SCS conceptualized the study design and methods. NSC developed the  
624 algorithms, performed the computations and generated results. NSC and SCS inter-  
625 preted the results. NSC developed the first draft of the manuscript and SCS further  
626 contributed to and revised the manuscript.

627 **Citation Diversity Statement**

628 Recent work in several fields of science has identified a bias in citation practices such  
629 that papers from women and other minority scholars are undercited relative to the  
630 number of papers in the field [35–37]. We recognize this bias and have worked diligently  
631 to ensure that we are referencing appropriate papers with fair gender and racial author  
632 inclusion.

633 **References**

- 634 [1] Mukherjee, D., Jani, N.D., Narvid, J., Shadden, S.C.: The Role of Circle of  
635 Willis Anatomy Variations in Cardio-embolic Stroke: A Patient-Specific Simu-  
636 lation Based Study. *Annals of Biomedical Engineering* **46**(8), 1128–1145 (2018)  
637 <https://doi.org/10.1007/s10439-018-2027-5>
- 638 [2] Mirramezani, M., Shadden, S.C.: A Distributed Lumped Parameter Model of  
639 Blood Flow. *Annals of Biomedical Engineering* **48**(12), 2870–2886 (2020) <https://doi.org/10.1007/s10439-020-02545-6>
- 640 [3] Arzani, A., Les, A.S., Dalman, R.L., Shadden, S.C.: Effect of exercise on patient  
641 specific abdominal aortic aneurysm flow topology and mixing. *International Jour-  
642 nal for Numerical Methods in Biomedical Engineering* **30**(2), 280–295 (2014)  
643 <https://doi.org/10.1002/cnm.2601>
- 644 [4] Sengupta, D., Kahn, A.M., Burns, J.C., Sankaran, S., Shadden, S.C., Mars-  
645 den, A.L.: Image-based modeling of hemodynamics in coronary artery aneurysms  
646 caused by Kawasaki disease. *Biomechanics and Modeling in Mechanobiology*  
647 **11**(6), 915–932 (2012) <https://doi.org/10.1007/s10237-011-0361-8>
- 648 [5] Updegrove, A., Wilson, N.M., Merkow, J., Lan, H., Marsden, A.L., Shadden, S.C.:  
649 SimVascular: An Open Source Pipeline for Cardiovascular Simulation. Springer  
650 (2017). <https://doi.org/10.1007/s10439-016-1762-8>
- 651 [6] Lan, H., Updegrove, A., Wilson, N.M., Maher, G.D., Shadden, S.C., Marsden,  
652 A.L.: A Re-Engineered Software Interface and Workflow for the Open-Source Sim-  
653 Vascular Cardiovascular Modeling Package. *Journal of Biomechanical Engineering*  
654 **140**(2), 1–11 (2018) <https://doi.org/10.1115/1.4038751>

656 [7] Arthurs, C.J., Khlebnikov, R., Melville, A., Marčan, M., Gomez, A., Dillon-  
 657 Murphy, D., Cuomo, F., Vieira, M.S., Schollenberger, J., Lynch, S.R., Tossas-  
 658 Betancourt, C., Iyer, K., Hopper, S., Livingston, E., Youssefi, P., Noorani, A.,  
 659 Ahmed, S.B., Nauta, F.J.H., Bakel, T.M.J., Ahmed, Y., Bakel, P.A.J., Mynard,  
 660 J., Di Achille, P., Gharahi, H., Lau, K.D., Filanova, V., Aguirre, M., Nama, N.,  
 661 Xiao, N., Baek, S., Garikipati, K., Sahni, O., Nordsletten, D., Alberto Figueroa,  
 662 C.: CRIMSON: An open-source software framework for cardiovascular inte-  
 663 grated modelling and simulation. *PLoS Computational Biology* **17**(5) (2021)  
 664 <https://doi.org/10.1371/journal.pcbi.1008881>

665 [8] Izzo, R., Steinman, D., Manini, S., Antiga, L.: The Vascular Modeling Toolkit: A  
 666 Python Library for the Analysis of Tubular Structures in Medical Images. *Journal*  
 667 of Open Source Software **3**(25), 745 (2018) <https://doi.org/10.21105/joss.00745>

668 [9] Moccia, S., De Momi, E., El Hadji, S., Mattos, L.S.: Blood vessel segmenta-  
 669 tion algorithms — Review of methods, datasets and evaluation metrics. Elsevier  
 670 Ireland Ltd (2018). <https://doi.org/10.1016/j.cmpb.2018.02.001>

671 [10] Nazir, A., Cheema, M.N., Sheng, B., Li, H., Li, P., Yang, P., Jung, Y., Qin, J.,  
 672 Kim, J., Feng, D.D.: OFF-eNET: An Optimally Fused Fully End-to-End Network  
 673 for Automatic Dense Volumetric 3D Intracranial Blood Vessels Segmentation.  
 IEEE Transactions on Image Processing **29**, 7192–7202 (2020) <https://doi.org/10.1109/TIP.2020.2999854>

676 [11] Chen, B., Kitasaka, T., Honma, H., Takabatake, H., Mori, M., Natori, H., Mori,  
 677 K.: Automatic segmentation of pulmonary blood vessels and nodules based on  
 678 local intensity structure analysis and surface propagation in 3D chest CT images.  
 International Journal of Computer Assisted Radiology and Surgery **7**(3), 465–482  
 679 (2012) <https://doi.org/10.1007/s11548-011-0638-5>

681 [12] Kong, F., Shadden, S.C.: Automating Model Generation for Image-Based Cardiac  
 682 Flow Simulation. *Journal of Biomechanical Engineering* **142**(11) (2020) <https://doi.org/10.1115/1.4048032>

684 [13] Kong, F., Wilson, N., Shadden, S.: A deep-learning approach for direct whole-  
 685 heart mesh reconstruction. *Medical Image Analysis* **74** (2021) <https://doi.org/10.1016/j.media.2021.102222>

687 [14] Maher, G., Wilson, N., Marsden, A.: Accelerating cardiovascular model building  
 688 with convolutional neural networks. *Medical and Biological Engineering and Com-  
 689 puting* **57**(10), 2319–2335 (2019) <https://doi.org/10.1007/s11517-019-02029-3>

690 [15] Maher, G., Parker, D., Wilson, N., Marsden, A.: Neural Network Vessel Lumen  
 691 Regression for Automated Lumen Cross-Section Segmentation in Cardiovascular  
 692 Image-Based Modeling. *Cardiovascular Engineering and Technology* **11**(6), 621–  
 693 635 (2020) <https://doi.org/10.1007/s13239-020-00497-5>

694 [16] Wolterink, J.M., Hamersveld, R.W.v., Viergever, M.A., Leiner, T., Išgum, I.:  
 695 Coronary artery centerline extraction in cardiac CT angiography using a CNN-  
 696 based orientation classifier. *Medical Image Analysis* **51**, 46–60 (2019) <https://doi.org/10.1016/j.media.2018.10.005>

698 [17] Li, Z., Xia, Q., Hu, Z., Wang, W., Xu, L., Zhang, S.: A Deep Reinforced Tree-  
 699 Traversal Agent for Coronary Artery Centerline Extraction. In: *Lecture Notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol. 12905 LNCS, pp. 418–428. Springer, Cham (2021). [https://link.springer.com/10.1007/978-3-030-87240-3\\_40](https://link.springer.com/10.1007/978-3-030-87240-3_40)

703 [18] Radl, L., Jin, Y., Pepe, A., Li, J., Gsaxner, C., Zhao, F., Egger, J.: AVT: Multi-center aortic vessel tree CTA dataset collection with ground truth segmentation masks. *SPIE Medical Imaging, Proceedings* **40**, 116001 (2022) <https://doi.org/10.6084/m9.figshare.14806362>

707 [19] Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnUNet: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021) <https://doi.org/10.1038/s41592-020-01008-z>

711 [20] Lorensen, W.E., Cline, H.E.: Marching Cubes: A High Resolution 3d Surface Construction Algorithm. *Technical Report* 4 (1987)

713 [21] Sabry Hassouna, M., Farag, A.A.: Robust Centerline Extraction Framework Using Level Sets. *Technical report* (2005)

715 [22] Taubin, G., Zhang, T., Golub, G.: Optimal surface smoothing as filter design. In: *Lecture Notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 1064, pp. 283–292. Springer, Heidelberg (1996). <https://doi.org/10.1007/bfb0015544>

719 [23] Rainio, O., Teuho, J., Klén, R.: Evaluation metrics and statistical tests for machine learning. *Scientific Reports* **14**(1) (2024) <https://doi.org/10.1038/s41598-024-56706-x>

722 [24] Nvidia, A.H., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Daguang, N., Nvidia, X.: UNETR: Transformers for 3D Medical Image Segmentation. *Technical report*. <https://monai.io/research/unetr>

725 [25] Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation. *Technical report*. <https://github.com/Beckschen/>

728 [26] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image Is Worth 16x16 Words: Transformers For Image Recognition At Scale.

731        Technical report. <https://github.com/>

732        [27] Milletari, F., Navab, N., Ahmadi, S.-A.: V-Net: Fully Convolutional Neural Net-  
733        works for Volumetric Medical Image Segmentation. In: 2016 Fourth International  
734        Conference on 3D Vision (3DV), pp. 565–571 (2016). <https://doi.org/10.1109/3DV.2016.79>

735

736        [28] Chen, S., Ma, K., Zheng, Y.: Med3D: Transfer Learning for 3D Medical Image  
737        Analysis (2019)

738        [29] Hatamizadeh, A., Xu, Z., Yang, D., Li, W., Roth, H., Xu, D.: UNetFormer: A  
739        Unified Vision Transformer Model and Pre-Training Framework for 3D Medical  
740        Image Segmentation (2022)

741        [30] Jia, D., Zhuang, X.: Learning-based algorithms for vessel tracking: A review.  
742        Elsevier Ltd (2021). <https://doi.org/10.1016/j.compmedimag.2020.101840>

743        [31] Abbasi-Sureshjani, S., Smit-Ockeloen, I., Bekkers, E., Dashtbozorg, B., Romeny,  
744        B.T.H.: Automatic detection of vascular bifurcations and crossings in retinal  
745        images using orientation scores. In: Proceedings - International Symposium  
746        on Biomedical Imaging, vol. 2016-June, pp. 189–192. IEEE Computer Society,  
747        Washington D.C. (2016). <https://doi.org/10.1109/ISBI.2016.7493241>

748        [32] Pratt, H., Williams, B.M., Ku, J.Y., Vas, C., McCann, E., Al-Bander, B., Zhao,  
749        Y., Coenen, F., Zheng, Y.: Automatic detection and distinction of retinal vessel  
750        bifurcations and crossings in colour fundus photography. Journal of Imaging **4**(1)  
751        (2018) <https://doi.org/10.3390/jimaging4010004>

752        [33] Li, Y., Ren, T., Li, J., Wang, H., Li, X., Li, A.: VBNet: An end-to-end 3D  
753        neural network for vessel bifurcation point detection in mesoscopic brain images.  
754        Computer Methods and Programs in Biomedicine **214** (2022) <https://doi.org/10.1016/j.cmpb.2021.106567>

755

756        [34] Dorobantu, A., Oorean, V., Brad, R.: Coronary centerline extraction from ccta  
757        using 3d-unet. Future Internet **13**(4) (2021) <https://doi.org/10.3390/fi13040101>

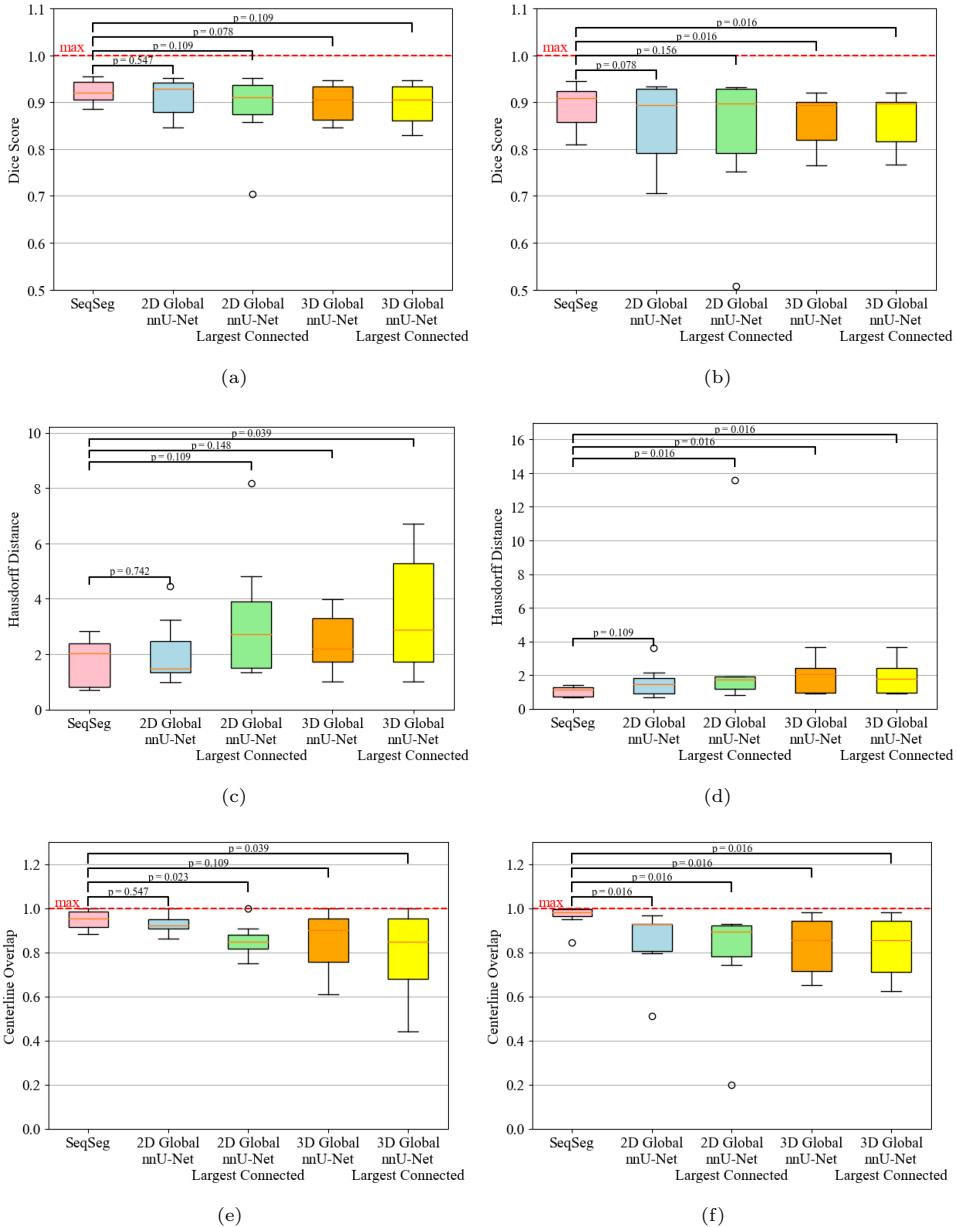
758        [35] Caplar, N., Tacchella, S., Birrer, S.: Quantitative Evaluation Of Gender Bias  
759        In Astronomical Publications From Citation Counts. Technical report (2016).  
760        <http://app.uio.no/ifi/texcount/>

761        [36] Maliniak, D., Powers, R., Walter, B.F.: The gender citation gap in international  
762        relations. International Organization **67**(4), 889–922 (2013) <https://doi.org/10.1017/S0020818313000209>

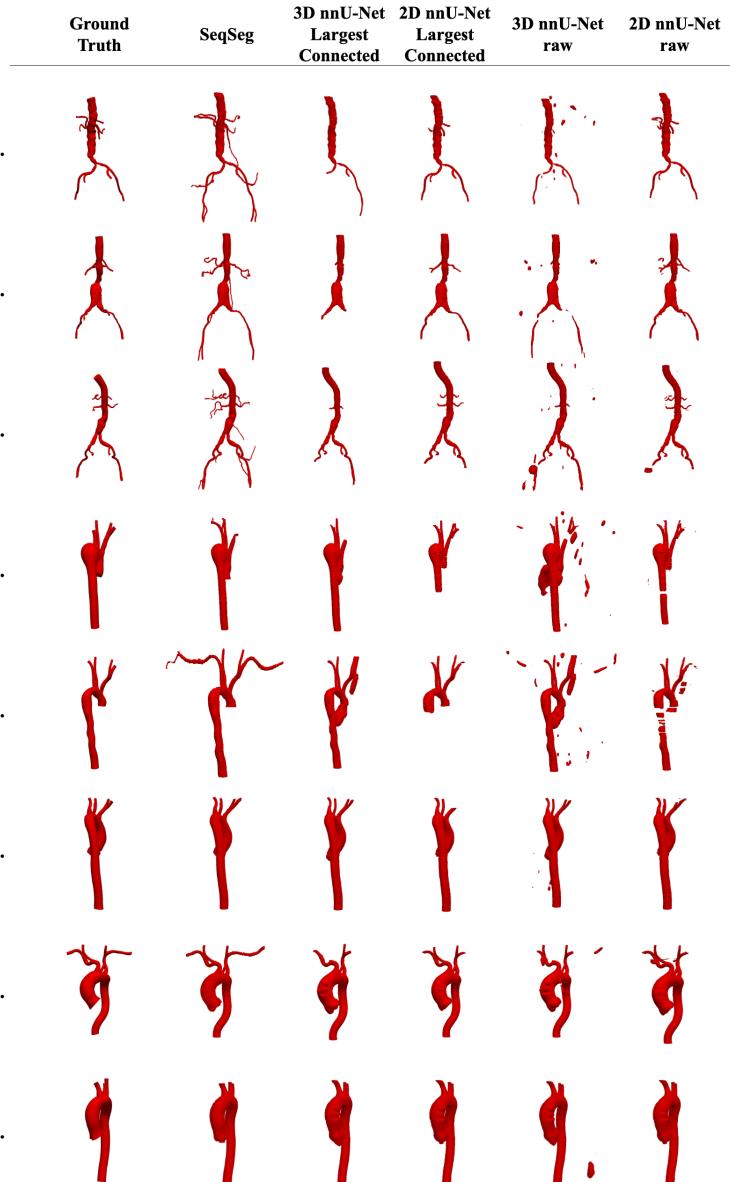
763

764        [37] Mitchell, S.M., Lange, S., Brus, H.: Gendered Citation Patterns in International  
765        Relations Journals. International Studies Perspectives **14**(4), 485–492 (2013)  
766        <https://doi.org/10.1111/insp.12026>

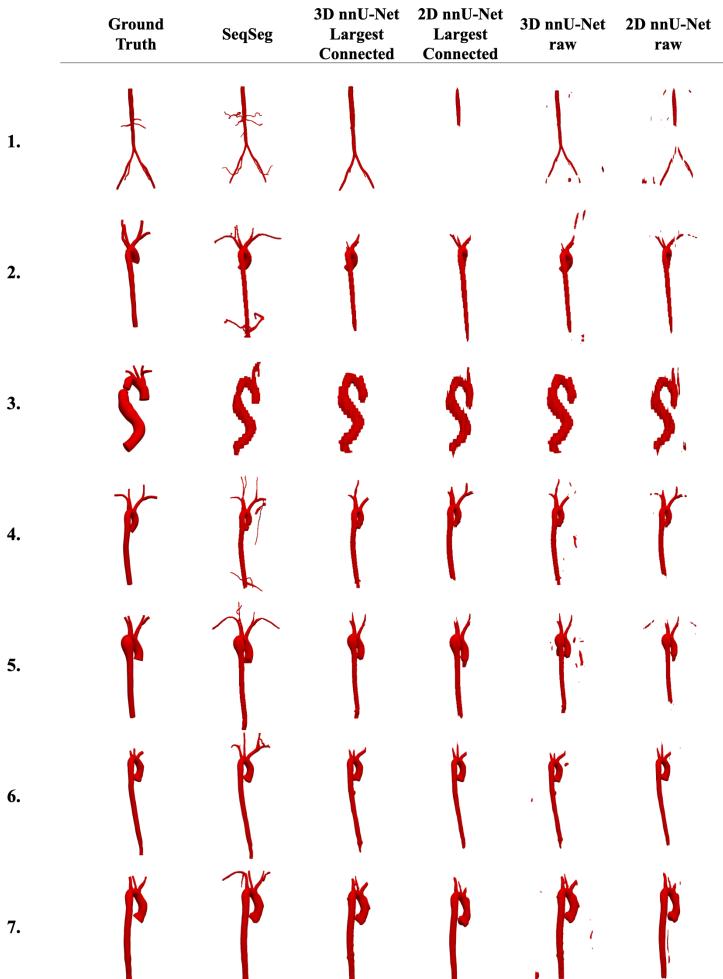
<sup>767</sup> **Appendix A**



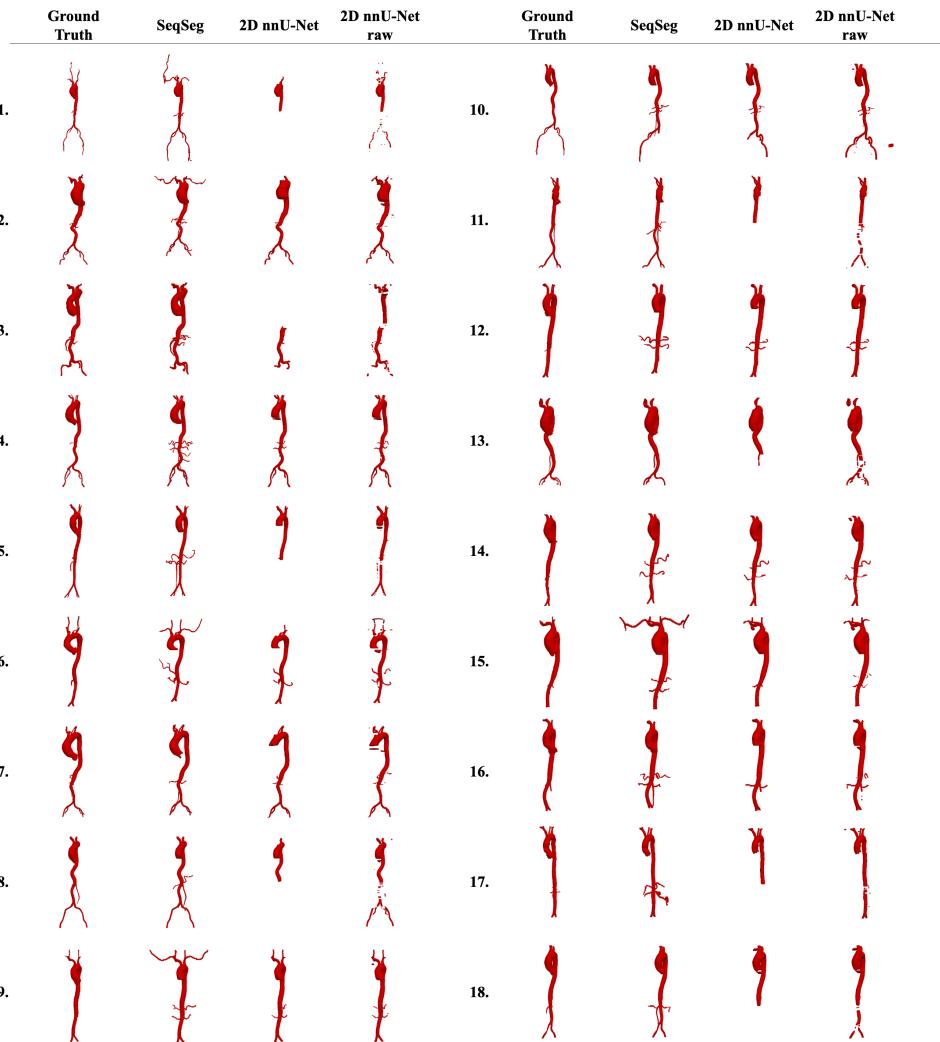
**Fig. A1** Quantitative metrics for VMR test data; (a) CT Dice (b) MR Dice (c) CT Hausdorff (in pixels) (d) MR Hausdorff (in pixels) (e) CT Centerline Overlap (f) MR Centerline Overlap. See Sec. 2.9 for definition of p-values.



**Fig. A2** Resulting meshes from complete VMR CT test dataset. Each row represents a different vascular model, labelled consistently with Tables 3, 4



**Fig. A3** Resulting meshes from complete VMR MR test dataset. Each row represents a different vascular model, labelled consistently with Tables 3, 4



**Fig. A4** Resulting meshes from complete AVT CT test dataset. Each row represents a different vascular model, labelled consistently with Table 5