

Geometric uncertainty in patient-specific cardiovascular modeling with convolutional dropout networks

Gabriel D. Maher^a, Casey M. Fleeter^a, Daniele E. Schiavazzi^b, Alison L. Marsden^{c,*}

^a *Institute for Computational and Mathematical Engineering, Stanford University, 475 Via Ortega, Stanford, CA, 94305, USA*

^b *Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, 102G Crowley Hall, Notre Dame, IN, 46556, USA*

^c *Departments of Pediatrics and Bioengineering, Stanford University, 318 Campus Dr, Clark Center E100b, Stanford, CA, 94305-5014, USA*

Received 16 September 2020; received in revised form 3 May 2021; accepted 5 July 2021

Available online 14 August 2021

Abstract

We propose a novel approach to generate samples from the conditional distribution of patient-specific cardiovascular models given a clinically acquired image volume. A convolutional neural network architecture with dropout layers is first trained for vessel lumen segmentation using a regression approach, to enable Bayesian estimation of vessel lumen surfaces. This network is then integrated into a path-planning patient-specific modeling pipeline to generate families of cardiovascular models. We demonstrate our approach by quantifying the effect of geometric uncertainty on the hemodynamics for three patient-specific anatomies, an aorto-iliac bifurcation, an abdominal aortic aneurysm and a sub-model of the left coronary arteries. A key innovation introduced in the proposed approach is the ability to learn geometric uncertainty directly from training data. The results show how geometric uncertainty produces coefficients of variation comparable to or larger than other sources of uncertainty for wall shear stress and velocity magnitude, but has limited impact on pressure. Specifically, this is true for anatomies characterized by small vessel sizes, and for local vessel lesions seen infrequently during network training.

© 2021 Elsevier B.V. All rights reserved.

Keywords: SimVascular; Cardiovascular model uncertainty; Uncertainty quantification; Neural networks; Deep learning; Computational fluid dynamics

1. Introduction

Results from cardiovascular models are affected by a number of uncertainty sources, including material properties, image-data resolution, and boundary condition selection to match clinical target data. A rigorous determination of simulation uncertainty and the development of numerical approaches to efficiently quantify its effects on patient-specific models are necessary to increase clinical adoption of simulation tools and improve their effectiveness for early treatment planning and non-invasive diagnostics.

Prior studies have investigated a range of methods for characterizing the influence of specific sources of uncertainty in cardiovascular modeling [1,2]. To reduce the computational burden with respect to Monte Carlo

* Corresponding author.

E-mail address: amarsden@stanford.edu (A.L. Marsden).

sampling, a stochastic collocation approach was proposed in [3], while a multi-fidelity approach was proposed by [4] in the context of mechanical stress analysis of abdominal aortic aneurysms. Additionally, a generalized polynomial chaos expansion is presented in [5] and applied to two pathological anatomies, i.e., an abdominal aortic aneurysm and an arteriovenous fistula. A generalized multiresolution expansion for uncertainty quantification was developed in [6] to better handle uncertainty in the presence of non-smooth stochastic responses, while mitigating the exponential complexity of multi-dimensional multi-wavelet refinement. Combined uncertainty in vessel wall material properties and hemodynamics are investigated in [7] for several patient-specific models of coronary artery bypass grafting, leveraging a novel submodeling approach to focus the analysis only on venous and arterial bypass grafts. This and other studies focusing on the coronary circulation, (see, e.g., [8]) contributed to show a loose coupling between hemodynamics and wall mechanics for such anatomies. One dimensional models have been used to better understand main pulmonary artery pressure uncertainty in mice due to material property and image segmentation uncertainty [9,10]. Generalized polynomial chaos simulations involving inlet waveform uncertainty were shown to induce significant wall shear stress variability in an aortic aneurysm model [11]. Multifidelity simulations based on approximate control variate variance reduction in Monte Carlo sampling, were thoroughly analyzed in the context of deformable cardiovascular models in [12]. Finally, uncertainty in combined cardiac–cardiovascular simulation due to input and model parameter variability has also been the subject of recent research [13–15].

While the above contributions focus on the propagation of uncertainty from model inputs to outputs, an end-to-end (or clinical data to simulation results) uncertainty analysis pipeline is proposed in [16] in the context of virtual stage II single ventricle palliation surgery. Additionally, the solution of inverse problems is discussed in [17], where automated Bayesian estimation is applied to tune close-loop boundary condition parameters for patient-specific multi-scale models of the coronary circulation, in order to match a number of non-invasive clinical measurements. The question of flow uncertainty due to mathematical modeling and noise approximation error was investigated using a mouse pulmonary model in [18], where a Bayesian inference approach was developed to reduce bias.

The vast majority of studies in the literature focus on uncertainty in the boundary conditions and mechanical properties of the vascular walls. A third major source is geometric uncertainty which results from errors and operator subjectivity in vessel segmentation from image data, which constitutes a fundamental step in the generation of cardiovascular models. Acquisition of medical image volumes is inherently noisy, has limitations related to the achievable resolution as well as artifact, motion, and aliasing errors. Construction of patient-specific model geometries from image volumes is therefore affected by image uncertainty. In the literature, analysis of the effects of geometrical uncertainty on the results of high-fidelity cardiovascular models has remained elusive due to the complexity of assembling end-to-end pipelines for automatic model generation and analysis. As discretization approaches invariably require the geometry to be represented through a discrete surface mesh, a popular technique within the biomechanics community is mesh morphing [19]. Other methods have focused on modeling the variation of geometry via segmentation approaches. For example, for a given input image the STAPLE algorithm generates a distribution of possible segmentations, but requires a set of ground-truth segmentations as input. Gaussian processes have also been used to model pre-existing segmentation variation in [20]. Segmentation priors and multivariate sensitivity analysis proved useful for segmentation variability estimation in [21], however it is unclear how to extend the method to multiple simultaneous images.

Only a few studies consider the effect of geometrical uncertainty for cardiovascular models. Sensitivity of hemodynamics to geometry variation in patient-specific cerebral aneurysms was investigated in [22], which considered two model samples generated using heuristic smoothing techniques. Manual segmentation uncertainty was shown to have varying influence on FFR-CT calculations in [23], where uncertainty depended on the mean FFR-CT value. In an aortic flow simulation, geometric uncertainty was shown to be a dominant factor when compared to computational fluid dynamics (CFD) model parameter uncertainty and boundary condition uncertainty [24]. Geometric uncertainty was also investigated in [25] for coronary artery simulations, and obtained through local perturbations of an idealized stenosis model. Effects on entire cardiovascular models were investigated in [26–28] by perturbing the area and surface points of selected vessel segments using a spatial Gaussian function with uniform parameterization. The variation in geometry was found to produce sensitivities of up to 10% in simulated FFR-CT measurements. We would like to point out that geometric uncertainty is assumed *a-priori* in the above studies, instead of being directly *learned* from the image data.

More recently, Bayesian Neural Networks, neural networks that are able to learn uncertainty from data, have been increasingly adopted in applications where it is crucial to quantify confidence in predictions [29,30]. In




Fig. 1. Proposed model building pipeline. (a) Image data and vessel pathline are supplied by the user. (b) Path information is used to extract local 2D cross-sectional images in the plane orthogonal to the vessel path. (c) Two-dimensional images are extracted along vessel pathlines and fed to the CNN as inputs. (d) and (e) The proposed CNN processes the cross-sectional images and directly outputs an array of point coordinates, characterizing a two-dimensional lumen segmentation. (f) The collection of two-dimensional points is transformed back to three-dimensional coordinates on the image volume. (g) The cross-sectional segmentations are lofted along the pathline to form the final lumen surface.

particular, [29] showed that augmenting neural networks with dropout layers enables them to learn uncertainty from the training dataset. In the medical imaging field dropout networks have been used to model segmentation uncertainty for MRI volumes [31]. In particular, the network's prediction uncertainty was found to be a useful marker for detecting human expert prediction error.

In this work, we use Bayesian deep learning to develop a cardiovascular model generation technique that learns the geometry distribution from a dataset of existing geometries and images. We then use this network along with Monte Carlo sampling and numerical blood flow simulation to characterize the change in model outputs due to geometric uncertainty.

In Section 2 we discuss our dropout network architecture and path-planning cardiovascular model generation process. Sections 3 and 4 provide an overview of the anatomical benchmarks we selected and the results we obtained. Finally, Section 5 contains a discussion and 6 presents our conclusions.

2. Methods

Given a medical image volume $\mathbf{X}$ and a set of vessel pathlines $\mathbf{V}$, our method produces samples from the distribution of patient-specific cardiovascular models, $\mathbf{Y} \sim P(\mathbf{Y}|\mathbf{X}, \mathbf{V}, \mathbf{z})$, compatible with both the image data, pathline and a collection of latent random variables $\mathbf{z}$. To better explain how this is accomplished, we summarize in Fig. 1 a typical *two-dimensional segmentation* or *path-planning* approach, i.e., a widely used method to generate anatomical surfaces developed in a prior work [32] and based on the cardiovascular model format developed for SimVascular [33] and the Vascular Model Repository (VMR) [34]. This requires to first define a vessel pathline (we will use the term *centerline* interchangeably) by connecting user-specified point locations inside the lumen of the vessel of interest (e.g., the aortic arch in Fig. 1). The tangent vector to this centerline is then used to generate a continuous collection of local 2D images slices. Two-dimensional vessel lumen segmentation on this slice is accomplished through a parametric estimator, trained using a large collection of 2D cross-sectional images and corresponding ground truth lumen boundary. The resulting two-dimensional segmentation, representing the intersection between the lumen wall surface and the cross-section plane, are then *lofted* into a three-dimensional lumen surface and the final model generated by boolean union of multiple vessels.

Note how our vessel lumen estimator also depends on a collection of latent variables $\mathbf{z}$, where two different realizations $\mathbf{z}^{(1)}, \mathbf{z}^{(2)} \sim p(\mathbf{z})$ will produce two distinct anatomical surfaces. This way, we can naturally generate families of cardiovascular models $\mathcal{Y} := \{\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(n)}\}$, where $\mathbf{Y}^i \sim P(\mathbf{Y}|\mathbf{X}, \mathbf{V}, \mathbf{z})$, $i = 1, \dots, n$. Given this ability to *sample* from distributions of cardiovascular anatomies, we adopt Monte Carlo sampling and a computational fluid dynamics (CFD) solver to estimate the changes in hemodynamics induced by geometrical uncertainty in the segmented anatomy. Specifically, each sample $\{\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(n)}\}$ provides a computational domain where we numerically solve the Navier–Stokes equations using the finite element method, as further discussed in Section 2.3. This Monte Carlo process is described in Algorithm 1 and outlined in Fig. 2.

Algorithm 1 Monte Carlo Sampling of Geometrically Uncertain Hemodynamic Solutions

```

medical image volume  $\mathbf{X}$ 
vessel pathlines  $\mathbf{V}$ 
element size  $h$ 
parametric estimator  $m_\theta$ 
 $\mathbf{x} := \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \leftarrow \text{extract}(\mathbf{X}, \mathbf{V})$ 
 $\mathcal{U} := \{\}$ 
for  $i = 1, \dots, K$  do
   $\mathbf{z}_i \sim P(\mathbf{z})$ 
   $\mathcal{Y} := \{\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n\} \leftarrow \{m_\theta(\mathbf{x}_1; \mathbf{z}_i), \dots, m_\theta(\mathbf{x}_n; \mathbf{z}_i)\}$ 
   $\mathbf{Y}_i \leftarrow \text{model}(\mathcal{Y}, \mathbf{V})$ 
   $\mathbf{Y}_i^h \leftarrow \text{mesh}(\mathbf{Y}_i, h)$ 
   $\mathbf{U}_i^h \leftarrow \text{simulate}(\mathbf{Y}_i^h)$ 
   $\mathcal{U} \leftarrow \mathcal{U} \cup \mathbf{U}_i^h$ 
end for
return  $\mathcal{U}$ 

```

extract($\mathbf{X}, \mathbf{V}$) creates a collection of cross-section images by slicing the image $\mathbf{X}$ orthogonal to the vessel pathlines $\mathbf{V}$.

model() generates a cardiovascular model by lofting the two-dimensional lumen segments along the vessel pathline, and merging multiple vessels together by boolean union.

mesh($\mathbf{Y}_i, h$) generates a tetrahedral mesh of the domain $\mathbf{Y}_i$, using an element size h (see details in [33,35]).

simulate($\mathbf{Y}^h$) computes the solution of the Navier–Stokes equations on $\mathbf{Y}^h$, with appropriately chosen initial and boundary conditions.




Fig. 2. Generation of geometrically uncertain cardiovascular model solutions following Algorithm 1.

2.1. Cardiovascular model construction using path-planning

In this section we provide a more formal description of the path-planning process (Fig. 3). The first input consists of a gray-scale medical image volume with H , W and D voxels in the axial, sagittal and coronal direction, respectively, i.e. $\mathbf{X} \in \mathbb{R}^{H \times W \times D}$. The second input is a single pathline which consists of a collection of spline




Fig. 3. Cardiovascular model construction workflow used in SimVascular [33]. Starting from (a) Image data, (b) pathlines are manually generated by the users, (c) two-dimensional lumen segmentations are generated at each cross section $\mathbf{x}_i$, $i = 1, \dots, N_s$, (d) the entire vessel lumen surface is reconstructed by lofting, (e) a Boolean union of multiple vessels is meshed to generate a 3D cardiovascular model.

segments $\mathbf{V} := \{v_1(s), \dots, v_{N_v}(s)\}$. Each segment, $v_i(s) : [0, 1] \rightarrow \mathbb{R}^3$, $i = 1, \dots, N_v$ is a function that maps an arc length parameter to three-dimensional locations within the image volume, obtained through spline interpolation from a collection of user-specified points locations [33].

Our approach is flexible with respect to the method selected to generate the points along the pathline. These can be manually specified by users as in the SimVascular workflow [33] or determined through libraries such as VMTK [36] or other pathline extraction algorithms (see, e.g., [37]). Manual pathline extraction may introduce user bias, whereas automated extraction algorithms are susceptible to approximation error. It is therefore likely that adoption of different pathline extraction methods will affect the resulting lumen segmentation. It is likely however, that automated pathline extraction methods could significantly accelerate the model construction workflow and could be used to analyze different QoI such as the maximum inscribed sphere radius, but the quantification of such effects is considered beyond the scope of this study. In this work, we use manually extracted pathlines from the Vascular Model Repository (VMR) [34].

A collection of N_s local two-dimensional cross-sectional images of the vessel lumen is then extracted for each pathline at the discrete image space locations $\mathcal{S} := \{v_{i_1}(s_1), \dots, v_{i_{N_s}}(s_{N_s})\}$ and the local tangent and normal vectors to the pathline at point $v_{i_j}(s_k)$ are used to construct a planar grid where the voxel intensities from $\mathbf{X}$ are interpolated, to create a gray-scale cross-sectional image of the vessel lumen $\mathbf{x}_i$, $i = 1, \dots, N_s$. Repeating this process for all selected N_s locations along the pathline produces a set of two-dimensional images $\mathbf{x} := \{\mathbf{x}_1, \dots, \mathbf{x}_{N_s}\}$.

Realizations from a Bernoulli random vector $\mathbf{z}$ and the images $\mathbf{x}_i$, $i = 1, \dots, N_s$ constitute the inputs to an artificial neural network designed to produce two-dimensional segmentations of the form

$$\mathbf{y} := \{\mathbf{y}_1, \dots, \mathbf{y}_{N_s}\} = \{m_\theta(\mathbf{x}_1; \mathbf{z}), \dots, m_\theta(\mathbf{x}_{N_s}; \mathbf{z})\}, \quad (1)$$

at all the N_s locations along the pathline. It is important to note how the *same* realization from the Bernoulli vector $\mathbf{z}$ is used to set the dropout layer in the network across all the N_s segmentation instances for each anatomical surface realization. In contrast to generating the segmentations with independent dropout vectors, our process ensures that the same network weights are used to segment all the $\mathbf{x}_i$, $i = 1, \dots, N_s$ images, leading to a consistent bias across the whole cardiovascular model geometry. Finally, the normal and tangent vectors to the pathline are used to re-orient the 2D lumen segments back to image space (Fig. 3c), where they are interpolated and joined together [35] to form a triangular surface mesh of the full cardiovascular model [33]. An unstructured tetrahedral mesh is then generated using TetGen [38] for finite element analysis.

2.2. Convolutional dropout networks for lumen segmentation

For the parametric vessel lumen estimator, $m_\theta(\mathbf{x}; \mathbf{z})$, we use a convolutional neural network. In particular, m_θ maps the input 2D gray-scale image slice, $\mathbf{x} \in \mathbb{R}^{H \times H}$, to a vector of K normalized radii $\mathbf{y} \in [0, 1]^K \subset \mathbb{R}^K$. The radii correspond to the distance of the vessel lumen from the center of the image along rays oriented according to

angular intervals $\phi := \{\phi_1, \dots, \phi_K\}$. This allows the radius $y_i^j \in \mathbf{y}_i$, $i = 1, \dots, N_s$, $j = 1, \dots, K$ to be converted to a single location in the cross-sectional slice $\mathbf{x}_i$ using the expression

$$\mathbf{p}_i^j = \left(y_i^j H \cos \phi_j, \hat{y}_i^j H \sin \phi_j \right), i = 1, \dots, N_s, j = 1, \dots, K, \quad (2)$$

and the full lumen segmentation from image slice $\mathbf{x}_i$ in the set of points $\mathbf{p}_i := \{\mathbf{p}_i^1, \dots, \mathbf{p}_i^K\}$. Even though the literature has witnessed an explosion in new layouts and arrangements in recent years [39,40], a CNN generally consists of a collection of layers, each applying a mathematical operation, such as a linear transformation or convolution, followed by an elementwise nonlinear activation. In our case, the transformation in layer l is expressed as

$$\mathbf{o}^{(l)} = m^{(l)}(\mathbf{a}^{(l-1)}; \boldsymbol{\Theta}^{(l)}), \mathbf{a}^{(l)} = g^{(l)}(\mathbf{o}^{(l)}), \quad (3)$$

where $m^{(l)}(\cdot)$ is the specific mathematical transformation occurring through layer l , $\mathbf{a}$ represents the generic input vector from the previous layer and output to the next one, and $g^{(l)}(\cdot)$ the selected non linear activation. The learnable parameters for the l th layer are denoted by $\boldsymbol{\Theta}^{(l)} \subset \boldsymbol{\theta}$. In this study, we employ a CNN combining dense and convolutional layers. Dense layers operate on vector inputs and outputs through the linear transformation

$$\mathbf{o}^{(l)} = \boldsymbol{\Theta}^{(l)} \mathbf{a}^{(l-1)} + \mathbf{b}^{(l)}, \quad (4)$$

where $\boldsymbol{\Theta}^{(l)}$ and $\mathbf{b}^{(l)}$ are a weight matrix and bias term, respectively. The convolutional layers instead transform a third order tensor input using

$$\mathbf{o}_{ijk}^{(l)} = \sum_o \sum_p \sum_q \boldsymbol{\Theta}_{opqk}^{(l)} \mathbf{a}_{i+o, j+p, q}^{(l-1)} + \mathbf{b}^{(l)}, \quad (5)$$

where $\boldsymbol{\Theta}^{(l)}$ is a fourth order tensor of trainable weights. Note how the outputs from convolutional layers are flattened into one-dimensional vectors before being fed to dense layers.

The activation functions $g^{(l)}(\cdot)$ allow the neural network to learn nonlinear relationships in the data [41] and determine the types of output it can produce. Since the output in our case is a vector of radii in $[0, 1]$ we use the elementwise sigmoid activation function

$$g^{(l)}(x) = \frac{1}{1 + e^{-x}}. \quad (6)$$

For the intermediate layers we instead use Leaky Rectified Linear Units (Leaky-RELU) because they avoid the problem with *vanishing gradients* when optimizing the network weights using gradient-descent [42]

$$g^{(l)}(x) = \begin{cases} x, & x > 0 \\ \alpha \cdot x, & x \leq 0. \end{cases} \quad (7)$$

We augment our convolutional network to sample from the distribution of vessel lumens for a given image by adding dropout layers to the network, which sets the outputs of the previous layer to zero through Hadamard (elementwise) products by a vector of Bernoulli random variables $\mathbf{z}^{(l)}$. In practice, this is implemented as

$$\mathbf{a}^{(l)} = m_{\theta}^{(l)}(\mathbf{o}^{(l-1)}) \quad (8)$$

$$\mathbf{o}^{(l)} = \mathbf{a}^{(l)} \odot \mathbf{z}^{(l)}, \mathbf{z}^{(l)} \sim \mathbf{B}(1 - p) \quad (9)$$

$$\mathbf{a}^{(l+1)} = m_{\theta}^{(l+1)}(\mathbf{o}^{(l)}), \quad (10)$$

where l denotes the layer number after which a dropout layer is applied, $\mathbf{B}$ is a multivariate distribution with independent Bernoulli components, and p the selected dropout probability. The inclusion of dropout layers induces stochasticity to the vessel lumen segmentation process, resulting in random collections of points $\mathbf{y}$ obtained as

$$\mathbf{y} = m_{\theta}(\mathbf{x}; \mathbf{z}), \mathbf{z} \sim \mathbf{B}(1 - p), \quad (11)$$

where $\mathbf{z}$ is a vector containing all Bernoulli dropout variables throughout the network. As discussed above, these variables are kept the same for every two-dimensional segmentation in a single cardiovascular model instance.

For the network architecture, we build on our previous work [32] and use the GoogleNet architecture [43] appropriately modified for vessel lumen regression, which consists of a CNN encoder followed by fully-connected




Fig. 4. CNN based vessel lumen regression with dropout sampling.

layers to transform the encoded vector into the vessel lumen space. A dropout layer is applied to the output of the penultimate layer in the network in order to inject stochasticity (see Fig. 4). A GoogleNet architecture was selected for computational efficiency, as the proposed algorithm is distributed as a SimVascular plug-in, targeting users without specialized hardware. The GoogleNet architecture is computationally efficient due to the use of convolutional and pooling layers with different dimensions to compress the input image while still retaining necessary input information. For more details the reader is referred to [43]. While more recent networks have been developed, earlier studies we conducted [32] showed that the GoogleNet network achieved accuracy comparable to human experts on a 2D vessel lumen segmentation task and so is sufficient for the purposes of this work.

2.2.1. Dropout lumen segmentation sampling

Dropout layers (or simply “dropout”) were first proposed as a regularization method for neural network training [44] and consist of augmenting a chosen layer in the network by randomly “dropping” its output components, that is, each output component is set to zero with probability p . This is achieved through elementwise multiplication with a vector of realizations from independent Bernoulli random variables, but use of independent $\mathcal{N}(1, \sigma^2)$ Gaussian variables has also shown similar performance [44]. Using dropout reduces the weight correlation in the augmented layer due to the possibility of every weight to be switched off during training, and thus promotes the importance of each individual output as an input to the downstream layers. Another key property is that dropout networks have non-deterministic outputs and have shown the ability to learn the conditional distribution of the training data [29]. Thus, neural networks trained with dropouts can be used to generate realistic samples from a given dataset.

While dropout can be added to each layer of the network, augmenting only a single layer with dropout is sufficient, in many cases, to learn the data distribution [29]. In our network, we selected to apply a dropout to the penultimate layer of the network (see Fig. 4). Thus, our network can be seen as a convolutional encoder that maps the input image to a linear combination of basis vectors learned by the final layer. Application of a dropout to the penultimate layer is thus equivalent to selectively including subsets of these basis functions to produce the output segmentation.

More concretely, an input sample $\mathbf{x}$ is processed by each layer of the network up to layer $L - 1$ to produce the hidden feature vector $\mathbf{a}^{(L-1)}$. The dropout layer then randomly sets to zero a subset of the elements of $\mathbf{a}^{(L-1)}$ before sending it to the next layer. The output of the final layer is

$$\mathbf{o}^{(L)}(\mathbf{x}) = \frac{1}{1-p} \Theta^{(L)}(\mathbf{a}^{(L-1)}(\mathbf{x}) \odot \mathbf{z}) + \mathbf{b}^{(L)}, \quad (12)$$

$$\mathbf{a}^{(L)}(\mathbf{x}) = g^{(L)}(\mathbf{o}^{(L)}(\mathbf{x})). \quad (13)$$

To highlight the effect of dropout we reformulate $\mathbf{o}^{(L)}$ as

$$o_i^{(L)}(\mathbf{x}) = \frac{1}{1-p} \left(\Theta^{(L)}(\mathbf{a}^{(L-1)}(\mathbf{x}) \odot \mathbf{z}) \right)_i + b_i^{(L)} = \begin{cases} \frac{1}{1-p} \left(\Theta^{(L)} \mathbf{a}^{(L-1)}(\mathbf{x}) \right)_i + b_i^{(L)} & \text{if } z_i = 1 \\ 0 & \text{if } z_i = 0, \end{cases} \quad (14)$$

which shows that dropout has the effect of generating an output from a random set of basis functions from the final layer. Training the network with samples $(\mathbf{x}, \mathbf{y})$ calibrates the network weights θ such that the output of the network is approximately distributed according to the conditional output distribution of the training data, i.e., $\mathbf{a}^{(L)}(\mathbf{x}) \sim P(\mathbf{y}|\mathbf{x})$.

2.2.2. Network training

During training, the network's weights are initialized according to the variance-scaling approach discussed in [45] and optimized using stochastic gradient-descent ADAM algorithm [46]. We apply the angular distance transform to each ground-truth lumen surface $\mathbf{p}_i^1, \dots, \mathbf{p}_i^K$ to transform it to a ground-truth vector $\mathbf{y}_i$ and create a training dataset of size N_d consisting of the following collection of image-radii pairs

$$\mathcal{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_{N_d}, \mathbf{y}_{N_d})\}. \quad (15)$$

Additionally, we employ a l_2 loss of the form

$$l(\mathbf{y}_i, \hat{\mathbf{y}}_i) = \sqrt{\sum_{j=1}^K (y_i^j - \hat{y}_i^j)^2}, \quad L(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}, \mathbf{z}, N_b) = \frac{1}{N_b} \sum_{i=1}^{N_b} l(\mathbf{y}_i, \hat{\mathbf{y}}_i) = \frac{1}{N_b} \sum_{i=1}^{N_b} l[\mathbf{y}_i, m_{\boldsymbol{\theta}}(\mathbf{x}_i; \mathbf{z})], \quad (16)$$

where $\mathbf{y}_i$ and $\hat{\mathbf{y}}_i = m_{\boldsymbol{\theta}}(\mathbf{x}_i; \mathbf{z})$ represent the i th ground-truth collection of normalized lumen radii and neural network prediction, respectively, while $L(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}, \mathbf{z}, N_b)$ represents the expected loss over a given *batch* of training examples. We also pre-process each input gray-scale image $\mathbf{x}_i \in \mathbb{R}^{H \times W}$, by computing a normalized image $\tilde{\mathbf{x}}_i$ having zero mean and unit variance pixel intensities expressed as

$$\tilde{x}_i^{jk} = \frac{x_i^{jk} - \mu_x}{\sigma_x}, \quad j = 1, \dots, H, \quad k = 1, \dots, W, \quad \mu_x = \frac{1}{H^2} \sum_{j=1}^H \sum_{k=1}^W x_i^{jk}, \quad (17)$$

$$\sigma_x = \left(\frac{1}{H^2} \sum_{j=1}^H \sum_{k=1}^W (x_i^{jk} - \mu_x)^2 \right)^{1/2},$$

where μ_x and σ_x are the mean and standard deviation pixel intensities, respectively. Finally, the training dataset $\mathcal{D}$ is augmented by randomly rotating and cropping each pair of image slice and 2D vessel lumen segmentation.

2.2.3. Dataset

Our dataset consists of 50 CT and 54 MR contrast-enhanced 3D medical image volumes, all publicly available from the Vascular Model Repository (VMR)¹ [34]. For each image volume, the VMR contains vessel pathlines, segmentations, 3D patient-specific models and hemodynamic simulation results (see Fig. 3) created in SimVascular by expert users, in many cases with supervision from a radiologist. To avoid anisotropic voxel spacing, all image volumes were re-sampled keeping an isotropic voxel spacing of 0.029 cm, which ensures the largest vessel diameter to be around 100 pixels, a relatively small window size which reduces the network computation and memory requirements. Specifically, we used a window size $H \times W$ of 160×160 pixels to allow the full range of vessel sizes to be represented with sufficient resolution by each two-dimensional slice. Finally, we split the data into training, validation and testing sets, of 86, 4 and 14 volumes, respectively. This resulted in 16 004, 239 and 6317 cross-sectional images and vessel lumen surface point labels for the training, validation and testing sets, respectively.

Trained with this dataset, our CNN produced 2D cross-sectional vessel lumen segmentations with accuracy comparable to expert SimVascular users across a range of anatomies [32]. An additional increase in accuracy is expected for larger training datasets.

2.3. Patient-specific hemodynamics simulations

The cardiovascular model generation process discussed above results in three-dimensional tetrahedral meshes which provide a domain $\Omega \subset \mathbb{R}^3$ where the incompressible Navier–Stokes equations are solved. These are the equations describing the evolution in time of a Newtonian fluid of constant density ρ , in a domain Ω with boundary $\partial\Omega = \Gamma_D \cup \Gamma_N$, partitioned according to the application of Dirichlet and Neumann boundary conditions, respectively

¹ <http://www.vascularmodel.com>.

$$\begin{cases} \rho \frac{\partial \mathbf{u}}{\partial t} + \rho(\mathbf{u} \cdot \nabla) \mathbf{u} - \nabla \cdot \boldsymbol{\tau} = \mathbf{f} & \text{in } \Omega \times [0, T] \\ \nabla \cdot \mathbf{u} = 0 & \text{in } \Omega \times [0, T] \\ \mathbf{u} = \mathbf{g} & \text{on } \Gamma_D \times [0, T] \\ \boldsymbol{\tau} \cdot \hat{\mathbf{n}} = \mathbf{h} & \text{on } \Gamma_N \times [0, T] \\ \mathbf{u}(0) = \mathbf{u}_0 & \text{in } \Omega \times \{0\}, \end{cases} \quad (18)$$

in which $\mathbf{u}$ is fluid velocity, p is the fluid pressure, $\mathbf{f}$ is a given forcing term, $\hat{\mathbf{n}}$ is the outward directed unit normal vector to Γ_N , and $\boldsymbol{\tau}$ the viscous stress tensor defined as $\boldsymbol{\sigma} = -p \mathbf{I} + 2\mu \boldsymbol{\epsilon}(\mathbf{u})$. Let μ be the dynamic viscosity of the fluid, $\mathbf{I}$ the second order identity tensor and $\boldsymbol{\epsilon}(\mathbf{u})$ the strain-rate tensor defined as $\boldsymbol{\epsilon}(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^T)$. The functions $\mathbf{g}$ and $\mathbf{h}$ are given Dirichlet and Neumann boundary data, while $\mathbf{u}_0$ is the initial condition.

We numerically solve the system (18) using a Streamline Upwind Petrov–Galerkin (SUPG) finite element method implemented in the SimVascular flow solver (svSolver) [47], which contains specialized routines for cardiovascular CFD such as backflow stabilization [48], algebraic system solvers and preconditioners [49] and a large collection of physiologic boundary conditions (see, e.g., [48–51]). The numerical solution is integrated in time using a second-order generalized- α method [52]. We also apply RCR boundary conditions for generic outlets and a coronary lumped parameter boundary condition for coronary artery outlets, respectively (see, e.g., [50,53,54]). Finally, we restrict our attention to simulations with rigid walls.

2.4. Monte Carlo sampling of cardiovascular flow solutions

Our model generation procedure generates a set of discrete meshes $\mathbf{Y}^h := \{\mathbf{Y}_1^h, \dots, \mathbf{Y}_{N_y}^h\}$ (here the superscript h is used to indicate the size of the discrete mesh). Numerical solution of the Navier–Stokes equations on each mesh subsequently produces a set of velocity fields $\mathbf{U}^h := \{\mathbf{U}_1^h, \dots, \mathbf{U}_{N_y}^h\}$ and pressure fields $\mathbf{P}^h := \{\mathbf{P}_1^h, \dots, \mathbf{P}_{N_y}^h\}$ with $\mathbf{U}_i^h : \mathbf{Y}_i \times [0, T] \rightarrow \mathbb{R}^3$ and $\mathbf{P}_i^h : \mathbf{Y}_i \times [0, T] \rightarrow \mathbb{R}$.

Our objective is to calculate relevant Monte Carlo (i.e., sample) statistics using the ensembles $\mathbf{U}^h$ and $\mathbf{P}^h$. However, the mesh geometry and hence the solution domain Ω is not constant for different realizations of the flow and pressure fields. This precludes us from considering quantities of interest defined at specific point locations in Ω , and we focus instead on output quantities that do hold meaning in the context of varying geometry. Consider a generic model result $\mathbf{r}_i(\boldsymbol{\omega}, t)$ for the i th geometry realization, and the cross-sectional area A_i^j corresponding to the j th slice location. We define the quantity

$$q_i^j(t) = \frac{1}{|A_i^j|} \int_{A_i^j} r_i(\boldsymbol{\omega}, t) d\Gamma, \text{ or the quantity } q_i^j(t) = \frac{1}{|\partial A_i^j|} \int_{\partial A_i^j(s)} r_i(\boldsymbol{\omega}, t) d\Gamma, \quad (19)$$

for situations where model outputs are only defined on the lumen surface ∂A_i^j (e.g., wall shear stress). The second type of quantities are time-average versions of the $q_i^j(t)$, such as,

$$q_i^j = \frac{1}{(T_2 - T_1)} \int_{T_1}^{T_2} q_i^j(t) dt. \quad (20)$$

In this study, we focus on quantities $\mathbf{r}_i(\boldsymbol{\omega}, t)$ such as the pressure as well as the wall shear stress and velocity magnitudes. Finally, for our Monte Carlo trials we choose to report a relative measure of variability, i.e., the coefficient of variation, defined as

$$\text{CoV} = \frac{\sigma}{\mu}, \quad (21)$$

where σ and μ are the sample mean and the standard deviation of the quantities of interest $q_i^j(t)$ or q_i^j , respectively.

We further report confidence intervals for our Monte Carlo estimates. By the Central Limit Theorem the sample mean tends to a Gaussian random variable with standard deviation

$$\tilde{\sigma} = \frac{\sigma}{\sqrt{N}}, \quad (22)$$

where N is the sample size. In many cases the absolute value of the sample mean is small, therefore, for the sake of clarity, we report the confidence interval as a percentage of the sample mean, that is

$$CI = \pm \alpha \frac{\tilde{\sigma}}{\mu}, \quad (23)$$

where α is a constant depending on the confidence level (e.g. $\alpha = 2$ for 95% confidence). The CoV is thus an estimate of the variability of a particular QOI and CI is a measure of the accuracy of our sample mean estimates.

2.5. Shape variability assessment through principal component analysis

To better understand the shape variation in our cardiovascular model samples we use Principal Component Analysis (PCA) to compute a low rank factorization of the matrix $Y \in \mathbb{R}^{M \times N}$ constructed from the collection of generated cardiovascular models. In particular we compute

$$Y - \bar{Y} = U_K \Sigma_K V_K^T, \quad (24)$$

where $\bar{Y}$ is the mean of the model sample, Σ_K is a diagonal matrix containing the singular values on its diagonal and the columns of $U \in \mathbb{R}^{M \times K}$ represent a reduced-order basis for the deviation of the cardiovascular models from the mean model. We use U to study the modes of variation of the generated cardiovascular models in our Monte Carlo samples.

By the properties of the PCA factorization, the columns of U_K are ordered starting with the modes that capture the most variance of Y . The first column thus represents the most significant mode, the second column the second most significant mode etc. revealing the dominant modes in which segmentation uncertainty is causing the cardiovascular model geometry to vary.

Furthermore the modes U_k also encode the correlation amongst the generated vessel lumen, since

$$C_Y = \frac{1}{N-1} (Y - \bar{Y})(Y - \bar{Y})^T = \frac{1}{N-1} U_K \Sigma_K V_K^T V_K \Sigma_K U_K^T = \frac{1}{N-1} U_K \Sigma_K^2 U_K^T, \quad (25)$$

where C_Y is the covariance matrix of Y . The singular values quantify how much each PCA mode contributes to explain the variance in the data. Modes associated with smaller singular values represent local perturbations that are present in a limited number of model geometry realizations. As such, the number of modes can be truncated once a sufficient fraction of the total variance is explained. In this work, we report the number of modes required to capture 99% of the variance.

3. Demonstration in selected cardiovascular anatomies

3.1. Aorto-iliac bifurcation model


The first anatomy we consider consists of the bifurcation of the abdominal aorta into the two iliac arteries, a model with one inlet and two outlets (see Fig. 5b). The inlet boundary condition is chosen to be a typical physiological waveform, corresponding to an average inflow of 6 L/min (see Fig. 5c), while outflow RCR boundary conditions are applied, where the resistance and compliance parameters were preliminarily tuned to produce a realistic outlet pressure range of 80–120 mmHg (see Table 5a). Initially, we conducted a mesh convergence study with meshes comprising 100,000, 250,000 and 1,500,000 tetrahedral elements and boundary layer mesh with 5 layers, and compared these to a reference mesh with 3,500,000 elements. Mesh convergence was assessed by first time-averaging the QOI and then using the mean absolute error

$$\epsilon^h = \frac{1}{|\Omega|} \int_{\Omega} |\bar{f}^h(\xi) - \bar{f}^*(\xi)| d\Omega, \quad (26)$$


where $\bar{f}^h$ and $\bar{f}^*$ are the time-averaged QOI on the investigated and reference mesh respectively. The 1,500,000 element mesh showed a less than 0.2%, 6% and 1.7% error for the pressure, WSS magnitude and velocity magnitude respectively (see Figs. 5d, 5e and 5f) and is employed in all additional experiments. The size of the Monte Carlo ensemble is finally selected equal to 150 for this anatomy.

Vessel	R_1	C	R_2
aorta	237.0	0.00115	2370
left iliac	376.0	0.00085	3760


(a) RCR boundary condition parameters. Resistance in $\text{dyne} \cdot \text{s}/\text{cm}^5$, capacitance in cm^5/dyne .




(b) Patient-specific model surface.




(c) Aortic inlet waveform



(d) Pressure convergence study.



(e) WSS convergence study.



(f) Velocity convergence study.

Fig. 5. Aorto-iliac bifurcation model with boundary conditions (a, c), lumen surface (b) and mesh convergence analysis (d, e, f).

3.2. Abdominal aortic aneurysm model

The second anatomy considered in this study includes the aorta and its main branches from an abdominal CT image of a patient with an abdominal aortic aneurysm (AAA, see Fig. 6a), subject to the same aortic inflow used for the previous anatomy (see Fig. 5c). RCR boundary condition parameters are reported in Table 6b. For this anatomy, a family of 110 geometries was generated through Monte Carlo sampling. The sample size was reduced compared to the previous model as 110 models were found sufficient for statistical convergence. We conducted a mesh convergence study, using the mean absolute error (26) and with meshes having roughly 500,000, 700,000 and 3,000,000 elements and boundary layer mesh with 5 layers and compared these to a reference mesh with 7,000,000 elements. The 3,000,000 elements mesh showed a less than 0.5%, 5% and 3% mean error for the pressure, WSS magnitude and velocity magnitude, respectively (see Fig. 6) and was subsequently used in all numerical experiments reported below.

3.3. Left coronary artery model

The third model we consider includes the left anterior descending (LAD) and left circumflex (LCx) coronary arteries extracted from a CT image volume, also studied in [49]. Coronary lumped parameter boundary condition values were selected to produce physiological pressure ranges (Fig. 7a). The coronary simulations use a pulsatile coronary inflow waveform (Fig. 7c). A sample size of 110 models was used for the Monte Carlo trials. We conducted a mesh convergence study, using the mean absolute error (26) and with meshes with roughly 500,000, 1,000,000 and 1,500,000 elements and boundary layers with 5 layers and compared these to a reference mesh with 3,500,000 elements. The 1,500,000 elements mesh showed a less than 0.075%, 3% and 1% mean absolute error for the pressure, WSS magnitude and velocity magnitude respectively (Fig. 7). Subsequently, this mesh was used for all further experiments.




Fig. 6. AAA model anatomy (a) with boundary conditions (b) and mesh convergence analysis (c, d, e).

4. Results


4.1. Comparison of dropout network lumen samples to human expert segmentation

Before commenting on the model results, we first investigate the statistical properties of the segmentations produced by our dropout network. In addition, we compare network samples and lumen segmentations produced by a number of expert SimVascular users. To do so, we selected four representative image volumes with vessel centerlines from the Vascular Model Repository [34], including a cerebrovascular anatomy imaged by MR, a coronary anatomy with aneurysms caused by Kawasaki disease imaged by CT, a coronary anatomy following bypass graft surgery imaged by CT, and a pulmonary anatomy imaged by MR. Slices were selected at discrete intervals along the vessel centerlines, and each location was segmented by three individual SimVascular experts (see Fig. 3), resulting in a total of 290 segmentations per expert. For the same slices, 50 neural network lumen samples were generated for various dropout probabilities, i.e., $p = 0.9$, $p = 0.7$ and $p = 0.4$, and used for statistical analysis.


Segmentation radius CoV observed for SimVascular expert users is separated into two classes, i.e., large vessels ($r > 0.4$ cm), where the radius CoV is typically less than 5%, and small vessels ($r \leq 0.4$ cm), where the radius CoV is larger, with values as large as 30% (see Fig. 8a). The inverse relationship between radius CoV and lumen radius is explained by the fact there is a minimum error produced by human segmentation, due to limits in image resolution, acquisition noise, and expert image interpretation. Similar trends are observed for neural network samples with varying dropout probabilities, where the radius CoV increases with the dropout probability (see Figs. 8b–8d). In addition, the network with dropout $p = 0.9$ produces similar CoV to expert SimVascular users, whereas a dropout

Vessel	R_1	C_1	R_2	C_2	R_3
LAD	125,342	$4.54 \cdot 10^{-7}$	203,681	$3.68 \cdot 10^{-6}$	62670
LAD - D_1	78,596	$7.26 \cdot 10^{-7}$	127,720	$5.87 \cdot 10^{-6}$	39,298
LCx	90,578	$6.30 \cdot 10^{-7}$	147,190	$5.09 \cdot 10^{-6}$	45,289
LCx - OM_1	38,155	$1.49 \cdot 10^{-6}$	62,002	$1.21 \cdot 10^{-5}$	19,077
LCx - OM_2	132,249	$4.31 \cdot 10^{-7}$	214,905	$3.49 \cdot 10^{-6}$	66,124
LCx - OM_3	254,268	$2.24 \cdot 10^{-7}$	413,186	$1.81 \cdot 10^{-6}$	127,134


(a) RCRCR boundary condition parameters for left coronary artery model, resistance in $\text{dyne} \cdot \text{s}/\text{cm}^5$, capacitance in cm^5/dyne .




(b) Left coronary artery model anatomy.




(c) Coronary inlet waveform



(d) Pressure convergence study.




(e) WSS convergence study.




(f) Velocity convergence study.


Fig. 7. Left coronary artery model with boundary conditions (a, c), lumen surface (b) and mesh convergence analysis (d, e, f).




(a) SimVascular experts



(b) CNN dropout $p = 0.9$



(c) CNN dropout $p = 0.7$



(d) CNN dropout $p = 0.4$

Fig. 8. Radius CoV against radius for SimVascular experts and GoogleNet network with different levels of dropout.

$p = 0.4$ produces more precise segmentations. Thus, the dropout probability can be tuned to modulate the variability in the resulting segmentation.

For large vessels, the network with dropout $p = 0.9$ produces a CoV distribution compatible to that produced by an expert SimVascular user (see Fig. 9a). However, for small vessels, the CoVs produced by the network are significantly larger (see Fig. 9b). Conversely, dropout probabilities equal to $p = 0.7$ and $p = 0.4$ produce significantly lower CoVs than an expert user for both large and small vessels. In particular, a dropout probability of $p = 0.4$ produces the smallest CoVs, showing direct proportionality between the dropout probability and the amount of segmentation uncertainty. The overall lumen shape is captured rather well by expert users, but their segmentations exhibit deviations in local vessel radii (see Figs. 10a, 10e). The same happens with our approach for a dropout probability equal to $p = 0.9$ (see Figs. 10b, 10f). For smaller p the variability in the segmentations is reduced, and they all converge closely to the mean lumen profile (see Figs. 10b–10d and 10f–10h).

Our results show how the dropout probability p can be tuned to represent a desired level of uncertainty in the network outputs. For example, choosing $p = 0.9$ would cause the network to generate samples with uncertainty similar to that generated by human experts, allowing the effects of human uncertainty to be investigated. Typical ranges for the dropout probability p from the literature are 0.5 to 0.8 for hidden layers and 0.8 for input layers [44].




Fig. 9. Comparison of mean radius CoV for large and small vessels between SimVascular experts and the proposed network, with varying dropout probability. The distributions of CoV shown above were obtained by collecting one CoV for each cross-section $\mathbf{x}_i$.






Fig. 10. Vessel lumen segmentation generated for large (top) and small (bottom) vessels by expert users and the proposed network with varying dropout probability.


Additional intuition on the relation between the loss from an ensemble of all possible sub-networks and the dropout loss is presented for a single linear unit in [55], where $p = 0.5$ is shown to be associated with maximum regularization. Our approach is however based on the GoogleNet network [43] where $p = 0.4$ is shown to produce optimal results. We choose to maintain this value, therefore focusing on the uncertainty produced by automated rather than manual segmentation.



(a) Aorta and Iliacs case - Left Iliac








(b) Radius CoV



(c) Radius standard deviation

Fig. 11. Lumen segmentation samples and radius CoV/standard deviation for aorto-iliac bifurcation test case, computed over cross-sectional slices $\mathbf{x}_i$, $i = 1, \dots, 156$.

(a) $\lambda_1 = 16.86$ (b) $\lambda_2 = 13.80$ (c) $\lambda_{19} = 3.00$ (d) $\lambda_{20} = 2.94$ 

(e) Singular values

Fig. 12. PCA modes overlayed on mean aorto-iliac bifurcation model geometry. 184 modes captures 99% of the variance.

4.2. Aorto-iliac bifurcation model

The lumen generated from our dropout network shows good qualitative agreement with the depicted vessel lumen for the left iliac artery (Fig. 11a) and is able to correctly identify the relevant main branch even in the presence of surrounding tissue noise and branching vessels.

Variation in the segmentation radii appears to be limited, with standard deviation σ_r between 0.005 cm and 0.01 cm (see Figs. 11b and 11c). In addition, variability in σ_r appears to increase with decreasing vessel size, likely due to the typically poorer resolution of smaller vessels. The roughly constant σ_r also results in a CoV that increases with decreasing vessel size, as typically seen for segmentations performed by expert operators [32], albeit with larger magnitude.

By considering the PCA modes in decreasing order of singular value, we can visualize the dominant ways in which the shape of individual models in the sample vary. The first two modes appear to be associated with the radial expansion of entire vessels, such as the right iliac or large segments of the aorta (see Figs. 12a–12d). Conversely, higher modes involve local geometrical perturbations such as the bulbous region at the proximal end of the left iliac.




Fig. 13. PCA modes overlaid on mean aorto-iliac bifurcation model geometry, when generating segmentations using independent dropout Bernoulli vectors for each vessel lumen.




Fig. 14. Monte Carlo moment traces for aorto-iliac bifurcation model QoIs.



These latter modes, however, contribute less, being associated with relatively small singular values. For example, the 19th and 20th modes have singular values that are a factor of approximately 5.7 smaller than the singular value for the first mode. Our network generates models through independent 2D segmentations along the vessel pathlines, as such correlated whole vessel variation is not guaranteed. Thus the fact that the dominant modes of the PCA result in whole vessel perturbations shows that the geometric variability produced by the proposed dropout network is distributed across the entire model in a correlated fashion, and increases with small vessel radii. We highlight this fact by generating a model sample while forcing the network to resample a new Bernoulli dropout vector at each vessel lumen location and then recompute the PCA modes (Fig. 13). The produced modes now appear to make more random perturbations and are less localized to specific features of the vessel geometry.

Convergence of Monte Carlo statistics such as the mean, standard deviation and CoV for the pressure, TAWSS and velocity magnitude integrated over the aorta appears to be satisfactory (Fig. 14), with 95% confidence intervals for the mean showing faster convergence for the pressure, then velocity magnitude and finally TAWSS.

Table 1

Monte Carlo sample mean, coefficient of variation (CoV) and 95% relative confidence interval for all QoIs in the aorto-iliac bifurcation model. n indicates the number of cross-sectional slices for the associated vessel.

Path	Aorta ($n = 95$)	Left iliac ($n = 61$)	Path	Aorta	Left iliac
Radius mean [cm]	0.84	0.61	TAWSS mean [dyne/cm ²]	40.24	46.55
Radius CoV	0.006	0.011	TAWSS CoV	0.027	0.030
Radius conf.	0.0012	0.0020	TAWSS conf.	0.0050	0.0056
Pressure mean [mmHg]	98.43	96.02	Velocity mean [cm/s]	42.84	41.78
Pressure CoV	0.003	0.002	Velocity CoV	0.014	0.019
Pressure conf.	0.0005	0.0004	Velocity conf.	0.0026	0.0035

**Fig. 15.** Outlet QoIs and $\pm 2\sigma$ interval for aorto-iliac bifurcation model.**Fig. 16.** Time averaged QoIs and $\pm 2\sigma$ interval for aorto-iliac bifurcation model, plotted along the vessel centerline.

The larger sensitivity of TAWSS to geometric uncertainty can be observed from their time histories over the last two cardiac cycles, as shown in Figs. 15a, 15b. Pressure variability tends to decrease towards the distal end of the vessel, whereas TAWSS and velocity magnitude show an opposite trend. This relates to an increase in the wall shear stress and velocity after the bifurcation which, for this model, amplifies the effect of the geometric uncertainty for these two QoIs. Pressure uncertainty instead depends on the variability of the vessel resistance which cumulates the contributions of each uncertain segmentation along the vessel (see Figs. 16 and 17).

CoVs were found to be approximately 0.4%, 1.5% and 3% for pressure, velocity magnitude and TAWSS, respectively. Thus, CoVs for TAWSS and velocity magnitude are roughly a factor of 10 and 5 larger compared to the pressure CoV, highlighting the increased sensitivity of TAWSS and velocity to geometry variation (see Table 1).

We note that changes to the mesh size or an alternative triangulation of a mesh with the same element size will slightly alter the results shown here. Our convergence study has shown these approximation errors to be small, on the order of less than 1–6% depending on the QoI, when using a mesh with a smaller element size than that used in our studies (e.g. 3,500,000 elements vs. 1,500,000 for the Aorto-Iliac bifurcation case). Thus, up to a relatively small error, our results are independent of the chosen mesh.

4.3. Abdominal aortic aneurysm model

The lumen generated by the network agrees well with the vessel lumen images, even in the presence of noise and bifurcations (see Fig. 18a). Values of σ_r from the dropout network are generally in the range of 0.005–0.01 cm, with some outliers, particularly for larger vessel lumens. Segmentation of small vessels is affected by increased uncertainty, with radius CoV of 3%, versus 1% for the largest vessels. PCA shape analysis shows the first two modes affecting large scale model features such as the aortic aneurysm as well as the entire aorta and iliac branches (see







Fig. 17. Nearest neighbor interpolation of cross-sectional time-averaged CoVs for aorto-iliac bifurcation model.



(a) Abdominal Aortic Aneurysm Case - Celiac Splenic



(b) Radius CoV



(c) Radius standard deviation

Fig. 18. Lumen segmentation samples and radius CoV/standard deviation for abdominal aortic aneurysm test case, computed over cross-sectional slices $\mathbf{x}_i$, $i = 1, \dots, 581$.

Fig. 19a). The 19th and 20th modes (with singular values a factor of four smaller than the first mode) are instead associated with local features, like celiac branches or by asymmetric aneurysm perturbations (see **Fig. 19c**).

The relative confidence intervals of the Monte Carlo estimates for the mean pressure, TAWSS and velocity magnitude were found to be within the ranges 0.07–0.8%, 0.44–0.79% and 0.28–0.65%, indicating a satisfactory convergence, particularly compared to the observed CoV for the same vessels (see **Table 2** and **Fig. 20**). Outlet profiles show increased variability in branch vessels for all QoIs with respect to the Aorta (see **Figs. 21a–21g**). Time average flow results show increasing TAWSS and velocity magnitude variability towards the distal end of the vessel (see **Figs. 22a–22g**).

As expected, radius CoV is inversely proportional to the vessel size (see **Fig. 23a**). Relative pressure variability is approximately uniform along the path of each vessel, and particularly elevated in the celiac hepatic branch, due to the fact that it branches off of the celiac splenic which itself branches off of the aorta (see **Fig. 23b**). The TAWSS CoV appears to be significant for small branches (see **Fig. 23c**), and is equal to 0.06 in the aneurysm region and 0.02 in the proximal regions of the aorta. Similarly, the velocity magnitude CoVs is found to be 0.04 within the aneurysm and 0.02 in the upstream aorta. This illustrates how diseased regions like aneurysms can lead to increased




Fig. 19. PCA modes overlaid on mean abdominal aortic aneurysm model geometry. 230 modes captures 99% of the variance.

geometric uncertainty, most likely due to increased ambiguity of the vessel lumen shape and increased surrounding noise sources in the input image volume, leading to higher neural network output variability. We note again that this is directly learned from the image volume by the neural network.

4.4. Left coronary artery model

Even for the smallest coronary arteries and in the presence of significant surrounding heart tissue, the vessel lumen shape was qualitatively well captured by the proposed dropout network (see Fig. 24a). A PCA quantification of segmentation uncertainty shows dominant modes distributed along major arteries, and higher modes inducing local changes to smaller branches (see Fig. 25a). While the relative variance is higher in the small vessel branches when compared to the larger branches, PCA determines modes based on absolute variance. The absolute variance is larger in the large vessels and explains their presence in the dominant PCA modes.

A satisfactory convergence is observed for the Monte Carlo statistical moments after 110 model evaluations, with asymptotic traces for more than 50 samples (see Fig. 26). The 95% relative confidence intervals for the Monte




Fig. 20. Monte Carlo convergence — Abdominal Aortic Aneurysm case.





Fig. 21. Outlet QoIs and $\pm 2\sigma$ interval for abdominal aortic aneurysm model.

Table 2

Monte Carlo sample mean, coefficient of variation (CoV) and 95% relative confidence interval for all QoIs in abdominal aortic aneurysm model. n indicates the number of cross-sectional slices for the associated vessel.

Path	Aorta ($n = 160$)	Celiac hepatic ($n = 30$)	Celiac splenic ($n = 69$)	Ext. iliac left ($n = 169$)	Renal left ($n = 51$)	Renal right ($n = 35$)	SMA ($n = 61$)
Radius mean [cm]	0.63	0.26	0.36	0.38	0.32	0.29	0.38
Radius CoV	0.007	0.018	0.016	0.019	0.018	0.021	0.015
Radius conf.	0.0012	0.0032	0.0028	0.0035	0.0033	0.0038	0.0027
Pressure mean [mmHg]	96.45	89.11	94.51	94.35	95.29	92.35	99.26
Pressure CoV	0.004	0.004	0.004	0.004	0.004	0.005	0.004
Pressure conf.	0.0007	0.0008	0.0008	0.0006	0.0007	0.0008	0.0008
TAWSS mean [dyne/cm ²]	47.66	125.70	85.11	58.47	87.28	117.50	35.35
TAWSS CoV	0.024	0.040	0.034	0.035	0.036	0.043	0.034
TAWSS conf.	0.0044	0.0073	0.0062	0.0064	0.0065	0.0079	0.0061
Velocity mean [cm/s]	38.05	79.84	52.21	43.88	51.43	66.67	31.98
Velocity CoV	0.015	0.036	0.020	0.027	0.019	0.028	0.024
Velocity conf.	0.0028	0.0065	0.0036	0.0048	0.0034	0.0052	0.0043

**Fig. 22.** Time averaged QoIs and $\pm 2\sigma$ interval for abdominal aortic aneurysm model, plotted along the vessel centerline.

Carlo estimates of the mean are approximately equal to 0.4%, 3% and 1.5% for pressure, TAWSS and velocity magnitude, respectively, and are well below the CoV computed for the same QoIs (see [Tables 3](#)).

The outlet time histories reflect the diastolic nature of the coronary flow (see [Figs. 27a–27f](#)), where the WSS exhibits the largest uncertainty followed by velocity magnitude and pressure. Unlike the other two anatomies considered in the previous sections, geometrical uncertainty significantly affects hemodynamic model outputs due to smaller vessel sizes, as previously suggested in the literature in the context of coronary artery disease [\[28\]](#). Time averaged quantities over the vessel length show a similar pattern (see [Figs. 28a](#) and [28f](#)). In particular, TAWSS and velocity uncertainty appear to be very similar and correlated with the vessel radius. Specifically, smaller radii (with







Fig. 23. Nearest neighbor interpolation of cross-sectional time-averaged CoVs for abdominal aortic aneurysm model.



(a) Left Coronary Artery case - $LAD - D_1$



(b) Radius CoV



(c) Radius standard deviation

Fig. 24. Lumen segmentation samples and radius CoV/standard deviation for left coronary artery test case, computed over cross-sectional slices $\mathbf{x}_i$, $i = 1, \dots, 222$.

larger radius variability) produce larger TAWSS and velocity uncertainty. A different behavior is instead observed for the $LCx - OM_3$ branch, where large TAWSS and velocity uncertainty are associated with a larger radius. However, this phenomenon is localized at the proximal end of the vessel and probably triggered by the bifurcation nearby.

Pressure, TAWSS and velocity magnitude CoVs were approximately equal to 2%, 10%–20% and 6%–15%, respectively (see Table 3). Notably, the $LAD - D_1$, $LCx - OM_2$, and $LCx - OM_3$ branches showed larger TAWSS and velocity CoVs, equal to 16%, 24.5% and 92.9%, still explained by the small range of vessel sizes




Fig. 25. PCA modes overlaid on mean left coronary artery model geometry. 209 modes captures 99% of the variance.




Fig. 26. Monte Carlo moment traces for left coronary artery model QoIs.

Table 3

Monte Carlo sample mean, coefficient of variation (CoV) and 95% relative confidence interval for all QoIs in left coronary artery model. n indicates the number of cross-sectional slices for the associated vessel.

Path	LCx ($n = 29$)	LCx-OM ₁ ($n = 58$)	LCx-OM ₂ ($n = 20$)	LCx-OM ₃ ($n = 12$)	LAD ($n = 48$)	LAD-D ₁ ($n = 55$)
Radius mean [cm]	0.14	0.12	0.09	0.10	0.13	0.11
Radius CoV	0.032	0.031	0.047	0.046	0.034	0.034
Radius conf.	0.0061	0.0059	0.0090	0.0088	0.0064	0.0064
Pressure mean [mmHg]	93.46	92.37	91.78	91.92	96.02	95.14
Pressure CoV	0.021	0.020	0.031	0.016	0.018	0.018
Pressure conf.	0.0040	0.0037	0.0059	0.0031	0.0034	0.0034
TAWSS mean [dyne/cm ²]	39.27	48.43	26.37	11.91	12.78	26.17
TAWSS CoV	0.106	0.108	0.142	0.220	0.114	0.133
TAWSS conf.	0.0201	0.0206	0.0269	0.0418	0.0217	0.0252
Velocity mean [cm/s]	21.91	28.04	13.69	6.66	8.56	15.50
Velocity CoV	0.069	0.067	0.091	0.158	0.075	0.084
Velocity conf.	0.0131	0.0126	0.0172	0.0299	0.0143	0.0160

present in the coronary anatomy, which amplifies the segmentation uncertainty, particularly towards the distal end of each vessel (see Table 3 and Fig. 29a). The pressure CoV appears to be larger in the LCx and its branches LCx – OM₁ and particularly LCx – OM₂. This is explained by the relatively small radius of such vessels, which




Fig. 27. Outlet QoIs and $\pm 2\sigma$ interval for left coronary artery model.




Fig. 28. Time averaged QoIs and $\pm 2\sigma$ interval for left coronary artery model, plotted along the vessel centerline.

increases resistance and amplifies geometric uncertainty and by the variations in flow split caused by the relatively smaller LCx branching off the LAD trunk (see Fig. 29b). Finally, the largest WSS and velocity CoVs are located near bifurcations, due to higher local flow variability and more ambiguity in the definition of the vessel lumen (see Figs. 29c and 29d).




Fig. 29. Nearest neighbor interpolation of cross-sectional time-averaged CoVs for left coronary artery model.

5. Discussion

Our experiments in the previous sections illustrate how the proposed Bayesian dropout network generates vessel lumen segmentations characterized by a σ_r between 0.005 cm and 0.01 cm. This translates into radius CoVs in the range 1%–5%, where smaller vessel sizes are associated with larger radius variability, as expected. Furthermore, our network generated models by producing independent 2D vessel lumen segmentations along vessel pathlines are thus not guaranteed to result in correlated variation of the vessel lumen in the models. However, PCA modes of the model samples showed that this resulted in correlated geometric perturbations along entire vessel for the dominant modes. Later modes adjusted finer-scale features of the model such as aneurysms regions of specific vessels. Thus our cardiovascular model sampling process resulted in realistic whole model variation as observed from the PCA modes across model realizations. In particular the realistic model variation generated by our dropout sampling process was highlighted when comparing to PCA modes of models generated without using the same Bernoulli dropout vector for all lumens. When using independent Bernoulli dropout vectors for each vessel lumen, the PCA modes were visibly more random in a nature and did not show localized effects to specific features of the vessel geometry.

Amongst the output QoIs considered in this study, wall shear stress was the most impacted by geometry uncertainty, followed by velocity magnitude and then pressure. Wall shear stress CoV ranged from 3% to 20%, whereas velocity magnitude and pressure resulted in CoV ranges equal to 1.4%–15% and 0.2%–3%, respectively. Larger variability is observed in the left coronary artery model, due to the prevalence of vessels with small radius.

To quantify the relative importance of geometric uncertainty with respect to other sources relevant in hemodynamic simulations, we compare the output variability found in our study to those found in studies investigating other sources of uncertainty. Uncertainty due to coronary pressure waveform, intramyocardial pressure, morphometry exponent and vascular wall Young's modulus was recently investigated in [49], in the context of coronary artery modeling. Coronary pressure waveform uncertainty resulted in a 7% CoV for the average pressure and <7% for TAWSS and velocity magnitude. Intramyocardial pressure uncertainty produced CoV of roughly 25% for TAWSS and velocity magnitude. Morphometry exponent uncertainty resulted in a 2% CoV for velocity magnitude and negligible impact on TAWSS magnitude and pressure. Finally, vascular wall Young's modulus uncertainty had negligible impact on the hemodynamics, as suggested in the literature [7]. The CoV of 1.4%–15% and 3%–20% for velocity magnitude and TAWSS produced by geometric uncertainty in this work is thus comparable to the CoV due to coronary pressure and intramyocardial pressure uncertainty, but larger than the CoV due to morphometry exponent and vascular wall Young's modulus. The pressure CoV of 0.2%–3% due to geometry uncertainty was smaller than that produced by intramyocardial pressure uncertainty.

A multifidelity uncertainty quantification approach was used in [12] to investigate simulation uncertainty due to material property and boundary condition uncertainty, for both healthy and diseased aortic and coronary anatomies. The uncertain parameters were uniformly distributed with $\pm 30\%$ variation around their means. Resulting CoV were $\sim 3\%$ for pressure, $< 1\%$ for velocity magnitude and $\sim 10\%$ – 20% for TAWSS, regardless of model and disease

Table 4

Comparison between the hemodynamic effect geometry uncertainty and other sources of uncertainty.

Model type	Coronary			
Uncertainty source	Pressure waveform	Intramyocardial pressure	Morphometry exponent	Wall Young's modulus
Input distribution	Uniform	Uniform	Uniform	Uniform
Reference	[49]	[49]	[49]	[49]
Pressure CoV	7%	Negligible	Negligible	Negligible
TAWSS CoV	<7%	25%	Negligible	Negligible
Velocity CoV	<7%	25%	2%	Negligible
Model type	Aorta/Coronary	Coronary bypass graft	Intracranial/Coronary	Aorta/AAA/Coronary
Uncertainty source	Boundary conditions	Multiple clinical targets	Boundary conditions	Full model geometry
Input distribution	Uniform $\pm 30\%$	Assimilated from clinical data	Gaussian 10% std	Dropout sampling
Reference	[12]	[17]	[56]	<i>This paper</i>
Pressure CoV	$\sim 3\%$	–	6%–9%	0.2–3%
TAWSS CoV	$\sim 10\%$ – 20%	5%–10%	5%–20%	3–20%
Velocity CoV	<1%	–	–	1.4%–15%

condition. Simulations in a subset of intracranial and coronary vessels, applying boundary conditions with Gaussian uncertainty and an efficient network based uncertainty propagation algorithm found a pressure CoV of 6%–9% and WSS CoV ranging from 5%–20% [56]. In particular uncertainty in flow parameters was found to decrease with further distance from the inlets and outlets of the models. These TAWSS CoV values were comparable to the CoVs produced by geometric uncertainty as simulated in this work. However the velocity magnitude CoV was smaller than the 1.4%–15% CoV produced by geometric uncertainty sampled from our dropout network, whereas the pressure CoV induced by the same geometric uncertainty was typically smaller than that produced by the aforementioned uniform and Gaussian boundary condition variability.

The effect of variability in closed-loop boundary conditions assimilated from uncertain clinical data including aortic pressure, cardiac output, pulmonary pressure, peak and total flow volumes is investigated in the context of coronary artery disease in [17]. Uncertainties in a number of independent clinical measurements were assumed to be normally distributed with standard deviation in the range 10%–40% of the corresponding (measured) mean value, and the boundary conditions parameters learned using a Bayesian parameter tuning framework, based on adaptive Markov chain Monte Carlo sampling. The results show CoVs for TAWSS in the range of 5% to 10%, which is comparable or smaller than the 3%–20% CoV observed for TAWSS in this work.

The results of the three test cases discussed in the previous sections seem to suggest, on the one hand, that geometrical uncertainty has a generally limited impact on hemodynamics compared to other sources of uncertainty, and, in practice, could be disregarded if pressure is the sole QoI. On the other hand, the velocity magnitude and TAWSS variability for the left coronary artery model were found to be approximately 3 to 5 times higher compared to the other two anatomies. This suggests that geometric uncertainty might play a dominant role for anatomies characterized by small vessel sizes (the coronary circulation is a particularly relevant case), especially for stenotic lesions, associated with substantial radius uncertainty (see Table 4).

Even though this is the first systematic study in the literature combining machine learning and high-fidelity cardiovascular models to study the effect of geometric uncertainty, we recognize several limitations. First, uncertainty propagation is performed using standard Monte Carlo sampling. Even though a number of approaches in the literature have shown promise to accelerate convergence to the true statistical moments (such as, e.g., stochastic collocation [3], or generalized multiresolution expansions [6] in the context of cardiovascular flow), the Bernoulli random vectors used in the dropout layer in this study have a dimension of around 10,000, which is extremely challenging for approaches based on stochastic spectral expansion. Second, our study includes only one diseased anatomy, and showed that geometrical uncertainty is amplified in the aneurysm region of the abdominal aorta. We also found such amplification in a healthy left coronary artery model due to the typically smaller vessel radii. This suggests how cases of stenosed or calcified coronary arteries and vascular lesions, that typically are a minority of network training samples, may offer new insights on the role of geometric uncertainty.

If uncertainty related to diseased anatomy is of particular interest it is likely the proportion of network training data containing pathological features would need to be increased. Additionally the network's segmentation accuracy in diseased regions would need to be compared that of human experts in those same regions.

Third, we assume that dropout networks are able to learn output uncertainty from their training data [29], while, in practice, they provide only an approximate representation of the true distribution of vessel lumen for a given image. Additionally, our segmentations are generated at discrete slices along the centerline path, and thus the training data might act as a filter on the whole geometric variability, leading the proposed algorithm to underestimate the true underlying geometric uncertainty. Removing these limitations would require new three-dimensional vessel segmentation paradigms, which is an active area of research. Fourth, the path planning approach intrinsically limits the uncertainty at the bifurcations, as it requires users to adjust pathlines so they originate within the parent vessel. This introduces user bias into the cardiovascular model samples and may constrain the underlying geometric uncertainty. Future work will be devoted to produce improved estimates for bifurcation uncertainty.

6. Conclusions

We have developed a Bayesian dropout network to generate families of two-dimensional lumen segmentations from slices of a clinically acquired image volume. Of particular note is the fact that our neural network learns lumen segmentation uncertainty directly from the image training data and is thus able to generate lumen samples with a realistic uncertainty distribution. This was combined with vessel centerlines and a path-planning model building workflow to create realizations of high-fidelity cardiovascular models with uncertain lumen surface. Finally, we characterized simulation output variability due to geometric uncertainty using Monte Carlo sampling.

We have also analyzed the principal components of the lumen surfaces we generated, showing how, despite segmenting slices independently and analyzing geometries characterized by a wide range of vessel radii, dominant modes appear to be equally distributed on the entire model, without amplifying any particular local feature. Additionally, our network generated vessel lumens with relatively constant radius standard deviation that were found to be independent of vessel size. This resulted in increasing relative uncertainty for smaller vessels, similar to manual segmentations generated by expert users.

Experiments on an aortic bifurcation model, an abdominal aortic aneurysm model and a left coronary artery model showed that geometry uncertainty primarily resulted in wall shear stress and velocity magnitude uncertainty. This was true in particular for the coronary anatomy, characterized by smaller vessel sizes. Moreover, while TAWSS and velocity magnitude were impacted by geometrical uncertainty, especially near the distal ends of small vessels and near bifurcations, pressure was only marginally affected. Compared to other sources of uncertainty, for example, the boundary conditions or material properties, the relative importance of geometry uncertainty was found to be determined by the particular patient-specific geometry being investigated and the vessel radius.

Our method still requires one to manually create the vessel centerlines, which may be laborious, time consuming and introduce additional uncertainty. Automated methods to predict vessel centerlines and corresponding uncertainty could therefore be combined with the proposed dropout network to further improve performance and increase model building efficiency. Finally, we have explored geometry uncertainty independently from other sources, disregarding their interaction. Future work will be devoted to combine all three sources of simulation uncertainty, i.e., boundary conditions, material properties and geometry.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was funded by the National Science Foundation (NSF), USA Software Infrastructure for Sustained Innovation (SSI) grants #1663671 (PI ALM) and #1562450 (PI ALM), Computational and Data-Enabled Science and Engineering (CDS&E), USA grant #1508794, NSF, USA CAREER award #1942662 (PI DES), National Institute of Health, USA grants R01HL123689 (PI ALM), R01EB018302 (PI ALM) and R01LM013120 (PI ALM) and the American Heart Association Precision Medicine Platform (PI ALM).

References

- [1] V.G. Eck, W.P. Donders, J. Sturdy, J. Feinberg, T. Delhaas, L.R. Hellevik, W. Huberts, A guide to uncertainty quantification and sensitivity analysis for cardiovascular applications, *Int. J. Numer. Methods Biomed. Eng.* 32 (2015).
- [2] D.A. Steinman, F. Migliavacca, Editorial: Special issue on verification, validation, and uncertainty quantification of cardiovascular models: Towards effective VVUQ for translating cardiovascular modelling to clinical utilities, *Cardiovasc. Eng. Technol.* 9 (2018).
- [3] S. Sankaran, A.L. Marsden, A stochastic collocation method for uncertainty quantification and propagation in cardiovascular simulations, *J. Biomech. Eng.* 133 (3) (2011) 1–12, <http://dx.doi.org/10.1115/1.4003259>.
- [4] J. Biehler, M.W. Gee, W.A. Wall, Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a Bayesian multi-fidelity scheme, *Biomech. Model. Mechanobiol.* 14 (2014).
- [5] S. Quicken, W.P. Donders, M.J. van Disseldorp, K. Gashi, B.M.E. Mees, F.N. van de Vosse, R.G.P. Lopata, T. Delhaas, W. Huberts, Application of an adaptive polynomial chaos expansion on computationally expensive three-dimensional cardiovascular models for uncertainty quantification and sensitivity analysis, *J. Biomech. Eng.* 138 (2016).
- [6] D.E. Schiavazzi, A. Doostan, G. Iaccarino, A.L. Marsden, A generalized multi-resolution expansion for uncertainty propagation with application to cardiovascular modeling, *Comput. Methods Appl. Mech. Engrg.* 314 (2017) 196–221, <http://dx.doi.org/10.1016/j.cma.2016.09.024>.
- [7] J.S. Tran, D.E. Schiavazzi, A.M. Kahn, A.L. Marsden, Uncertainty quantification of simulated biomechanical stimuli in coronary artery bypass grafts, *Comput. Methods Appl. Mech. Engrg.* 345 (2019) 402–428, <http://dx.doi.org/10.1016/j.cma.2018.10.024>.
- [8] J. Seo, D. Schiavazzi, A. Kahn, A. Marsden, The effects of clinically-derived parametric data uncertainty in patient-specific coronary simulations with deformable walls, *Int. J. Numer. Methods Biomed. Eng.* n/a (n/a) (2020) e3351, <http://dx.doi.org/10.1002/cnm.3351>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/cnm.3351> URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.3351>.
- [9] M.J. Colebank, L.M. Paun, M.U. Qureshi, M. Chesler, D. Husmeier, M.S. Olufsen, L.E. Fix, Influence of image segmentation on one-dimensional fluid dynamics predictions in the mouse pulmonary arteries, *R. Soc. Interface* (2019).
- [10] M.J. Colebank, M.U. Qureshi, M.S. Olufsen, Sensitivity analysis and uncertainty quantification of 1-d models of pulmonary hemodynamics in mice under control and hypertensive conditions, *Int. J. Numer. Methods Biomed. Eng.* (2019).
- [11] A. Boccadifuoco, A. Mariotti, K. Capellini, S. Celi, M.V. Salvetti, Uncertainty quantification applied to hemodynamic simulations of thoracic aorta aneurysms: Sensitivity to inlet conditions, *Quantif. Uncertain.: Improv. Effic. Technol.* (2020).
- [12] C. Fleeter, G. Geraci, D. Schiavazzi, A. Kahn, A. Marsden, Multilevel and multifidelity uncertainty quantification for cardiovascular hemodynamics, *Comput. Methods Appl. Mech. Engrg.* 365 (2020) 113030.
- [13] G.R. Mirams, S.A. Niederer, R.H. Clayton, The fickle heart: Uncertainty quantification in cardiac and cardiovascular modeling and simulation, *Philos. Trans. R. Soc.* (2020).
- [14] J.O. Campos, J. Sundnes, R.W. dos Santos, B.M. Rocha, Uncertainty quantification and sensitivity analysis of left ventricular function during the full cardiac cycle, *Philos. Trans. R. Soc.* (2020).
- [15] R.H. Clayton, Y. Aboelkassem, C.D. Cantwell, C. Corrado, T. Delhaas, W. Huberts, C.L. Lei, H. Ni, A.V. Panfilov, C. Roney, R.W. dos Santos, An audit of uncertainty in multi-scale cardiac electrophysiology models, *Philos. Trans. R. Soc.* (2020).
- [16] D. Schiavazzi, G. Arbia, C. Baker, A. Hlavacek, T.-Y. Hsia, A. Marsden, I. Vignon-Clementel, T.M. of Congenital Hearts Alliance (MOCHA) Investigators, Uncertainty quantification in virtual surgery hemodynamics predictions for single ventricle palliation, *Int. J. Numer. Methods Biomed. Eng.* 32 (3) (2016) e02737.
- [17] J. Tran, D. Schiavazzi, A. Ramachandra Bangalore, A. Kahn, A. Marsden, Automated tuning for parameter identification and uncertainty quantification in multi-scale coronary simulations, *Comput. & Fluids* 142 (2017) 128–138, <http://dx.doi.org/10.1016/j.compfluid.2016.05.015>.
- [18] L.M. Paun, M.J. Colebank, M.S. Olufsen, N.A. Hill, D. Husmeier, Assessing model mismatch and model selection in a bayesian uncertainty quantification analysis of a fluid-dynamics model of pulmonary blood circulation, *J. R. Soc. Interface* (2020).
- [19] I.A. Sigal, M.R. Hardisty, C.M. Whyne, Mesh-morphing algorithms for specimen-specific finite element modeling, *J. Biomech.* 41 (2008).
- [20] M. Le, J. Unkelbach, N. Ayache, H. Delingette, Sampling image segmentations for uncertainty quantification, *Med. Image Anal.* 34 (2016).
- [21] L. Joskowicz, D. Cohen, N. Caplan, J. Sosna, Automatic segmentation variability estimation with segmentation priors, *Med. Image Anal.* 50 (2018).
- [22] A.M. Gambaruto, J. Janela, A. Moura, A. Sequeira, Sensitivity of hemodynamics in a patient specific cerebral aneurysm to vascular geometry and blood rheology, *Math. Biosci. Eng.* (2011).
- [23] P. Venugopal, X. Li, L. Cheng, R. Bhagalia, P.M. Edic, Sensitivity of FFR-CT to manual segmentation, *Proc. SPIE* (2019).
- [24] H. Xu, D. Baroli, A. Veneziani, Global sensitivity analysis for patient-specific aortic simulations: the role of geometry, boundary condition and LES modeling parameters, *Biomech. Eng.* (2020).
- [25] S. Sankaran, L. Grady, C.A. Taylor, Fast computation of hemodynamic sensitivity to lumen segmentation uncertainty, *IEEE Trans. Med. Imaging* 34 (2015).
- [26] S. Sankaran, L.J. Grad, C.A. Taylor, Real-time sensitivity analysis of blood flow simulations to lumen segmentation uncertainty, in: *Conference on Medical Imaging and Computer Assisted Intervention*, 2014.
- [27] S. Sankaran, L. Grady, C.A. Taylor, Impact of geometric uncertainty on hemodynamic simulations using machine learning, *Comput. Methods Appl. Mech. Eng.* 297 (2015).
- [28] S. Sankaran, H.J. Kim, G. Choi, C.A. Taylor, Uncertainty quantification in coronary blood flow simulations: impact of geometry, boundary conditions and blood viscosity, *J. Biomech.* 49 (2016).

- [29] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: representing model uncertainty in deep learning, in: International Conference on Machine Learning, 2016.
- [30] R. Jena, S.P. Awate, A Bayesian neural net to segment images with uncertainty estimates and good calibration, in: International Conference on Information Processing in Imaging, 2019.
- [31] T. Nair, D. Precup, D.L. Arnold, T. Arbel, Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation, *Med. Image Anal.* 59 (2020).
- [32] G. Maher, D. Parker, N. Wilson, A. Marsden, Neural network vessel lumen regression for automated segmentation in cardiovascular image-based modeling, *Cardiovasc. Eng. Technol.* (2020).
- [33] A. Updegrove, N. Wilson, J. Merkow, H. Lan, A. Marsden, S. Shadden, SimVascular: an open source pipeline for cardiovascular simulation, *Ann. Biomed. Eng.* 61 (2013).
- [34] N. Wilson, A. Ortiz, A. Johnson, The vascular model repository: a public resource of medical imaging data and blood flow simulation results, *J. Med. Devices* 7 (4) (2013).
- [35] A. Updegrove, N.M. Wilson, S.C. Shadden, Boolean and smoothing of discrete surfaces, *Adv. Eng. Softw.* 95 (2016).
- [36] L. Antiga, M. Piccinelli, L. Botti, B. Ene-Iordache, A. Remuzzi, D.A. Steinman, An image-based modeling framework for patient-specific computational hemodynamics, *Med. Biol. Eng. Comput.* 46 (2008).
- [37] D. Lesage, E.D. Angelini, I. Bloch, G. Funka-Lea, A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes, *Med. Image Anal.* 13 (2009).
- [38] H. Si, TetGen, a Delaunay-based quality tetrahedral mesh generator, *ACM Trans. Math. Softw.* 41 (2015).
- [39] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444, <http://dx.doi.org/10.1038/nature14539>.
- [40] J. Schmidhuber, Deep learning in neural networks: an overview, *Neural Netw.* 61 (2015) 85–117, <http://dx.doi.org/10.1016/j.neunet.2014.09.003>.
- [41] K. Hornik, Approximation capabilities of multilayer feedforward networks, *Neural Netw.* 4 (1991).
- [42] A.L. Maas, A.Y. Hannun, A.Y. Ng, Rectifier nonlinearities improve neural network acoustic models, ICML (2013).
- [43] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 1–9.
- [44] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* (2014).
- [45] K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: surpassing human-level performance on ImageNet classification, in: International Conference on Computer Vision, 2015.
- [46] D. Kingma, J. Ba, Adam: A method for stochastic optimization, in: International Conference on Learning Representations, 2015.
- [47] A.N. Brooks, T.J.R. Hughes, Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations, *Comput. Methods Appl. Mech. Eng.* (1982).
- [48] M.E. Moghadam, Y. Bazilevs, T.Y. Hsia, I.E. Vignon-Clemente, A.L. Marsden, M. Alliance for the Modeling of Congenital Hearts Alliance, A comparison of outlet boundary treatments for prevention of backflow divergence with relevance to blood flow simulations, *Comput. Mech.* 48 (2011).
- [49] J. Seo, D. Schiavazzi, A. Marsden, Performance of preconditioned iterative linear solvers for cardiovascular simulations in rigid and deformable vessels, *Comput. Mech.* 64 (3) (2019) 717–739.
- [50] M.E. Moghadam, I.E. Vignon-Clementel, R. Figliola, A.L. Marsden, M. Alliance for the Modeling of Congenital Hearts Alliance, A modular numerical method for implicit 0D/3D coupling in cardiovascular finite element simulations, *J. Comput. Phys.* 244 (2013).
- [51] M.E. Moghadam, Y. Bazilevs, A.L. Marsden, A bi-partitioned iterative algorithm for solving linear systems arising from incompressible flow problems, *Comput. Methods Appl. Mech. Eng.* 286 (2015).
- [52] K.E. Jansen, C.H. Whiting, G.M. Hulbert, A generalized- α method for integrating the filtered Navier-Stokes equations with a stabilized finite element method, *Comput. Methods Appl. Mech. Eng.* 190 (2000).
- [53] I.E. Vignon-Clementel, C.A. Figueroa, K.E. Jansen, C.A. Taylor, Outflow boundary conditions for 3D simulations of non-periodic blood flow and pressure fields in deformable arteries, *Comput. Methods Biomech. Biomed. Eng.* 13 (2009).
- [54] H.J. Kim, I.E. Vignon-Clementel, J.S. Coogan, C.A. Figueroa, K.E. Jansen, C.A. Taylor, Patient-specific modeling of blood flow and pressure in human coronary arteries, *Ann. Biomed. Eng.* 38 (2010).
- [55] P. Baldi, P. Sadowski, Understanding dropout, *Adv. Neural Inf. Process. Syst.* 26 (2013) 2814–2822.
- [56] S. Guzzetti, L.A.M. Alvarez, P.J. Blanco, K.T. Carlberg, A. Veneziani, Propagating uncertainties in large-scale hemodynamics models via network uncertainty quantification and reduced-order modeling, *Comput. Methods Appl. Mech. Eng.* 358 (2020).