

HIERARCHICAL LOG BAYESIAN NEURAL NETWORK FOR ENHANCED AORTA SEGMENTATION

*Delin An*¹ *Pan Du*² *Pengfei Gu*³ *Jian-Xun Wang*^{2,4} *Chaoli Wang*¹

¹ Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN

² Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, IN

³ Department of Computer Science, University of Texas Rio Grande Valley, Edinburg, TX

⁴ Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY

ABSTRACT

Accurate segmentation of the aorta and its associated arch branches is crucial for diagnosing aortic diseases. While deep learning techniques have significantly improved aorta segmentation, they remain challenging due to the intricate multiscale structure and the complexity of the surrounding tissues. This paper presents a novel approach for enhancing aorta segmentation using a Bayesian neural network-based hierarchical Laplacian of Gaussian (LoG) model. Our model consists of a 3D U-Net stream and a hierarchical LoG stream: the former provides an initial aorta segmentation, and the latter enhances blood vessel detection across varying scales by learning suitable LoG kernels, enabling self-adaptive handling of different parts of the aorta vessels with significant scale differences. We employ a Bayesian method to parameterize the LoG stream and provide confidence intervals for the segmentation results, ensuring robustness and reliability of the prediction for vascular medical image analysts. Experimental results show that our model can accurately segment main and supra-aortic vessels, yielding at least a 3% gain in the Dice coefficient over state-of-the-art methods across multiple volumes drawn from two aorta datasets, and can provide reliable confidence intervals for different parts of the aorta. The code is available at <https://github.com/adlsn/LoGBNet>.

1. INTRODUCTION

Semantic segmentation is a crucial research direction in deep learning technologies, extensively applied in medical image processing. It finds wide-ranging applications in organ segmentation, lesion detection, tumor segmentation, and vessel segmentation. Vessel segmentation includes coronary artery, aorta, and pulmonary artery segmentation. Among these, aorta segmentation presents unique challenges. While U-Net [1] has successfully segmented other blood vessels, the aorta poses specific difficulties due to its characteristics. Other methods, like U-Net 3D [2], are tailored for general abdominal organs but fall short of effectively segmenting

the aorta and arch branches. Moreover, providing reliable confidence intervals for aorta segmentation results is challenging. Uncertainty quantification (UQ) offers analysts a basis for assessing the reliability of current predictions and significantly impacting downstream tasks involving the aorta. For instance, in computational fluid dynamics (CFD) simulations of the aorta, UQ helps verify and calibrate CFD models, ensuring they accurately reflect physiological conditions.

The aorta is conformed by the aortic arch, ascending aorta, descending aorta, and numerous branches. These components exhibit significant variations in size and shape across individuals. Following [3], the main arch and branches of the aorta are referred to as the main aorta (MA) and supra-aorta branches (SA), respectively. Additionally, aorta segmentation requires higher computational power due to the need for high-resolution imaging. Traditional aorta segmentation methods involve manual centerline tracking and boundary filtering [4]. Deep learning models like U-Net reduce manual work, using encoders such as 3D CNNs, transformers, and attention-based models [5]. Some use labeled centerlines to improve performance [6], but obtaining them is challenging. Thus, automatic centerline extraction is preferred [7]. These methods perform well on the MA but have limitations: they neglect the SA, and incorporating centerlines increases model complexity and training difficulty.

We propose LoGB-Net, a Laplacian of Gaussian (LoG)-based Bayesian network, for aorta segmentation, enhancing multiscale vessel recognition (especially for SA) and improving model automation. LoGB-Net features two streams: a U-Net-based regular stream for overall aorta segmentation and a LoG stream with a Bayesian LoG module and balanced gate to refine local details. Our key contributions include

- A hierarchical LoG module is introduced to improve the performance of aorta segmentation. This design enhances features of aortic vessels of various sizes in a self-adaptive manner, effectively emphasizing multiscale features and strengthening the model's segmentation capability.
- An imbalance problem between the aorta's foreground and background is addressed by designing a balanced gate

in the LoG stream. This gate balances the ratio between multiscale blood vessels and prevents the model from overfitting, ensuring improved segmentation accuracy.

- The LoG module is parameterized to obtain the UQ for the segmentation results. This approach provides confidence intervals for the segmentation results, allowing for measuring the uncertainty in the model’s predictions, reflecting the prediction’s reliability in cases where the edges of the aorta are ambiguous, and offering UQ for downstream tasks in aortic analysis.

2. RELATED WORK

Medical image semantic segmentation methods are semi-automatic and automatic [8]. Semi-automatic methods involve manual lumen boundary and centerline annotation, achieving good performance but are labor-intensive. SimVascular [4] reduces manual effort by combining centerline and lumen edge for segmentation but still requires annotations for new data. To improve efficiency, machine learning methods, such as [9], use Gaussian filters and Fourier transforms to identify regions of interest. However, they are noise-sensitive and require parameter tuning for different object sizes.

Deep learning has shifted segmentation towards automation, with methods categorized based on centerline usage: with and without centerline. While centerlines improve performance, annotating them is labor-intensive, leading to a focus on the direct segmentation of raw images. Encoder-decoder models like U-Net, PSPNet [10], FC-DenseNet [11], and kCBAC-Net [12] are widely used. U-Net excels in medical image processing tasks. Transformer-based models like UNETR [13], UNETR++ [14], Swin-UNETR [15], and ConVFormer [16] integrate attention mechanisms, performing well in blood vessel segmentation with consistent scale but struggle with the aorta and supra branches due to significant scale variation. As a dual-stream model, Gated-SCNN [17] combines a regular and a shape stream leveraging image gradients, achieving favorable results in small object detection. This model guides SA-enhancing architectures for aorta segmentation. Furthermore, UQ provides confidence intervals for segmentation predictions, enhancing reliability in aortic research. Existing UQ methods often increase complexity by integrating VAE encoders into U-Net architectures. [18] evaluated these approaches on brain tumor datasets.

3. LOGB-NET

Multiscale aortic vessels require attention to the overall structure and local details, and the imbalanced foreground-background ratio complicates training. We propose a hierarchical LoG module to address these issues and enhance aorta segmentation. Our framework (Figure 1) comprises regular and LoG streams. The regular stream employs a 3D U-Net [2] to segment the overall aorta structure. The LoG




Fig. 1. The top shows the flow diagram of LoGB-Net. The Bayesian segmentation results’ blue, red, and green lines represent the segmentation’s upper bound, mean, and lower bound. The bottom displays the application of UQ in aorta segmentation for downstream tasks, such as CFD simulation.

stream consists of a Bayesian hierarchical LoG module and a balanced gate, augmenting the segmentation with local details. The balanced gate contains two queues Q_1 and Q_2 to store multiscale vessel images. When these queues are full, the gate will open for the regular and LoG streams.

Hierarchical LoG module. Encoder-decoder models effectively segment vessels with manageable complexity [14, 19], achieving high-quality results for retinal [20] and liver [21] vessels. However, further enhancement is required for the SA of the aorta to achieve satisfactory segmentation results. The LoG stream can strengthen aorta features by self-adapting to multiscale vessel characteristics.

The LoG model, used by [19], is a second-order derivative operator for edge detection in images. Specifically, the 2D LoG model is defined as $\nabla^2 G(x, y) = \frac{\partial^2 G(x, y)}{\partial x^2} + \frac{\partial^2 G(x, y)}{\partial y^2}$, where $G(x, y)$ is the Gaussian function: $G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$. Since the Laplacian is sensitive to noise, applying Gaussian smoothing is necessary. In 3D CTA images, we adopt the 3D LoG model, defined as $\nabla^2 G(x, y, z) = \frac{\partial^2 G(x, y, z)}{\partial x^2} + \frac{\partial^2 G(x, y, z)}{\partial y^2} + \frac{\partial^2 G(x, y, z)}{\partial z^2}$. This continuous function is transformed into a 3D convolutional kernel through discretization. The standard deviation σ and kernel size influence the detection size of target objects. [9] tested various parameters to fit retinal vessels, increasing manual operations.

To leverage the LoG model’s self-adaptive capability for main and supra aorta detection, we design a hierarchical LoG stream, as shown in Figure 1. This module comprises five 3D convolutional layers with kernel sizes 3, 5, 7, 9, and 11, initialized with 3D LoG kernels at different trainable σ values (0.5 to 2.5). Unlike the original LoG method with fixed parameters, we make them trainable in our approach. The initial σ values are optimized to achieve the best lumen boundaries. The module’s multiple layers facilitate the self-adaptive enhancement of multiscale aortic vessels.

Bayesian method and balanced gate. The low resolu-

Table 1. Quantitative results (reported as mean $\pm$ standard deviation) of different methods for the testing set. The best results are highlighted in bold. The experiment was conducted with varying random seeds five times to evaluate the outcomes.

Type	Metric	CNN-based				Attention-based					Ours	
		FPN	U-Net 3D	PSPNet	nnUNet	Attention-UNet	MISSFormer	Swin-UNETR	TransUNet	UNETR	UNETR++	LoGB-Net
SA	Dice $\uparrow$	0.722 $\pm$ 0.023	0.764 $\pm$ 0.042	0.775 $\pm$ 0.037	0.766 $\pm$ 0.068	0.824 $\pm$ 0.025	0.882 $\pm$ 0.021	0.896 $\pm$ 0.019	0.823 $\pm$ 0.022	0.878 $\pm$ 0.031	0.892 $\pm$ 0.037	0.927$\pm$0.011
	ASD $\downarrow$	1.523 $\pm$ 0.323	1.376 $\pm$ 0.311	1.329 $\pm$ 0.304	1.214 $\pm$ 0.322	0.929 $\pm$ 0.213	0.813 $\pm$ 0.196	0.786 $\pm$ 0.249	0.928 $\pm$ 0.284	0.702 $\pm$ 0.167	0.682 $\pm$ 0.156	0.678$\pm$0.168
	Hausdorff $\downarrow$	8.490 $\pm$ 3.252	8.323 $\pm$ 4.022	9.765 $\pm$ 3.773	8.941 $\pm$ 2.547	6.725 $\pm$ 0.294	6.899 $\pm$ 0.077	6.324 $\pm$ 0.133	6.452 $\pm$ 0.236	6.139 $\pm$ 0.698	6.034$\pm$0.356	6.225 $\pm$ 0.957
MA	Dice $\uparrow$	0.734 $\pm$ 0.016	0.773 $\pm$ 0.026	0.778 $\pm$ 0.034	0.775 $\pm$ 0.028	0.847 $\pm$ 0.021	0.899 $\pm$ 0.016	0.909 $\pm$ 0.033	0.904 $\pm$ 0.023	0.904 $\pm$ 0.036	0.907 $\pm$ 0.023	0.937$\pm$0.006
	ASD $\downarrow$	1.454 $\pm$ 0.364	1.342 $\pm$ 0.344	1.335 $\pm$ 0.313	1.397 $\pm$ 0.311	0.879 $\pm$ 0.225	0.923 $\pm$ 0.233	0.804 $\pm$ 0.218	0.931 $\pm$ 0.314	0.688 $\pm$ 0.199	0.744 $\pm$ 0.194	0.682$\pm$0.197
	Hausdorff $\downarrow$	10.245 $\pm$ 3.779	11.556 $\pm$ 3.929	9.987 $\pm$ 4.201	9.294 $\pm$ 2.326	6.942 $\pm$ 0.309	6.125$\pm$0.191	7.021 $\pm$ 0.065	6.878 $\pm$ 0.254	6.422 $\pm$ 0.453	6.321 $\pm$ 0.241	6.322 $\pm$ 1.485

tion of CTA images makes lumen edge segmentation challenging due to blurry boundaries. We address this by using a Bayesian method [22] to parameterize the LoG module, treating the LoG kernel as the prior distribution and the module’s output as the posterior, defined as: $p(\theta|D) = \frac{p(D|\theta)p(\theta)}{p(D)}$, where θ are the LoG module parameters and D is the image. The optimization aims to maximize the Evidence Lower Bound (ELBO): $\theta = \text{argmax}(\text{ELBO})$. Our overall loss function combines Dice loss and ELBO loss: $\mathbb{L} = \text{Dice} - \text{ELBO}$. We decrease $\mathbb{L}$ to maximize the ELBO value. We compared the LoG module with and without Bayesian parameterization in the ablation study. Our results demonstrate that the Bayesian parameterized LoG module produces more stable predictions.

Since the aorta occupies only a small section of the CTA image, there is an imbalance between the foreground (aorta) and the background. To address this, we propose a balanced gate to achieve foreground-background balance, defined as: $T = \frac{\sum_{i=1}^{N_F} f_i}{\sum_{j=1}^{N_B} b_j}$, where T denotes the foreground-to-background ratio. In Figure 1, the threshold μ is determined statistically by calculating the mean of the foreground-to-background ratio for the SA. Then, we introduce two counters, c_1 and c_2 , associated with two queues, Q_1 and Q_2 , each with a capacity of C . When the foreground area occupied by blood vessels is small, indicated by $T < \mu$ and $c_1 < C$, we enqueue the input image into Q_1 . Conversely, when the foreground area is large, indicated by $T > \mu$ and $c_2 < C$, we enqueue the input image into Q_2 . Once both c_1 and c_2 reach the capacity C , we concatenate the $2C$ images along the channel dimension and initiate the regular stream. In our experiments, we set $\mu = 0.15$. This balanced gate ensures a more balanced ratio between large and small foreground regions, enhancing the model’s robustness.

Algorithm Description. The overall pipeline follows the steps below. (1) Input the CTA image into the LoG stream and calculate the mean of the output from the five layers. (2) Balance the input data ratio using the balanced gate. When counters c_1 and c_2 reach C , concatenate the $2C$ images along the channel dimension. (3) The outputs from the regular and LoG streams are concatenated and fed into a four-layer atrous spatial pyramid pooling (ASPP) module for further processing. (4) Process the concatenated result using a 3D CNN with a $1 \times 1 \times 1$ kernel.

4. EXPERIMENTS

Datasets, network training, and metrics. Two datasets were used: the first is from [23], and the second is the Aortic Vessel Tree (AVT) CTA dataset [24]. The training set contains 34 volumes (24 from the first, 10 from the second), and the testing set has 16 volumes (8 from each). Volumes were re-sampled using MONAI’s Spacingd function to $0.8 \times 0.8 \times 0.3$ spacing, cropped into $64 \times 64 \times 64$ blocks, and normalized to $[0, 1]$. Data augmentation included RandFlipd and RandRotate90d (probability 0.1), applied only to training data. We used the Adam optimizer. The hyperparameters are set as $\beta_1 = 0.8$ and $\beta_2 = 0.999$, while the penalty weight is set to 10^{-5} . The learning rate is 10^{-4} . The model is trained for 5000 epochs on an NVIDIA RTX 3080 GPU. We adopted three metrics: Dice coefficient, average surface distance (ASD), and symmetric Hausdorff distance.

UQ. We assessed uncertainty due to varying boundary clarity and annotation errors by calculating confidence intervals from ten Bayesian inferences (Figure 2). Blurred boundaries led to larger confidence intervals.




Fig. 2. Confidence intervals for aortic predictions based on boundary clarity variations.

Evaluation results. We compared LoGB-Net with CNN-based (FPN [25], U-Net 3D [2], PSPNet [10], nnUNet [26]) and attention-based (Attention-UNet [5], MISSFormer [27], Swin-UNETR [15], TransUNet [28], UNETR [13], UNETR++ [14]) methods. LoGB-Net’s computational cost is 48.88 GFLOPS, modestly above UNETR’s 41.19 GFLOPS but significantly lower than nnUNet’s 412.65 GFLOPS. Using the same training and data augmentation for all models, Table 1 shows that LoGB-Net outperforms all baselines across various metrics, particularly achieving approximately 3% higher Dice coefficient than Swin-UNETR.

As illustrated in Figure 3, while most methods perform similarly on MA vessels, LoGB-Net outperforms others on SA vessels, detecting small-radius vessels and producing

Table 2. Quantitative results of the LoGB-Net’s ablation study. $\mathcal{L}_{LoG}^-$, $\mathcal{L}_{Bay}^-$, and $\mathcal{L}_{Gate}^-$ represent LoGB-Net without LoG module, Bayesian optimization, and balanced gate, respectively. $\mathcal{L}(1)$, $\mathcal{L}(2)$, $\mathcal{L}(3)$, $\mathcal{L}(4)$, and $\mathcal{L}(5)$ represent LoGB-Net with 1, 2, 3, 4, and 5 LoG layers, respectively.

	Metric	$\mathcal{L}_{LoG}^- (\mathcal{L}(0))$	$\mathcal{L}_{Bay}^-$	$\mathcal{L}_{Gate}^-$	$\mathcal{L}(1)$	$\mathcal{L}(2)$	$\mathcal{L}(3)$	$\mathcal{L}(4)$	$\mathcal{L}(5)$ (Ours)
SA	Dice $\uparrow$	0.735 $\pm$ 0.033	0.863 $\pm$ 0.031	0.896 $\pm$ 0.012	0.852 $\pm$ 0.033	0.863 $\pm$ 0.018	0.872 $\pm$ 0.011	0.893 $\pm$ 0.033	0.927$\pm$0.011
	ASD $\downarrow$	1.384 $\pm$ 0.334	0.810 $\pm$ 0.214	0.801 $\pm$ 0.179	1.291 $\pm$ 0.123	1.226 $\pm$ 0.202	0.922 $\pm$ 0.244	0.772 $\pm$ 0.202	0.678$\pm$0.168
	Hausdorff $\downarrow$	8.271 $\pm$ 2.276	7.993 $\pm$ 2.112	7.909 $\pm$ 1.093	8.203 $\pm$ 1.903	8.014 $\pm$ 1.221	7.027 $\pm$ 0.721	6.626 $\pm$ 0.121	6.225$\pm$0.957
MA	Dice $\uparrow$	0.741 $\pm$ 0.022	0.871 $\pm$ 0.033	0.882 $\pm$ 0.021	0.866 $\pm$ 0.014	0.865 $\pm$ 0.009	0.877 $\pm$ 0.036	0.923 $\pm$ 0.003	0.937$\pm$0.006
	ASD $\downarrow$	1.399 $\pm$ 0.123	0.961 $\pm$ 0.120	0.933 $\pm$ 0.200	1.105 $\pm$ 0.239	0.979 $\pm$ 0.103	0.933 $\pm$ 0.210	0.727 $\pm$ 0.121	0.682$\pm$0.197
	Hausdorff $\downarrow$	8.335 $\pm$ 2.996	8.220 $\pm$ 2.003	8.011 $\pm$ 1.898	8.277 $\pm$ 2.516	8.100 $\pm$ 2.274	7.965 $\pm$ 1.001	7.482 $\pm$ 1.112	6.322$\pm$1.485




Fig. 3. Qualitative results of different methods for one volume in the testing set (drawn from the first dataset). The top row displays the 3D segmentation prediction, where selected slices (S1 to S2) are marked on the ground-truth result.

smoother boundaries. The Bayesian segmentation results (Figure 1) further demonstrate LoGB-Net’s robustness, showing minimal deviation in the Dice metric (Table 1).




Fig. 4. CFD simulation results for the aortic vessel’s velocity, pressure, and wall shear stress. Each case represents a different UQ realization.

CFD simulation. We evaluated segmentation effectiveness by conducting CFD simulations on nine aorta shapes sampled from the posterior distribution. Using vmtk [29] for mesh generation and OpenFOAM [30] to solve the Navier-Stokes equations, Figure 4 compares simulation results among different UQ realizations. Results show significant flow variations, demonstrating sensitivity to segmented geometry and underscoring our model’s UQ capability.

Ablation study. An ablation study assessed each component’s effectiveness. Table 2 shows that the LoG module, Bayesian optimization, and balanced gate each contribute to performance enhancement. Figure 5 presents the qualitative results of the ablation study.




Fig. 5. Qualitative results of the LoGB-Net’s ablation study.

5. CONCLUSIONS AND FUTURE WORK

We have presented LoGB-Net, a hierarchical, Bayesian-optimized LoG segmentation model for enhanced aorta segmentation. Combining a regular stream for overall structure with a LoG stream for local details, LoGB-Net enhances SA detection, addresses foreground-background imbalance, and improves robustness by providing the UQ confidence intervals, outperforming baseline methods as demonstrated in experiments. The hierarchical structure and trainable parameters enable the detection of vessels with significant scale differences, particularly for SA. We will assess generalizability for future work by applying our method to other blood vessel datasets beyond the aorta.

6. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by two publicly available datasets [23, 24]. Ethical approval was not required, as confirmed by the licenses attached to the open-access datasets.

7. ACKNOWLEDGMENTS

This research was supported in part by NSF grants IIS-1955395, OAC-2047127, IIS-2101696, OAC-2104158, and IIS-2401144, and NIH grant 1R01HL177814-01. The authors would like to thank the anonymous reviewers for their insightful comments.

8. REFERENCES

- [1] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [2] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in *Proc. MICCAI*, 2016, pp. 424–432.
- [3] M. A. Scarpolini, M. Mazzoli, and S. Celi, "Enabling supra-aortic vessels inclusion in statistical shape models of the aorta: A novel non-rigid registration method," *Front. Physiol.*, vol. 14, pp. 1211461, 2023.
- [4] A. Updegrove, N. M. Wilson, J. Merkow, H. Lan, A. L. Marsden, et al., "SimVascular: An open source pipeline for cardiovascular simulation," *Ann. Biomed. Eng.*, vol. 45, no. 3, pp. 525–541, 2017.
- [5] O. Oktay, J. Schlemper, L. L. Folgoc, M. C. H. Lee, M. P. Heinrich, et al., "Attention U-Net: Learning where to look for the pancreas," *arXiv:1804.03999*, 2018.
- [6] G. D. Maher, C. M. Fleeter, D. E. Schiavazzi, and A. L. Marsden, "Geometric uncertainty in patient-specific cardiovascular modeling with convolutional dropout networks," *Comput. Methods Appl. Mech. Eng.*, vol. 386, pp. 114038, 2021.
- [7] J. Zhao and Q. Feng, "Automatic aortic dissection center-line extraction via morphology-guided CRN tracker," *IEEE J. Biomed. Health Informatics*, vol. 25, no. 9, pp. 3473–3485, 2021.
- [8] Y. Jin, A. Pepe, J. Li, C. Gsaxner, F.-H. Zhao, et al., "AI-based aortic vessel tree segmentation for cardiovascular diseases treatment: Status quo," *arXiv:2108.02998*, 2021.
- [9] K. A. Vermeer, F. M. Vos, H. G. Lemij, and A. M. Vossepoel, "A model based method for retinal blood vessel detection," *Comput. Biol. Med.*, vol. 34, no. 3, pp. 209–219, 2004.
- [10] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. CVPR*, 2017, pp. 2881–2890.
- [11] S. Jegou, M. Drozdal, D. Vazquez, A. Romero, and Y. Bengio, "The one hundred layers tiramisu: Fully convolutional DenseNets for semantic segmentation," in *Proc. CVPRW*, 2017, pp. 11–19.
- [12] P. Gu, H. Zheng, Y. Zhang, C. Wang, and D. Chen, "kCBAC-Net: Deeply supervised complete bipartite networks with asymmetric convolutions for medical image segmentation," in *Proc. MICCAI*, 2021, pp. 337–347.
- [13] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, et al., "UNETR: Transformers for 3D medical image segmentation," in *Proc. WACV*, 2022, pp. 574–584.
- [14] A. Shaker, M. Maaz, H. Rasheed, S. Khan, M. Yang, et al., "UNETR++: Delving into efficient and accurate 3D medical image segmentation," *arXiv:2212.04497*, 2022.
- [15] A. Hatamizadeh, V. Nath, Y. Tang, D. Yang, H. R. Roth, et al., "Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images," in *Proc. BrainLes@MICCAI*, 2022, pp. 272–284.
- [16] P. Gu, Y. Zhang, C. Wang, and D. Chen, "ConvFormer: Combining CNN and transformer for medical image segmentation," in *Proc. ISBI*, 2023, pp. 1–5.
- [17] T. Takikawa, D. Acuna, V. Jampani, and S. Fidler, "Gated-SCNN: Gated shape CNNs for semantic segmentation," in *Proc. ICCV*, 2019, pp. 5229–5238.
- [18] Abhinav Sagar, "Uncertainty quantification using variational inference for biomedical image segmentation," in *Proc. WACVW*, 2022, pp. 44–51.
- [19] S. Moccia, E. De Momi, S. El Hadji, and L. S. Mattos, "Blood vessel segmentation algorithms - review of methods, datasets and evaluation metrics," *Comput. Methods Programs Biomed.*, vol. 158, pp. 71–91, 2018.
- [20] X. Yang, Z. Li, Y. Guo, and D. Zhou, "DCU-Net: A deformable convolutional neural network based on cascade U-Net for retinal vessel segmentation," *Multim. Tools Appl.*, vol. 81, no. 11, pp. 15593–15607, 2022.
- [21] Y.-Z. Zeng, S.-H. Liao, P. Tang, Y.-Q. Zhao, M. Liao, et al., "Automatic liver vessel segmentation using 3D region growing and hybrid active contour model," *Comput. Biol. Medicine*, vol. 97, pp. 63–73, 2018.
- [22] P. Thiagarajan, P. Khairnar, and S. Ghosh, "Explanation and use of uncertainty quantified by Bayesian neural network classifiers for breast histopathology images," *IEEE Trans. Medical Imaging*, vol. 41, no. 4, pp. 815–825, 2022.
- [23] N. M. Wilson, A. K. Ortiz, and A. B. Johnson, "The vascular model repository: A public resource of medical imaging data and blood flow simulation results," *J. Med. Devices*, vol. 7, no. 4, pp. 0409231, 2013.
- [24] L. Radl, Y. Jin, A. Pepe, J. Li, C. Gsaxner, et al., "AVT: Multicenter aortic vessel tree CTA dataset collection with ground truth segmentation masks," *Data in Brief*, vol. 40, pp. 107801, 2022.
- [25] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, et al., "Feature pyramid networks for object detection," in *Proc. CVPR*, 2017, pp. 936–944.
- [26] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," *Nat. Methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [27] X. Huang, Z. Deng, D. Li, and X. Yuan, "MISSFormer: An effective medical image segmentation transformer," *arXiv:2109.07162*, 2021.
- [28] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, et al., "TransUNet: Transformers make strong encoders for medical image segmentation," *arXiv:2102.04306*, 2021.
- [29] L. Antiga, M. Piccinelli, L. Botti, B. Ene-Iordache, A. Remuzzi, and D. A. Steinman, "An image-based modeling framework for patient-specific computational hemodynamics," *Medical Biol. Eng. Comput.*, vol. 46, pp. 1097–1112, 2008.
- [30] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A tensorial approach to computational continuum mechanics using object-oriented techniques," *Comput. Phys.*, vol. 12, no. 6, pp. 620–631, 1998.