Automatic Segmentation of the Abdominal Aorta and Stent-Grafts

Bertram Sabrowsky-Hirsch, Stefan Thumfart, Wolfgang Fenz, and Richard Hofer
Unit for Medical Informatics, RISC Software GmbH, Hagenberg, Austria
Email: {bertram.sabrowsky-hirsch, stefan.thumfart, wolfgang.fenz, richard.hofer}@risc-software.at

Pierre Schmit and Franz Fellner
Central Radiology Institute, Kepler University Hospital, Linz, Austria
Email: {pierre.schmit, franz.fellner}@kepleruniklinikum.at

***Abstract*—Understanding and monitoring changes of the treated vessel after Endovascular aneurysm repair is crucial for the prediction of complications and risk assessment to facilitate timely intervention. Due to the complexity of the stent-graft wire frame enveloping the aortic blood lumen and the inherent artifacts caused by the metal wire, segmenting the structures required for simulation and further analysis is a non-trivial task. In this paper we present a fully automatic segmentation architecture combining two 3D U-Nets in a novel patching approach leveraging knowledge of the target anatomy. We evaluated our approach on a real world clinical dataset against a competitive baseline, yielding results that surpass the baseline in both accuracy and computation time. On our data we achieve a median Dice similarity coefficient of 0.97 for the blood lumen and 0.88 for the stent-graft segmentation. We point out two common flaws in current segmentation models: undersampling and indiscriminate patching. By addressing them appropriately, our approach gains an advantage that may benefit a multitude of segmentation tasks.**

***Index Terms*—segmentation, patch-based, centerline, U-net, stent graft, abdominal aneurysm**

I. Introduction

Endovascular Aneurysm Repair (EVAR) was chosen for 65% of surgical interventions of Abdominal Aortic Aneurysms (AAA) between 2010 and 2013 [1], and has therefore found its place as a minimally-invasive alternative to open surgery for suited patients. EVAR greatly reduces the intraoperative stress on patients and shortens the period of convalescence. However, the procedure also entails a high reintervention rate of 20% [2], rendering postoperative monitoring indispensable. In an endeavour to improve postoperative risk assessment by predicting complications, we plan to automatically analyse blood-flow simulations based on segmentations of the treated abdominal aorta and stent-graft prosthesis (i.e., the blood lumen and wire frame). The main obstacle in streamlining and deploying such an analysis to clinical practice is the dependence on said segmentations.

Manuscript received January 18, 2021; revised July 23, 2021.

Computed Tomography Angiography (CTA) is acquired within the standard clinical monitoring procedure of AAA patients [3]. A practically viable workflow must therefore rely on this imaging modality for the segmentations. Creating these segmentations semi-automatically is, however, a time-consuming task due to the complex structure of the stent-graft wire frame and the imaging artifacts it introduces. Segmenting the target structures in one scan takes a trained expert between 25 and 40 minutes. In this paper, we present a method to automatically create combined segmentations of both the blood lumen and the stent-graft wire frame.

A. Related Work

There are several publications on the segmentation of the abdominal aorta blood lumen and stent-graft wire frame, respectively, and some of them describe fully automated methods. We are, however, not aware of any approach that encloses both segmentation tasks. As generalized approaches for blood lumen segmentation struggle due to the unique challenges introduced by the aneurysm thrombus and stent-graft wire frame, specialized methods for segmentation of AAAs and aortic dissections are better suited for the first segmentation task. While purely intensity-based methods fail due to indistinct boundaries and strong imaging artifacts, these methods often rely on graph-based techniques or deformable models.

Graph-based techniques [4]-[7] utilize shape constraints to prevent leakage into neighbouring structures. The methods rely on a rough blood lumen segmentation (or centerline information [5]) that is acquired in a semi-automatic manner (e.g., using a graph-cut technique [6]) and subsequently refined. Approaches based on deformable models [8], [9] try to automatically fit contours to the target structures, but depend on seed points for the determination of the initial contour. While Kovács *et al*. [8] describe an automatic calculation of these seed points, their method suffers from a general lack in accuracy, especially for postoperative scans. More exotic approaches make use of radial models [10], level-set methods [11] and tracking [12], again depending on manual selection of seed points for initialization. Of the above methods, [6], [8] and [9] are the only ones tested

doi: 10.18178/joig.9.3.67-73

with postoperative scans. Another recurring issue is the dependence on highly optimized parameters based on undisclosed datasets.

As for learning-based methods, Maiora *et al.* [13] employ a random forest classifier to segment lumen, thrombus and bone simultaneously in an active learning approach. While it reports high accuracy and offers a comparison with a set of other classifiers, no evaluation against ground truth segmentations is provided. López-Linares *et al.* [14] achieve a mean Dice Similarity Coefficient (DSC) score of 0.80 and 0.88, respectively, using a 2D and 3D CNN for the segmentation of the thrombus volume in postoperative scans. Zheng *et al.* [15] yield a DSC of 0.82 with a U-Net for thrombus segmentation given a dataset of pre-operative scans. The model used by Li *et al.* [16] for the segmentation of Type-B aortic dissection is based on the 3D U-Net and achieves a DSC of 0.92 for the aorta.

Regarding the segmentation of stent-graft wire frames, the literature is much more limited. Klein *et al.* [17] present a graph-based method to segment the stent-graft wire frame, which focuses, however, on the known topology of two stent-types and disregards the blood lumen altogether. As the structures of blood lumen and stent-graft are highly entangled, segmenting them separately would introduce another challenge when trying to consistently merge them. Instead, in this work we focus on the 3D U-Net, which is capable of segmenting multiple structures simultaneously and has performed well in similar segmentation tasks.

As of today, the most crucial limitation of 3D U-Nets remain their hardware requirements, calling for a trade-off between the extent of the region of interest and segmentation accuracy. Our segmentation task challenges both aspects as it requires the extraction of large abdominal structures (the aortic blood lumen) as well as minute details (the stent-graft wire frame with a diameter as small as 0.4mm [18]). This dilemma is generally approached by using patch-based models which in turn are prone to segmentation artifacts at the patch borders and incapable of capturing relations beyond the scale of the patches. There are methods to reduce these issues (such as overlapping patches, test-time augmentation and multi-scale architectures; all used by Isensee *et al.* [19]), however, these methods also exponentially increase the computational costs and fail to address the underlying problem of ill-placed patches. Another often neglected issue of U-Nets is the prerequisite of a common spacing within the input image data, necessitating intermediate resampling that may significantly distort the information if not considering the *Nyquist-Shannon sampling theorem*. Nevertheless, a common approach is to resample to the median voxel spacing within the training dataset (see Isensee *et al.* [19]).

In this paper, we propose an approach using a combination of two 3D U-Nets to reliably and efficiently create high-resolution segmentations of the blood lumen and stent-graft wire frame. Specifically, we present an architecture using a novel patching method to overcome some of the limitations of U-Nets. We validated both our approach and the state of the art segmentation framework *nnU-Net* [19] based on our dataset of 76 abdominal CTA scans and manual segmentations. Our approach yielded significantly better results for the segmentation of aortic blood lumen and stent-graft wire frame than the *nnU-Net*.

II. The Dataset

A dataset consisting of 76 postoperative scans of 36 AAA patients (4 female, 31 male, mean age of 71 years) treated with EVAR was provided by the Kepler University Hospital Linz. Five different stent-types are present in the dataset: Anaconda, Gore EXCLUDER, Medtronic ENDURANT, Medtronic ENDURANT II and Ovation TRIVASCULAR. The scans originate from Siemens Somatom hardware (Force, Sensation Cardiac 64 and Sensation Open) and vary in size and voxel spacing as outlined in Table I.

TABLE I. Statistics of Spacing, Resolution and Size over the 76 Scans in the Dataset. F: Frontal, S: Sagittal, L: Longitudinal

	Spacing (mm)		Resolution (voxel)		Size (mm)	
	F/S	L	F/S	L	F/S	L
Mean	0.677	1.361	512	411.39	346.49	508.59
Median	0.695	1.500	512	341.50	356.00	473.25
Minimum	0.404	0.800	512	155.00	207.00	390.00
Maximum	0.977	3.000	512	873.00	500.00	732.00

The segmentations have been created semi-automatically by trained experts using the protocol outlined in Fig. 1: The Active Contour Segmentation [20] featured in the software ITK-Snap 3.8.0 [21] was used to create the initial blood lumen segmentations (a) from below the heart to the second bifurcation of the iliac arteries. The stent-graft wire frame was then segmented (b) by thresholding a region of interest around the blood lumen segmentation. The final segmentations (c) were then obtained by manually correcting the segmentation labels using ITK-Snap. To the bottom of Fig. 1 are examples of the manual segmentations in our dataset featuring different stent-types like Anaconda (d), Medtronic Endurant (e) and Medtronic Endurant II (f).

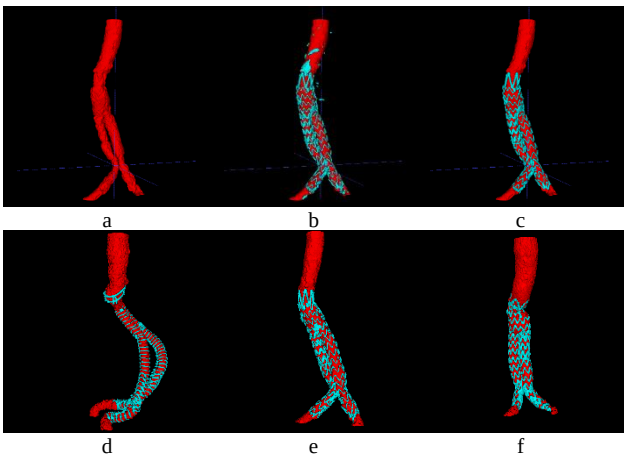

Figure 1. Steps for obtaining the ground truth segmentations (a, b, c; see text) and exemplary segmentations of our dataset (d, e, f).

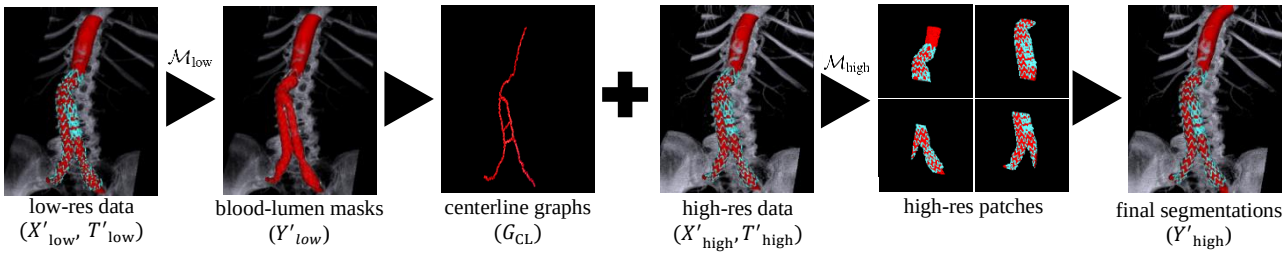

Figure 2. Overview of the method's workflow.

We conducted a 5-fold cross-validation for both our approach and the *nnU-Net*, using the same fold configuration. For some patients, several postoperative scans exist, which has been accounted for via a grouping criterion in our fold-stratification to avoid a configuration where scans of the same patient occur in both the training and the validation set.

III. Method

The general architecture of our 3D U-Net models closely resembles the one used by Isensee *et al.* for the *BRATS 2017* challenge [22]. Distinctive features of our approach are the combination of two models and a novel patching method in a process that is outlined in Fig. 2.

A. Definitions

Let X be the set of CTA scans and T the corresponding ground truth segmentations in our dataset: Our method can be described as a function $f: X \to Y$, where X is the set of input scans and Y the set of resulting segmentations. We use two models, $\mathcal{M}_{\text{low}}$ and $\mathcal{M}_{\text{high}}$, that are trained on preprocessed data X' and T' (see the following section) and can be used for inference in the form $\mathcal{M}: X' \to Y$.

$\mathcal{M}_{\text{low}}$ is trained to segment the blood lumen using instances of the large-area low resolution scans X'_{low} and truth T'_{low}. The resulting low-resolution blood lumen masks Y_{low} are used to extract the centerline graphs G_{CL}.

$\mathcal{M}_{\text{high}}$ is trained to segment blood lumen and stent-graft in patches extracted from the high resolution scans X'_{high} and truth T'_{high}. Patches are extracted at positions along the entire span of the respective centerline graph $g \in G_{\text{CL}}$ and finally merged into a high-resolution segmentation $y \in Y$.

B. Data Preprocessing

With the CTA scans of our dataset varying considerably in physical extent and voxel spacing as outlined in Table I, the first step is to resample and crop the scans and ground truth for each segmentation model.

For $\mathcal{M}_{\text{low}}$ a voxel spacing of 1×1×3mm (frontal, sagittal, longitudinal) proved sufficient to capture enough detail in the resulting blood lumen segmentation to facilitate the extraction of centerline information. We then crop the scans to a large Volume of Interest (VOI) of 192×192×128 voxels (i.e., 192×192×384mm).

For $\mathcal{M}_{\text{high}}$ we do not use the optimal voxel spacing (according to the *Nyquist rate*, i.e., half of the minimum within the dataset: 0.202×0.202×0.4mm) but rather a more feasible resolution of 0.35×0.35×0.75mm, which is considerably higher than the median voxel spacing of 0.695×0.695×1.5mm. We then crop the images to VOIs of 580×580×512 voxels (i.e., 203×203×384mm).

We use third order B-spline interpolation to resample each scan in the set of scan volumes X, clip the intensity values to the interval between the 0.5th and 99.5th percentile to remove outliers and further normalize by subtracting the mean and dividing through the standard deviation of the clipped intensity values. These steps are done individually for each segmentation model over the set of all scans to yield the preprocessed sets of scans X'_{low} and X'_{high}. The ground truth segmentations T are equally resampled using the label-linear interpolation suggested by Schaerer *et al.* [23] to yield the preprocessed sets T'_{low} and T'_{high}.

C. Model Architecture

Based on David G. Ellis' implementation [24] of the Isensee *et al.* model [22], we adjusted some configurable parameters such as the model-depth (number of layers), number of segmentation levels (used to create secondary segmentation maps for deep supervision) and base-filters (number of output filters for the first convolution kernel). The parameter changes are based on a grid-search for each model.

We configured $\mathcal{M}_{\text{low}}$ with 3 segmentation levels, a model-depth of 5 and base-filters set to 8. The input size was set to the full scan size of 192×192×128 (i.e., 203×203×384mm), and the target label is *aorta*.

$\mathcal{M}_{\text{high}}$ was configured with 5 segmentation levels, a model-depth of 7 and base-filters set to 8. The patch size was set to 256×256×192 (i.e., 90×90×144mm), and the target labels are aorta and stent.

The models each yield a volume where every voxel is assigned a set of softmax values corresponding to the target labels. The output is processed into a label-map volume y by assigning the label corresponding to the largest softmax value to each voxel or the background label if that value is below a threshold of 0.5.

D. Training Procedure

We train the models using a weighted multi-class Dice loss as used by Isensee *et al.* [22] computed for the respective target labels. The setup consists of an Adam optimizer with an initial learning rate $\eta_0 = 5 \cdot 10^{-4}$, a learning rate drop criterion set to 10 epochs with a drop factor of $\lambda_\eta = 0.5$ and an early stopping criterion set to 50 epochs, with the training continuing for 70 to 120 epochs of 200 training samples per epoch. With $\mathcal{M}_{\text{low}}$

being trained on the full scans X'_{low} before $\mathcal{M}_{\text{high}}$, the resulting segmentations Y_{low} are used to locate the aorta centerline for the training of $\mathcal{M}_{\text{high}}$. The largest connected region of voxels with the *aorta* label is masked and skeletonized (according to Lee *et al.* [25]) to extract the centerline graph G_{CL}. The center coordinates of the patches extracted from X'_{high} and used for the training of $\mathcal{M}_{\text{high}}$ are retrieved by randomly sampling equally distributed positions along the edges of the respective centerline graph $g \in G_{\text{CL}}$.

E. Inference

During inference, instead of sampling random coordinates along the centerline graph $g \in G_{\text{CL}}$ we equally distribute n coordinates along all combined edges to center the patches at. For our experiments $n = 20$ patches were sufficient. The softmax outputs for each patch are further combined into a volume spanning the full extent of the preprocessed scan $x' \in X'_{\text{high}}$. The softmax values for each target label are interpolated using weights from a Gaussian kernel to attenuate values at the patch borders. This was done to account for the inherent uncertainty at the patch borders. The resulting softmax volume is resampled to the original spacing of the respective input volume $x \in X$ using a first order B-spline interpolation before computing the final label-map volume y. The resampling step is necessary since we compare the results to the original ground truth T for our evaluation against the *nnU-Net*. By resampling the softmax volume rather than the label-map, fine structures like the stent-graft are preserved more accurately. If used for mesh reconstruction, one may skip downsampling and leverage the high resolution segmentations instead.

IV. Evaluation

We evaluated both our approach and the *nnU-Net* on our dataset in a 5-fold cross-validation. The DSC scores were calculated based on the respective label-map output and the original ground truth. The results are a median DSC (aorta and stent-graft respectively) of 0.969 and 0.879 for our approach and 0.964 and 0.866 for the *nnU-Net*, with our approach yielding slightly better results. This is particularly interesting as the *nnU-Net* uses an improved U-Net architecture and additional techniques – such as test-time augmentation – compared to our relatively simple implementation. Adjusting our architecture and adapting techniques used by the *nnU-Net* might further improve the performance of our approach.

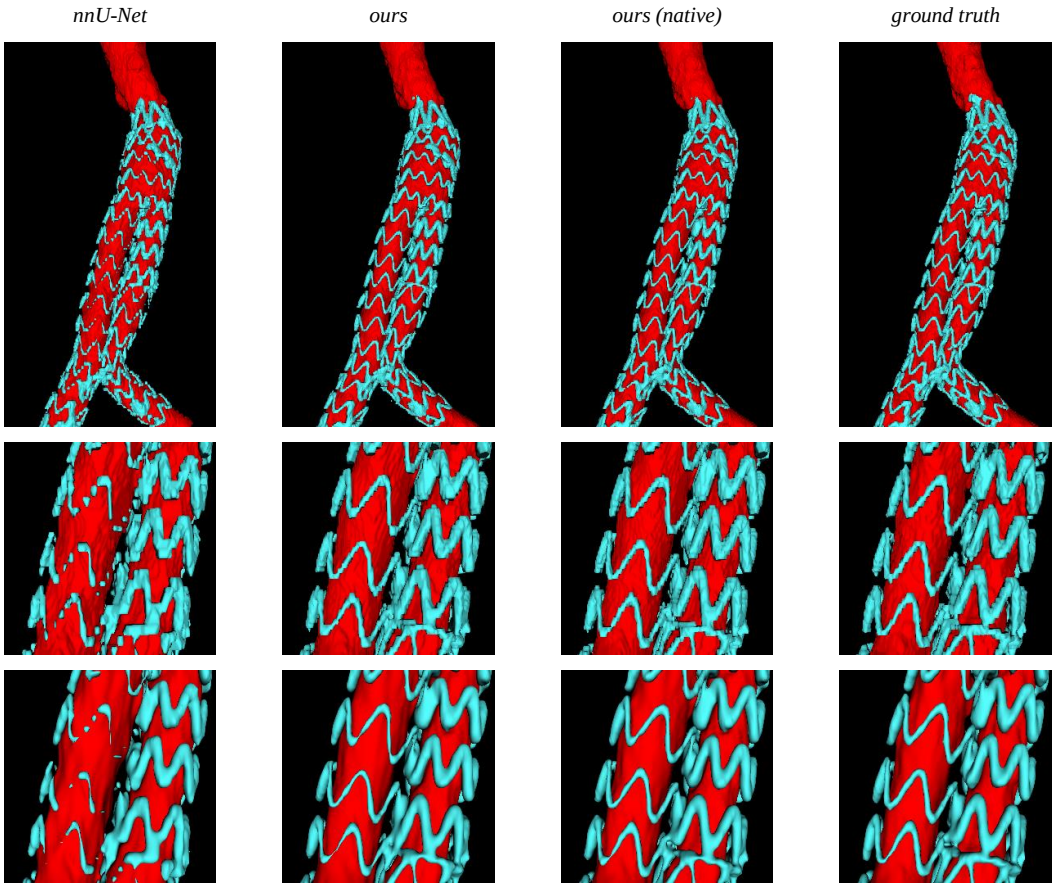

Figure 3. Comparison of meshes created from different segmentations of the same CTA scan. 3DSlicer [26] was used to create the meshes with a smoothing factor of either 0 (first two rows) or 0.3 (bottom row). The columns show (from left to right): the results of the nnU-Net, our approach resampled to the original voxel-spacing, our approach in the model's native spacing, and the original ground truth.

The main advantage of our method, however, is its ability to natively segment at a high resolution, while taking up significantly less computation time. In contrast to the *nnU-Net*, which uses the dataset's median voxel spacing of 0.695×0.695×1.5mm, our model segments at 0.35×0.35× 0.75mmm - increasing the resolution by a factor of almost 7. Both experiments were run on the same workstation featuring an *NVIDIA Titan RTX* GPU (24 GB VRAM). The resulting segmentations are optically clearly distinguishable, with our approach yielding clean reconstructions of the stent-graft structure whereas the structure is fragmented in the results of the *nnU-Net* as shown in Fig. 3. Yet, by resampling to the original voxel spacing information is lost or distorted due to resampling, especially at the fine structure of the stent-graft, making segmentations in our models high native resolution the preferable source for mesh reconstruction.

Other than the segmentation quality, we also compared the processing time for both approaches. While *nnU-Net* took 10 minutes and 24 seconds on average (and up to 21 minutes), our approach performed considerably faster as it took 1 minute and 53 seconds on average (and up to 2 minutes 58 seconds) per volume, which makes a huge difference in terms of clinical practicability. Table II shows a comprehensive comparison of our approach, which features median values for the FNE (False Negative Error), FPE (False Positive Error), JSC (Jaccard Similarity Coefficient) and AHD (Average Hausdorff Distance) as defined in [27] as additional evaluation metrics. Fig. 4 to Fig. 8 outline the statistical distribution of the results for each metric. We explain the discrepancy in Fig. 5 with the fact that our approach, though it is more accurate (as apparent from the DSC metric), is picking up calcifications along the aorta more frequently. These can be located distant from the stent-graft and affect the AHD disproportionally. Apart from that, our approach yields better or equivalent results, while outperforming the nnU-Net in terms of processing time.

TABLE II. DETAILED EVALUATION OF OUR APPROACH AGAINST THE NNU-NET (MEDIAN SCORES FOR EACH METRIC)

Approach	Target	DSC	FNE	FPE	JSC	AHD
Ours	Blood lumen	0.969	0.018	0.040	0.939	0.027
	Stent-Graft	0.879	0.117	0.108	0.783	0.198
	Time	1min 53secs				
nnU-Net	Blood lumen	0.964	0.025	0.039	0.931	0.043
	Stent-Graft	0.866	0.121	0.141	0.764	0.127
	Time	10min 24secs				

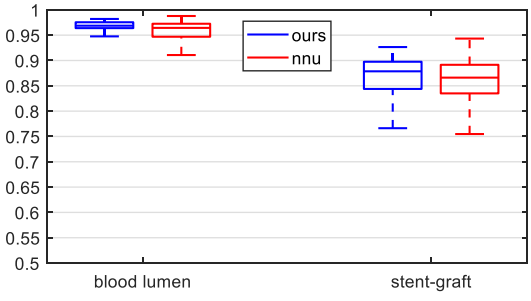

Figure 4. Dice Similarity Coefficient (DSC).

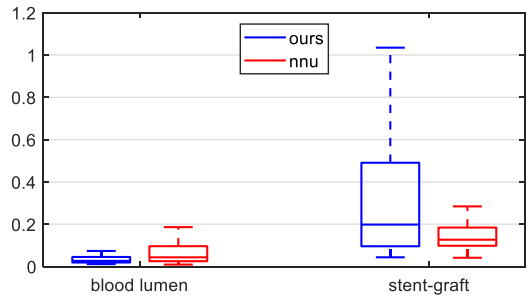

Figure 5. Average Hausdorff Distance (AHD).

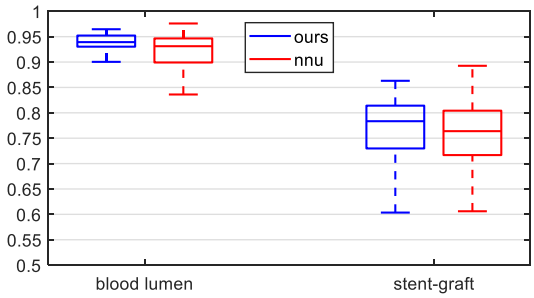

Figure 6. Jaccard Similarity Coefficient (JSC).

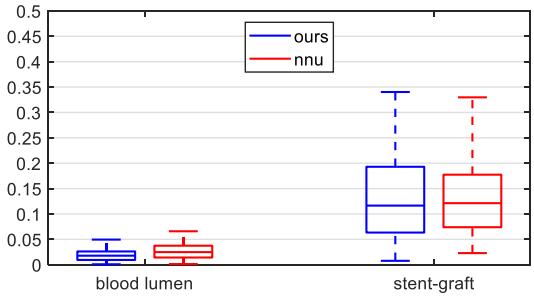

Figure 7. False Negative Error (FNE).

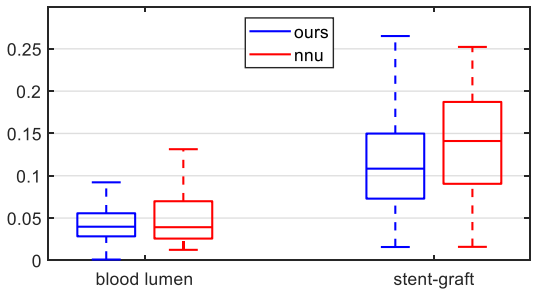

Figure 8. False Positive Error (FPE).

V. CONCLUSION AND FUTURE WORK

We have proposed a fully automated segmentation architecture combining two 3D U-Nets with a novel patching approach in order to segment both the aortic blood lumen and the stent-graft wire frame in abdominal CTA scans. The resulting meshes can be used for blood-flow simulations, facilitating a fully automated pipeline for post-EVAR risk assessment and prediction of complications based on CTA scans.

We showed that our approach is able to surpass the state of the art segmentation framework *nnU-Net*, despite lacking advanced techniques like test-time augmentation, while taking a fraction of the time on the same hardware. We attribute this to the high native resolution of our segmentation model as well as the strategic placement of the patches during training and inference. By choosing a small voxel spacing, we can reduce the distortion introduced by resampling. Using topological knowledge of the target geometry allowed us to drastically reduce the amount of patches needed, while still maintaining superior segmentation quality.

Apart from the direct application to other elongated structures like various vessels and bones, our method can benefit any segmentation task where the target structure covers the input volume only partially, as it facilitates more efficient training/inference and thus segmentation at a higher resolution. For the future, we plan to apply our approach to other anatomical structures and evaluate it on public datasets.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

AUTHOR CONTRIBUTIONS

Bertram Sabrowsky-Hirsch conducted the research and wrote the paper in collaboration with Stefan Thumfart and Wolfgang Fenz; Franz Fellner and Pierre Schmit provided and evaluated the medical data and results; Richard Hofer analyzed and annotated data; all authors had approved the final version.

ACKNOWLEDGEMENTS

This work was funded by the FFG (Austrian Research Promotion Agency) under the grants 851461 (EndoPredictor), 872604 (MEDUSA) and 867536 (vizARd). This project was supported by the strategic economic and research programme “Innovatives OÖ - 2020” of the province of Upper Austria. RISC Software GmbH is a Member of UAR (Upper Austrian Research) Innovation Network.

REFERENCES

[1] A. W. Beck, *et al.*, “Variations in abdominal aortic aneurysm care: A report from the international consortium of vascular registries,” *Circulation*, vol. 134, pp. 1948-1958, December 2016.

[2] M. Malina, “Reinterventions after open and endovascular AAA repair,” *J. Cardiovasc Surg.*, vol. 56, pp. 257-268, April 2015.

[3] M. S. Patel and J. P. Carpenter, “The value of the initial post-EVAR computed tomography angiography scan in predicting future secondary procedures using the Powerlink stent graft,” *Journal of Vascular Surgery*, vol. 52, pp. 1135-1139, November 2010.

[4] A. A. Duquette, P. M. Jodoin, O. Bouchot, and A. Lalande, “3D segmentation of abdominal aorta from CT-scan and MR images,” *Computerized Medical Imaging and Graphics*, vol. 36, pp. 294-303, June 2012.

[5] J. Egger, *et al.*, “Aorta segmentation for stent simulation,” unpublished, arXiv:1103.1773, March 2011.

[6] M. Freiman, S. J. Esses, L. Joskowicz, and J. Sosna, “An iterative model-constrained graph-cut algorithm for Abdominal Aortic Aneurysm thrombus segmentation,” in *Proc. IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Rotterdam, 2010, pp. 672-675.

[7] T. Siriapisith, W. Kusakunniran, and P. Haddawy, “Outer wall segmentation of abdominal aortic aneurysm by variable neighborhood search through intensity and gradient spaces,” *J. Digit Imaging*, vol. 31, pp. 490-504, August 2018.

[8] T. Kovács, P. Cattin, H. Alkadhi, S. Wildermuth, and G. Székely, “Automatic segmentation of the vessel lumen from 3D CTA images of aortic dissection,” in *Image Processing for Medicine (In German)*, Berlin/Heidelberg: Springer-Verlag, 2016, pp. 161-165, 2006.

[9] F. Lalys, V. Yan, A. Kaladji, A. Lucas, and S. Esneault, “Generic thrombus segmentation from pre- and post-operative CTA,” *Int. J. CARS*, vol. 12, pp. 1501-1510, September 2017.

[10] I. Macía, *et al.*, “Segmentation of abdominal aortic aneurysms in CT images using a radial model approach,” in *Proc. International Conference on Intelligent Data Engineering and Automated Learning*, 2009, pp. 664-671.

[11] C. Zohios, G. Kossioris, and Y. Papaharilaou, “Geometrical methods for level set based abdominal aortic aneurysm thrombus and outer wall 2D image segmentation,” *Computer Methods and Programs in Biomedicine*, vol. 107, pp. 202-217, August 2012.

[12] Y. Xie, J. Padgett, A. M. Biancardi, and A. P. Reeves, “Automated aorta segmentation in low-dose chest CT images,” *Int. J. CARS*, vol. 9, pp. 211-219, March 2014.

[13] J. Maiora and M. Graña, “A hybrid segmentation of abdominal CT images,” in *Proc. International Conference on Hybrid Artificial Intelligent Systems*, 2012, pp. 416-423.

[14] J. Li, L. Cao, W. Cheng, M. Bowen, and W. Guo, “Towards automatic measurement of type B aortic dissection parameters: Methods, applications and perspective,” in *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, Cham: Springer, 2018, pp. 64-72.

[15] J. Q. Zheng, X. Y. Zhou, Q. B. Li, C. Riga, and G. Z. Yang, “Abdominal aortic aneurysm segmentation with a small number of training subjects,” unpublished, arXiv:1804.02943, April 2018.

[16] K. López-Linares, I. García, A. García-Familiar, I. Macía, and M. A. G. Ballester, “3D convolutional neural network for abdominal aortic aneurysm segmentation,” unpublished, arXiv:1903.00879, March 2019.

[17] A. Klein, *et al.*, “Automatic segmentation of the wire frame of stent grafts from CT data,” *Medical Image Analysis*, vol. 16, pp. 127-139. January 2012.

[18] R. J. Winder, Z. Sun, B. Kelly, P. K. Ellis, and D. Hirst, “Abdominal aortic aneurysm and stent graft phantom manufactured by medical rapid prototyping,” *Journal of Medical Engineering & Technology*, vol. 26, pp. 75-78, January 2002.

[19] F. Isensee, P. Kickingereder, W. Wick, M. Bendszus, and K. H. Maier-Hein, “Brain tumor segmentation and radiomics survival prediction: Contribution to the BRATS 2017 challenge,” in MICCAI *Brainlesion Workshop*, 2018, pp. 287-297.

[20] V. Caselles, R. Kimmel, and G. Sapiro, “Geodesic active contours,” in *Proc. IEEE International Conference on Computer Vision*, Cambridge, MA, USA, 1995, pp. 694-699.

[21] P. A. Yushkevich, *et al.*, “User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability,” *NeuroImage*, vol. 31, pp. 1116-1128, July 2006.

[22] F. Isensee, P. F. Jäger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, “Automated design of deep learning methods for biomedical image segmentation,” unpublished, arXiv:1904.08128, April 2020.

[23] J. Schaerer, F. Roche, and B. Belaroussi. (December 2014). A generic interpolator for multi-label images. *The Insight Journal*. [Online]. vol. 950. Available: http://hdl.handle.net/10380/3506

[24] D. G. Ellis. 3DUnetCNN. [Online]. Available: github.com/ellisdg/3DUnetCNN

[25] K. Lee, *et al.*, “Three-Dimensional thrombus segmentation in abdominal aortic aneurysms using graph search based on a triangular mesh,” *Computers in Biology and Medicine*, vol. 40, pp. 271-278, March 2010.

[26] A. Fedorov, *et al.*, “3D slicer as an image computing platform for the quantitative imaging network,” *Magnetic Resonance Imaging*, vol. 30, pp. 1323-1341, November 2012.

[27] N. Tustison and J. Gee. (December 2009). Introducing Dice, Jaccard, and Other Label Overlap Measures to ITK. *The Insight Journal*. [Online]. vol. 707. Available: http://hdl.handle.net/10380/3141

Bertram Sabrowsky-Hirsch received the MSc degree from the University of Applied Sciences Upper Austria, in 2016. He is currently pursuing the PhD Degree from the Institute of Computational Perception, Johannes Kepler University Linz while being employed at the Unit for Medical Informatics at RISC Software GmbH. His research is centered around Deep Learning for Medical Image Analysis.