DOI: 10.1111/cgf.14808
Eurographics Conference on Visualization (EuroVis) 2023
D. Archambault, R. Bujack, and T. Schreck
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 42 (2023), Number 3

A Fully Integrated Pipeline for Visual Carotid Morphology Analysis

P. Eulzer[1], F. von Deylen[1], W.-C. Hsu[1], R. Wickenhöfer[2], C. M. Klingner[3] and K. Lawonn[1]

[1]University of Jena, Faculty of Mathematics and Computer Science, Germany
[2] Herz-Jesu Hospital Dernbach, Clinic for Radiology, Germany
[3] University Hospital Jena, Clinic for Neurology, Germany

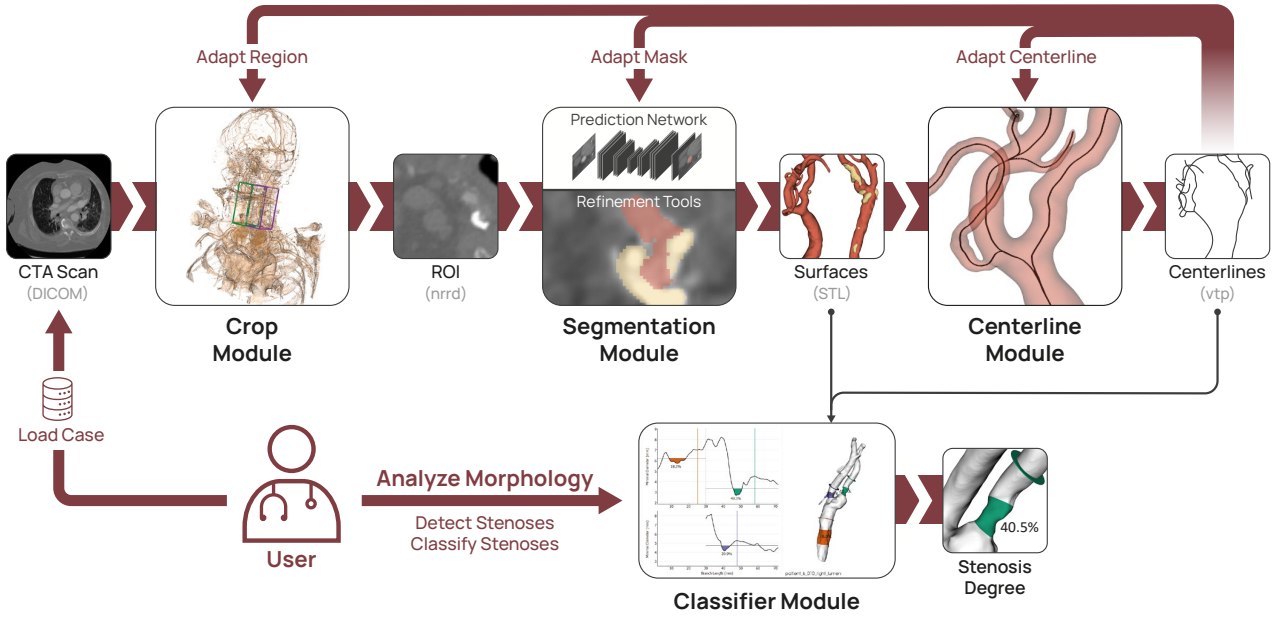

Figure 1: *The integrated pipeline is split into three pre-processing modules (cropping, segmentation, and centerline generation) and one visualization module (the classifier module) to analyze vessel stenoses. At any time, the user can inspect and, if necessary, make corrections to any of the processing steps. The results are automatically propagated.*

Abstract

Analyzing stenoses of the internal carotids – local constrictions of the artery – is a critical clinical task in cardiovascular disease treatment and prevention. For this purpose, we propose a self-contained pipeline for the visual analysis of carotid artery geometries. The only inputs are computed tomography angiography (CTA) scans, which are already recorded in clinical routine. We show how integrated model extraction and visualization can help to efficiently detect stenoses and we provide means for automatic, highly accurate stenosis degree computation. We directly connect multiple sophisticated processing stages, including a neural prediction network for lumen and plaque segmentation and automatic global diameter computation. We enable interactive and retrospective user control over the processing stages. Our aims are to increase user trust by making the underlying data validatable on the fly, to decrease adoption costs by minimizing external dependencies, and to optimize scalability by streamlining the data processing. We use interactive visualizations for data inspection and adaption to guide the user through the processing stages. The framework was developed and evaluated in close collaboration with radiologists and neurologists. It has been used to extract and analyze over 100 carotid bifurcation geometries and is built with a modular architecture, available as an extendable open-source platform.

CCS Concepts

• ***Human-centered computing*** → ***Scientific visualization;*** • ***Applied computing*** → *Life and medical sciences;*

1. Introduction

The two internal carotid arteries provide oxygen supply to the majority of the brain. If their blood flow is constricted, stroke is imminent. Stroke is a major type of cardiovascular disease – the leading cause of mortality and morbidity worldwide [GBD19]. Full recovery from a stroke is a rare exception. In non-fatal cases, permanent neurological impairments are the norm and post-stroke care represents a tremendous societal and economic burden [RGB*19]. Therefore, effective stroke prevention strategies are of fundamental importance. Due to the buildup of atherosclerotic plaque on its vessel walls, *stenosis* of the internal carotid may occur, where a localized narrowing restricts the blood flow (see Figure 2). A typical predilection site for stenosis is closely after the carotid bifurcation, located in the upper neck area, where the internal and external carotid split off the common carotid artery. A developing stenosis, if detected at the right stage, can be treated through surgical removal of plaque or stent insertion. As with most cardiovascular diseases, however, surgical intervention is an option but carries its own risks [HTW*09]. Therefore, treatment decisions need to be made on an individual basis, considering a large number of factors involved. In general, better comprehension of these factors and a more accurate assessment of objective descriptors of the disease are crucial for ideal treatment and monitoring of a patient's condition. In the case of carotid stenosis, essential objectives are the proper analysis of the stenosis location, degree, composition, and shape, i.e., the vessel morphology. In this work, we introduce a fully integrated pipeline to efficiently extract and visualize 3D vascular models from computed tomography angiography (CTA) images. We pair AI-assisted segmentation and automatic feature extraction with carefully selected user interaction options in a processing pipeline with step-wise visual feedback. Based on the resulting vascular models, we introduce a specialized tool to immediately detect, assess, and accurately classify stenoses of the internal carotids. The pipeline is outlined in Figure 1.

Considering the prevalence of cardiovascular disease, it is unsurprising that abundant proposals have been made to improve monitoring, treatment planning, and medical research using techniques from the visualization toolbox. This includes volume renderings [KGNP12], image reformations [MMV*13, RFK*07], map-like depictions [EMML22], and visualizations of hemodynamics (blood flow) [OJMN*19]. In spite of the positive feedback of potential users for these algorithms and applications and decades of development, we find that the majority of techniques are hardly on their way to being used in clinical practice or medical research. We identify four major bottlenecks for the practical transfer of methods that we aim to address with our approach: the users' trust in what is visualized, the efficiency of using a visualization with new data, the compatibility of the used data formats, and the scale at which such systems are validated. An integrated pipeline, as delineated in the following, targets these bottlenecks by streamlining data processing. Giving the user carefully chosen control over the pipeline then allows for uncomplicated validation of data and even on-the-fly correction, ultimately improving understanding and trust in what is shown. Additionally, a modularized pipeline design supports the implementation of new processing steps and visualization modules, where the database does not need to be recreated. We elaborate on the design of the pipeline for the analysis of stenoses of the internal carotids. However, the concept could be transferred to other domains like cranial circulation [MMNG16], aneurysm risk-assessment [MGB*19, MPL21], liver surgery planning [LL20], or the analysis of heart valves [EEL*19]. In summary, our contributions are:

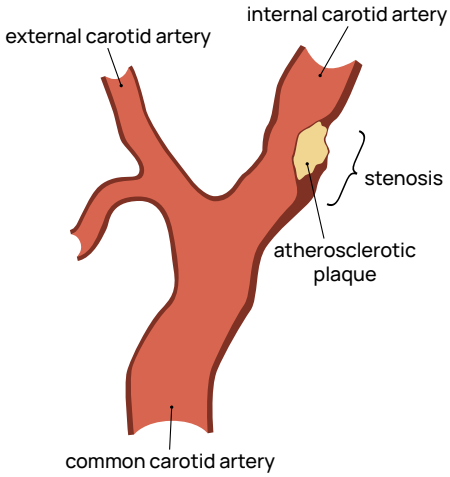

Figure 2: *Anatomy of the carotid bifurcation. A stenosis caused by plaque on the arterial wall restricts blood flow in the internal carotid.*

- A fully self-contained processing and visualization pipeline for carotid bifurcation geometries.
- A modularized and open-source framework design, which allows easy integration of processing and visualization modules interfacing with a shared patient database.
- A novel visualization for highly efficient carotid stenosis detection and classification implemented as a pipeline module.

To evaluate the effectiveness of our approach, the pipeline has been used to generate over 100 carotid bifurcation models from real patients. We hope that this will provide a basis for further research to validate vessel visualization methods more broadly. The open-source code of all integrated modules and the carotid model database are available online [EL23a, EL23b].

2. Related Work

Related work encompasses all pipeline processing and visualization stages. We highlight key works in each field and provide relevant surveys for more in-depth discussions of each stage.

Vessel Segmentation. From medical volume images, like CTA, vascular structures can be segmented manually or automatically. The voxels of interest are marked in the volume. This is often followed by a surface reconstruction step, where a geometric model, e.g., a triangulation, is extracted. Sometimes, multiple models are created, for instance, to differentiate the vessel inner wall (lumen) and atherosclerotic plaque [JN*18]. There are two principal types of segmentation methods, "traditional" and machine learning-based. Traditional methods to segment blood vessels are reviewed by Lesage et al. [LABFL09]. These include, for example, active contours [MVN06] or graph-cuts [BPS*10, ELD10]. As

many of these methods are computationally slow and not as robust against noise and artifacts as desired, recent years have seen a surge in machine learning methods [LK16, PHM*16, WXG*16]. Particularly convolutional neural networks (CNNs) have been found advantageous for the automatic segmentation of vessels. Specialized networks for carotid lumen and plaque segmentation have been proposed [BKS18, CVOA08], where good performing networks, with a Dice similarity of 82%, were based on as low as 15 training data sets [ZTP*21]. For further reading, we refer to the surveys by Moccia et al. [MMHM18] and Zhao et al. [ZCHH17].

Centerline Extraction. A typical way to describe vascular topology is to derive centerlines passing through the geometric center of the tubular structures. Centerlines can be directly computed from image volumes through topological thinning [SSZZ01, WL08] or graph search [BSB*00, BKS01]. They can also be computed from triangulated vessel models through thinning [ATC*08] or using the medial axis transform. The medial axis passes through the origins of the maximally inscribed spheres inside the model. A distinct advantage of the latter approach is that it computes the minimal radii inside the vessel as a byproduct. The medial axis can be derived from bisectors [CKM04], which is computationally complex. A faster method is to first construct the 3D Voronoi diagram from the surface points and compute the medial axis as a subset using the cells that lie inside the geometry. Antiga et al. [AEIR03] implemented this approach and Wang et al. [WCH*10] provide a comparison of different methods.

Blood Vessel Visualization. Various visualization frameworks for blood vessels exist that are designed to solve specific tasks of clinical practitioners [MKP*16, MVB*17, MOJB*19]. A common approach is the combination of multiple coordinated data views with different purposes, for example, slice-based renderings, volume renderings, 3D surfaces, views of integral lines, focus lenses, map-like depictions, and also data graphs or charts. Generally, a differentiation is made between model-free and model-based techniques [PO08]. The latter rely on a processing pipeline that extracts geometric models from volume images. These can then be further used in advanced visualizations showing, e.g., blood flow features [OJMN*19] or map-like depictions with strategically reduced visual complexity [EMML22]. Eulzer et al. [EMKL21] specifically target the visualization of hemodynamics in carotid stenoses. Wall-related parameters like shear stress or plaque distribution at the carotid bifurcation have been visualized using parameterizations of the vessel surface geometry [CLC13, CCR20, ERM*21]. Lawonn et al. [LMW*19] propose a technique for detecting and segmenting aneurysms (pathological vessel bulges). Similarly, in this work, we introduce a technique for the extraction of stenosis geometries.

There are a number of data processing frameworks designed for researchers working with vascular models. The vascular modeling toolkit (VMTK) [APB*08, ISMA18] is a suite of command-line scripts for segmentation, centerline creation, vessel mesh processing, and volume mesh generation. There are also generalized, multi-purpose software frameworks for processing medical volume data, most notably 3D Slicer [KPV13] and ImageJ [SRE12]. Due to their high versatility, they are broadly used. However, their steep learning curve and vast range of functions means, in terms of usability for specialized (clinical) tasks, they cannot compete with individualized tools specifically designed for these purposes. Notably, in almost all related medical applications, data processing happens separately from data visualization. The processing pipeline is usually present but not fully integrated into the final application. From the user's perspective, this means there is no insight into the correctness of the data and no easy way to fix errors or adapt the presented models.

3. Bottlenecks in Clinical Transfer

The traction of vessel visualization techniques and frameworks in clinical practice has been notoriously challenging when compared to other, non-medical, application domains. These challenges are partly attributed to the difficulty of properly developing and evaluating medical visualizations [PRI18] – user studies with a large number of highly specialized medical experts are not a conceivable option – and also to the integration of new algorithms into clinical workflows, which requires going through elaborate legal procedures. These obstacles mean that, for most visualization researchers, it is simply not lucrative to sustainably determine the benefits of new vessel visualizations. From the literature on vessel visualization systems and our discussions with clinicians we do, however, repeatedly get the impression that advanced visualizations of vascular features, including morphology, connectivity, wall parameters, and blood flow, do posses a high *potential* value for diagnosis and treatment. So why are they only sparsely adopted? While various probable reasons may exist, in our discussions we kept iterating on a particular aspect: commonly, in prototypical development, we do not focus on the streamlined integration of data processing. This leads to the fact that many solutions for pre-processing, segmentation, feature extraction, and visualization exist, yet most of them are individual and scattered. At the same time, these solutions are becoming more and more complex, as they tackle increasingly sophisticated problems, making it harder to functionally connect them. From this lack of integration, we distill four aspects that are dominant bottlenecks for the clinical transfer of vessel visualization techniques:

- **B1 Trust.** With highly processed data, how can a user assure the correctness of what they see? This is an exceptionally critical factor in clinical environments. Medical personnel has to be able to fully trust the algorithms and their results, partly without understanding how they are computed. In medical decision-making, the liability question is always a focal consideration. Physicians bear immediate responsibility over their patients' lives and, thus, need to know to which extent they can trust a data representation.
- **B2 Efficiency.** Many advanced medical visualization frameworks have a high adoption cost, as they depend on various fragmented tools for data curation and processing. The result is that small corrections early in a processing pipeline often require extensive effort to propagate through later stages.
- **B3 Compatibility.** Additionally, these tools sometimes suffer from high interface costs, as specific data formats with varying degrees of standardization are used [GSG*21].
- **B4 Validation Scale.** Typically, only a handful of example data sets are used for the development and validation of vessel visualizations, as processing steps like model extraction or simulation are profoundly time-intensive. This reliance on selected

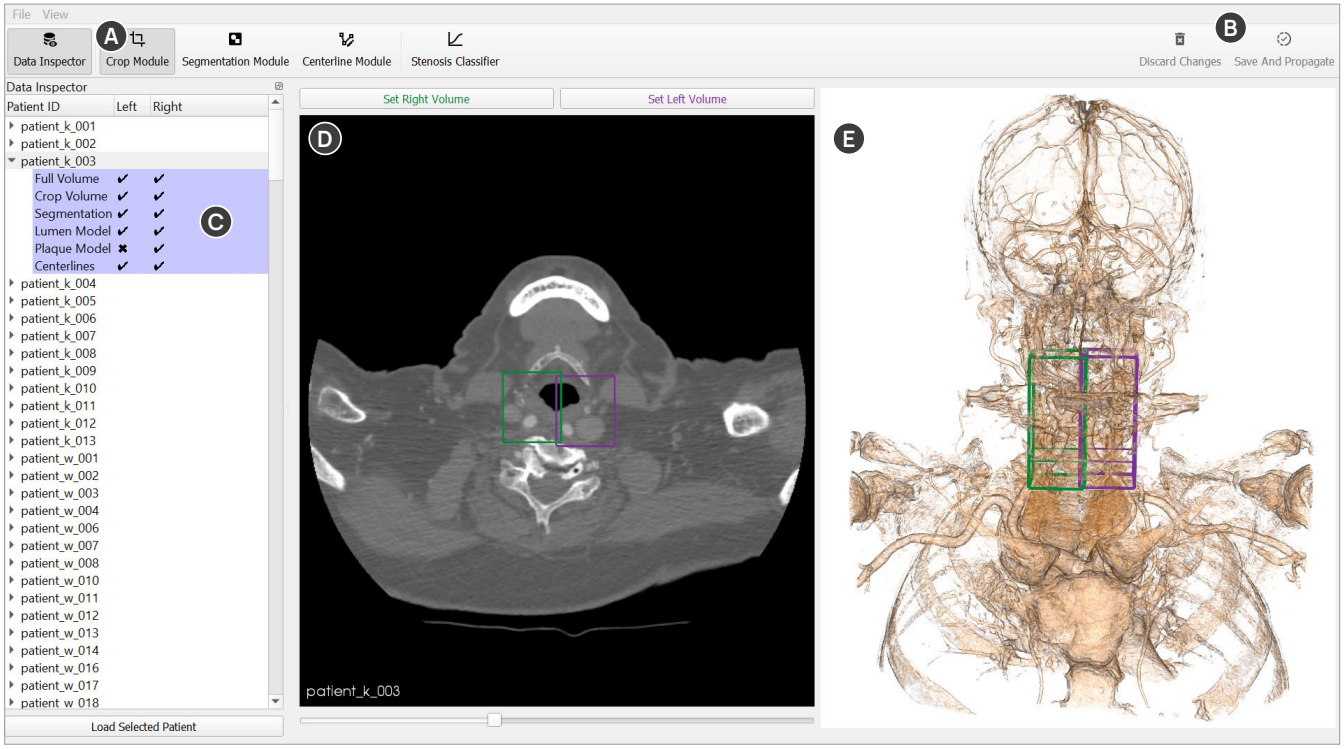

Figure 3: *User interface of the analysis framework. (A) Module selection bar. The crop module is currently active. (B) Global discard and save buttons. (C) The data inspector shows the list of available data sets and which files exist for the left and right carotid. (D) Slice view of the crop module. The left and right volumes can be set here. (E) Volume rendering of the CTA scan with the selected crop regions highlighted.*

data sets, however, may lead to unwanted side effects, such as algorithms and visualization techniques that are biased toward certain data attributes. Repeating the performed processing steps with new data might be more challenging than initially anticipated. In the worst case, this results in low generalizability and transferability.

4. System Design

In this work, we propose a first approach to address the described bottlenecks using the example of visual carotid morphology analysis. Advancing the detection and classification of carotid stenoses has a high potential value due to the frequency and severity of the affliction. In routine diagnostics, a suspected carotid stenosis is screened using angiographic volume imaging, typically a head and neck CTA. A crucial factor for deciding if surgery is necessary is the stenosis degree S. Following the guidelines from the North American Symptomatic Carotid Endarterectomy Trial (NASCET) [FEB*99], a high indication for surgical intervention is given for symptomatic patients with $S \geq 70\%$. To derive the stenosis degree, the smallest diameter inside the suspected stenosis d_{min} and the poststenotic "normal" diameter of the internal carotid d_{normal} need to be measured:

$$S = \frac{d_{normal} - d_{min}}{d_{normal}} \cdot 100\%. \quad (1)$$

These values are usually taken from CTA or sonography imaging. In practice, the exact positions for d_{min} and d_{normal} are not known and are subject to the estimate of the attending physician. The stenosis degree then needs to be considered in light of all factors, including the general anatomy, progression, and composition of the stenotic region.

In discussions with two collaborating physicians who treat stroke patients, one radiologist and one neurologist, we found that a tailored visualization based on an extracted model of the carotid could aid in these tasks. The efficiency could be improved if the identification and classification of stenoses could be performed faster. Also, the classification accuracy would benefit from a computed diameter that can be used to derive the stenosis degree objectively and automatically. Additionally, general qualitative aspects could be enhanced, like the assessment of the overall anatomy, insight into the 3D distribution of plaque, and communication between professionals could be fostered. To meet these goals, however, we require a 3D reconstruction of the vessel lumen, as well as a diameter model, e.g., based on a centerline. The necessary processing of the image data automatically runs into the bottlenecks discussed above.

As a way to alleviate the low transferability that usually results from complex data processing, we propose a fully integrated, adaptive, and modular pipeline. Integrated means that all tasks around data processing *and* visualization can be performed within a single software framework. A minimal data interface that only requires

the CTA image data in standardized DICOM format as input ensures maximal compatibility (**B3 Compatibility**). By adaptive, we mean that manual edits at any pipeline stage will propagate to the later stages automatically, where only the necessary parts of the data are updated. This greatly simplifies validating and correcting the models (**B2 Efficiency**) and fosters a better understanding of the processing steps (**B1 Trust**). Modular means that the pipeline is compartmentalized into secluded processing and visualization stages with minimal data interfaces between them. Processing modules may be optional and new modules can be easily integrated. The database resulting from the processing modules is shared between distinct visualization modules. This means new visualizations that provide a different view of the data can simply be appended, allowing direct comparisons and evaluations.

The overarching objective of our work is to increase the efficiency, accuracy, and quality of clinical workflows. However, we noticed that more concern should fall on intermediary goals, primarily increasing the trust medical users have in the data views. We consider this the only way to ultimately increase the adoption rate of promising vessel visualizations. With the integrated pipeline, we target to give the user insight into the processing stages directly. We believe these stages should be automated where sensible, but at the same time control should be enabled where required. We found that this balancing act of controllability versus ease of use is a considerable challenge on its own.

5. System Implementation

The pipeline takes as input only the raw CTA images of patients as they are exported by clinical workstations (DICOM format). As shown in Figure 1, three data processing modules are integrated for cropping the original volume, segmenting the carotid lumen and plaque, and computing the vessel centerlines and diameter information. In each stage, the results are saved in standard formats. Volumes and masks are written as nearly raw raster data (nrrd) and triangulated surface models as STL files. This enables easy integration of further processing modules if desired. Similarly, plain visualization modules can be appended, which do not modify the data but only visualize it. We propose one such module for detecting and classifying stenoses. The system user interface is shown in Figure 3. The pipeline stages are layered on top in left-to-right order. In each stage, the user can save and propagate or discard changes they made. The selected module is shown in the window center. A database view can be toggled that displays the imported patients and shows which files were created per case. All files can also be externally used if desired. We append additional header information in the files to achieve cross-compatibility with typical volume image processing tools like 3D Slicer.

The implementation is Python-based with backend C/C++ libraries to allow an efficient setup and development while reducing any performance compromises. The user interface is developed in PyQt for platform independency. New modules can be easily integrated as Qt widgets. The segmentation module uses PyTorch for machine learning with a U-Net architecture from MONAI [CLBe22]. The 3D visualizations are based on the Python VTK wrapper and 2D graphs are created with PyQtGraph [Cam22] to enable smooth real-time interaction.

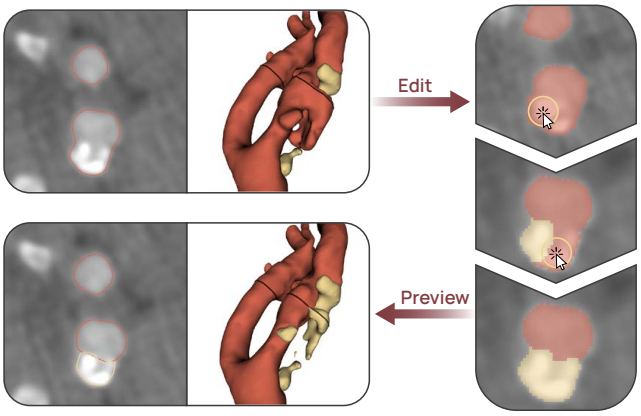

Figure 4: *In the segmentation module, the result of the automatic labeling can be reviewed in a slice-based view (left). The resulting 3D model is shown in a second view. If adjustments are required, the user can switch into edit mode (right), where multiple drawing tools are available to correct the segmentation mask.*

5.1. Crop Module

Before the vessel can be segmented from the input image, the regions of interest need to be cropped. The subsequent automatic labeling stage uses a convolutional neural network (CNN), whose quality can be significantly improved by consistently providing a small target domain with a fixed resolution. To crop the image, we first show an axial slice view of the volume (see Figure 3 left), which the users are familiar with, as it is commonly used to analyze the carotids in CTA. The user can set the desired windowing by dragging on the image, which is a common interaction in radiology workstations. Selecting any point close to the carotid bifurcation automatically positions the bounding box of the crop region around that point. We include voxels in a fixed region of $35 \times 45 \times 80$ millimeters, which describes a reasonable section around the carotid bifurcation in adults to cover any possible stenoses. The region of interest is also indicated in a volume rendering (see Figure 3 right) to facilitate the quick assessment of the 3D box position. We use a transfer function preset that highlights vascular structures, but the transfer function may also be modified. The user can place one crop region each for the left and right sides. When saved, the volume is cropped and interpolated to a fixed resolution of $120 \times 144 \times 248$ voxels using the windowed sinc kernel, as it provides favorable passband characteristics, preserving local contrasts [LGS99]. If a segmentation already exists, for example, if the user decides to adjust the region of interest retrospectively, the segmentation label map is not discarded but only cropped to align with the new region extends.

5.2. Convolutional Neural Network

Deep neural networks, CNNs in particular, have rapidly evolved and been widely used in medical image analysis [LKB*17]. Inspired by the biological visual system [Lin21], CNNs were designed to process data with grid patterns through convolution operation with kernels, and thus are ideal for image processing and

analysis. One of the main components of a CNN is a group of convolutional layers, which can automatically extract spatial features from an input image in a hierarchical manner. For image segmentation, one of the most popular candidates is the U-Net [RFB15], which is a variant of CNN proposed for biomedical image segmentation. A U-Net consists of a contracting (encoding) path for feature extraction and a symmetric expanding (decoding) path for precise localization and image dimension restoration. U-Nets have been extensively studied and improved in recent years [ÇAL*16,DCL*20,SPED21]. To segment the carotids in CTA volumes, we chose a state-of-the-art 3D U-Net implemented in PyTorch from MONAI [CLBe22]. The input to the network is the cropped CTA dataset. The output is the predicted segmentation mask with three channels: background, lumen, and plaque. We manually segmented 37 datasets, which we then split into training (25), validation (6), and test (6) sets. The preprocessing steps include windowing (Hounsfield units [180, 650]) to remove irrelevant structures, as well as normalization to the intensity range of [0, 1]. The data was augmented by random horizontal and vertical flipping with a probability of 0.5. The network was optimized using Dice loss, which is robust against class imbalance [SLV*17] and trained for 100 epochs with a learning rate of 0.001 and a batch size of 4. The network performance was evaluated on the test datasets, reaching a Dice score of 0.776. Note that this score contains the plaque label whose shape and proportions vary extensively between datasets and is thus hard to train for.

5.3. Segmentation Module

While the automatic labeling of the carotid vessel lumen and plaque produces highly useful results and the trained models have substantially improved in recent years, the segmentation masks do sometimes contain errors. Especially patients with atypical anatomy, e.g., with complex plaque distributions, are difficult to train for. To cover all possible input data, a correction mechanism must be in place that ideally allows the user to quickly inspect and adjust the segmentation mask without leaving the application. In the segmentation module, we show an axial view of the cropped CTA region with an outline of the segmentation mask rendered on top (see Figure 4). To make adjustments of the segmentation directly visible, next to the slice view we show the extracted 3D models of the lumen and plaque. The model is continuously and automatically generated by applying discrete marching cubes [Gro16] to the segmentation label map. It is smoothed with a windowed sinc filter [TZG96] (20 iterations, passband 0.005).

When the user decides to switch to edit mode, we show a direct overlay of the discrete segmentation mask in which pixel-precise changes with a brush and eraser are possible. With the brush, new pixels are labeled as plaque or lumen while the eraser removes the label of the current volume and inscribes the pixels as background. The size of the brush can be adjusted with a slider. The user can choose to draw or erase either two-dimensionally, i.e., in the individual slices, for precise corrections, or three-dimensionally with a spherical volume to cover more area faster. A marker shows the current drawing position to facilitate the navigation of the brush both in the slice and the 3D view. To further simplify the manual editing, a threshold can be set which restricts drawing to the image values above. This threshold allows locally adjusting the image value range that is considered lumen or plaque and quickly drawing over larger areas without overwriting the background. The finished segmentation can then be saved and the generated models are propagated to the subsequent modules. A fundamental advantage of the integrated pipeline is that edits can also be made retrospectively, for example, if during the analysis of the vessel morphology an inaccuracy in the model is suspected. Retrospective changes are also automatically forwarded and the data is updated in all consecutive modules.

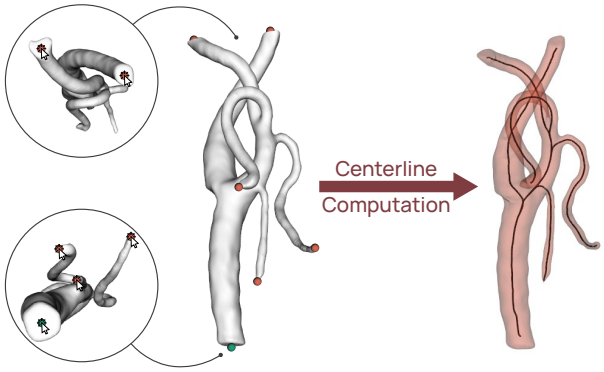

Figure 5: *In the centerline module, the user can select the approximate start and endpoints of the desired centerline tree by clicking on the model. The resulting lines can be reviewed before they are saved.*

5.4. Centerline Module

To enable analyzing the change in vessel width, we compute a centerline tree using the medial axis method, which gives us the radius of the maximally inscribed sphere, i.e., the minimal internal radius, at every centerline point. The minimal lumen diameter is used by clinicians to determine the stenosis degree, which means, with the radius information, we can effectively specify the stenosis degree at any point of the vessel. To compute the centerlines, we use the VMTK implementation of the method by Antiga et al. [AEIR03]. By clicking on the model, the user can choose the centerline source and target points. This provides flexibility regarding the number of branches to be analyzed. The user may select the tip of the internal carotid as a target, but they can also add the external carotid and any potential further outgoing branches that are of interest. Then, the centerline tree is automatically computed and displayed with the vessel hull as shown in Figure 5, such that the user can quickly visually inspect the results.

5.5. Stenosis Classifier Module

We integrated an interactive visualization module to detect and classify stenoses of the internal carotid as an example of how the results of the processing pipeline can be directly used by physicians. The module consists of two linked views, one that shows diameter plots and a 3D view of the vessel model (see Figure 6 right). The most important information is the vessel diameter, for

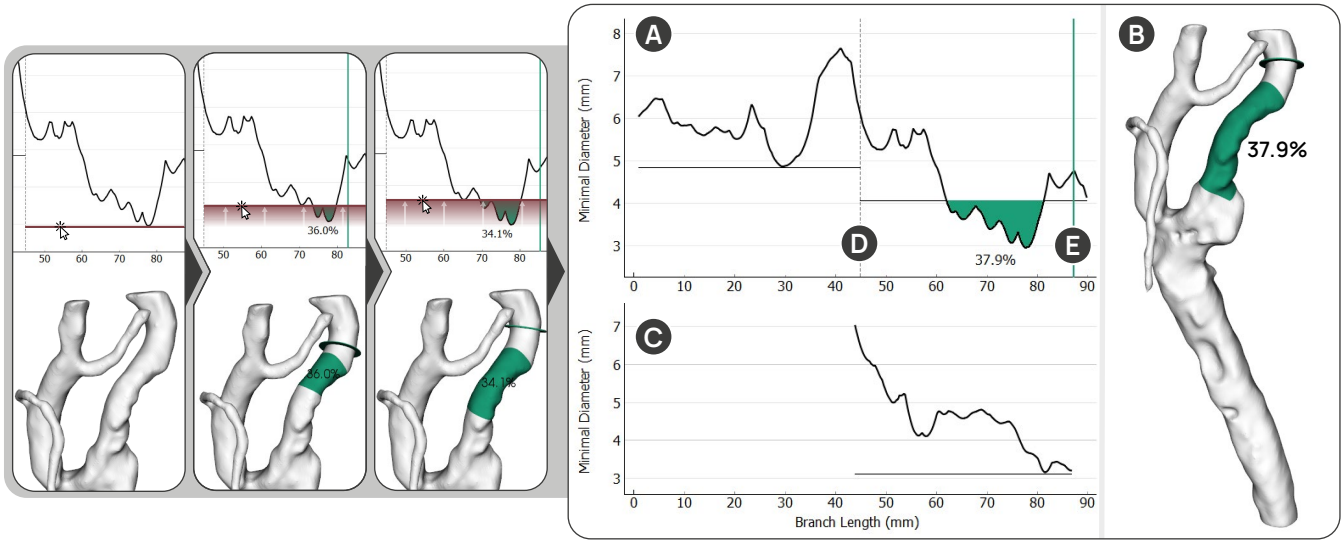

Figure 6: *The classifier module is shown on the right. It is composed of a diameter graph for each branch modeled by a centerline and a 3D view of the vessel lumen. As shown on the left, dragging the branch-specific threshold, indicated by a horizontal line, segments the corresponding stenosis candidate in the surface model. Based on the diameter model, the stenosis degree is automatically computed. (A) Diameter graph of the common and internal carotid. (B) 3D view of the vessel model. (C) Diameter graph of the external carotid. (D) The dashed line shows the bifurcation point. (E) The vertical line can be dragged to set the sampling point for the poststenotic reference diameter. This point is also shown on the 3D model by a small disc.*

which we chose a simple line graph as a direct form of representation. We plot the vessel length (x-axis) against the computed diameter (y-axis). As centerlines for multiple branches might exist, we show one graph for each branch. We considered superposing the line graphs to enable comparisons of the changing diameter after branches, but we later rejected this approach as it makes the individual diameters less readable. Instead, we juxtapose the graphs vertically with a shared x-axis and indicate the branching points with dashed lines (visible in Figure 6D). Hovering over a graph highlights the respective vessel segment in the 3D model to intuitively connect the two visualizations.

Candidates for stenoses can be easily spotted as local minima of the line graphs. To analyze such a candidate, the user can drag a horizontal line over the plot, which interactively sets a local diameter threshold. Everything below the threshold is classified as part of a potential stenosis. This interaction is shown in Figure 6 left. The selected region is filled with a color in the line graph as well as on the 3D model surface to visually link the two representations. If multiple stenotic regions are selected, which is a common occurrence in patients with increasing atherosclerosis, they automatically receive different colors to make them discernible. We apply a qualitative color map, which we generated using the ColorBrewer [BH22].

The central benefit of the interactive region selection is that stenoses can be instantly and automatically classified, removing the need for approximating diameters in slice views of the CTA volume. The stenosis degree, following the NASCET guidelines, is derived from the smallest diameter d_{min} inside the stenosis and the poststenotic diameter d_{normal} of the vessel. We use the local minimum in each user-selected region for d_{min}. As the precise location of d_{normal} is not defined, we first default to a diameter at a position that lies ½ of the length of the stenosis behind it. The user can interactively change this location by dragging a vertical line on the line chart. One such line is given for each detected stenosis and is visually linked using the same color. The position d_{normal} is also marked with a circle in the 3D view, providing the user with an intuitive way to quickly chose a location where they consider the diameter to be normal. The stenosis degree is displayed at d_{min} under the diameter curve and on top of the 3D model. Hovering over a stenosis reveals additional information, such as the stenosis length, and the exact values of d_{min} and d_{normal}. Exemplary results of the automatic stenosis classification with different degrees of severity can be seen in Figure 7.

6. System Evaluation

During the course of the development cycle, we internally evaluated the system's applicability by processing a wide range of input CTA volumes. Our first goal was to validate the generalizability regarding heterogeneous input (**B3**). The second goal was to establish a database of carotid bifurcation models that is publicly accessible for use in further research. We imported head and neck CTA scans from two different clinics. In total, we processed 60 patients, resulting in about 120 carotid bifurcation models, one each for the left and right sides. Then, to assess the usefulness of the integrated pipeline for carotid morphology analysis, we performed a user study with four physicians who routinely treat stroke patients. Our goal was to gather impressions regarding the encodings we use for interactive stenosis classification and also regarding the concept

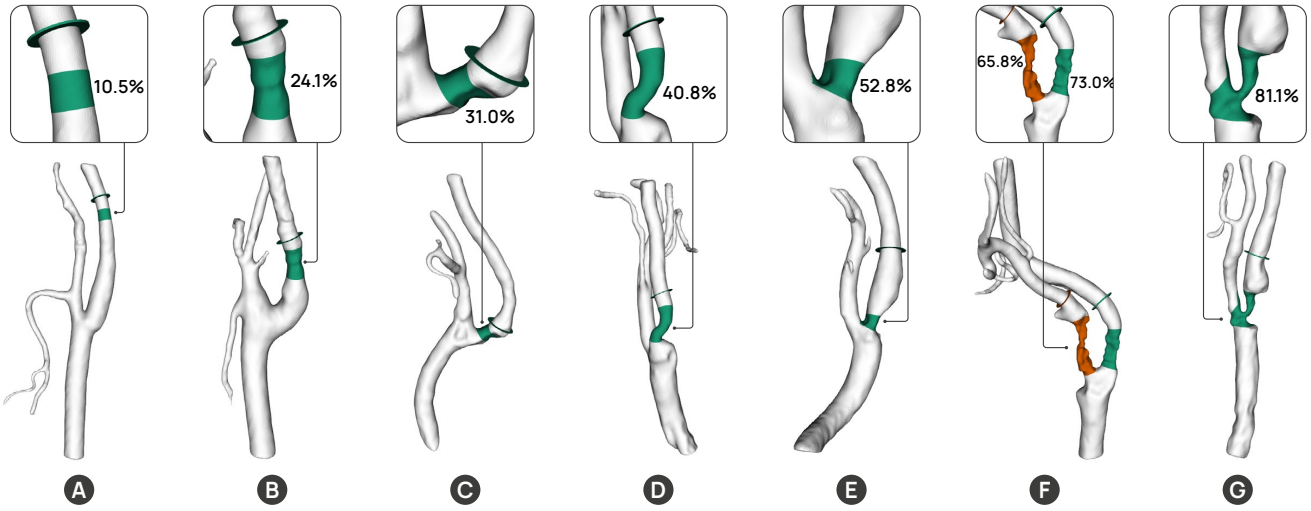

Figure 7: *Exemplary results of the stenosis classifier with increasing stenosis degree from left to right. The disc glyph marks the sampling point for the poststenotic reference diameter. If multiple stenoses are selected, they are differentiated by color, as shown in (F).*

of the integrated and adaptive pipeline in general. In particular, we wanted to know if the clinicians' trust (**B1**) is aided by the option to inspect the processing steps and whether they find the workflow efficient enough (**B2**) to be applicable in practice.

6.1. Method

We recruited four neurologists by word of mouth (P1-P4; two female, two male; ages 34-42). They have 17, 13, 8, and 15 years of working experience in clinics for neurology, respectively. P1 is a direct collaborator and co-author of this paper, the others are not associated with the project. None have used the developed software framework before. We conducted individual interviews lasting about 60 minutes. We first explained the pipeline concept and goals and, subsequently, demonstrated the execution of all processing steps using an example case. Then, we asked the physicians to evaluate the carotid morphology and classify possible stenoses of three patients using the framework. The data sets were not preprocessed, so they had to perform the cropping, segmentation, and centerline creation themselves. We chose data sets with high variation in the carotid geometries and stenosis degrees. One contained no stenosis, one represented an edge case with mild stenosis, and one had strong calcifications and multiple stenoses. We deliberately included a case where we found the segmentation prediction to often exhibit errors, to see if the participants would find and correct them.

After they interacted with the framework, the participants were given a questionnaire to rate their impressions of the pre-processing workflow, the carotid stenosis classification module, and their overall trust in the application. We used a series of statements, summarized in Figure 8, regarding the workflow and modules, which the participants were to score on a five-point Likert scale $(--, -, \circ, +, ++)$. We further discussed multiple open questions, including whether the data views in the analysis framework would present new possibilities for clinical practice and if they found the insights into the processing pipeline necessary and helpful. We also compared the results of the stenosis classification performed by the participants with the stenosis degree from the radiology report of each case. If the values differed notably, we asked the physicians how they believed the difference came about.

6.2. Findings

We report on frequent comments and feedback that we gathered from the interviews. Overall, the framework was very well received. Participants repeatedly mentioned the advantage of the interactive stenosis analysis based on the 3D reconstruction, as compared to the image-based representation they currently use:

> "*Understanding the location, the shape, the length [of a stenosis] is immensely important when we discuss surgery if we should do an endarterectomy [open plaque removal] or stenting [insertion of a wire mesh tube]. The 3D model would be much better to discuss this. I particularly like that I can also see the plaque. We often need to estimate if the stenosis is stable or if emboli can break off.*" (P4)

All physicians stated they would find the tool highly useful to support the screening of the carotids. P3 and P4 immediately asked how they could import new patients. P1 and P4 requested details about the implementation and how complicated the setup of the framework would be on their local system.

The comments the participants made while using the framework were also reflected in the questionnaire. The results are shown in detail in Figure 8. All participants found executing the pipeline straightforward and the processing steps and interactions comprehensible. Only when correcting errors in the segmentation mask, did P2 and P4 note that the interaction was occasionally tricky and would require some training to be smooth and efficient. The physicians agreed that the stenosis classification module enables highly

efficient identification and classification of stenoses. P3 stated, "*Seeing the dips in the graph, I can very quickly highlight the stenosis. That it instantly gives me the stenosis degree is very impressive.*" P4 said, "*Computing the NASCET value with this is super efficient. If the model has no [segmentation] errors, doing all the steps before is also quite fast, so I think it could be integrated into practice.*" All four noted that the ability to smoothly change the reference diameter sampling position is highly advantageous. P1 said, "*It is always different where the poststenotic diameter should be measured. With this, I can tune the position very accurately and get immediate feedback.*"

Regarding their trust in the visualizations, the participants responded very positively. The general impression was that performing or simply being able to inspect the processing steps fosters understanding the model extraction and preventing errors:

> "*I would not want to use this [the classifier module] without being able to check if the model is correct. I can adjust what I need to see and make quick corrections, that is what makes this tool actually useful. The automatic NASCET computation is really neat but it would be pointless if I couldn't see where the model comes from.*" (P1)

> "*Seeing the overlay on the CT is really helpful. I can immediately see if something is not right and correct it. I could even imagine merging the CT with the last tab [the classifier module], so I don't have to switch between the views.*" (P4)

One participant (P2), however, argued that she would rather use the classifier module only, as inspecting the processing stages would be too time-intensive in urgent situations. Nevertheless, she agreed that as long as the processing stages were not widely established and tested for accuracy, there should be an option to manually verify the correctness of the models.

When we checked the stenosis degree computed with the classifier module against the degree given in the radiology report, we sometimes found differences in the values. The participants made the same observation on the example data sets they explored with the framework. Remarkably, when asked which value they thought to be more accurate, all four claimed they would rather trust the application than the report:

> "*It might be more effort to first create the 3D model and make sure it is correct, but the upside is, I think then it gives us the most accurate results. Getting the diameters from the CT can be challenging, it is up to the observer where to measure the two values. With the model, the minimum can be found numerically. And if I can see that the model fits with the CT, I would rather trust the computed value.*" (P4)

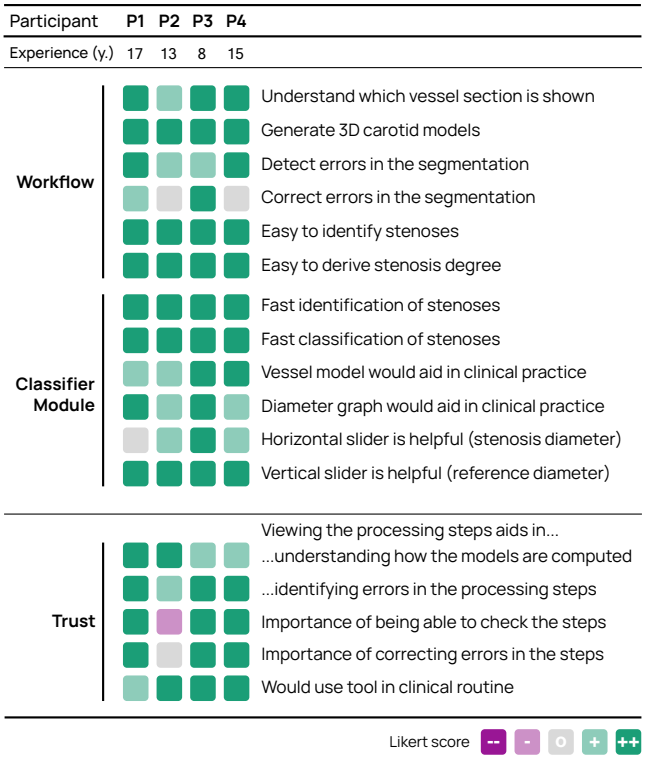

Figure 8: *Results of the questionnaire with color-encoded Likert scores. Each box represents the answer of one physician.*

7. Discussion

Our discussions during the interviews showed that the visualizations of the extracted carotid bifurcation models would benefit the understanding of the morphology and aid in clinical decision-making. The physicians repeatedly stated that the demonstrated visual encodings allowed them to intuitively assess the lumen shape, the location of stenoses, and the distribution of plaque. The automatic stenosis labeling based on coordinated views, as implemented in the classifier module, was considered highly promising. Here, the users particularly cherished real-time interactions, which are directly reflected in the visualizations. These enable experimenting with suspected stenosis sites considerably more efficiently than when relying on slice-based representations only. These observations are underlined by Figure 7, where it becomes evident that stenoses, even reduced to those at the carotid bifurcation, vary in their location, shape, and size, making individual assessments indispensable. If the models are correctly extracted, the resulting stenosis degrees are also highly accurate, as the minimum internal diameter can be precisely computed.

With this work, we attempted to address typical bottlenecks in the clinical transfer that arise for visualizations of processed medical data. In the interviews, we found that the concept of the fully integrated pipeline allows users to better understand the different processing stages and assure the correctness of what they see (**B1**). To enable these insights, tailored visualizations are key, as they allow uncomplicated inspection of each processing step. From the feedback we gathered, we determine that this approach is an effective way to lower the adoption costs (**B2**) since the combined framework can be set up once as a self-contained unit. The only inputs are the already existing CTA volumes in the format they are exported from clinical workstations. This results in increased transferability, which we noticed during the interviews, as participants were eager

to try the framework on their own data, for which they would only need to import existing image volumes (**B3**). Furthermore, to combat algorithmic bias (**B4**), we leveraged the streamlined processing to create a comprehensive database of models extracted from 60 different patients. This effort was only viable due to the tightly intertwined processing modules, which allow new image volumes to be swiftly processed and later edits to be directly propagated. As a result, the modules were developed and tested with a large variety of geometries, increasing their robustness against different inputs.

It should be noted that even with **B1**-**B4** addressed, the transfer of methods into clinical everyday use still presents a major challenge. While the experts agreed that the pipeline would be a useful tool, many more certification and integration steps are necessary to actually deploy such a framework. For these to be feasible, any change in a clinical procedure must prove to provide a substantial benefit that would justify integration costs. Therefore, in a next step, we intent to deploy the pipeline in clinical research, gathering further feedback and testing the applicability in the real world.

7.1. Controllability vs. Ease of Use

When asked about the importance of being able to inspect and potentially correct the processing steps, one of the interviewed experts argued that the required time to do this properly would be too high in an applied clinical context. The others disagreed and two even highlighted the ability to control the pipeline themselves as the most important feature for enabling adoption. This is indicative of a core challenge. In an integrated application, like the presented pipeline, should we favor controllability or ease of use from the users' perspective? Allowing more control will likely produce more accurate results. Furthermore, our evaluation indicates that pipeline insight and control do increase user trust and understanding of the processed data. However, increasing control options makes the whole process inevitably more complex, which might ultimately backfire and hinder adoption. On the other hand, favoring automated processing would benefit efficiency, making the application more suited for time-critical scenarios. Automating and hiding processing stages, though, comes at the cost of less fidelity, which could lead to decreased validity of the results.

The challenge is to reduce control options up to the point where they have the maximum benefit. For example, in the proposed pipeline, the user may select the endpoints for the centerline computation. The points could be computed automatically in the background [LL20], but the selection can be done quickly, as only a handful of points need to be clicked, and this short task greatly improves the results. Automatic endpoint selection can be unstable and the user might only be interested in a subset of vascular branches. We found a promising rule of thumb is to automate processes where reasonably possible, but always let the user inspect stages where automation is prone to fail. For example, the machine learning models for vessel segmentation have become impressively capable of delivering fast and "good enough" results. Still, to cover unusual cases, the user must be able to correct the geometries. If inaccuracies in the processing are sometimes observed, in a medical treatment context, the whole application becomes useless if these errors cannot be reliably detected and rectified.

7.2. Opportunities and Future Work

The integrated pipeline proposed in this work, together with the database of already processed and tuned carotid bifurcation models, lays the foundation for the efficient development of further extensions. The framework is implemented using a modularized architecture, where additional processing or visualization stages can simply be inserted. Conceivable are, for example, modules geared for hemodynamic simulation analysis or epidemiological cohort studies. Existing processes could also be improved, for instance, it would be beneficial if the segmentation prediction network would continuously and automatically learn from the user-corrected masks. In clinical practice, the carotid stenosis degree is currently estimated from diameters measured in CTA scans, as well as from blood flow velocities recorded with Doppler sonography. We plan to evaluate the automatic stenosis classification in a clinical comparative study to investigate the differences between these methods.

8. Conclusion

We presented a system for the visual analysis of the carotid morphology to improve cardiovascular disease screening and treatment. We implemented our methods in a modular and fully integrated processing and visualization pipeline. By combining state-of-the-art techniques, like pairing a neural network for segmentation mask prediction with graphics-based user control, we make the pre-processing required for advanced visualization applicable to practical scenarios. We show how the processed data can be used to interactively detect and classify stenoses quickly and reliably.

A general takeaway from this work is that the integration of medical data processing and visualization is a challenging task. Implementing and testing the presented system was a considerable collaborative and time-intensive effort, not least because of the difficulty of connecting previously detached solutions for model extraction and interactive visualization. Yet, we found that this type of integration may make the adoption of methods feasible that could improve medical imaging analysis tasks. A key question is where the user should be able to interrupt and adapt automatic processing. Favoring higher accuracy in the results might be desirable but is also more time-intensive for the user and adjusting complicated algorithms may also be hard to learn. Our impression from this line of work is that if intensive processing is a requirement for novel medical visualization tools, for example, if segmenting images is necessary, any required interaction must be fast enough such that introducing a new way to analyze the data does not significantly disrupt the clinical workflow. An interesting opportunity arises from the fact that machine learning models, like CNNs, are just becoming good enough that they can substitute tasks in this process that would otherwise be too time-intensive. We hope that the insights gathered in this work will spawn further endeavors to connect fragmented processing and visualization techniques in medical applications, which may ultimately pave the way for applying advanced visual analysis tools in medical practice.

Acknowledgements: This project was funded by the BMBF Joint Project 05M20SJA-MLgSA. Open Access funding enabled and organized by Projekt DEAL.

References

[AEIR03] ANTIGA L., ENE-IORDACHE B., REMUZZI A.: Computational geometry for patient-specific reconstruction and meshing of blood vessels from MR and CT angiography. *IEEE Transactions on Medical Imaging 22*, 5 (may 2003), 674–684. doi:10.1109/tmi.2003.812261. 3, 6

[APB*08] ANTIGA L., PICCINELLI M., BOTTI L., ENE-IORDACHE B., REMUZZI A., STEINMAN D. A.: An image-based modeling framework for patient-specific computational hemodynamics. *Medical & Biological Engineering & Computing 46*, 11 (2008), 1097. doi:10.1007/s11517-008-0420-1. 3

[ATC*08] AU O. K.-C., TAI C.-L., CHU H.-K., COHEN-OR D., LEE T.-Y.: Skeleton extraction by mesh contraction. *ACM Transactions on Graphics 27*, 3 (aug 2008), 1–10. doi:10.1145/1360612.1360643. 3

[BH22] BREWER C. A., HARROWER M.: ColorBrewer, 2022. URL: https://colorbrewer2.org/. 7

[BKS01] BITTER I., KAUFMAN A., SATO M.: Penalized-distance volumetric skeleton algorithm. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 195–206. doi:10.1109/2945.942688. 3

[BKS18] BOZKURT F., KÖSE C., SARI A.: An inverse approach for automatic segmentation of carotid and vertebral arteries in CTA. *Expert Systems with Applications 93* (mar 2018), 358–375. doi:10.1016/j.eswa.2017.10.041. 3

[BPS*10] BAUER C., POCK T., SORANTIN E., BISCHOF H., BEICHEL R.: Segmentation of interwoven 3D tubular tree structures utilizing shape priors and graph cuts. *Medical Image Analysis 14*, 2 (apr 2010), 172–184. doi:10.1016/j.media.2009.11.003. 2

[BSB*00] BITTER I., SATO M., BENDER M., MCDONNELL K., KAUFMAN A., WAN M.: CEASAR: a smooth, accurate and robust centerline extraction algorithm. In *Proceedings Visualization 2000. VIS 2000 (Cat. No.00CH37145)* (2000), IEEE. doi:10.1109/visual.2000.885675. 3

[ÇAL*16] ÇIÇEK Ö., ABDULKADIR A., LIENKAMP S. S., BROX T., RONNEBERGER O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2016), Springer, pp. 424–432. 6

[Cam22] CAMPAGNOLA L.: PyQtGraph, 2022. URL: https://www.pyqtgraph.org/. 5

[CCR20] CHOI G. P. T., CHIU B., RYCROFT C. H.: Area-preserving mapping of 3D carotid ultrasound images using density-equalizing reference map. *IEEE Transactions on Biomedical Engineering 67*, 9 (sep 2020), 2507–2517. doi:10.1109/tbme.2019.2963783. 3

[CKM04] CULVER T., KEYSER J., MANOCHA D.: Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design 21*, 1 (jan 2004), 65–98. doi:10.1016/j.cagd.2003.07.008. 3

[CLBe22] CARDOSO M. J., LI W., BROWN R., ET AL.: MONAI: An open-source framework for deep learning in healthcare. doi:https://doi.org/10.48550/arXiv.2211.02701. 5, 6

[CLC13] CHIU B., LI B., CHOW T. W. S.: Novel 3D ultrasound image-based biomarkers based on a feature selection from a 2D standardized vessel wall thickness map: a tool for sensitive assessment of therapies for carotid atherosclerosis. *Physics in Medicine and Biology 58*, 17 (aug 2013), 5959–5982. doi:10.1088/0031-9155/58/17/5959. 3

[CVOA08] CUISENAIRE O., VIRMANI S., OLSZEWSKI M. E., ARDON R.: Fully automated segmentation of carotid and vertebral arteries from contrast enhanced CTA. In *SPIE Proceedings* (mar 2008), Reinhardt J. M., Pluim J. P. W., (Eds.), SPIE. doi:10.1117/12.770481. 3

[DCL*20] DU G., CAO X., LIANG J., CHEN X., ZHAN Y.: Medical image segmentation based on U-Net: A review. *Journal of Imaging Science and Technology 64*, 2 (mar 2020), 20508–1–20508–12. doi:10.2352/j.imagingsci.technol.2020.64.2.020508. 6

[EEL*19] EULZER P., ENGELHARDT S., LICHTENBERG N., SIMONE R. D., LAWONN K.: Temporal views of flattened mitral valve geometries. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. doi:10.1109/tvcg.2019.2934337. 2

[EL23a] EULZER P., LAWONN K.: Carotid Analyzer, 2023. URL: https://github.com/PepeEulzer/CarotidAnalyzer. 2

[EL23b] EULZER P., LAWONN K.: A dataset of reconstructed carotid bifurcation lumen and plaque models with centerline tree, 2023. doi:10.5281/ZENODO.7634643. 2

[ELD10] ESNEAULT S., LAFON C., DILLENSEGER J.-L.: Liver vessels segmentation using a hybrid geometrical moments/graph cuts method. *IEEE Transactions on Biomedical Engineering 57*, 2 (feb 2010), 276–283. doi:10.1109/tbme.2009.2032161. 2

[EMKL21] EULZER P., MEUSCHKE M., KLINGNER C. M., LAWONN K.: Visualizing carotid blood flow simulations for stroke prevention. *Computer Graphics Forum 40*, 3 (jun 2021), 435–446. doi:10.1111/cgf.14319. 3

[EMML22] EULZER P., MEUSCHKE M., MISTELBAUER G., LAWONN K.: Vessel maps: A survey of map-like visualizations of the cardiovascular system. *Computer Graphics Forum 41*, 3 (jun 2022), 645–673. doi:10.1111/cgf.14576. 2, 3

[ERM*21] EULZER P., RICHTER K., MEUSCHKE M., HUNDERTMARK A., LAWONN K.: Automatic cutting and flattening of carotid artery geometries. *Eurographics Workshop on Visual Computing for Biology and Medicine* (2021). doi:10.2312/VCBM.20211347. 3

[FEB*99] FERGUSON G. G., ELIASZIW M., BARR H. W. K., CLAGETT G. P., BARNES R. W., WALLACE M. C., TAYLOR D. W., HAYNES R. B., FINAN J. W., HACHINSKI V. C., BARNETT H. J. M.: The north american symptomatic carotid endarterectomy trial. *Stroke 30*, 9 (sep 1999), 1751–1758. doi:10.1161/01.str.30.9.1751. 4

[GBD19] GBD 2016 STROKE COLLABORATORS: Global, regional, and national burden of stroke, 1990–2016: a systematic analysis for the global burden of disease study 2016. *The Lancet Neurology 18*, 5 (2019), 439 – 458. doi:10.1016/S1474-4422(19)30034-1. 2

[Gro16] GROTHAUSMANN R.: Providing values of adjacent voxel with vtkDiscreteMarchingCubes. *The VTK Journal* (jul 2016). doi:10.54294/2aeqx3. 6

[GSG*21] GILLMANN C., SMIT N. N., GROLLER E., PREIM B., VILANOVA A., WISCHGOLL T.: Ten open challenges in medical visualization. *IEEE Computer Graphics and Applications 41*, 5 (sep 2021), 7–15. doi:10.1109/mcg.2021.3094858. 3

[HTW*09] HALM E. A., TUHRIM S., WANG J. J., ROCKMAN C., RILES T. S., CHASSIN M. R.: Risk factors for perioperative death and stroke after carotid endarterectomy: results of the New York carotid artery surgery study. *Stroke 40*, 1 (2009), 221–229. doi:10.1161/STROKEAHA.108.524785. 2

[ISMA18] IZZO R., STEINMAN D., MANINI S., ANTIGA L.: The vascular modeling toolkit: a python library for the analysis of tubular structures in medical images. *Journal of Open Source Software 3*, 25 (2018), 745. 3

[JN*18] JAWAID M. M., , NAREJO S., QURESHI I. A., PIRZADA N.: A review of the state-of-the-art methods for non-calcified plaque detection in cardiac CT angiography. *International Journal of Computer Theory and Engineering 10*, 3 (2018), 84–92. doi:10.7763/ijcte.2018.v10.1204. 2

[KGNP12] KUBISCH C., GLASSER S., NEUGEBAUER M., PREIM B.: Vessel visualization with volume rendering. In *Visualization in Medicine and Life Sciences II. Mathematics and Visualization*, Linsen L., Hagen H., Hamann B., Hege H. C., (Eds.). Springer, Berlin, Heidelberg, 2012, pp. 109–132. doi:10.1007/978-3-642-21608-4_7. 2

[KPV13] KIKINIS R., PIEPER S. D., VOSBURGH K. G.: 3D Slicer: A platform for subject-specific image analysis, visualization, and clinical support. In *Intraoperative Imaging and Image-Guided Therapy*. Springer New York, nov 2013, pp. 277–289. doi:10.1007/978-1-4614-7657-3_19. 3

[LABFL09] LESAGE D., ANGELINI E. D., BLOCH I., FUNKA-LEA G.: A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical Image Analysis 13*, 6 (2009), 819–845. 2

[LGS99] LEHMANN T., GONNER C., SPITZER K.: Survey: interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging 18*, 11 (1999), 1049–1075. doi:10.1109/42.816070. 5

[Lin21] LINDSAY G. W.: Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of Cognitive Neuroscience 33*, 10 (2021), 2017–2031. 5

[LK16] LISKOWSKI P., KRAWIEC K.: Segmenting retinal blood vessels with deep neural networks. *IEEE Transactions on Medical Imaging 35*, 11 (nov 2016), 2369–2380. doi:10.1109/tmi.2016.2546227. 3

[LKB*17] LITJENS G., KOOI T., BEJNORDI B. E., SETIO A. A. A., CIOMPI F., GHAFOORIAN M., VAN DER LAAK J. A., VAN GINNEKEN B., SÁNCHEZ C. I.: A survey on deep learning in medical image analysis. *Medical Image Analysis 42* (2017), 60–88. 5

[LL20] LICHTENBERG N., LAWONN K.: Parameterization, feature extraction and binary encoding for the visualization of tree-like structures. *Computer Graphics Forum 39*, 1 (nov 2020), 497–510. doi:10.1111/cgf.13888. 2, 10

[LMW*19] LAWONN K., MEUSCHKE M., WICKENHÖFER R., PREIM B., HILDEBRANDT K.: A geometric optimization approach for the detection and segmentation of multiple aneurysms. *Computer Graphics Forum 38*, 3 (jun 2019), 413–425. doi:10.1111/cgf.13699. 3

[MGB*19] MEUSCHKE M., GUNTHER T., BERG P., WICKENHOFER R., PREIM B., LAWONN K.: Visual analysis of aneurysm data using statistical graphics. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (jan 2019), 997–1007. doi:10.1109/tvcg.2018.2864509. 2

[MKP*16] MEUSCHKE M., KÖHLER B., PREIM U., PREIM B., LAWONN K.: Semi-automatic vortex flow classification in 4d pc-mri data of the aorta. 351–360. doi:https://doi.org/10.1111/cgf.12911. 3

[MMHM18] MOCCIA S., MOMI E. D., HADJI S. E., MATTOS L. S.: Blood vessel segmentation algorithms — review of methods, datasets and evaluation metrics. *Computer Methods and Programs in Biomedicine 158* (may 2018), 71–91. doi:10.1016/j.cmpb.2018.02.001. 3

[MMNG16] MIAO H., MISTELBAUER G., NAŠEL C., GRÖLLER M. E.: Visual quantification of the circle of willis: An automated identification and standardized representation. *Computer Graphics Forum 36*, 6 (jul 2016), 393–404. doi:10.1111/cgf.12988. 2

[MMV*13] MISTELBAUER G., MORAR A., VARCHOLA A., SCHERNTHANER R., BACLIJA I., KÖCHL A., KANITSAR A., BRUCKNER S., GRÖLLER E.: Vessel visualization using curvicircular feature aggregation. In *Computer Graphics Forum* (2013), vol. 32, pp. 231–240. doi:10.1111/cgf.12110. 2

[MOJB*19] MEUSCHKE M., OELTZE-JAFRA S., BEUING O., PREIM B., LAWONN K.: Classification of blood flow patterns in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics 25*, 7 (2019), 2404–2418. doi:10.1109/TVCG.2018.2834923. 3

[MPL21] MEUSCHKE M., PREIM B., LAWONN K.: Aneulysis – a system for the visual analysis of aneurysm data. *Computers & Graphics 98* (aug 2021), 197–209. doi:10.1016/j.cag.2021.06.001. 2

[MVB*17] MEUSCHKE M., VOSS S., BEUING O., PREIM B., LAWONN K.: Glyph-based comparative stress tensor visualization in cerebral aneurysms. In *Computer Graphics Forum* (2017), vol. 36, pp. 99–108. doi:10.1111/cgf.13171. 3

[MVN06] MANNIESING R., VIERGEVER M. A., NIESSEN W. J.: Vessel enhancing diffusionA scale space representation of vessel structures. *Medical Image Analysis 10*, 6 (dec 2006), 815–825. doi:10.1016/j.media.2006.06.003. 2

[OJMN*19] OELTZE-JAFRA S., MEUSCHKE M., NEUGEBAUER M., SAALFELD S., LAWONN K., JANIGA G., HEGE H.-C., ZACHOW S., PREIM B.: Generation and visual exploration of medical flow data: Survey, research trends and future challenges. In *Computer Graphics Forum* (2019), vol. 38, pp. 87–125. doi:10.1111/cgf.13394. 2, 3

[PHM*16] PRENTAŠIĆ P., HEISLER M., MAMMO Z., LEE S., MERKUR A., NAVAJAS E., BEG M. F., ŠARUNIC M., LONCARIC S.: Segmentation of the foveal microvasculature using deep learning networks. *Journal of Biomedical Optics 21*, 7 (jul 2016), 075008. doi:10.1117/1.jbo.21.7.075008. 3

[PO08] PREIM B., OELTZE S.: 3D visualization of vasculature: An overview. In *Visualization in Medicine and Life Sciences. Mathematics and Visualization.*, Linsen L., Hagen H., Hamann B., (Eds.). Springer Berlin Heidelberg, 2008, pp. 39–59. doi:10.1007/978-3-540-72630-2_3. 3

[PRI18] PREIM B., ROPINSKI T., ISENBERG P.: A critical analysis of the evaluation practice in medical visualization. *Eurographics Workshop on Visual Computing for Biology and Medicine* (2018). doi:10.2312/VCBM.20181228. 3

[RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2015), Springer, pp. 234–241. 6

[RFK*07] ROOS J. E., FLEISCHMANN D., KOECHL A., RAKSHE T., STRAKA M., NAPOLI A., KANITSAR A., SRAMEK M., GROELLER E.: Multipath curved planar reformation of the peripheral arterial tree in CT angiography. *Radiology 244*, 1 (jul 2007), 281–290. doi:10.1148/radiol.2441060976. 2

[RGB*19] RAJSIC S., GOTHE H., BORBA H., SROCZYNSKI G., VUJICIC J., TOELL T., SIEBERT U.: Economic burden of stroke: a systematic review on post-stroke care. *The European Journal of Health Economics 20*, 1 (2019), 107–134. doi:10.1007/s10198-018-0984-0. 2

[SLV*17] SUDRE C. H., LI W., VERCAUTEREN T., OURSELIN S., JORGE CARDOSO M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*. Springer, 2017, pp. 240–248. 6

[SPED21] SIDDIQUE N., PAHEDING S., ELKIN C. P., DEVABHAKTUNI V.: U-Net and its variants for medical image segmentation: A review of theory and applications. *IEEE Access 9* (2021), 82031–82057. doi:10.1109/access.2021.3086020. 6

[SRE12] SCHNEIDER C. A., RASBAND W. S., ELICEIRI K. W.: NIH image to ImageJ: 25 years of image analysis. *Nature Methods 9*, 7 (jun 2012), 671–675. doi:10.1038/nmeth.2089. 3

[SSZZ01] SHAHROKNI A., SOLTANIAN-ZADEH H., ZOROOFI R. A.: Fast skeletonization algorithm for 3D elongated objects. In *Proc. SPIE 4322, Medical Imaging 2001: Image Processing* (jul 2001), Sonka M., Hanson K. M., (Eds.), SPIE. doi:10.1117/12.431102. 3

[TZG96] TAUBIN G., ZHANG T., GOLUB G.: Optimal surface smoothing as filter design. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1996, pp. 283–292. doi:10.1007/bfb0015544. 6

[WCH*10] WANG Z., CHI Y., HUANG W., VENKATESH S. K., TIAN Q., OO T., ZHOU J., XIONG W., LIU J.: Comparisons of centerline extraction methods for liver blood vessels in ImageJ and 3D slicer. *APSIPA ASC* (2010), 276–279. 3

[WL08] WANG Y.-S., LEE T.-Y.: Curve-skeleton extraction using iterative least squares optimization. *IEEE Transactions on Visualization and Computer Graphics 14*, 4 (jul 2008), 926–936. doi:10.1109/tvcg.2008.38. 3

[WXG*16] WU A., XU Z., GAO M., BUTY M., MOLLURA D. J.: Deep vessel tracking: A generalized probabilistic approach via deep learning. In *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)* (apr 2016), IEEE. doi:10.1109/isbi.2016.7493520. 3

[ZCHH17] ZHAO F., CHEN Y., HOU Y., HE X.: Segmentation of blood vessels using rule-based and machine-learning-based methods: a review. *Multimedia Systems 25*, 2 (dec 2017), 109–118. doi:10.1007/s00530-017-0580-7. 3

[ZTP*21] ZHOU T., TAN T., PAN X., TANG H., LI J.: Fully automatic deep learning trained on limited data for carotid artery segmentation from large image volumes. *Quantitative Imaging in Medicine and Surgery 11*, 1 (jan 2021), 67–83. doi:10.21037/qims-20-286. 3