


PEPE EULZER

DATA AND SIMULATION-BASED EXPLORATION, ANALYSIS,
AND TREATMENT OF STENOSES FOR PREVENTION OF
ISCHEMIC STROKES





**FRIEDRICH-SCHILLER-
UNIVERSITÄT
JENA**

**DATA AND SIMULATION-BASED EXPLORATION,
ANALYSIS, AND TREATMENT OF STENOSES FOR
PREVENTION OF ISCHEMIC STROKES**

DISSERTATION

zur Erlangung des akademischen Grades

Doktor-Ingenieur (Dr.-Ing.)

vorgelegt dem Rat der Fakultät für Mathematik und Informatik
der Friedrich-Schiller-Universität Jena

von **M.Sc. Pepe Eulzer**
geboren am 24.07.1995 in Trier

Gutachter

1. Prof. Dr.-Ing. habil. Kai Lawonn
Friedrich-Schiller-Universität Jena
2. Prof. Dr.-Ing. habil. Bernhard Preim
Otto-von-Guericke-Universität Magdeburg
3. Univ.Prof. Dr.techn. Eduard Gröller
Technische Universität Wien

Tag der öffentlichen Verteidigung: 05. August 2024

ABSTRACT

Ischemic stroke is one of the primary causes of death and disability worldwide. Common origins of cerebral ischemia are the two carotid arteries that supply most of the intracranial circulation. Therefore, detecting and analyzing carotid stenoses, i.e., potentially lethal constrictions of these arteries, is a critical task in clinical stroke prevention and treatment. Carotid surgery can reduce the stroke risk associated with stenoses, however, the procedure entails risks itself. Therefore, a thorough evaluation of each case is necessary. Clinical evaluation requires assessing the vessel shape (morphology) and the properties of the internal blood flow (hemodynamics). The latter is typically assessed using Doppler ultrasonography, which requires specialized equipment and personnel to be carried out.

With the emergence of computational fluid dynamics (CFD), the simulation of intricate biological transport processes became possible, including flow in actual physiological structures, such as blood vessels. These simulations of the internal blood flow field could facilitate the visual integration of hemodynamic and morphological information, provide a higher resolution on relevant parameters than current screening methods, and enable the analysis of new flow-derived parameters associated with stenosis progression. While hemodynamic simulations offer groundbreaking opportunities to enhance clinical decision-making, a successful translation of CFD to the clinical forefront is highly challenging.

In this thesis, a self-contained pipeline is built, which integrates streamlined model extraction and pre-processing along with a collection of visual exploration tools into a single framework. Methods are proposed and evaluated to analyze the internal blood flow, anatomical context, and vessel wall composition, and to automatically and reliably classify stenosis candidates. Through retrospective user control over the processing stages, error detection and correction are greatly simplified. As hemodynamic simulations are intrinsically complex, time-consuming, and resource-intensive, their application conflicts with the time-sensitive nature of clinical workflows. To address these transfer challenges, a novel visualization system is integrated into the pipeline, which enables instant flow exploration without performing on-site simulation. The framework was developed and evaluated in multiple iterative user studies, involving a group of eight radiologists and neurologists working in stroke care. It is publicly available, along with a database of 152 carotid bifurcation geometries with simulated flow that were extracted with the framework from computed tomography data.

ZUSAMMENFASSUNG

Der ischämische Schlaganfall ist eine der führenden Invaliditäts- und Todesursachen. Ein häufiger Ursprung für die Entstehung zerebraler Ischämien sind die beiden internen Karotisarterien, die den größten Teil des intrakraniellen Kreislaufs versorgen. Daher ist die Erkennung und Analyse von Karotisstenosen, also potenziell tödlichen Verengungen dieser Arterien, eine entscheidende Aufgabe der klinischen Schlaganfallprävention und -behandlung. Ein chirurgischer Eingriff kann das mit Stenosen verbundene Schlaganfallrisiko verringern, doch dieser selbst ist mit Risiken verbunden. Daher ist in jedem Einzelfall eine gründliche Risiko-Nutzen-Abwägung eines möglichen Eingriffs erforderlich. Die klinische Bewertung erfordert eine Beurteilung der Gefäßform (Morphologie) und der Eigenschaften des internen Blutflusses (Hämodynamik). Letzteres wird in der Regel mittels Doppler-Ultraschalluntersuchung beurteilt, für deren Durchführung entsprechende Geräte und spezialisiertes Personal erforderlich sind.

Mit dem Aufkommen der numerischen Strömungsmechanik (engl. computational fluid dynamics, Abk. CFD) wurde die Simulation komplizierter biologischer Transportprozesse möglich, einschließlich der Strömung in tatsächlichen physiologischen Strukturen, wie zum Beispiel Blutgefäßen. Diese Simulationen des internen Blutflussfeldes können die visuelle Integration hämodynamischer und morphologischer Informationen erleichtern, eine höhere Auflösung relevanter Parameter als die derzeitigen Screening-Methoden bieten und die Analyse neuer Parameter im Zusammenhang mit dem Fortschreiten von Stenosen ermöglichen. Während hämodynamische Simulationen neue Möglichkeiten zur Verbesserung der klinischen Entscheidungsfindung bieten, ist eine erfolgreiche Umsetzung von CFD in die klinische Praxis eine große Herausforderung.

In dieser Arbeit wird eine in sich geschlossene Pipeline entwickelt, die eine optimierte Modellextraktion und -vorverarbeitung zusammen mit einer Sammlung von visuellen Untersuchungswerkzeugen in ein einziges Framework integriert. Es werden Methoden zur Analyse des internen Blutflusses, des anatomischen Kontexts und der Gefäßwandzusammensetzung vorgeschlagen und bewertet, um Stenosekandidaten automatisch und zuverlässig zu klassifizieren. Durch die retrospektive Kontrolle des Benutzers über die Verarbeitungsschritte werden Fehlererkennung und -korrektur erheblich vereinfacht. Da hämodynamische Simulationen von Natur aus komplex, zeitaufwändig und ressourcenintensiv sind, steht ihre Anwendung im Konflikt mit dem zeitkritischen Charakter klinischer Arbeitsabläufe. Um diesen Herausforderungen zu begegnen, wurde ein neuartiges Visualisierungssystem in die Pipeline integriert, das eine sofortige Untersu-

chung der Strömung ermöglicht, ohne dass eine Simulation vor Ort durchgeführt werden muss. Das Framework wurde in mehreren iterativen Nutzerstudien, mit einer Gruppe von acht Radiologen und Neurologen, entwickelt und evaluiert. Das Framework ist öffentlich verfügbar, zusammen mit einer Datenbank von 152 Geometrien von Karotisbifurkationen mit simuliertem Blutfluss, die mit dem System aus Computertomographiedaten extrahiert wurden.

DANKSAGUNGEN

Traditionsgemäß möchte ich an erster Stelle meinem Betreuer und Doktorvater Kai Lawonn danken, ohne dessen Unterstützung es diese Arbeit nie gegeben hätte. Kai hat mir die Möglichkeit gegeben wissenschaftlich zu arbeiten als ich noch grün hinter den Ohren war und gerade mal mit dem Bachelor fertig wurde. Nicht nur, dass diese frühe Förderung an sich schon ein großes Privileg war, Kai hat es zusätzlich geschafft, mir von Anfang an das Gefühl zu vermitteln, dass wir auf Augenhöhe sprechen können, was absolut keine Selbstverständlichkeit ist. Kai war es auch, der mir die Promotion ans Herz legte und mir ohne zu zögern angeboten hat mitzukommen als der Ruf nach Jena für ihn Realität wurde. Ich weiß es sehr zu schätzen, dass das Vertrauen in mich so groß war und ist, dass ich mich nie fragen musste was passiert, wenn mal was nicht funktioniert wie geplant (Kai nimmt sich die Zeit und wir suchen zusammen nach einer Lösung), wenn mal eine Deadline nicht machbar ist (nicht schlimm, mal schauen was als nächstes kommt) oder wenn eine Finanzierung ausläuft (Kai hat einen Backup-Plan, klar). Kai versteht es auch, nicht nur auf wissenschaftlicher Ebene hilfreich unter die Arme zu greifen, sondern auch die Brücke zu schlagen zu einer Büroatmosphäre in der ich mich ausnahmslos jeden Tag wohl gefühlt habe. Ich hätte mir in der gesamten Zeit, vom Studium bis heute, keine bessere Betreuung wünschen können, Danke Kai.

Ich bedanke mich bei meinen Kollegen, die in Jena schnell zu meinen engsten Freunden geworden sind. Simon "Jazen" Lieb, danke für echte Gespräche, fürs Gedanken machen, quality-time und da sein, für die frühen Stunden im Gym (ohne dich wäre ich nicht mal ansatzweise so motiviert gewesen) und für den seidenmatten und feinporigen Milchschaum. Jan Hombeck, danke dass ich immer auf dich zählen konnte. Man denke an die Aktion mit dem Laufband, tut mir echt leid dass es dir auch noch auf den Fuß gefallen ist, autsch. Danke auch für die vielen Abende an denen du mich motiviert hast die Work-Life-Balance zumindest mal ein bisschen Richtung Life zu verschieben. Selbiges gilt für Brian Zahoransky, danke dafür dass du mir das Einleben in Jena so leicht gemacht hast. Ich danke auch meiner Büronachbarin Anna Sterzik, dass du dir immer die Zeit genommen hast Ideen durchzudiskutieren und, dass wir uns zusammen über unnötige Bürokratie und Reviewer 2 echauffieren können (geteiltes Leid ist bekanntlich halbes Leid). Henrik Voigt, danke für deinen bedingungslosen Einsatz in ML-Projekten, ohne dich wäre ich oft aufgeschmissen. Für ihre Unterstützung bei allen organisatorischen Herausforderungen danke ich Silvia Blaser.

Ich bedanke mich weiterhin bei allen Partnern, ohne die diese Arbeit nicht möglich gewesen wäre. Besonderer Dank gilt Kevin Richter und Anna Hundertmark, die mit ihrer Expertise in numerischer Simulation und Strömungsmechanik die Qualität und den Umfang dieses Projektes maßgeblich gesteigert haben. Ein weiterer Dank gilt Carsten Klingner, der mir jederzeit bei medizinischen Fragen ausgeholfen und viele Denkanstöße für diese Dissertation gegeben hat. Danke auch dafür, dass du mir tatkräftig klinische Evaluationen ermöglicht hast. In diesem Sinne bedanke ich mich auch bei Ralf Wickenhöfer, der die ursprüngliche Idee für dieses Projekt mit ausgearbeitet hat und mir ebenfalls mit seiner umfangreichen Erfahrung bei radiologischen Fragestellungen weitergeholfen hat.

Nicht zuletzt möchte ich meiner Familie danken. Bei meinen Eltern, dafür dass ihr mir diesen Weg ermöglicht habt und mir alle Freiheiten eingeräumt habt, ohne Erwartungen oder je an mir zu Zweifeln. Und bei Wanda fast-Eulzer, für deine liebevolle Unterstützung, das Aufmuntern, das Herumalbern und das gemeinsame Durchstehen von Allem. Dafür, dass ich weiß, immer deine ehrliche Meinung zu bekommen und deine Gabe Dingen neue Perspektiven zu geben.

Pepe Eulzer, 05. August 2024

CONTENTS

I Preliminaries	
1 Introduction	5
1.1 Motivation	6
1.2 Contributions and Structure	7
II Background	
2 Medical Background	15
2.1 Carotid Artery Stenosis	16
2.2 Stenosis Screening, Diagnostics, and Treatment	17
3 Vessel Extraction, Flow Simulation, and Visualization	23
3.1 Vessel Extraction and Pre-Processing	23
3.1.1 Vessel enhancement	24
3.1.2 Model extraction	24
3.1.3 Centerline extraction	26
3.2 Computational Hemodynamics	27
3.2.1 Blood Flow Simulation	28
3.3 Vessel and Blood Flow Visualization	31
3.3.1 Overview	32
3.3.2 Temporal analysis	33
3.3.3 Probing	34
3.3.4 Filtering	34
3.3.5 Contextualization	34
3.3.6 Comparison	35
3.3.7 Robustness Evaluation	36
4 Vessel Maps	39
4.1 Vessel Map Requirements	42
4.2 Data Structures	45
4.3 Taxonomy	46
4.3.1 Layout Generation	47
4.3.2 Geometry Mapping	49
4.3.3 Attribute Mapping	50
4.3.4 Classification	51
4.4 Image Volume Maps	51
4.5 Flow Volume Maps	58
4.6 Vessel Wall Maps	60
4.6.1 Maps of Cerebral Aneurysms	61
4.6.2 Stenosis Predilection Sites	65
4.6.3 Maps of the Heart	67
4.6.4 Arbitrary Vessel Trees	72
4.7 Vessel Network Maps	72
4.7.1 Connectivity Maps of Large Arteries	73

4.7.2	Cerebral Connectivity Maps	77
4.7.3	Maps of Microvasculature Networks	79
4.8	Recommendations for Vessel Map Creation	81
4.8.1	Identifying Domain Tasks	81
4.8.2	Choosing Data Type and Region Complexity	82
4.8.3	Preservation versus Standardization	82
4.8.4	Layout, Geometry, and Attribute Mapping	83
4.9	Open Challenges	84
4.10	Conclusion	85
III Visual Exploration of Hemodynamic Simulations in Stenoses		
5	Automatic Vessel Wall Maps	91
5.1	Objectives	92
5.2	Cutting	93
5.2.1	Sorting boundaries	94
5.2.2	Aligning curvature	95
5.2.3	Finding optimal paths	96
5.3	Flattening	98
5.4	Application	103
5.4.1	Implementation	104
5.5	Results and Discussion	104
5.5.1	Performance	106
5.5.2	Generalizability	107
5.5.3	Limitations	108
5.6	Conclusion	109
6	Visualizing Blood Flow in Carotid Stenoses	113
6.1	Domain Characterization	114
6.1.1	Data Abstraction	114
6.1.2	Stakeholder Analysis	115
6.1.3	Task Abstraction	116
6.2	Pre-Processing	117
6.3	Methods	118
6.3.1	Domain Overview	119
6.3.2	Vessel Map	120
6.3.3	Flow Rate Graph	122
6.3.4	Navigation and Probing	122
6.3.5	Contextualization	124
6.4	Evaluation	127
6.5	Results and Discussion	128
6.6	Conclusion and Future Work	130
IV Identifying and Solving Clinical Transfer Challenges		
7	A Fully Integrated Pipeline Architecture	137
7.1	Introduction	137
7.2	Bottlenecks in Clinical Transfer	140
7.3	System Design	141
7.4	System Implementation	143

7.4.1	Crop Module	143
7.4.2	Convolutional Neural Network	144
7.4.3	Segmentation Module	145
7.4.4	Centerline Module	146
7.4.5	Stenosis Classifier Module	147
7.5	System Evaluation	149
7.5.1	Method	151
7.5.2	Findings	151
7.6	Discussion	154
7.6.1	Controllability vs. Ease of Use	155
7.6.2	Opportunities and Future Work	156
7.7	Conclusion	156
8	Similarity-based Hemodynamics Analysis	161
8.1	Task Analysis	163
8.2	Streamline Clustering	165
8.3	Methods	165
8.3.1	Creating a Database of High-Resolution Carotid Flow Models	167
8.3.2	Extracting Carotid Bifurcation Geometries	169
8.3.3	Aligning the Carotid Data	170
8.3.4	Visualizing the Alignment Space Neighborhood	171
8.3.5	Comparative Visualization of Multiple Vascular Flow Models	175
8.4	Evaluation	177
8.4.1	Performance Evaluation	177
8.4.2	Numerical Evaluation	178
8.4.3	Structured Interviews	180
8.5	Results and Discussion	180
8.5.1	Interview Results	183
8.5.2	Limitations	185
8.5.3	Opportunities and Future Work	187
8.6	Conclusion	187
v	Conclusion	
9	Conclusion	193
9.1	Summary	193
9.2	Future Work	196
vi	Appendix	
	Bibliography	201
	List of Publications	239
	Declaration	241

ACRONYMS

2D	two-dimensional
3D	three-dimensional
AHC	agglomerative hierarchical clustering
ARAP	As-Rigid-As-Possible
BEP	Bull's Eye Plot
CCA	common carotid artery
CFA	curvicircular feature aggregation
CFD	computational fluid dynamics
CNN	convolutional neural network
CPR	curved planar reformation
CT	computed tomography
CTA	computed tomography angiography
CVD	cardiovascular disease
DICOM	Digital Imaging and Communications in Medicine
DVR	direct volume rendering
ECA	external carotid artery
FEM	finite element method
GBD	Global Burden of Disease
GUI	graphical user interface
ICA	internal carotid artery
LSCM	least squares conformal maps
MIP	maximum intensity projection
MRA	magnetic resonance angiography
MRI	magnetic resonance imaging
NASCET	North American Symptomatic Carotid Endarterectomy Trial
NWS	normal wall stress
OSI	oscillatory shear index
PC-MRI	phase contrast magnetic resonance imaging
PET	positron emission tomography
RT-MRI	real-time magnetic resonance imaging
TAVI	transcatheter aortic valve implantation
VMTK	vascular modeling toolkit
WSS	wall shear stress

Part I

PRELIMINARIES

1

Introduction

INTRODUCTION

The effective treatment and prevention of stroke are of high societal importance [107]. Globally, stroke has remained the second-leading cause of death for several years and it also is the third-leading cause of death and disability combined [100]. Post-stroke care is associated with considerable economic expenses [257]. The annual number of deaths from stroke is increasing, which is partly attributed to the shift towards an older population and also to an increase in risk factors like high blood pressure and high body-mass index. The Global Burden of Disease (GBD)¹ 2019 stroke collaborators conclude that “*without urgent implementation of effective primary prevention strategies, the stroke burden will probably continue to grow across the world*” [100].

The majority of strokes, about 80–90%, are *ischemic* [2, 80, 301], which means they are caused by lack of blood flow to the brain. Ischemic stroke is typically caused by a constricted or blocked blood vessel. A crucial component of clinical stroke care and prevention is the detection, classification, and possibly treatment of arterial stenoses – local vessel constrictions – which are often due to atherosclerotic plaque. A predilection site for stenosis development is the carotid bifurcation, where the internal carotid splits off the common carotid artery. The two internal carotids supply the largest part of the brain’s circulation and a large portion of ischemic strokes are directly linked to carotid stenosis [80]. Therefore, sophisticated clinical workflows are in place for screening the carotids [76]. Typically, a computed tomography angiography (CTA) or magnetic resonance angiography (MRA) and an ultrasonographic assessment are used to determine the morphological and hemodynamic properties of potential stenoses. Physicians face multiple challenges when analyzing the resulting images. The derivation of geometrical properties of a stenosis, e.g., the shape, length, minimal diameter, and plaque distribution from slice-based depictions is not intuitive, as multiple successive images need to be mentally combined. Volume renderings are sometimes also used but can be problematic to interpret due to visual clutter from the surrounding blood vessels, connective tissue, and artifacts from tooth fillings. Yet, a comprehensive understanding of the stenosis properties is central to determining the ideal treatment approach. Furthermore, integrating the results from multiple modalities, like CTA and ultrasonography, is challenging. Important hemodynamic properties, such as peak velocity and the extent of poststenotic turbulence, are only visible using

¹ The GBD study is the largest global observational epidemiological study to date. See also: <https://www.thelancet.com/gbd>.

Doppler ultrasonography. However, many morphological parameters are better derived from angiographic volume imaging. A central value is the *stenosis degree*, since it is used as a threshold determining the need for surgical intervention [88]. It can be attained from both CTA and ultrasonography, however, considerable fluctuations of the stenosis degree exist depending on the observer and modality [209].

The clinical decisions made in acute stroke care and stroke prevention are highly case-specific. Determining the appropriate management strategy for carotid stenosis depends not only on the degree of stenosis but also on individual patient factors such as age, comorbidities, and overall cardiovascular risk profile. The risk of stroke always needs to be weighed against the risk of complications from intervention and different surgical strategies exist, for instance, plaque removal or stent insertion. The pathological progression of developing stenoses and the interplay of anatomical and hemodynamic factors are not yet fully understood. Due to the prevalence of stroke, stenosis formation and stenosis risk assessment remain major research areas [107].

1.1 MOTIVATION

In the clinical detection and evaluation of arterial stenoses, many challenges persist. While several imaging techniques are available for detecting carotid stenosis, each has its limitations. Common methods include carotid ultrasound, MRA, CTA, and digital subtraction angiography. Each modality has its strengths and weaknesses, and interpreting results accurately requires expertise. Often, multiple assessments are carried out by different experts at several points in time and sometimes different clinics. To derive the best treatment strategy, physicians then need to piece together the results to get the full picture. Stenoses can progress gradually over time, and the severity can vary from mild to severe. Monitoring disease progression and accurately assessing the degree of stenosis require regular follow-up and precise measurements, which can be challenging. Determining the appropriate management strategy for carotid stenoses necessitates a comprehensive evaluation and often multidisciplinary decision-making involving neurologists, vascular surgeons, and other specialists.

Blood flow simulations, typically based on computational fluid dynamics (CFD), can complement and improve the processes involved in detecting and managing carotid stenosis in several ways. Through visualization of hemodynamics, CFD simulations can provide insights into blood flow patterns within the carotid arteries. These can help clinicians better understand the complex flow dynamics associated with stenotic lesions, including areas of turbulence, recirculation, and low shear stress, which are important factors in the progression and rupture of atherosclerotic plaques [186]. Visualizing simulated blood flow enables a direct integration of morphology and hemodynamics.

CFD further allows for the quantification of hemodynamic parameters such as wall shear stress (WSS), pressure gradients, flow velocities, and flow rates. These parameters can help identify regions of potential plaque vulnerability or increased risk of thrombosis [124, 229].

Despite their potential benefits, integrating CFD-based methods into clinical workflows poses several challenges, due to the often time-consuming, resource-intensive, and complex nature of the simulations. Executing and analyzing CFD simulations requires expertise in fluid dynamics, numerical methods, and specialized software tools. Clinicians and healthcare professionals may lack the necessary training and experience to perform or interpret CFD analyses, making it difficult to incorporate these methods into routine clinical practice. Effectively communicating the results of CFD simulations to clinicians requires clear visualization and interpretation of complex hemodynamic data. Furthermore, performing CFD simulations often requires access to high-performance computing resources, which may be expensive and not readily available in clinical settings, limiting the feasibility of conducting real-time simulations or large-scale analyses.

In visualization research, a variety of techniques were developed that enable tailored exploration and analysis of simulated blood flow data. Techniques were proposed for both, conducting medical research and decision-making in clinical workflows. Existing methods primarily target flow features in cerebral aneurysms and the aorta [234]. Visualizing blood flow in stenoses and analyzing risk-factors for ischemic stroke has only been sparsely covered.

1.2 CONTRIBUTIONS AND STRUCTURE

In this work, to address the challenges outlined above, methods from blood vessel and flow visualization are refined and adapted for the exploration of carotid hemodynamics and morphology [342]. On this basis, novel methods are developed for the spatio-temporal exploration of vascular models employing map-like abstractions [343, 346]. Further, techniques are introduced to efficiently detect stenoses and perform automatic and accurate stenosis degree classification [339, 344]. To target the challenge of the applicability of blood flow simulations in direct patient care, a novel approach is implemented to circumvent complex individual flow simulations [345].

At the core of this work lies a multi-stage, unified software framework, which combines pre-processing and the data exploration tools listed above, making them readily available to clinical practitioners and researchers. Interactive visualizations are used for data inspection and adaption to guide the user through the processing stages. Streamlining the data processing and enabling retrospective user control seem to have a major impact on user trust in medical visualizations and increase the data validity, ultimately making sophisticated vi-

sualization systems more readily usable in routine. All techniques developed in this thesis were evaluated by an interdisciplinary group of visualization experts, numeric simulation researchers, neurologists, and radiologists. The framework has a modular and extendable architecture and is available as an open-source platform [48]. It was used to procure a database of, at the time of writing, 152 carotid bifurcation geometries. The database was also extended to include simulated blood flow and is publicly available [347].

This thesis is structured as follows. The initial part provides a broad review of the medical and application background:

- [Chapter 2](#) gives an overview of the relevant medical background. The prevalence and epidemiology of carotid stenosis are outlined, as well as current clinical workflows regarding carotid screening, diagnostic capabilities, and possible treatment strategies.
- [Chapter 3](#) complements the medical background with the technical state-of-the-art in computerized medical imaging, hemodynamic simulation, and vessel visualization. The particular focus lies on the carotids and carotid stenosis, however, all techniques relevant to the development of the methods presented in later chapters are reviewed. [Section 3.1](#) discusses how vascular models can be extracted from image data, [Section 3.2](#) provides an overview of hemodynamic simulation techniques, and [Section 3.3](#) summarizes how the resulting vascular models and possibly blood flow data can be visualized.
- [Chapter 4](#) provides a classification of *vessel maps* – map-like visualizations of cardiovascular structures [343]. A comprehensive review of vessel maps is given, yielding an overview of the vast landscape of these techniques, which have been broadly applied to simplify and abstract complex three-dimensional (3D) vascular structures. Vessel maps are applied on several occasions in this thesis, where the investigations from [Chapter 4](#) serve as a basis.

In the next part, techniques for the visual analysis of carotid stenoses are developed. The proposed methods extend existing procedures for visualizing hemodynamics. The contributions of this thesis are as follows:

- [Chapter 5](#) targets the exploration of hemodynamic parameters that affect the vessel wall. An algorithm for creating a vessel surface map is developed, which cuts and flattens vessel tree surfaces fully automatically [346]. The advantage over previous approaches is that only the vessel mesh is required as an input. No centerlines, manual cuts, or landmarks are needed. Both the cutting and flattening stages are designed to facilitate an intuitive and accurate assessment of parameter distributions on the surface. *Natural vessel cuts* are introduced to solve the tendency

of most cutting algorithms to wrap cuts around elongated structures. The flattening develops a free boundary with minimized area distortions while keeping the layout close to the original structure.

- [Chapter 6](#) targets the exploration of flow parameters inside stenoses and their integration into the patient-specific anatomical context. A comprehensive visualization framework is developed that allows for exploring and analyzing time-dependent hemodynamic simulations of the carotids [342]. A vessel network map is used as an overview depiction of the vessel diameter to quickly identify stenosis candidates. The map doubles as a navigational aid to efficiently and unambiguously probe any region in the simulation domain.

During the development and testing of the above frameworks, several challenges became apparent, which hinder the clinical adoption of advanced vessel visualization and patient-specific hemodynamic simulation. These challenges are not specific to the analysis of carotid stenosis but similarly occur in most cardiovascular simulation and visualization applications. In the following part of this thesis, targeted methods were developed that directly address these transfer and application bottlenecks:

- [Chapter 7](#) introduces a fully integrated pipeline for the extraction, exploration, and analysis of carotid bifurcation geometries from clinical image data [339]. The pipeline directly connects multiple sophisticated processing stages, including a neural prediction network for lumen and plaque segmentation and automatic global diameter computation. Visualization modules are provided, which allow exploring the processed data. The frameworks developed in [Chapter 5](#) and [Chapter 6](#) are accessible as visualization modules. The integration of processing and visualization stages enables interactive and retrospective user control over the processing steps without leaving the application. Providing the users with guided insight into the data processing showed to increase their trust in the resulting visualizations and helped to make the application of the pipeline scalable. The pipeline was used to extract 152 carotid bifurcation models, which were made publicly available [347].
- [Chapter 8](#) targets a second overarching issue: the translation of hemodynamic simulation into clinical decision-making and disease monitoring is highly challenging. Hemodynamic simulations are intrinsically complex, highly time-consuming, and resource-intensive, which conflicts with the time-sensitive nature of clinical workflows and the fact that hospitals usually do not have the necessary resources or infrastructure to support

flow simulations. The chapter describes the development of a similarity-based visualization system to address these transfer challenges, which enables instant flow exploration without performing on-site simulation [345]. A database of high-resolution carotid blood flow models is created from the models extracted in [Chapter 7](#). Then, similarity metrics are used to place a new carotid surface model into a neighborhood of simulated cases with the highest geometric similarity. The neighborhood can be immediately explored. A core visualization that is used here is the surface map developed in [Chapter 5](#). The geometries, simulated flow, and wall-related parameters can be comparatively analyzed. If the models are similar enough in the regions of interest, a new simulation leads to coinciding results, allowing the user to circumvent complex individual flow simulations.

Part II

BACKGROUND

2

Medical Background

MEDICAL BACKGROUND

Medical conditions that affect the heart or blood vessels are grouped as cardiovascular disease (CVD) [211]. With more than 30% of global deaths attributed to CVD, it is the leading cause of fatality worldwide [99]. It is estimated that up to 90% of CVD may be preventable [207, 231] by reducing risk factors like smoking, changes in diet and exercise, and timely identification and treatment of developing conditions. The most prominent complication from CVD is myocardial infarction, commonly known as a heart attack, where blood flow to the coronary arteries of the heart decreases, damaging the heart muscle. The second most relevant complication is stroke, a condition in which a circulatory disorder limits blood flow to the brain, ultimately causing cell death. Other diseases include venous thrombosis, which can impact lung function, peripheral artery disease, which may necessitate limb amputation, heart valve insufficiency, which impairs overall blood circulation, and many more [211]. The underlying mechanism of many CVDs is atherosclerosis, in which lesions in the inner wall of an artery occur that, over time, lead to the buildup of fatty deposits called plaques. The progression of atherosclerosis is usually asymptomatic for decades, as the arteries compensate for the plaque growth by enlarging [265]. Symptoms only occur after severe narrowing of an artery reduces the blood flow enough to affect an organ or if particles break off a lesion site and block a smaller vessel upstream (embolization). The buildup of atherosclerotic plaque is particularly alarming if it affects the arteries supplying the heart or brain.

This work focuses on the prevention of stroke and stroke-like attacks, in particular ischemic stroke. Stroke can either be hemorrhagic if it is caused by internal bleeding or ischemic if it is due to reduced blood flow. Ischemic stroke accounts for 80–90% of all cases [2, 80, 301]. In 2015, stroke led to 6.3 million deaths worldwide, about 11% of the total, making it the second most frequent cause of death after coronary artery disease and before cancer [98]. In non-fatal cases, stroke often results in permanent neurological disabilities [107] and post-stroke care is associated with substantial economic costs [257]. About 50% of people who survived an acute stroke live less than a year [80]. In non-fatal strokes, common ensuing symptoms include paralysis, sensory malfunction, speech disorder, and impaired vision, coordination, and cognition [200]. Known causal risk factors for ischemic stroke are high blood pressure, hypercholesterolaemia¹, atrial fibrillation², and carotid

¹ A lipid disorder in which the low density lipoprotein (“bad cholesterol”) is too high.

² A rapid and irregular beating of the atrial heart chambers.

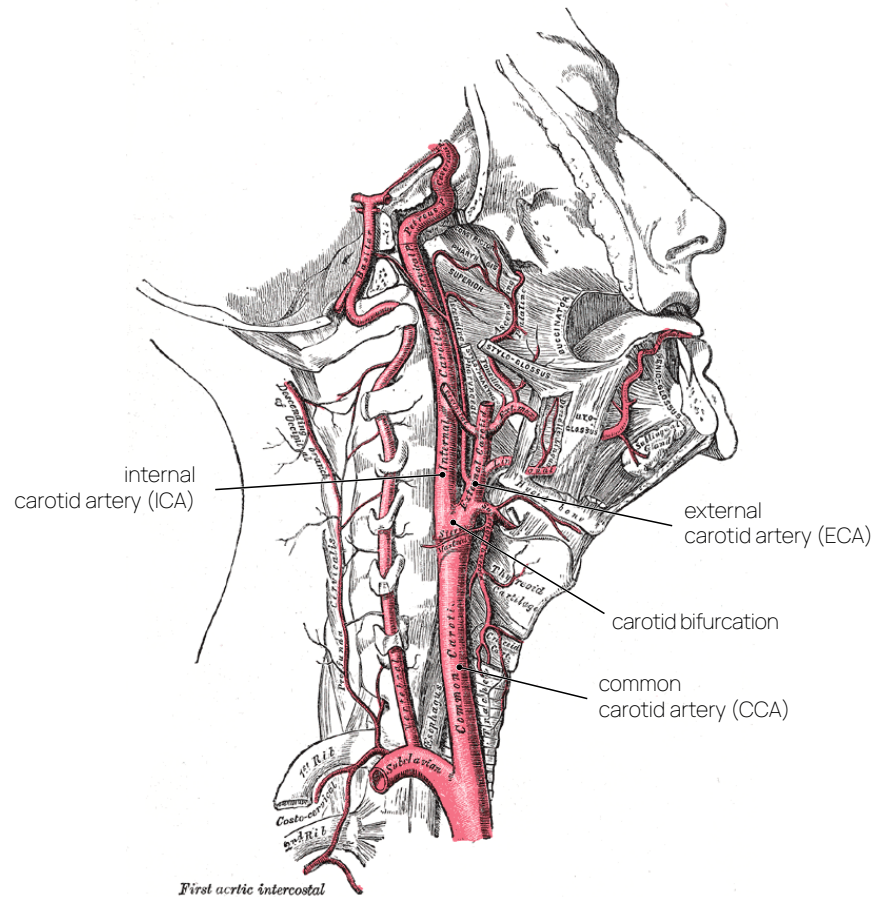


Figure 1: Arteries of the neck. The relevant parts of the right carotid artery are annotated. Image from [113] (public domain).

stenosis [123]. Carotid stenosis describes a localized narrowing of the (internal) carotid arteries, which is usually caused by atherosclerosis.

2.1 CAROTID ARTERY STENOSIS

The carotids are two large arteries that run on both sides of the neck. Together with the two vertebral arteries, they provide blood supply to the brain. The anatomical layout of the carotid arteries is shown in Figure 1. The right and left common carotid arteries (CCAs) originate from the aorta. Each CCA splits into an internal carotid artery (ICA) and an external carotid artery (ECA), at about half the length of the neck. The location where these vessels split is called the carotid bifurcation. While the ECA provides blood flow for the face and neck, the ICA supplies the brain and eyes. After the bifurcation, the ECA immediately branches into further appended vessels, whereas the ICA travels along the remaining length of the neck and into the skull without any side branches. Only after emerging from the base of the skull, the first side branches can be found. Inside the skull, the ICA forms a small arc,

often S- or U-shaped, which is also called the carotid siphon. After the carotid siphon, the vessel's diameter quickly diminishes, as the ICA branches into several smaller arteries responsible for brain blood supply.

Carotid stenosis is usually caused by the buildup of atherosclerotic plaques in the inner lining of the artery. As these plaques grow, they can constrict the artery and reduce blood flow. A stenosis of the ICA can result in a stroke in three ways. (1) A stenosis can grow such that it severely impedes blood flow to the brain or entirely blocks the artery. Reduced blood flow in one ICA may be partially substituted by other arteries. A severe stenosis or multiple stenoses in the left and right ICA increase the stroke risk. (2) Plaque deposits can rupture and emboli can break off and travel to the cranial blood vessels, where they block a smaller artery. (3) Blood clots can form on the irregular surface of the stenosed artery wall, which may occlude the artery at the stenosis site or break off and block a smaller vessel upstream. A typical predilection site for stenosis development is shortly after the carotid bifurcation, as shown in Figure 2. The prevalence of extracranial³ carotid stenosis is about 4% in adults, but rises to 6–15% after the age of 65 [76]. Considering the demographic shift towards an older population, the healthcare system needs to be prepared for an increased risk of stroke due to carotid stenosis. Already, in Germany alone, about one million patients are living with extracranial carotid stenosis and approximately 30,000 strokes every year are directly linked to this condition [76]. As such, the prevention of carotid stenosis-induced stroke is of fundamental societal importance.

2.2 STENOSIS SCREENING, DIAGNOSTICS, AND TREATMENT

Patients who suffered a stroke or show symptoms that may indicate possible stenoses are typically subjected to ultrasonography of the carotids [76]. With Doppler ultrasonography, the blood flow velocity in the CCA and ICA can be measured [279]. Based on the Doppler effect, the frequency shift between a transmitted and received ultrasonic signal is measured. This shift is proportional to the speed of the blood which reflected the signal, allowing quantitative measurements of blood flow. The constricted vessel walls of a stenosis lead to an increase in flow velocity, which abruptly drops behind the stenosis. Using this pattern, stenoses can be detected and classified according to measured flow speeds. Clinically relevant are especially the systolic⁴ peak velocities within the stenosis maximum and after the stenosis. These two values can be used to grade a stenosis, for example, a systolic peak velocity of 350–400 $\frac{\text{cm}}{\text{s}}$ inside the stenosis and a poststenotic peak velocity of $< 50 \frac{\text{cm}}{\text{s}}$ would classify an 80% ICA stenosis [76]. The

³ Outside of the skull.

⁴ Part of the cardiac cycle where blood is pumped through the body.




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

sudden reduction in velocity leads to turbulent flow after the stenosis that sometimes also becomes visible, especially in severe stenoses. Similarly, the end-diastolic⁵ velocity within the stenosis maximum is sometimes used as an additional gradation criterium; it often exceeds $100 \frac{\text{cm}}{\text{s}}$ in severe stenoses [76].

The hemodynamic information gained with Doppler ultrasonography is highly relevant for neurologists treating stroke patients. The method is also widely adopted as it is fast, inexpensive, and non-invasive. A limitation, however, is the susceptibility of Doppler ultrasonography to artifacts. For instance, image quality is considerably worse in the case of obese patients. Additionally, the results are highly dependent on the orientation of the ultrasound transducer held by the physician, resulting in possible inter-observer variability [269]. Not only does this affect which cross-section of the examined vessel is visible, but also the measured speeds [289]. The Doppler effect is dependent on the angle between the flow and measurement direction. The most accurate signal is found at 0° difference between these directions, i.e., the blood flows directly towards or away from the probe. This is problematic, as the carotid arteries do not necessarily have a straight layout, meaning the direction of flow varies. Consequently, carotid ultrasound screenings require a highly qualified physician to execute. Smaller hospitals often do not have the proper equipment or personnel to perform these types of tests. Further, stenoses might occur after the ICA enters the skull. Although transcranial⁶ Doppler ul-

⁵ Part of the cardiac cycle where the heart fills with blood.

⁶ Through the skull.

trasonography does exist, it is even more difficult to properly execute, as the relevant arteries are occluded by the skull.

In many cases, an angiography is also performed, often in the form of a CTA or MRA, to reveal the intracranial⁷ vessels and give a more accurate representation of the morphology. The injection of a contrast agent highlights the *lumen*, i.e., the region inside a vessel where blood actively flows, in the resulting volume images. This allows for determining the position and size of a stenosis, based on which its severity can be derived. Besides the evaluation of individual slices, procedures used in practice include maximum intensity projection [226], shaded surface display [187], and volume rendering [46]. Current clinical guidelines base the treatment recommendations largely on the stenosis degree. Following the results from the North American Symptomatic Carotid Endarterectomy Trial (NASCET) [88], a high indication for surgical intervention is given for symptomatic patients with above 70% ICA stenosis. On angiographic imaging, the stenosis degree S_{NASCET} is computed using the NASCET criteria. The diameters of the vessel lumen inside the stenosis d_{min} and after the stenosis d_{normal} are compared:

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

However, CTA and MRA reveal no hemodynamic flow features. As a result, in carotid artery diagnostics, clinicians need to mentally combine the outcomes of multiple examinations before deciding on viable treatment options. Using the stenosis degree as a threshold value alone does not include other properties, such as the individual shape of the stenosis, the wall composition, or the plaque type and distribution. Comprehending these factors and an accurate assessment of objective descriptors of the disease are crucial for ideal treatment and monitoring of a patient's condition. Furthermore, treatment decisions need to be made on an individual basis, considering a large number of factors involved, including the age, background, gender, and medical history of a patient.

Conservative treatments using medication to control vascular risk factors, e.g., lowering blood pressure or antithrombotic medication, are recommended for almost all patients with carotid stenosis [76]. If developing stenoses are identified at the right stage, endovascular⁸ or operative approaches can be taken to avoid total closure and drastically minimize the risk of embolization. It has been repeatedly shown that timely surgical intervention can prevent stroke as a result of carotid stenosis [120, 258]. The most common procedure is carotid endarterectomy, where plaque is removed from the carotid bifurcation in open surgery. Current evidence suggests that carotid endarterectomy has some benefit for patients with 50–60% symptomatic ICA

⁷ Inside the skull.

⁸ A minimally invasive procedure, where a catheter is used to access the artery.

stenosis and is highly beneficial for those with 70–99% stenosis [259]. Alternatively, a stent, i.e., a small tube-like mesh, can be inserted in an endovascular procedure to stabilize the arterial wall and keep the passageway open. However, carotid surgery has its own risks that must be carefully considered against the likelihood of a stroke [121]. For example, the stroke risk during carotid endarterectomy is between 1–5% [37]. Selecting the right moment for surgery and the best treatment strategy is critical, which requires careful evaluation and an interdisciplinary discussion of each case.

3

Vessel Extraction, Flow Simulation, and Visualization

VESSEL EXTRACTION, FLOW SIMULATION, AND VISUALIZATION

Various frameworks and techniques for visualizing blood vessels were developed to solve specific tasks of clinical practitioners [343, 234, 251, 252]. Generally, a differentiation is made between model-free and model-based techniques [251]. Model-free techniques work directly on the image data, for example, using slice-based depictions, image reformations [142, 221], or volume rendering [167, 349]. Model-based techniques rely on a processing pipeline that extracts geometric models from volume images. Typically, the vessels of interest are segmented on the image, after which a surface is reconstructed. In many clinical and research applications, such models are used to describe the bounds of the target structure, for instance, an organ [75, 108], bones [338, 348], or the surface of a vessel tree [183, 223]. The model usually is a triangulated surface mesh or quad mesh [283]. An advanced application of patient-specific vascular models is their use in computational hemodynamics, i.e., the simulation of the internal blood flow.

The ability to simulate and analyze flow in actual physiological structures, such as patient-specific blood vessels, was made possible by complementary progress in image processing, high-performance computing, and the creation of advanced clinical visualization methods [27]. In computational hemodynamics, a volumetric mesh is built from the target domain, and then a CFD simulation is carried out, where predefined boundary conditions dictate the flow behavior at vascular walls, inlets, and outlets. A large body of visualization methods has been developed that enable analyzing the resulting flow field and any derived parameters like WSS, the pressure gradient, or the volumetric flow rate [234]. A common approach is the combination of multiple coordinated data views with different purposes, for example, slice-based renderings, volume renderings [167], 3D surfaces, views of integral lines [40], focus lenses [96], map-like depictions [343], and also data graphs or charts [213].

3.1 VESSEL EXTRACTION AND PRE-PROCESSING

The two principal source modalities for non-invasive capturing of a patient's vasculature are CTA and MRA. Both yield volume images of the scanned region where vessels are contrasted against other structures. They are typically not temporally resolved, i.e., they record a static image. Exceptions are specialized techniques, like real-time

magnetic resonance imaging (RT-MRI) [64] and phase contrast magnetic resonance imaging (PC-MRI) [288], which have a temporal dimension. PC-MRI in particular is used to determine not only the occurrence of blood flow but also the flow velocities in large arteries. In clinical routine, cardiovascular structures are also often examined using ultrasound. Ultrasound generally has a lower resolution and more artifacts compared to CTA or MRA, which is why in diagnostics it is often used in addition to, e.g., a CTA scan of the same region. In routine examinations, mostly traditional two-dimensional (2D) ultrasound is used but 3D ultrasound is an emerging alternative [133]. Further imaging modalities that have been applied in some cases are rotational angiography [118] and digital subtraction angiography [140]. Also, invasive capturing techniques exist, such as intravascular ultrasound [95] and intravascular optical coherence tomography [24], where a probe is inserted into the vasculature via a catheter. For model-based vessel visualization and to simulate blood flow, a processing pipeline as shown in Figure 3 is typically used.

3.1.1 *Vessel enhancement*

The recorded volume image is often filtered to enhance the contrast of the target structure. A comprehensive overview of vessel enhancement procedures can be found in the survey of Fraz et al. [91]. Prominent are vesselness filters, originally introduced by Frangi et al. [90]. A variety of similar and improved filters have been proposed [171]. Other approaches include Wavelets [285] and diffusion filtering [166]. In some cases, like PC-MRI processing, artifacts need to be explicitly removed [156].

3.1.2 *Model extraction*

Model extraction can be split into segmentation, where a target structure is marked in the volume, and surface reconstruction, where the geometric model is created. Sometimes, multiple models are defined, for example, to differentiate the inner and outer walls of a vessel or to extract regions with atherosclerotic plaque [139]. Due to their simplicity, intensity-based approaches like thresholding and region growing are sometimes used to segment vessels. While they might apply to image data where vessels are delineated, they often do not produce the expected results, as they are not robust against noise and artifacts. Deformable models have also been applied to vascular structures, particularly active contours [199, 204]. They match an initial model to fit the specific image, e.g., based on its gradient. Similarly, graph-based methods, especially graph cuts, have been adapted to vessel tracking [21, 84]. Temporally resolved volumes, such as from PC-MRI or ultrasound, require a segmentation of every time step to adjust for the




Figure 3: A typical processing pipeline used for model-based vessel visualization and hemodynamic simulation.

moving morphology [41]. For example, Köhler et al. [159] proposed a modified graph cut to segment the aorta in PC-MRI data. Many model extraction techniques require some form of manual input, such as landmarks or outlined contours. A survey of conventional techniques for vessel lumen segmentation is provided by Lesage et al. [183].

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-based methods [194, 254, 321]. Especially convolutional neural networks (CNNs) have proven to be highly successful tools for image analysis. Their advantage lies in the fact that image features are automatically learned, i.e., no specific descriptor needs to be provided. This allows complex automatic segmentations of data with noise, image artifacts, or different types of atherosclerotic plaque, which is often difficult to differentiate from the vessel lumen. CNNs have been applied to segment various types of vessels [194,

224, 254, 321]. They generally require a basis of manually segmented image volumes for training. Specialized networks for carotid lumen and plaque segmentation have been proposed [34, 72], where good performing networks, with a Dice similarity¹ of 82%, were based on as low as 15 training data sets [332]. For further reading, we recommend the surveys of Moccia et al. [223] and Zhao et al. [330]. Extracted surfaces can also be further processed and given additional information, such as relevant landmarks or morphological features. For example, aneurysms can be detected and segmented in models of vessel trees [176]. Particularly relevant are also centerline extraction methods that produce a topological model of the vessel tree structure.

3.1.3 Centerline extraction

Many mesh processing and visualization algorithms that work with vascular models require a centerline. The centerline is the geometric vessel skeleton. More generally speaking, it is a graph that defines the tree or network topology of the vascular structure it lies in. It can be reconstructed either from a surface model or directly from the volume image. An overview of skeletal representations for surface meshes is provided by Tagliasacchi [294] and a survey on skeletonization algorithms is given by Saha et al. [270].

Various approaches to computing the skeleton of voxel volumes, like medical volume images or voxelized surface meshes, have been proposed. A common idea is topological thinning, for instance, by peeling one layer of voxels at a time [101, 238, 243, 281] or iteratively contracting the volume until a linear graph is retained [309]. If the centerline is to be derived from a volume image, these techniques typically require a binarized volume, where the voxels belonging to the target structure are separated [180]. A second widely studied approach is to use distance metrics on the image graph. By connecting neighboring voxels in a graph structure, individual skeleton paths can be defined through the volume [333]. After a start node is selected, the end node can be determined by maximal distance and an ideal path can be found using Dijkstra's algorithm [78]. For the centerline to adhere to the center of the volume, passing through edges closer to the border of the structure can be penalized [25, 26, 273].

If a surface geometry is given, a common strategy for finding the centerline of a vessel tree is to use the medial axis transform of the geometry [28]. For a 3D shape, the medial axis tracks the positions of the maximally inscribed spheres. Each sphere is described by three points on the surface. The medial skeleton is then defined by a set of sheets, each described by three points [102]. Connecting the centers

¹ The Dice similarity coefficient computes a similarity measure between two samples. It is a common measure for comparing algorithm output against reference masks in medical image segmentation.

of the spheres retrieves a curve skeleton – the vessel centerline. An advantage of this method is that the minimum radii of the tubular structure are computed as a byproduct. For surfaces made of polygons, the medial axis can be derived by computing the bisectors [73]. This method, however, is computationally complex. Faster strategies usually approach this as a Voronoi diagram problem [235]. First, the boundary points are sampled, then a 3D Voronoi diagram is computed on them. The medial axis can then be derived by choosing a fitting subset of the diagram, i.e., cells that lie inside the geometry. Antiga et al. [6] demonstrate this approach with surface models of vessel trees. Other popular approaches for the skeletonization of surface meshes include, similar to the thinning of volumes, a topology-based thinning of the surface [12] and retrieving the centerline from a clustering of the points composing the object [87].

Implementations of vessel centerline algorithms have been included in commonly used medical image analysis toolkits. Examples are the skeletonization plug-in in ImageJ [276], which is based on the method by Lee et al. [180] and the vascular modeling toolkit (VMTK) [7, 135], which has been implemented as an extension for 3D Slicer [151, 247]. A comparison of these methods is provided by Wang et al. [310].

Centerlines and lumen segmentations have also been used for semi-automatic stenosis grading. Tang et al. [295] used level sets with a centerline prior to segment the carotid artery lumen and propose an automated stenosis quantification based on cross-section diameters or areas. Kaufhold et al. [149] developed a tool to quantify lumen diameters on different image data sets, yielding an estimate of the quantification uncertainties. There has also been a documented challenge for carotid lumen segmentation and stenosis grading at the 2009 MICCAI conference [122].

3.2 COMPUTATIONAL HEMODYNAMICS

In an extensive processing step, numerical methods can be used to compute a prediction for the blood flow. Overviews of techniques for the generation and visualization of blood flow data can be found in the reviews of Caballero and Laín [44], Vilanova et al. [307], Köhler et al. [156], and Oeltze-Jafra et al. [234]. Usually, CFD is used to solve the flow field based on boundary conditions on the vessel wall and domain inlets and outlets at the caps of the vessel tree [172]. CFD requires a volumetric mesh that discretizes the full domain where the simulation should occur. After the simulation, each mesh cell stores information about the predicted flow, which can also be time-resolved. Further quantities, like the WSS or normal wall stress (NWS), can be derived from the flow field.

CFD involves simulating the movement of a fluid and its interaction with other elements. The behavior of the fluid is determined by a

set of governing equations that describe the conservation of mass, momentum, and energy [275]. Blood, as a fluid, can be characterized by properties such as density and viscosity, making it suitable for analysis through CFD simulations. To solve fluid dynamics problems, a defined problem space is necessary. In the case of hemodynamics, i.e., the dynamics of blood flow, this space typically consists of an internal flow field containing flow properties like velocity and pressure, surrounded by spatial boundaries such as arterial walls. Flow enters the domain through predetermined inlets and exits through designated outlets. By providing initial conditions for variables and accounting for potential input or boundary changes over time, the internal fields can be solved. Although analytical solutions exist theoretically, they are only applicable to small and often artificial CFD problems. Various numerical approaches are available, requiring the domain to be divided into discrete volumes called cells. Each cell holds information about internal variables and interfaces with one or more boundary layers (inlets, outlets, or walls). The numerical solver updates these values iteratively, taking into account neighboring cells and boundary conditions.

Another approach, not discussed in this work, is smoothed-particle hydrodynamics, which models the problem domain using freely flowing particles. Smoothed-particle hydrodynamics methods have gained popularity due to their parallelizability and sparse characteristics [71, 255]. However, they have drawbacks when it comes to handling boundary conditions [280], particularly in computing particle displacements near walls. For this reason, hemodynamic simulations are typically performed using cell grids.

3.2.1 Blood Flow Simulation

Hemodynamic simulation has been a widely explored area of interest in biomedical research for several decades. As a result, numerous studies have incorporated CFD techniques to improve patient-specific evaluation and treatment planning in medicine. Given a vessel wall model, an internal flow field can be simulated. The process typically involves many additional steps, including capping the ends of the model, building a volumetric mesh from the surface, and creating profiles for incoming and outgoing flow. One major research focus is the study of aneurysms.² Typically, mechanical risk factors associated with the development and formation of aneurysms are identified. Moreover, analyzing flow characteristics before and after surgical treatment can aid in clinical decision-making [61]. Other fruitful areas of research include CFD investigations of large arteries such as the aorta [249, 328] and the carotids [5, 66, 245, 248], as well as advanced

² Enlargements of arterial walls that can potentially rupture.

simulations of internal hemodynamics in the human heart and the interaction of blood flow with heart valves [74].

The main objective is to detect and analyze flow patterns that may indicate pathological vessel deformations, offering additional information to clinicians. This involves deriving numerical attributes to classify cases or using visualization techniques to enhance the intuitive understanding of flow patterns suggestive of pathologies. The information obtained can then be used alongside other factors to determine the most promising treatment approach [250]. Additionally, the results of surgical interventions can be evaluated post-treatment [66].

Performing a CFD simulation requires the predefined geometry of the target vessel. Various modalities, such as CTA, MRA, and ultrasonography, can be used to acquire this geometry. However, sonography is less suitable due to the high noise associated with the technique [324]. The segmentation of the domain can be done manually, automatically, or semi-automatically [183, 326], as described in Section 3.1.2. Typically, the lumen is identified in each image slice, and it may be beneficial to mark the outer vessel walls or other relevant structures. The segmentation is then interpolated across image slices to generate a 3D surface representation, which is used to create a volumetric mesh for the simulation [193]. In a CFD challenge conducted in 2013 to estimate the rupture probability of aneurysms, the majority of participating groups used commercial solutions to compute hemodynamics [23, 137]. ANSYS Fluent³ is one commonly used tool in this context [131]. Open-source solvers like OpenFOAM⁴ or specific in-house solutions are also alternatives [260].

Boundary conditions need to be defined for CFD simulations. Regarding the carotid arteries, similar conditions and assumptions are often employed [324]. Lopes et al. [198] review hemodynamic simulation approaches applied to carotid artery models. Typically, the flow inlet is set at the CCA, and two outlets are assigned, one at the ICA and another at the ECA. For simulations covering one or multiple cardiac cycles, a pulsatile waveform is defined at the inlet to specify the inflow volume. The waveform can be standardized based on average flow values or customized using patient-specific data if available. Blood viscosity is usually assumed to be constant throughout the domain, and the internal arterial walls are considered rigid. The outlets can have different boundary conditions based on the desired results. Options include a zero pressure boundary, where the fluid exits outlets without resistance, explicit flow splitting on the ICA and ECA, where a predefined volume exits each outlet, or the more accurate Windkessel condition, which models the flow resistance of downstream microvasculature [137]. For instance, Conti et al. [66] implemented the Windkessel condition in their work. An example for setting up such

³ <https://www.ansys.com/products/fluids/ansys-fluent>

⁴ <https://www.openfoam.org>

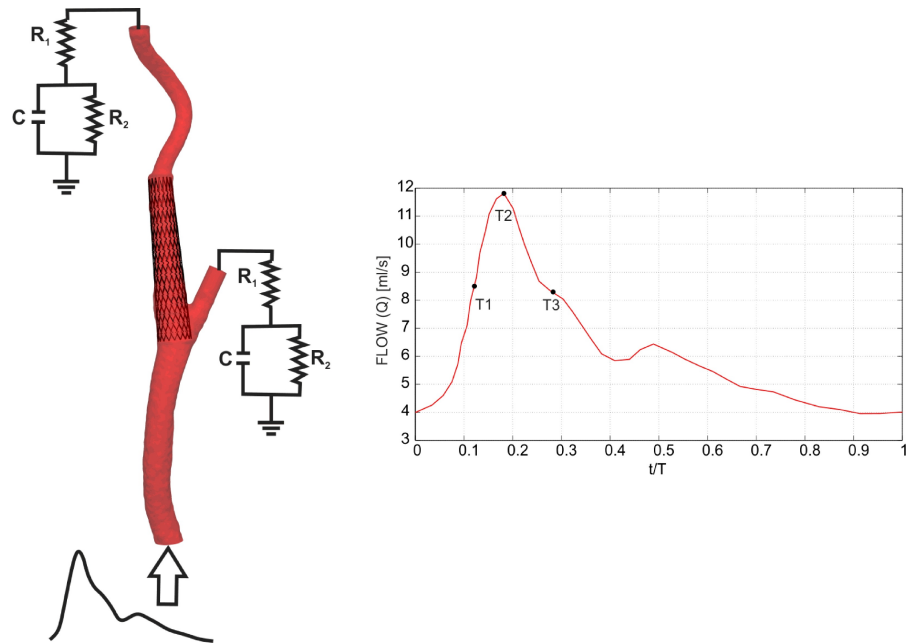


Figure 4: Schema of boundary conditions used to simulate carotid hemodynamics. A post-operative (stented) case is shown with a cyclical flow rate waveform at the inlet and two resistance/compliance models at the outlets. Reprinted from [66] with permission from Elsevier.

boundary conditions is shown in Figure 4. After generating the mesh and defining the boundary conditions, the simulation is performed over several time steps, the number of which depends on the solver configuration requirements. The literature reports values ranging from hundreds [79] to thousands [131] of time steps per cardiac cycle. Sousa et al. [287] and Guerciotti et al. [116] compared the simulated velocities in carotid arteries with measurements from Doppler ultrasonography. Both show good agreements between the measurement and simulation. Sousa et al. report a maximum velocity deviation of 5% in the ICA, demonstrating the applicability of CFD to model physiological blood flow in the carotids.

The resulting velocity fields are sometimes saved for specific time points, although this can be memory-intensive and may not always be done. Wall-related attributes are often computed in post-processing, serving as potential predictors or indicators of vascular wall dysfunction. Low WSS, in particular, has been repeatedly linked to atherosclerotic plaque development [124]. These attributes require fewer memory resources and can be saved for multiple time steps or as time averages. Examples of such parameters include WSS, the oscillatory shear index (OSI), and relative residence time [179]. Probing the velocity field also allows the calculation of the volumetric flow rate, which measures

the amount of blood volume passing through a defined surface region per unit of time:

$$Q = \|\mathbf{u}\|A \cos(\theta), \quad (2)$$

where \mathbf{u} is the velocity vector, θ is the angle between the flow direction and the surface, and A is the surface area.

It is important to note that there are several limitations and challenges in carotid hemodynamic simulations. Many studies rely on a limited number of datasets, sometimes as few as one or two, which raises concerns about the generalizability of the results [61]. Modeling the carotid artery with only two outlets introduces further inaccuracies since the ECA has multiple branches that split immediately after the bifurcation. Therefore, the computation of blood volume may be inaccurate if only the main branch is considered as an outlet. Defining flow input conditions is challenging due to waveform variations among patients and changes based on factors such as physical activity or body position. Almost all works in the literature assume rigid walls [198], although technically the carotid walls are elastic and stretch depending on blood pressure. Modeling this elasticity is difficult as it depends on various factors, including patient age. These limitations emphasize the need for method refinement and overcoming persistent challenges in simulating carotid hemodynamics. Tailored visualizations could potentially aid research in this field by providing better access to the simulation results and enabling effective comparisons, leading to more accurate simulations and improved communication of results.

3.3 VESSEL AND BLOOD FLOW VISUALIZATION

The emergence of more accurate CFD methods to simulate hemodynamics with higher resolutions and new parameters has led to a necessity for specialized visualizations to explore and analyze the resulting flow data. Furthermore, these visualizations provide a means of transfer into medical practice, where clinicians can benefit from depictions tailored to specific data exploration tasks. Specialized techniques have been proposed for three areas in particular: hemodynamics of aneurysms, blood flow in the aorta, and nasal airflow, as outlined in the state-of-the-art report by Oeltze-Jafra et al. [234]. One focus of this thesis lies in the visual analysis of blood flow in stenoses, which has only been sparsely covered before. However, it will be assumed that advances made in the other application domains for blood flow visualization can be transferred with some modulation to be applicable for carotid stenosis analysis. Therefore, this section refers to aspects of related works, from which methods and findings are then generalized and adapted to match specific requirements regarding carotid stenosis. It is based on the state-of-the-art reports by Vilanova et al. [307], Köhler et al. [156], and Oeltze-Jafra et al. [234].




Figure 5: Animated pathlines of blood flow in a vessel with aneurysm. The pathlines are visualized using arrow glyphs. Image from [175], © 2016 IEEE, reprinted with permission.

Analyzing biomedical flow fields is a complex task that can benefit greatly from tailored visualization methods [307]. Oeltze-Jafra et al. [234] identify which types of exploration different approaches facilitate, such as overview representation, probing, and contextualization. Often, multiple coordinated views of the underlying data are combined to facilitate different purposes. These may include 3D renderings of vessel surfaces, integral lines of the internal flow field, tools for data probing, and also graphs or plots of the data. Using these multiple-view frameworks, focus-and-context visualizations are an effective strategy to analyze complex flow data. For example, a flow lens can be used to probe a user-selected focus region, while spatial context is provided by the vessel wall model [96, 299]. Most approaches focus on exploring flow in a single vascular model, for instance, in aneurysms [213, 215, 216, 233] and the aorta [4, 31, 83, 212].

3.3.1 Overview

Generally, users first require an overview of the properties across the spatio-temporal domain. These can refer to the hemodynamics, wall-related attributes, or both [215]. A useful overview depiction guides towards points of interest. Wall-related parameters can be shown using colormaps, textures, and opacity [10]. The visualization of 3D flow fields generally relies on either integral curves [52], arrow glyphs [40], or particles [322]. Integral curves track the flow over a certain distance, starting at a defined point in the field. Integration techniques are used to follow the velocity vector field, resulting in paths, which can be shown as lines. Arrow glyphs are simple glyphs, which primarily encode direction and need to be seeded at defined points of the flow field. Arrow glyphs may also be used to trace along integral curves, such as pathlines. An example is shown in Figure 5. Particles, similar to integral curves, track the flow field from a starting position and display the fluid dynamics through animation.

More advanced approaches take these ideas further and perform, e.g., clustering of flow curves [216, 233, 303]. An example of stream-line clustering is shown in Figure 6. Similarly, the simulation space can be partitioned using glyphs scattered across the domain [169].




Figure 6: Streamlines (left) and streamline clusters with nine representatives (right) visualizing flow in an aneurysm. Image from [233], © 2014 IEEE, reprinted with permission.

Information on vessel walls can be combined with internal flow field data through dynamic cutaway views, which show both surface and flow-related parameters [175]. This effect can be seen in Figure 5. Straightened representations, for instance of the aorta [4], result in simplified depictions that allow at-a-glance assessments [157, 233]. Such reduced depictions of the domain may also be used in combination with statistical data, resulting in 2.5D visualizations that register abstract and spatial data [215]. Tree-like structures like vasculature can be transferred to graphs or maps that display the underlying connectivity [14, 30, 189]. In general, dimensional reduction techniques have shown great effect when creating overview visualizations in the medical domain [343, 163]. In Chapter 4 the wide range of techniques that create map-like depictions from vascular structures will be classified as *vessel maps*.

3.3.2 Temporal analysis

Blood flow is intrinsically time-dependent. This aspect of the data often needs to be shown, e.g., to allow investigation of stability or to assess the flow behavior over the cardiac cycle [218]. The general goal is to display the temporal evolution of the data and to highlight important spatio-temporal features. Typically, consecutive time steps of a simulation are shown in a juxtaposition [19, 86, 320]. However, the interpretation of such depictions can be time-intensive and exhausting. Therefore, advanced approaches facilitate animation or integrated displays of multiple time steps. For example, integral lines can be highlighted [304] or animated [174] according to the point in time. Abstract representations may also help. For instance, a circular plot was proposed to show the spatio-temporal flow behavior of vortices using the angle-to-clock analogy [158].

3.3.3 Probing

After familiarizing themselves with the domain, users often need to further explore specific regions of interest and quantify values. For this purpose, many visualization tools offer the ability to probe flow fields at specified locations, resulting in the output of numerical values or more detailed views. Common approaches are cut planes, which show 2D cross-sections of the 3D domain [19]. Also, integral curves are often dynamically seeded at such user-defined planes [315]. More advanced ideas include flexible probing geometries that can be arbitrarily placed or dragged through the 3D object. Such *slice widgets* can, for instance, show the contour of walls at the cross-section of vessels [105, 206]. Similarly, the use of a *flow lens* was proposed to separate the domain into a focus (inside) and context region (outside) [96].

3.3.4 Filtering

Exploration can be simplified if physically meaningful patterns are automatically extracted. Hiding uninteresting parts of the data, the visualization is then restricted to such features or highlights them. This is related to the concept of overview followed by probing. The key difference is the automatic filtering of the domain, which can be based on known importance measures for specific application scenarios. For instance, vortices are often the relevant parts of a flow field. They can be extracted and shown as isosurfaces [42, 43]. Some works have also proposed the automatic classification of vortices based on pattern matching [244]. Further examples include view-dependent filtering, such as automatic opacity modulation [117] and line-density control [146]. If features are known, integral curves and surfaces can also be seeded accordingly, achieving a more focused flow visualization with less distracting lines [36, 165].

3.3.5 Contextualization

Context rendering of the anatomy surrounding a focused structure, like a vascular tree, can support spatial orientation and correlation of flow and morphology [234]. The objective of the underlying task is to understand the initiation and progression of a possible pathology. A typical approach to display context information is to use semi-transparent surfaces, e.g., of the near vasculature [3, 93]. For example, cerebral aneurysms can be embedded within the arterial system [125]. Advanced visualization methods are often specific regarding either general structures (far context) or directly connected tissue (near context) [234].

Far context might refer to bones, vessels, or tissue that could help to understand the location and orientation of the focus structure [128,




Figure 7: Slice-based representations of the scalar fields of two data sets directly on one component: surface mesh (top), volume (middle), and path lines (bottom). The pressure and velocity values are compared before and after the surgical treatment of laryngeal cancer. Image from [351], under Creative Commons license [70].

[129]. This type of information can be found directly in the medical volume images and can be filtered before display, to avoid clutter and occlusion [150, 228]. It can also be displayed indirectly using volume rendering with pre-defined transfer functions [128, 129]. Sophisticated transfer functions take the gradient of the volume into account, in addition to the intensity values [128, 129]. This allows visual differentiation between vessels, bones, and other types of structure.

Near context is often associated with vessel walls, a display of which is useful to provide context for the internal flow field [97]. Methods that proved successful to display vessel walls with this aim are front face removal, view angle-dependent transparency, local shadows, and distance-dependent desaturation [97]. Crossing into the area of illustrative visualization, other approaches rely on contour rendering [138, 304], contour enhanced rendering of shaded surfaces [127], and silhouette rendering with half-toning [31, 32].

3.3.6 Comparison

If multiple instances of data need to be compared, e.g., to contrast treatment approaches, to study differences across patients or cohorts, or to validate a simulation, techniques for explicit comparison are required. Such methods target to highlight differences in data sets

and allow exploration of these disparities at various levels of detail. In CFD research, again, juxtapositions are usually used to accomplish this task [3, 94, 173]. Dedicated visualizations, on the other hand, utilize integrated comparative views. Such displays can be based on glyphs [216] or flattening techniques [110]. Another example is the work by van Pelt et al. [305], who allow the concurrent investigation of an untreated aneurysm's hemodynamics and up to four virtual treatment approaches. This is made possible by limiting the comparison to selected important parameters. Recently, Meuschke et al. [351] investigated different methods for the integrated comparison of two data sets with internal flow fields. An example of the approach is shown in Figure 7.

3.3.7 Robustness Evaluation

Relying on CFD to derive hemodynamics requires evaluation and validation of the computed data. This is especially important if the objective is clinical use. Multiple studies have found significant correlations between data measured in-vivo and computed through CFD [11, 51, 132], concluding that fluid simulations might be a non-invasive alternative to current methods [109]. However, it has also been demonstrated that uncertainties persist in the prediction of flow patterns [274]. These can arise from image artifacts, incorrect segmentation, and oversimplified boundary conditions. Also, a multitude of solving procedures exists, which can produce varying results. For this reason, visualizations have been developed that encode uncertainty directly [29, 35, 160].

4

Vessel Maps

This chapter is partly based on:

Pepe Eulzer, Monique Meuschke, Gabriel Mistelbauer, and Kai Lawonn. "Vessel Maps: A Survey of Map-Like Visualizations of the Cardiovascular System."

In: *Computer Graphics Forum* 41.3 (2022), pp. 645–673.

Maps of the cardiovascular system have existed for a long time for education and reference purposes. Their function is to simplify complex 3D vascular structures in 2D abstractions to allow efficient identification of principal vessels and branching points and ultimately foster an understanding of blood flow supply areas. Take, for instance, Weber's schemata of the human pathways, a collection of abstract depictions of the circulatory and nervous system designed in the 1970s for medical education [312]. Figure 8 shows the human arterial system, including all major arteries, their topological connections, and supply areas on a single page. Notably, a high level of abstraction is employed to reduce the visual complexity. The arteries are flattened to a 2D plane and straightened links are used to increase legibility and achieve a very high information density. Line thickness is utilized to encode the size of different arteries and make important pathways more prominent. The anatomical context is provided by outlines of the head and respective organs. Due to these abstractions, the resulting illustration bears similarity to a complex street map.

The advance of imaging techniques that can non-invasively capture patient-specific vasculature has spawned the creation of numerous techniques for map-like visualization of cardiovascular structures. The objectives of these *vessel maps* follow a similar principle as the illustration discussed above: they provide a single uncluttered overview of anatomical features that can be interpreted at a glance. These views can be used to detect and evaluate features by abstracting information and presenting it in a comprehensible layout. Occlusion-free maps of vascular structures can guide medical practitioners to efficiently understand a patient's state, assess possible vessel pathologies, and navigate complex vessel trees. As opposed to other rendering techniques, such as direct volume rendering, map-like depictions generally require more preprocessing of the underlying data but yield representations that drastically reduce the required user interactions and already filter key features. If the map creation is standardized, it can further enable comparisons in cohort studies and help quantify the effects of treatments, ultimately leading to better and increasingly individualized treatments. Due to the prevalence of cardiovascular disease [99], techniques that capture clinically relevant features from angiographic imaging in clear and expressive map-like depictions are of fundamental importance. Creating a vessel map, however, is a twofold challenge. First, in the medical application domain, a broad spectrum of requirements exists that depends on specific tasks in di-

E. M. W. Weber:
Schemata der Leitungsbahnen des Menschen
© Springer-Verlag Berlin Heidelberg 1978

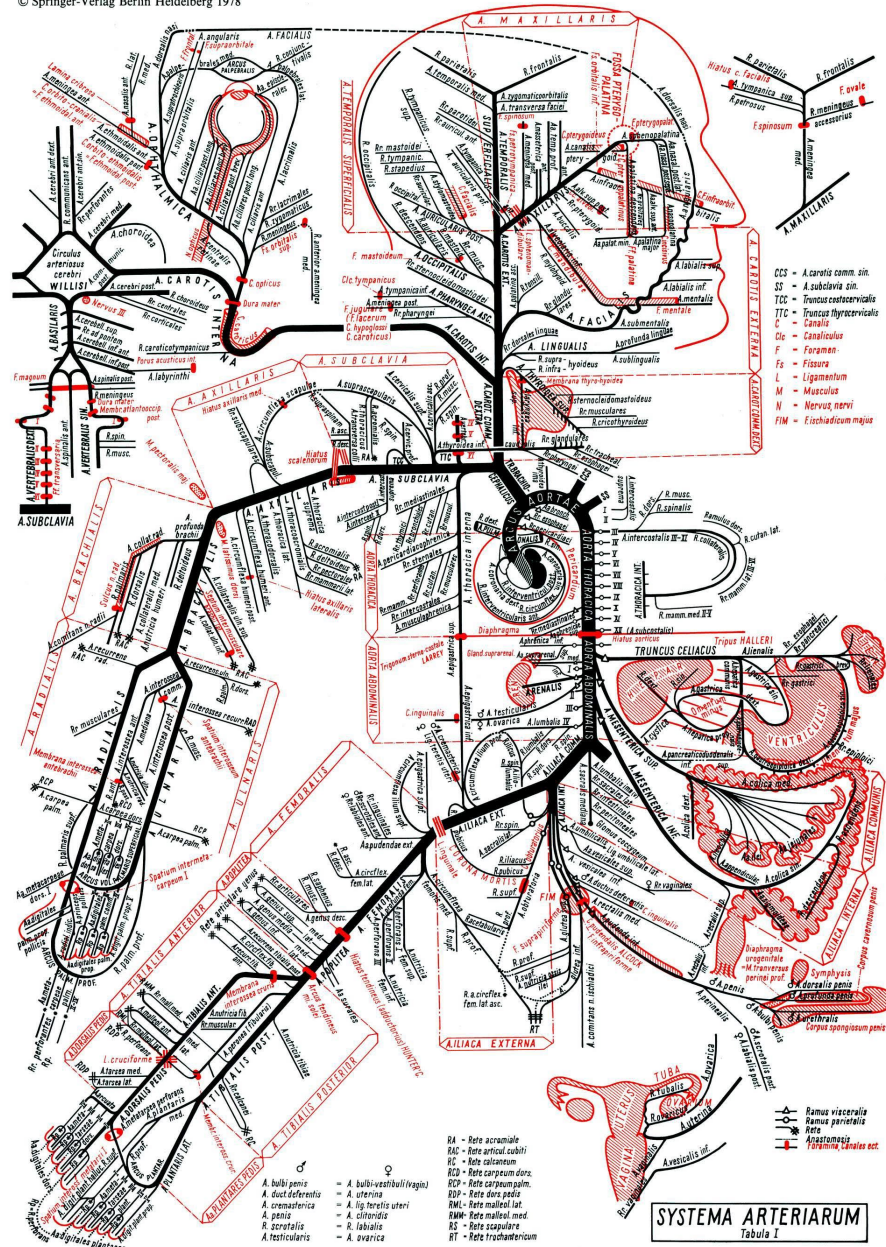


Figure 8: Map-like illustration of the human arterial system from [312].
© Springer-Verlag Berlin Heidelberg 1978.

agnosis, treatment planning, or research. There is no single depiction of a patient’s cardiovascular system that can cover all possible applications. Second, numerous data types can be visualized in a vessel map, necessitating different approaches for creating map layouts, transferring the geometry of the data to that layout, and visualizing the underlying attributes (cf. Figure 9). To make the challenge of creating vessel maps more approachable, this chapter reviews techniques to

compute patient-specific vessel maps that emerged at the cross-section of visualization, computerized medical imaging, and radiology.

Despite the relevance of the topic and the extensive amount of map-like depictions proposed to visualize cardiovascular structures, no taxonomy or classification of vessel maps has been proposed yet. Most existing surveys on vessel segmentation and visualization do not cover map-like depictions of vessels but focus on 3D techniques [39, 183, 251]. The recent state-of-the-art report on map-like visualization by Hogräfer et al. [130] discusses abstract map creation in different visualization contexts, however, it does not cover medical data or applications. Only in their survey on flattening-based medical visualization techniques, Kreiser et al. [163] review a subset of the methods proposed to create 2D overviews of vascular structures under the umbrella of “circulatory system flattenings”. However, the selected literature is incomplete, as the focus lies on mesh parameterization, missing other techniques for the generation of map-like depictions. Also, no further differentiation of methods is provided. In the following, we will show that a diverse assortment of such methods exists, which are based on heterogeneous data spaces. For instance, the curved planar reformation of an image volume is an entirely different approach than the mesh parameterization of a vessel surface or the radial graph embedding of a centerline. Yet, all techniques intend to facilitate similar user tasks centered around overview, exploration, navigation, and/or comparison. This chapter aims to fill this gap in the literature and provide a consistent classification and complete overview of vessel map techniques that have been proposed. Furthermore, we connect the techniques to the domain-specific tasks, from which we ultimately derive suggestions regarding the usefulness and applicability of various vessel map approaches.

Vessel maps are also different from map-like depictions of other anatomical structures. They usually need to transfer geometry from the \mathbb{R}^3 to the \mathbb{R}^2 space, while simultaneously creating a readable network layout. The core challenge lies in creating an optimal layout, mapping the data geometry into the layout, and then encoding the data attributes in the resulting map-like visualization. For each technique, we further filtered relevant attributes, including their algorithmic dependencies (e.g., if a centerline or view direction is required), which type of graphs the technique can handle, whether it is generalizable to different vascular structures and if properties like the vessel length are preserved. In summary, this chapter has the following core contributions:

- A review of domain requirements for vessel maps.
- A classification of the literature on vessel maps according to their mapping technique.

- Recommendations for the creation of vessel maps, depending on user task and data source for the relevant cardiovascular structure.

Selection Criteria. The literature included in this chapter comes from a range of different journals and venues, covering interdisciplinary contributions to the topic. In all selected cases, a map-like visualization of some part of the cardiovascular system is employed to address a medical domain task. The vessel maps are created from patient-specific data and are designed to facilitate diagnostic, treatment, or research purposes. Primarily, the search engines from Google Scholar¹, the IEEE Xplore Digital Library², the ACM Digital Library³, the Vispubdata data set⁴ [134], and the Eurographics Digital Library⁵ were used. We searched for keywords that are a combination of (1) vessel, vascular, or cardiovascular with (2) map, planar visualization, unfolding, untangling, straightening, projection, flattening, parameterization, or reformation.

Outline. First, an overview of the medical requirements for vessel maps is provided, followed by a review of the data spaces from which vessel maps are created. Then, we introduce the taxonomy of vessel maps, structuring different approaches and outlining the design considerations that must be made. Next, the existing literature is categorized by the taxonomy. We use the geometry mapping type (cf. Figure 9) for the top-level categorization of techniques because we aim to provide an overview for readers who know which data they will work with and are looking for related techniques. Finally, we conclude with recommendations on how to create a vessel map, which are derived from the domain requirements, the taxonomy, and the challenges solved by existing techniques.

4.1 VESSEL MAP REQUIREMENTS

As described in Chapter 2, an expansive range of symptoms, causes, risk factors, and underlying mechanisms for CVD exist. This leads to the fact that diagnostics, treatment, prevention, and research are often handled by interdisciplinary teams. Within the medical domain, experts from cardiology, hematology, pulmonology, neurology, radiology, vascular surgery, cardiac surgery, and neurosurgery need to cooperate to combat the complexity of CVD. By simplifying the visual output of vascular imaging and increasing recognizability and comparability, vessel maps can aid the cross-communication of different experts. A useful vessel map needs to distill the important details relevant to the

¹ <https://scholar.google.com/>

² <https://ieeexplore.ieee.org/>

³ <https://dl.acm.org/>

⁴ <https://vispubdata.org/>

⁵ <https://diglib.eg.org/>




Figure 9: We describe the creation of a vessel map as a three-step process: the layout of the map must be derived, the geometry or spatiality of the data must be mapped, and the data attributes must be mapped. We differentiate between image volume maps, flow volume maps, vessel wall maps, and vessel network maps depending on the primary data structure the map-like visualization is based on.

clinical task and present them effectively. Therefore, it is crucial to determine which features can be targeted and which tasks exist.

CVDs that include visible morphological changes to cardiovascular structures, i.e., vascular malformations that are discernible in imaging data, can be addressed by the use of a vessel map. Three major types of vascular malformation exist, stenosis, aneurysm, and dissection. A *stenosis* is a localized vessel narrowing, often due to atherosclerotic plaque that builds up on vessel walls. It can cause a stroke if the blood supply to areas of the brain is restricted and heart disease if coronary arteries are affected. An *aneurysm* is a localized bulging of a vessel, often due to a weak spot on the vessel wall. Aneurysms come in various shapes and sizes but generally increase in size over time.

They may rupture, causing uncontrolled internal bleeding. A ruptured aneurysm in the brain can also cause a stroke. A *dissection* is a tear in the wall of a blood vessel, leading to a cavity or pouch of blood that forms within the wall. It may also cause a stroke or heart attack if the blood supply to the brain or heart is reduced as a result of the dissection.

In diagnostics, treatment planning, and treatment evaluation of CVD, multiple tasks exist that clinicians need to perform when analyzing a (potential) vascular malformation. First, patient-specific anatomy needs to be assessed. This visual search task requires clinicians to localize pathologies by spotting irregularities. Concrete examples include the detection of stenoses, aneurysms, or missing arteries. If any candidates are found, they need to be contextualized to judge their severity, e.g., their spatial location must be known. Next, the distribution of a measured or simulated attribute on the patient's anatomy may need to be assessed. For example, the distribution of atherosclerotic plaque at a stenosis or the WSS on an aneurysm wall can be of interest. At this point, the integration of multiple features is often required, for instance, by combining the morphology, hemodynamics, and vascular connectivity in the vicinity of the targeted segment. For standardized treatments, many medical guidelines require the classification of cases by measuring predetermined properties, e.g., width, length, volume, shape, or blood flow velocity. An example is the classification of the stenosis degree by measuring its diameter inside versus behind the stenosis [88], see also Section 2.2. Lastly, if surgical intervention is deemed necessary, the accessibility of the target region and the fitness of different approaches need to be determined. For instance, the topology of a vascular tree is analyzed before a minimally invasive procedure, during which a surgical instrument needs to be inserted and traverse the inside of the vasculature.

Concurrent with tasks in medical practice, many objectives in CVD research exist that can also benefit from vessel maps. For one, medical researchers are trying to assess correlations of different attributes to determine new and more accurate markers for classifying CVD. Often, recurring patterns are sought in multiple data sets, i.e., cohort studies, to generate new hypotheses. In later-stage clinical trials, usually, the effects of a particular treatment are studied to test these hypotheses and create a predictor for future cases. Ultimately, the findings of many studies are then condensed into clinical guidelines, for instance, by incorporating a new measurable quantity into the decision-making process for a particular disease treatment. In all of these tasks, visualizations can be effective tools. Map-like visualizations in particular can help to find correlations by providing better overviews that require less interaction. Standardized vessel maps can aid in comparative tasks, make patterns visually discernible, and provide a way to quantify attributes.




Figure 10: The typical pipeline used to derive the data structures from which vessel maps can be created. Other possibilities also exist, for example, some modalities allow measuring flow instead of a simulation.

From the tasks in medical practice and research, abstract visualization tasks can be derived that map-like depictions can facilitate. We differentiate five categories of abstract tasks:

OVERVIEW: Identifying features, detecting outliers, exploring attribute distributions, and browsing the topology.

CONTEXTUALIZATION: Analyzing the spatial and/or topological context of features.

QUANTIFICATION: Measuring attributes on the local coordinate space of the map and classifying features.

NAVIGATION: Using the map to navigate other representations of the domain.

COMPARISON: Comparing features, distributions, and/or the topology of multiple domain instances.

4.2 DATA STRUCTURES

Vessel maps can be created from divergent data structures, yet they are all based on similar imaging techniques. How these structures are typically derived is shown in Figure 10. A complete description of the

processing pipeline can be found in [Section 3.1](#). The underlying data can be encoded in four spatial structures:

Voxel Volume. The cells of a voxel volume are a regular 3D grid. This can be the original or filtered medical volume image. Data is typically expressed as intensity, i.e., scalar values, which are specific to the imaging modality used. The volume not only covers the cardiovascular structure but also the surrounding context. Usually, there is no temporal resolution, although exceptions exist. Often, the location of a target structure is automatically or semi-automatically marked during preprocessing.

Flow Volume. The cells of a flow volume typically are an irregular 3D grid that describes the blood flow inside a vessel. It can originate from flow simulation or [PC-MRI](#) and is normally confined to the vessel lumen (the volume encased by the inner wall of a vessel, where blood is flowing). A flow volume is often temporally resolved. Common attributes stored in the volume are flow velocity (vector), pressure (scalar), and derivations of the flow field.

Surface. Data can be located on vertices or polygons of a surface geometry. These are usually one or multiple surfaces embedded in \mathbb{R}^3 , like the inner and outer walls of a vessel. Wall surfaces can be temporally resolved, this is typically the case if they are combined with a flow volume. Common attributes stored on vessel wall geometry include wall thickness (scalar), [WSS](#) (scalar), wall normal stress (scalar), wall displacement (vector), and the occurrence or thickness of plaque (scalar).

Graph. Lastly, data can be located on the vertices of a centerline graph, which is also embedded in \mathbb{R}^3 . It is usually not temporally resolved and encodes properties like the vessel radius (scalar), cross-section area (scalar), or branch label.

4.3 TAXONOMY

We introduce the notion of vessel maps as a collective for map-like visualizations of the cardiovascular system. Vessel maps have been proposed for heterogeneous data structures, like volumes, surfaces, or trees, and can be based on a divergent range of algorithmic approaches, for example, untangling, straightening, unfolding, flattening, or reformation. However, they share a common purpose in the simplification of complex cardiovascular structures for visual interpretation, they map attribute distributions and/or connectivity information, and they are designed to aid one or multiple of the tasks described in [Section 4.1](#). We use the term map-like to allude to the properties of the described visualizations, which are similar to maps in the traditional sense. First, they are 2D depictions of a spatial domain and aim to preserve the locality and the relative positioning of features. Second, they abstract the attributes relevant to the observer to increase legibility. Third, they

conform to the uses of maps, which include getting an overview of a region, navigating a domain, exploring distributions and connections, and comparing different areas. Following the definition of Hogräfer et al. [130], map-like visualizations exhibit traits of both (cartographic) maps and charts/plots. In this sense, vessel maps are schematizations of cardiovascular structures instead of geospatial data.

Vessel maps are also an interesting case to study from a pure visualization point of view. Their spectrum ranges from visualizations of networks and trees to the depiction of surface fields. Interestingly, many ideas proposed for vessel maps do not fall completely into either category but are situated somewhere in between. This range requires otherwise distinct techniques to be merged, like network and surface visualization. For the creation of a vessel map, three aspects must be considered:

1. What is the layout of the resulting map?
2. How is the geometric component of the data mapped to this layout?
3. How is the attribute component of the data mapped to this layout?

These steps are also shown in [Figure 9](#). Steps one and two are sometimes solved interdependently [55, 59, 334], sometimes independently [30, 188, 205, 319]. Depending on the user task, the branching topology, object geometry, or data attributes can be of interest. It should be kept in mind, however, that preserving a specific property in the mapping process might require sacrificing another.

4.3.1 *Layout Generation*

Commonly, a vessel map needs to create a flat layout from an input graph in \mathbb{R}^3 . Most techniques use constraints to build the layout. Typical constraints include preventing self-intersections and preserving original properties, such as angles, overall shape, or the relative length of segments. As illustrated in [Figure 9](#), the layout directly depends on the input graph of the vascular structure. Also, it indirectly depends on the user task, as it needs to visualize the appropriate region.

In the simplest situation, the region is a single segment, in which case no considerations regarding branching must be made. This is easy to solve, as no actual graph layout must be determined. In practice, a map of a specific vascular segment or surface patch is created, like an individual aneurysm [110, 215]. If the region is a tree of depth 1, singular branch points must be considered. We list this as an individual case, as sometimes a specific vessel segment is focused but branches exist in the evaluated region [142, 201, 264]. Often, solving the layout for a tree of depth 1 does not require sophisticated algorithms. An




Figure 11: Straightening of a mouse aorta centerline tree with depth 1 for curved planar reformation (CPR). (1) The projected centerline, (2) the straightened centerline, (3) the straightened centerline with 90° bifurcation angles, and (4) the straightened centerline with normalized segments. On the right, a polygonal mesh derived from the centerline is shown. Image from [264], © 2009 IEEE, reprinted with permission.

example of a straightening for a depth 1 tree is shown in Figure 11. If the region is a tree of depth n , a variable number of successive sub-branches must be considered. Usually, a generalized solution is used that can be applied to different trees. For the most part, on the macroscopic level, the vascular system is a tree and many layout approaches make use of this property. If the region is a cyclical graph, sub-branches *and loops* must be considered, which is a harder problem to solve. This is the case, for instance, for the circle of Willis, a central arterial structure in the cranium.

Depending on which properties the final visualization should have, different 2D layouts are suitable. In simple cases, projection has been used to create a 2D from a 3D tree structure [33, 143, 227, 298]. The obvious flaws of this approach are possible self-intersections that may lead to visual ambiguity and/or occlusion of data. Still, simple projection may be viable if, e.g., only a single segment is mapped. Many techniques attempt to create a readable layout that follows certain rules. A *straightening* lays out the branches of a tree as straight lines, which can be connected [342] or disconnected [104]. The advantage of straightening is that the individual segments can be easily visually followed. Often, the arc length of segments is preserved. An *untangling* removes self-intersections in the layout, while simultaneously preserving attributes like the overall curve or relative positioning of

segments [205]. Untangled views are closer to the original layout, improving familiarity, but can also create more complex depictions than other techniques. Some techniques also use a fully pre-determined layout. The input graph is then fitted to align with this blueprint [239]. With this type of procedure, high comparability between multiple vessel maps can be achieved and it is often used for standardization purposes. However, the original proportions of the input graph will be lost.

4.3.2 Geometry Mapping

The layout determines where the data should appear in the visualization. How the data is mapped from \mathbb{R}^3 to \mathbb{R}^2 is determined by the *geometry mapping*, see Figure 9. Sometimes, layout and geometry mapping are solved as the same step in an algorithm [218] but they can also be solved independently [205]. Common constraints used for geometry mapping techniques are the prevention of overlaps and the minimization of the area or angle distortions, i.e., the original proportions are attempted to be preserved.

The geometry mapping directly depends on the data structure of the geometry – not on *what* the data is but *where* it is. The data structure can be a regular 3D grid, which typically is the voxel volume of an angiographic imaging modality. Hence, we group the vessel map visualizations that directly build on this type of data under the term *image volume maps*. These can either be created by cutting the volume with a surface and reformatting the cut in a planar layout or by showing voxel values that were aggregated in the vicinity of the vessel. The geometry can also be an irregular grid, which is normally the case for flow volumes, where each cell holds properties of a field that describes the blood flow. Therefore, we call the resulting visualizations *flow volume maps*. Flow volume maps can be created by reforming the data space, e.g., by straightening the flow volume, or by deriving a graph layout from the flow features. Often, data is situated not in a volume but on a surface embedded in \mathbb{R}^3 . Focusing the analysis on wall properties makes sense, as CVD generally develops on or inside the walls of the cardiovascular system. If the data shown in the resulting map-like visualization is associated with the vertices or faces of a wall geometry, we call it a *vessel wall map*. Such maps were created with projection techniques or surface parameterizations. Lastly, data can be located on the vertices of a centerline graph. If the data is encoded exclusively on this type of graph, the layout step already defines the geometry mapping. As the resulting vessel maps focus on the vascular tree or network, we call them *vessel network maps*. They can either keep some geometric properties by straightening or untangling the graph structure or be abstract representations that only convey the topology.

4.3.3 *Attribute Mapping*

In the final step, a visual encoding for the data attributes needs to be chosen. The attribute mapping is notably distinct from the layout and geometry mapping, as it only defines how the data is shown, not where. We encourage thinking of this step as an individual part of a vessel map visualization. Some data structures have a strong association with certain attribute mappings, e.g., visualizations of image volumes tend to use a grayscale colormap. Defaulting to such a mapping can have benefits regarding recognizability but might also mean that other possibilities are overlooked and remain untested.

The suitability of an attribute mapping must be determined by the user task that it should facilitate. For instance, if the user attempts to find certain features in an attribute range, they should be highlighted by the chosen encoding. For vessel maps, five types of attribute mappings have proven to be useful:

1. Size encodings, which are often used as a way to show the thickness of a vessel by varying the width of rendered segments. The width can be quickly visually read and is, naturally, an intuitive representation of vascular morphology.
2. Windowing encodings that map intensity values to gray values while providing the necessary interaction to explore the entire data range (usually 12 bits). Interactions include changing brightness and contrast. These encodings can be considered a special case of colormaps and are predominantly applied to image volume maps.
3. Color encodings are widely utilized to make full use of the properties of colormaps. They are particularly beneficial to show scalar field distributions.
4. Some approaches exist that use illustrative encodings [213, 215]. Techniques like hatching can be used to identify certain regions.
5. If multiple attributes must be visualized, glyphs can be used as an additional way to encode information [112, 219]. Symbols that encode attributes through their shape, size, and color are common in traditional maps, as they can add explorable layers of information to a spatial domain without the need for direct interaction. The advantages of glyphs can just as well be utilized for vessel maps.

Attribute mappings can also be combined to show multiple attributes at once. For example, a size encoding can easily be used in combination with a colormap. Combining attribute mappings to exploit multiple visual channels makes map-like depictions a powerful tool for communicating information efficiently and exploring correlations.

4.3.4 Classification

In the following, a classification of the literature on vessel maps is provided. We discuss the application domains, which structures are visualized, and why the respective data space is chosen. The primary questions we answer follow the taxonomy:

- What is the input graph? (segment, tree of depth 1 or n , cyclical graph)
- What is the layout of the resulting map? (projection, straightening, untangling, fitting)
- How is the geometric component mapped? (cut surface, voxel aggregation, flow straightening, flow graph, surface projection, mesh parameterization, network straightening/untangling, topology map)
- How is the attribute component mapped? Which encodings are used? (size, windowing, color, illustrative encodings, glyphs)

If applicable, the following secondary questions are answered:

- Which tasks are facilitated? (overview/exploration, contextualization, quantification, navigation, comparison)
- Which dependencies exist in addition to the data structure? (centerline, view direction, landmarks, manual cut)
- Is the approach generalizable to different structures? If yes, has this been demonstrated?
- Which properties are preserved? (arc length, diameter, surface area, angles)

4.4 IMAGE VOLUME MAPS

This section covers all techniques that work on the voxel level. Mostly, these are derivations of curved planar reformation (CPR), i.e., techniques that use centerlines to align non-planar cuts and then perform image reformation to display selected vascular structures in volume images within a single 2D depiction. The results are map-like visualizations of the vasculature, which are commonly used to identify and assess calcified plaque and potential vessel stenosis. Typical challenges are diameter preservation, how to incorporate vessel wall features, and how to preserve or display the surrounding context. We differentiate between techniques that display volume cuts [1, 13, 104, 126, 142–145, 164, 177, 201, 263, 272, 290] and techniques that aggregate voxel values to display certain features [45, 77, 221, 222, 264].

Achenbach et al. [1] generated CPRs to study stenotic regions of coronary arteries in electron beam computed tomography (CT). They manually defined the cross-sections of each coronary artery in axial slices and use a reference plane (axial, sagittal, or possibly coronal) to create the reformations. The results show an image for each vessel branch in stretched form, allowing distance measurements. Context visualizations outside the vessel lumen were not presented. This approach can be applied to any tubular structure, since only a curve reflecting the course of the vessel is needed, together with a reconstruction or reference plane.

Kanitsar et al. [143] describe a workflow to examine peripheral arteries for stenosis or occlusion. To provide an unobstructed view of the blood vessels in CTA data, the bones must be segmented without holes and masked out. To this end, they use a threshold-based procedure that consists of two steps, the identification of outer cortical bone and the subsequent extension to the softer marrow. In the next step, vessels are tracked between user-defined start and endpoints. The shortest path between the respective two points is calculated using a cost function. Eventually, the obtained vessel centerlines are centered in perpendicular cross-sections using a ray-casting approach. After that, CPRs can be generated and the vessels examined. This workflow is generalizable to arbitrary tubular structures, but some thresholds have to be adjusted accordingly.

Medial axis reformation was introduced by He et al. [126] to visualize the interior of blood vessels. Their approach extracts the medial axis of each vessel branch using a centroid-based skeletonization approach and then improves the centering based on their medialness. These centerlines are then spread segment-wise in image space. The direction of the spread can be controlled by the user. At branching points, the image is split and the process continues for each sub-image. Since each vessel branch is considered a cylinder, it can be projected into image space without any deformation, allowing the user to measure the length and diameter of the branch. This is similar to a stretched CPR as described by Kanitsar et al. [142]. The technique was applied to coronary, carotid, and iliac arteries imaged by electron beam CT.

Kanitsar et al. [142, 145] discuss the properties of projected, stretched, and straightened CPR. The projected CPR offers good spatial orientation, while the straightened CPR does not. The stretched CPR still offers good spatial orientation, because the centerline is stretched only in local areas. Stretching and straightening retain the arc length of the centerline, while straightening also preserves the vessel diameter horizontally. The authors also introduce a rotation of the ruling vector, which is referred to as the vector of interest in this work. This allows the entire vessel lumen to be inspected from different viewing angles. However, the axis of rotation is fixed, which causes problems when the vessel path is nearly perpendicular to it. To visualize the interior




Figure 12: Untangled CPR of a peripheral CTA data set showing the entire vessel tree in a single image. Image from [144], © 2003 IEEE, reprinted with permission.

of multiple vessels simultaneously in a single image, the multipath CPR method is presented. To improve accuracy when inspecting small vessels, a thick CPR can be rendered. Here, instead of a thin surface, a small slab is resampled and rendered using averaging, maximum intensity projection (MIP), or minimum intensity projection. The presented techniques are demonstrated on a phantom data set and mainly discussed on peripheral CTA data sets, except for a bronchial example. All presented methods are generalizable to any tubular structure but require a tree as input. The generated 2D layout of peripheral vessels is created by partitioning the image according to the projected branch points.

Although (rotated) CPRs require only a few images to examine the entire vessel lumen, they can still add up to a considerable amount. To reduce the number of images, Kanitsar et al. [144, 145] proposed two approaches, namely a helical and an untangled CPR. In the former, a sampling helix is used instead of a ruling vector. This allows the entire lumen to be acquired at once and eliminates the need for rotation. However, only a single vessel can be displayed straightened using this method. The other approach, the untangled CPR (see Figure 12), is composed of stretched CPRs of each tree branch without overlaps. For each vessel branch, even if the ruling vector is rotated, a so-called vessel hull (circle sector) is constructed from bottom to top, i.e., starting from the leaves of the tree, in the image space. Then these envelopes are arranged without overlap and the final image is partitioned similarly to the multipath CPR. Results are presented using peripheral CTA data sets, and both methods can be generalized to tree-like tubular structures, where the spiral CPR can visualize only one branch at a time.

To improve visual discontinuities when depicting multiple vascular branches in a single image using ruled surfaces as in medial axis reformation or CPR and to represent cut surfaces of irregular structures such as the jaw or human pelvis, Saroul et al. [272] used free-form surfaces. First, they extend CPR by defining the orientation of the ruling vector based on the principal direction of the analyzed 3D centerline. The ruled surface is then flattened into image space by

resampling the rectangular facets spanned by the vector between two consecutive centerline points and the ruling vector. Flattening this ruled surface produces no angular or metric distortions. They then propose free-form Coons surfaces [67] for analysis of the human aortic arch, including its branches, and other structures such as the jaw. By defining curves along the centerline of a vessel that represent the maximum diameter, the resulting surface interpolates these curves and displays the variation in aortic diameter in a single image. Since Coons surfaces cannot be flattened without distortion, the user can choose a direction along which distances are preserved, as demonstrated on the aortic arch and its branches. The authors also show all teeth of a jaw that was flattened with Coon surfaces in a single image. Specifying free-form surfaces by interpolating user-defined boundary curves does not require centerlines of vascular structures.

To this time, CPR has only represented a cut surface along a curve or tubular structure, without considering alternative visual representations outside the vessel lumen. Straka et al. [290] presented several ways in which CPR (focus) can be combined with direct volume rendering (DVR) or MIP (context). They differ mainly in the transitions between the focus and context regions and in the way obstructed vessels are represented, e.g., by occlusion lines. Results are presented using peripheral CTA data sets, and the layout of the visualization resembles a multipath CPR. The approach is generalizable to other focus structures and applies not only to vascular structures but to importance-driven rendering in general.

Lee and Rasch [177] focused on improving curved sections through generally-oriented vessel trees by aligning the ruling vectors perpendicular to the vessel centerlines. The projected vessel tree is arranged from top to bottom with the largest medialness node at the top. The reason for this decision is that the branch diameter decreases with increasing depth of the tree and the largest node should be at the top. Since the projected y-direction corresponds to the parameterized arc length, the approach resembles a stretched CPR. Results are presented for coronary and peripheral arterial trees.

Roos et al. [263] evaluated multipath CPR in a clinical prospective study involving ten patients with peripheral arterial disease. The result showed that multipath CPR produces artifacts when vessels are coincident or nearly collinear with the ruling vector and when multiple vessels positioned behind each other are examined as they overlap in image space. Other than that, it is a viable solution for examining peripheral arteries in a single image but does not replace either single-path CPR or MIP.

Since a single CPR is usually insufficient to assess the entire vessel lumen, several such sections must be examined from different angles (about 20–30°). Cai [45] attempted to reduce the number of images to be inspected by enclosing the entire 3D vessel lumen with a so-

called biconvex slab rendered with MIP or X-ray. This approach can be viewed as an inverse VesselGlyph [290] with enhanced contextual representation of the lumen and a thin section through the surrounding anatomy. In this way, the entire lumen of a vessel is represented in a single image. However, a MIP of the biconvex slab shows a stenosis or complete occlusion in case of concentric calcification, because the rays are cast along the viewing direction. This is the reason why CPR needs to be rotated in the first place. Furthermore, this approach presents only a single-path stretched-like CPR on coronary and carotid arteries from CTA data sets.

To obtain centerlines of vascular structures, Lv et al. [201] used active contours or snakes. Since active contours depend strongly on their initial curve, the authors proposed an improved initial contour determined by casting rays radially from the center to each point on the contour and taking the point with the largest slope. Since this initial contour is close to the true boundary, the snakes quickly converge to the true contour. The midpoints of the resulting contours in the chosen slices are then used to create a straightened CPR of the vessel branch. Results are presented on several individual vessel branches from a head aneurysm CTA data set.

Multimodal vascular reformation was explored by Ropinski et al. [264] on positron emission tomography (PET)/CT data of the mouse aortic arch. The aim was to study the development of atherosclerotic lesions at the morphological, functional, and molecular levels. The PET and CT data sets were registered using three artificially inserted ceramic markers. PET captures inflammatory activity, whereas CT captures structure. The vessel centerlines were extracted using a curve skeletonization approach. To compare several different mice or a single mouse over time, a normalization of the aortic arch with its outgoing arteries was introduced. For this purpose, a modified straightened multipath CPR was used. Vessels are flattened outward from the centers using a ray-casting approach to allow comparative assessment. Two options are proposed: the first preserves and displays the diameter of the aorta, while the second flattens the entire aorta into a rectangle. Since this leads to undesirable distortions, especially when normalizing the distances of the aortic arch branches between different individuals, and degradation of spatial orientation, they use multiple linked views to obtain a spatial overview and a comparable detailed view.

Mistelbauer et al. [222] proposed centerline reformation, which uses wavefronts instead of a ruling vector to render a CPR. Their approach resembles a projected multipath CPR and renders the cut surface through the vessel lumen of arbitrarily oriented vascular structures. To ensure proper visibility of multiple overlapping CPRs in image space, their approach uses a depth buffer and parameterizes the vessel tree according to the length of its branches. Depending on




Figure 13: Anatomical layout of a peripheral CTA data set (right in 3D) showing several CFA visualizations arranged from the center to the left and right. Image from [220], courtesy of Gabriel Mistelbauer.

the viewing direction, the lumen that is closer to the viewer and has a minimum distance from the current lumen in the buffer, along the graph (arc length), is chosen. To provide additional depth information, halos [85] can optionally be added around the lumen visualization and a volume rendering can be displayed as context. Vessel centerlines were extracted from CTA data sets using multiscale vesselness. The results show the vessels of a human head, pulmonary arteries, and a human abdominal aorta. As demonstrated on a phantom data set consisting of differently-sized helices, the technique is generalizable and extensible.

To reduce the number of CPR images when examining a vessel lumen, Mistelbauer et al. [221] use nonlinear ray-casting along concentric circles perpendicular to the centerline of the vessel. Samples along these circular rays are then combined (or aggregated) into a single value using the minimum or maximum. By straightening the centerline of the vessel in the final image, it is divided into left and right sides. Each side can then show either the same aggregation method or two different methods. A combination of MIP and minimum intensity projection allows the radiologist to view either calcified or soft plaque in CTA data. Results are presented for stenoses of the abdominal aorta and peripheral arteries. For the latter, an anatomic layout was proposed because multiple vessels were inspected, see Figure 13. This technique is generalizable to all tubular structures, but the anatomic layout of the multiple straightened curvicircular feature aggregations (CFAs) is case-specific.

At the same time as CFA, Diepenbrock et al. [77] introduced normalized circular projection. They acquired PET/CT data to analyze the development of atherosclerotic lesions in mouse carotid arteries. Analogous to CFA, normalized circular projection uses a sampling

scheme around the centerline of a vessel but starts outside the vessel wall. Within cross-sections perpendicular to the centerline, rays are defined and sampled outward from the centerline point. The vessel wall is estimated in CTA using the maximum gradient along these rays. Once the wall is found, the PET data are sampled. By inflating the vessel wall into a cylinder, PET activity outside the vessel is visualized using normalized circular projection with MIP aggregation and in a straightened way. In a multiple linked view application, users can examine the left and right carotid arteries of mice in a standardized and normalized manner.

Previously, CPR rotation was mainly constrained by the ruling vector, with the main challenge to ensure a continuous and smooth section along the centerline of the vessel and through its surrounding parts. Curved surface reformation, introduced by Auzinger et al. [13], allows unrestricted rotation of a vessel tree while continuously examining the lumen and surrounding parts of multiple vessels with sufficient visibility. The approach is based on ray-casting, where each ray cuts small strips generated by extruding the small line segments of the sampled vessel branching curves perpendicular to their direction and the viewing direction. A cost function decides which intersection point to use for the final result, favoring points that are closer to the viewer and the vessel centerline. The resulting images correspond to a projected multipath CPR, but with unrestricted rotation. Vessels of a human head and peripheral vessels, both acquired with CTA, are shown and also presented with maximum intensity difference accumulation [38] as context. The technique requires only a graph consisting of curves (vessel branches) and is therefore generalizable.

Kretschmer et al. [164] propose an improvement to most CPR approaches. By filtering the depth image of the cut surface immediately before sampling the data set, small discontinuities are removed. To ensure that the surface still passes through the vessel lumen, the projected pixels of the centerline are marked as fixed and are not considered in the filtering; this could be extended to the entire lumen if diameter information were available. To preserve large discontinuities caused by distant vessel branches along the vascular tree, but to smooth vessels that are close to each other, a bilateral filter is used. The results are demonstrated on a CTA data set.

Gillmann et al. [104] present a visualization for coronary artery analysis, especially for surgical preparation, intervention, and restoration. Coronary artery centerlines are obtained from coronary CT data sets. The layout of their visualization is oriented from left to right and is similar to a treemap and the visualization presented by Borkin et al. [30]. The individual branches show a CFA with maximum aggregation at the top and average aggregation at the bottom. They use constant arc length sampling along the curved rays or concentric sampling circles. Calculations or intensity values that are above a user-defined isovalue

are highlighted in red to draw the user's attention to that region and prompt further analysis.

4.5 FLOW VOLUME MAPS

Maps of flow volumes are a less explored area that is nonetheless distinct from other techniques. Blood flow information can be acquired from [PC-MRI](#), Doppler ultrasonography, or simulated with [CFD](#). The results are often visualized with [3D](#) techniques like integral lines, particle animations, volume rendering, or flow profiles of cross-sections. Some visualization techniques have been proposed to create map-like depictions of complex flow volumes, which are easier to compare and can be quickly assessed. One approach is to straighten vascular structures, including the flow field [[4](#), [22](#), [278](#)], another is to create a [2D](#) graph visualization of the flow [[296](#), [329](#)]. The challenge here is not only to straighten the vessel structure but also attributes derived from the surface area. Information such as scalar fields on the domain and especially flow data represented as vector information within the surface domain also needs to be transformed consistently. Therefore, it is not sufficient to just straighten the surface; [3D](#) flow information must also be straightened consistently.

Flow volume straightening. Angelelli and Hauser [[4](#)] present a straightening approach to simplify the aorta with its simulated blood flow. Using volumetric data, they first extract the centerline of the aorta and use it to construct a curvilinear grid, which is then straightened preserving the length of the centerline. The curved grid can also be used to reformat the blood flow into the straightened space, giving a simplified overview of the aorta and blood flow simulation. The blood flow is then represented with streamlines or path lines using a color map. Behrendt et al. [[22](#)] use a 2.5D representation of the aorta with its blood flow. As in the approach of Angelelli and Hauser [[4](#)], a centerline is needed to transform the vertices of the surface mesh as well as the points of the pathlines into a straightened space. In addition, the centerline has six landmarks that are used to further simplify the straightened aorta. The parts of the centerline between the landmarks are either stretched or compressed to ensure equal spacing, and the radius is also normalized such that the simplified aorta becomes a cylinder, see [Figure 14](#). During this simplification process, length and radius are intentionally distorted to ensure a consistent model for the comparison of different data sets. Seifert et al. [[278](#)] applied a straightening approach to [2D](#) fluorescence reflectance imaging. Using their standardized layout they aim to detect patterns in the development of atherosclerosis. To achieve this, multiple boundary points on the aorta must be defined by the user. Then, the boundary is smoothed and connected to form a polygonal net. In the next step, the centerline is straightened with the polygonal net, which also results in straightened




Figure 14: Normalized map of the aorta with flow, created with the technique of Behrendt et al. [22]. From top to bottom: 3D rendering of the domain, straightening, even distribution of segments, radii normalization, flattening along the circumference.

fluorescence reflectance imaging. To ensure comparability, the image is normalized over the radii of the vessel.

Flow graphs. Another approach to mapping the flow volume while preserving the perceptual structure of a vessel was proposed by Tao et al. [296]. The volume of the vessel is divided into blocks, which are then transformed into 2D points. The formation of the 2D vessel structure is formulated as a graph layout problem. Both Euclidean distances and geodesic distances of the volume are integrated to preserve the perceptual structure and to ensure that there is no self-

occlusion in the final result. After the 2D map is created, the user can further explore the underlying blood flow data with heatmaps, graphs, histograms, and color-coded matrices. Brushing-and-linking techniques on the map are used for the navigation of the domain. Recently, Zhang et al. [329] proposed a simplification of flow data based on a 2D map-like depiction of stream surfaces. Their approach can be applied to a variety of data, including vascular structures with flow fields. To achieve a simplified 2D representation of the flow, the problem is considered as a graph layout optimization based on three energy terms. These energy terms are used to order the simplified flow so that parts of the flow are closer together when they are close in time. This applies to the x - and y -axes and it avoids temporal flipping, i.e., when the temporal ordering changes in the x -direction. The simplification yields a 2D overview of stream surfaces, where color is used to differentiate multiple seeding curves. Additional ellipsoidal glyphs are used to represent velocity or vorticity. A stacked layout of the resulting curves facilitates overview and comparison tasks.

4.6 VESSEL WALL MAPS

As vascular pathologies develop on the vessel walls, numerous techniques focus on creating maps of these walls, on which properties like thickness, plaque occurrence, normal stress, or shear stress can be displayed. There are two principal approaches to generating map-based visualizations of vessel walls: employing *mesh parameterizations* or *projections*. Approaches that rely on mesh parameterizations create one-to-one maps of surface meshes embedded in 3D. Mesh parameterizations are used for mapping different vascular structures comprising: specific structures like aneurysm surfaces [110, 213, 215, 217, 218], stenosis predilection sites [8, 9, 53–59, 331], or heart valves and cavities [340, 147, 350, 202, 230, 242, 261, 316] but also arbitrary vessel tree walls [346, 334–337]. For techniques based on surface *projections*, the properties are projected onto a parametric structure like a cylinder, disk, or sphere [33, 168, 227, 232, 282, 298]. The approach is similar to mesh parameterization, however, the topology is not retained, which means bijectivity cannot be ensured. The advantage of these approaches is that they usually can be computed on the fly for local structures.

All of these techniques result in maps where surface parameters can be assessed without rotation. Some are additionally used for navigation in visualization frameworks and some are specifically standardized maps for anatomical correspondence and statistical evaluations in medical studies. Typical challenges are how to handle branching, how to preserve properties like the area of features, how to standardize layouts, and how to cut vascular structures into topological disks.

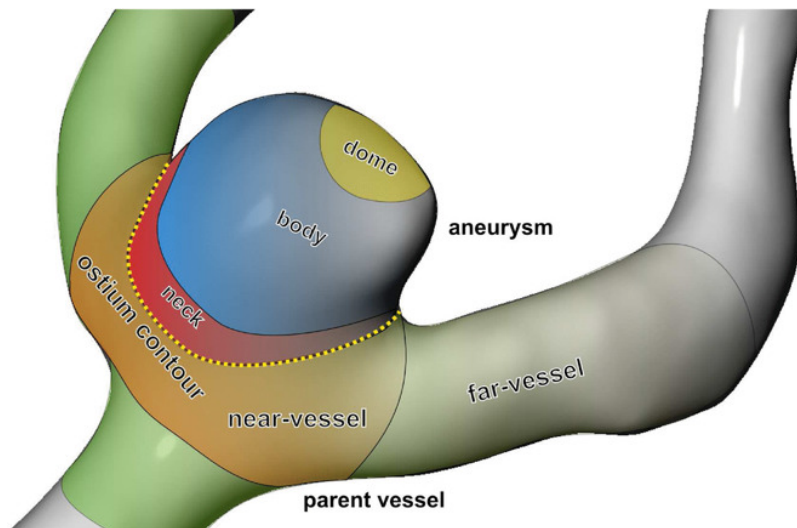


Figure 15: Morphology of a cerebral aneurysm. Image from [234], © 2018 The Eurographics Association and John Wiley & Sons Ltd.

4.6.1 Maps of Cerebral Aneurysms

Several works dealt with the generation of map-based visualizations of cerebral aneurysms. A patient-specific assessment of the rupture risk is important to decide whether treatment, which is also associated with risks, is necessary or whether the aneurysm can continue to be observed. Because the estimation of rupture risk is based solely on morphological characteristics, such as the size and shape of the aneurysm, is not reliable, numerous papers are concerned with simulating cerebral blood flow in aneurysms to derive meaningful factors regarding rupture risk. For visual analysis of the simulation data, 3D surface models of the aneurysm and adjacent vessels are usually reconstructed from the clinical image data. An example of a reconstructed cerebral aneurysm is shown in Figure 15. On these 3D models, various simulated scalar values, such as WSS, pressure, or wall thickness, are displayed using standard techniques such as color coding. The goal is to find wall regions where rupture-prone correlations of the scalar fields occur, such as regions with a low wall thickness and high WSS. Due to the time-dependence of the scalar fields over a cardiac cycle as well as complex shapes of the aneurysms, the visual analysis of the 3D models is complicated. The user must rotate the 3D model, as well as change the encoding of the currently displayed scalar field, to explore the complete domain. However, since the scalar fields usually change via animation on the surface to reflect the cardiac cycle, it is often difficult for the user to spot prominent wall regions based on 3D models. The creation of 2D aneurysm maps solves the rotation problem. In addition, the 2D maps were combined with techniques to explore more than one scalar field at a time to more easily identify conspicuous wall regions. Currently, all aneurysm maps based on the

mapping of the vessel wall take a segment as input, since only the aneurysm, without adjacent vessels, is mapped.

Neugebauer et al. [227] developed the only approach in which the geometry of the aneurysm was mapped using a multi-perspective projection. The 3D vessel surface is centered within an invisible cube. Afterward, the perspective projections are used to map the 3D surface to the cube sides. The resulting map consists of five regions representing the flow information of the left, right, top, bottom, and back sides of the vessel surface, giving an overview of the whole 3D surface. Besides the surface mesh, no other dependencies are needed for mapping. However, the map does not preserve any properties such as arc length, diameter, area, or angles of the 3D surface. The map is linked to a 3D view of the aneurysm, where a scalar field can be color-coded. The 3D vascular surface is shown in the center of the illustration surrounded by the map, where a spatial relation between both views is created. With bidirectional interaction between the two views, the map supports navigation on the 3D surface. In principle, their approach could be applied to aneurysms on arbitrary vessels.

While there is only one projection-based approach to generating an aneurysm map, several works used mesh parameterization techniques for aneurysm mapping. One of the first concepts in this area is the work of Goubergrits et al. [110], who generated an aneurysm map to analyze statistical WSS distributions. The basic idea is to transfer the aneurysm wall to a unit circle. For this purpose, first, the vertices of the aneurysm surface are moved towards the center of mass of the aneurysm, ensuring that the distance to the center is equal for all vertices. This results in a unit sphere, where an azimuthal equidistant projection is then used for mapping the geometric component, where angle information is preserved. Besides the 3D surface mesh of the aneurysm, no further dependencies such as a centerline are needed for the map generation. To assign the attribute component (the WSS distribution), a simple color coding based on the rainbow color map was used. With the generated map, an overview is given about the WSS on the aneurysm wall, where different maps can easily be compared due to the uniform shape of the map. However, reasonable results are just produced for convex structures, where the center of mass lies inside the surface. For irregularly shaped aneurysms, where the center of mass is outside the surface, this method results in area distortions in the map, especially for regions farther from the center of the map. Furthermore, this dependence on regular shapes limits the generalizability of the method to other anatomical structures.

Meuschke et al. generated aneurysm maps with different parameterization techniques to support the visual analysis of multiple scalar fields simultaneously [213, 215, 217, 218]. Since only the aneurysm without adjacent vessels is to be mapped, the aneurysm has to be separated from the healthy vessel part first. Therefore, the user can

click on the 3D surface mesh to define landmarks that are automatically connected to a cut line by applying the Dijkstra algorithm [78] that determines the shortest path based on the Euclidean distances. The cut line follows the border of the aneurysm ostium, which is an imaginary surface that separates the aneurysm from the healthy vessel part. To provide anatomical context information, the cut line along the ostium is color-coded on the aneurysm map as well as the 3D surface.

In their first work [215], they employed least squares conformal maps (LSCM) [185] to generate angle-preserving maps. In addition to the cut line, two points are needed as constraints for the parameterization, which are also set by the user by clicking on the mesh. For anatomical context information, these landmarks are also shown on the resulting map. On the map, two scalar fields can be visualized simultaneously. The first attribute is mapped to a cool-to-warm color scale, while the second attribute is encoded by an image-based hatching scheme. Dark red, strongly cross-hatched regions represent high scalar values. Additionally, a 3D bar chart is used above the map to visually encode a third attribute through the height and color of the bars. The resulting 2D map provides an occlusion-free overview of the aneurysm domain and supports the navigation of the corresponding 3D surface. By clicking on an interesting region in the 2D map, the virtual camera automatically moves to the region on the 3D surface, facilitating navigation in the 3D space. Regarding generalizability, the authors mentioned that their method could also be applied to other aneurysm types, such as aortic aneurysms, but they did not show exemplary results for other structures.

Later Meuschke et al. [218] used the previous aneurysm map based on LSCM parameterization again to visually explore the interplay between scalar fields defined on the aneurysm wall and scalar fields defined on path lines that represent the internal blood flow. On the aneurysm map, a scalar field is color-coded using a gray-to-red color scale. Moreover, the internal blood flow is grouped by hierarchical clustering, where the user can explore the resulting clusters. The path line points of a selected cluster are projected onto the aneurysm map and are visualized as circles by applying depth-dependent halos. For this projection, for each path line point, the nearest surface point of the aneurysm part is determined using the Euclidean distance. On the resulting contour-based rendering of the path lines, a blue-to-yellow color map is applied to encode a flow-related attribute such as flow velocity. With this, the context and interplay of the wall- and flow-related attributes can be analyzed to further improve the rupture risk assessment. To reduce the manual effort needed for generating aneurysm maps, Meuschke et al. [213] later used the spectral conformal parameterization [225], which is also an angle-preserving approach. In contrast to LSCM, spectral conformal parameterization can be applied without constraints. It tries to optimize the boundary of the mesh such




Figure 16: Simultaneous analysis of two scalar fields on an aneurysm map using a checkerboard visualization. The continuous value γ provides a smooth transition between both fields. Image from [213], © 2019 IEEE, reprinted with permission.

that the mapping leads to appropriate results. Again, the map was used to explore two scalar fields simultaneously. For this purpose, the authors introduced a checkerboard visualization to create a smooth transition between two selected scalar fields, see Figure 16. Both scalar fields are represented by two different color scales and the user has the possibility to transition between them with a slider. Moreover, by clicking on the map, the corresponding scalar field values are shown.

To enable a simultaneous exploration of more than two scalar fields, Meuschke et al. [217] introduced an aneurysm map combined with a glyph-based depiction of the corresponding scalar values, called “Skyline visualization”, see Figure 17. Since the size of conspicuous wall regions on the map should be preserved, the area-preserving As-Rigid-As-Possible (ARAP) parameterization [196, 286] was chosen. On the resulting map, a 2D regular grid is defined. For each grid cell, a 3D bar is computed, which is separated into floors along its height, where each floor represents a simulated time step. On the side surfaces of the floors, the scalar fields are color-coded. This gives the user an overview of the temporal behavior of the simulated data on the aneurysm wall. In addition, a threshold can be set for each scalar field. If the associated scalar fields within a floor fall below the thresholds, the floor is hidden. Thus, the user can interactively find regions where conspicuous correlations of scalar fields occur. Besides




Figure 17: Skyline visualization to visually analyze multiple scalar fields on surfaces simultaneously. On the aneurysm map, 3D bar charts are generated, where the color coding of the floors encodes different scalar values. Reprinted from [217], with permission from Elsevier.

aneurysms, the authors have applied their method to other examples, such as simulated inducer data and iterative smoothed surfaces of organs, to show the generalizability of their approach.

4.6.2 Stenosis Predilection Sites

Several works explored the visual analysis of vessel bifurcations, as these highly curved regions favor the initiation and development of atherosclerosis. Atherosclerotic plaque deposits can lead to vessel stenosis or stroke in case the vessel is fully blocked or deposits detach from the vessel wall. Moreover, there are different drug treatments where it is necessary to analyze if plaque deposits are receding. As the assessment of plaque distribution and thickness based on 3D imaging is challenging, much work has been done to create 2D maps of bifurcating regions, i.e., trees of depth 1. The main applications are the carotid arteries and their bifurcations, where stenoses often develop and have a high risk of causing a stroke.

Antiga et al. [8, 9] introduced an automated mesh parameterization to straighten bifurcations to a 2D parametric plane. For this purpose, first, the centerlines are computed automatically. Based on the centerlines, a reference system is defined at the bifurcation. It consists of four points, which are determined by computing the intersections between the centerline branches and the mesh tubes corresponding to the individual branches. With this, a bijective mapping can be generated to preserve the location of the parameterized surface points relative to the reference system. Using the reference points, the bifurcating vessel is then split into its three constituent vessel branches. Afterward, each branch is parameterized in the longitudinal direction by solving a Laplacian equation over the circumferential direction of the surface based on the tortuosity of the centerline. Each parameterized branch is mapped onto a rectangular parametric space, where the individual branches are patched together. Due to the bijective mapping on the parametric plane, their algorithm allows a quantitative comparison of

anatomically different surface distributions. On the resulting patches, a scalar field is color-coded to observe the results of blood flow simulations, e.g., wall regions with high *WSS* values which might be an indicator for the generation of atherosclerotic plaque. The authors applied their technique to different bifurcating vessel geometries which shows its robustness and generalizability.

Chiu et al. [55] introduced a map-based representation of carotid arteries to visualize the point-wise change of wall thickness due to plaque deposits. First, the outer vessel wall and vessel lumen were segmented in high-resolution 3D ultrasonography images for two points in time, at baseline and a second scanning session. This results in stacked contour-based representations of the lumen and vessel wall which can be transformed into 3D surface meshes. Then, the wall thickness is computed for both time points by determining the point-wise distance between the lumen and outer vessel wall. Afterward, the change in wall thickness must be calculated. For this purpose, the two 3D meshes of the vessel wall are registered to each other using the iterative closest point algorithm. Subsequently, the point-wise change of wall thickness can be computed using the Euclidean distance between the registered walls. To straighten the baseline wall mesh to the 2D domain an arc length parameterization is used that maps the contours in each slice to straight lines. The parameterization was applied to each of the three branches of the bifurcation. For this, three planes are defined to cut open the mesh for mapping. The cut planes are described by the bifurcation apex as well as the centerlines and corresponding tangent vectors of the left and right carotid branches. Since an area-preserving map is needed to assess the change in the amount of plaque within a selected region, the area-preserving method by Zhu [336] is applied in a post-processing step. On the resulting map, the change of wall thickness is color-coded, where the user can select specific regions on the map to quantitatively assess the change in the amount of plaque. However, the carotid map by Chiu et al. [55] is less suited to compare carotid arteries of different patients, since, the map shape is influenced by the shape of the corresponding 3D vessel wall mesh. To overcome this problem, the authors presented an extension of their carotid map by introducing a standardized carotid map [56, 57]. While the segmentation, wall reconstruction, and thickness computations stay the same, the straightening of the carotid branches changed. To generate a standardized map, the bifurcation apex is mapped to the origin of a Cartesian coordinate system. Moreover, the left and right branches of the carotid artery are mapped to the negative and positive *x*-axis, where their centerlines are straightened to the *y*-axis. Thus, the 2D standardized map allows quantitative group-wise comparison between distributions of the thickness changes. Choi et al. [58, 59] extend this line of work and propose an L-shaped continuous template to map the carotid bifurcation as a single 2D patch. They

present a conformal mapping approach that retains the local geometry [58] and an angle-preserving technique [59]. The standardized vessel wall thickness change maps were used in clinical studies to investigate biomarkers that show the effect of different therapies for atherosclerosis [53, 54, 56, 331].

4.6.3 *Maps of the Heart*

The heart is the center of the cardiovascular system. In the following, we summarize techniques to generate map visualization of different anatomical structures of the heart comprising maps of its wall muscle (see Section 4.6.3.1) as well as its valves and cavities (see Section 4.6.3.2) together with maps of inserted stents to treat cardiac diseases.

4.6.3.1 *Heart Wall Muscle*

Abnormal movement of the myocardium⁶ is an important indicator of many heart diseases. A crucial role falls on the left ventricle, the thickest of the heart's four chambers, as it supplies the body with oxygenated blood. With the help of magnetic resonance and CT imaging, the complex motion of the left ventricle as well as its perfusion can be analyzed to diagnose cardiac pathologies. To simplify the analysis of cardiac function parameters on the complex left ventricle geometry, several works dealt with the generation of map-based visualizations. All of them used the cardiac Bull's Eye Plot (BEP) as a basic map design. The BEP is an abstract 2D plot that represents a standardized segmentation of the myocardium into 17 regions. These regions are arranged in a circular pattern, with the apex of the heart being projected into the center of the plot and more distant regions being projected into peripheral plot segments. This corresponds to an anatomical unfolding of the cone-shaped myocardium, where only the myocardial surface is needed as input to generate the BEP. Cardiologists are familiar with this segmentation, therefore, they can immediately understand the information presented and mentally reconstruct the original geometry.

Kuehnel et al. [168] use the BEP to display perfusion data that are relevant for the diagnosis of coronary artery disease, where plaque buildup in the wall of the arteries that supply blood to the heart. The acquisition of perfusion data occurs frequently at rest and under stress, induced by drugs. Within the plot segments, a perfusion parameter is color-coded to provide an overview about myocardial perfusion which simplifies the detection of regions with poor blood supply. Oeltze et al. [232] extend the BEP-based visualization of myocardial perfusion data. For the simultaneous representation of rest and stress perfusion, the segments of the plot are divided in two, resulting in 34 segments, with 17 each color-encoding rest and stress perfusion, respectively,

⁶ The heart wall muscle.




Figure 18: Simultaneous visualization of two perfusion parameters, in rest and stress state, based on the BEP to identify ischemic areas. Image courtesy of Anja Kuß.

see Figure 18. The ring-wise subdivision of the segments ensures that adjacent segments in the plot are also adjacent on the heart wall. In this way, it is possible to clearly show the influence of rest and stress on perfusion in the respective area.

To visually analyze regional motion patterns of the myocardium Sheharyar et al. [282] applied novel glyph based techniques to the BEP. For visual encoding of the myocardial motion, different techniques are provided, i.e., warped BEP segments, pins, and warped lines, where a blue-to-red color scale is used to depict the motion strengths. This supports the detection of abnormal myocardial motion, which is an important indicator for multiple cardiac pathologies. Since the standard BEP does not provide any anatomical information and does not show the amount of scar tissue induced by cardiac diseases, Teermeer et al. [298] introduced the volumetric BEP. While the basic projection of the myocardium is the same as before, in addition, the thickness of the myocardial wall is mapped to the height of the plot segments. The resulting volume shows the amount and position of scar tissue. Furthermore, anatomical information is added to the volumetric BEP by projecting the coronary arteries and displaying two dots that mark the locations where the right ventricle connects to the left ventricle. While all BEP-based visualizations simplify the analysis of complex myocardial data, the depictions do not preserve any properties of the 3D myocardium surface. Moreover, the generalizability is limited due to the underlying standardized segmentation of the myocardium.

4.6.3.2 Heart Valves and Cavities

Several approaches have been proposed to generate maps of the components of the heart, such as the mitral valve.⁷ A possible strategy

⁷ One of the four heart valves.




Figure 19: Mitral valves with two color maps quantifying local surface properties indicative of pathologies. A 3D and map-like 2D depiction is shown. Regions of possible blood backflow in (d) are occluded by geometry in the 3D model but become visible on the map. Image from [340], © 2019 IEEE, reprinted with permission.

to provide an overview of the valve surface is to flatten its structure, which can also facilitate exploring time-dependent data on the resulting map [340]. The motivation behind this technique is to enhance the clinical assessment of heart valve defects. The surface mesh of the mitral valve is rolled out in the 2D domain to provide a simplified overview [341, 350], see Figure 19. Layering the map-like depictions allows full exploration of the time-dependent nature of these data. For further exploration, different color maps can be used to encode scalar field information, such as the probable closure zone of the heart valve. Later, Casademunt et al. [49] also employed the method to encode additional information, e.g., contact distribution, curvature, and the principal stress, as a colormap on the flattened map.

Karim et al. [147] proposed a map of the left atrium.⁸ The left atrium has multiple incoming vessel connections and can thus be seen as a branched structure of depth 1. Their goal was to simplify this structure to study patients with atrial fibrillation, which is an irregular and rapid heart rhythm. The used surface parameterization attempts to find a trade-off between an area-preserving and angle-preserving parameterization to balance distortions of the relative locations and sizes of features. It uses a fixed rectangular layout to enhance comparability. To apply the surface parameterization, an initial manual cut must be defined. The resulting map can then be used to color-code blood flow information, such as delayed enhanced magnetic resonance imaging (MRI) data.

Ma et al. [202] presented two methods to obtain an overview of cardiac structures. First, a BEP of the left ventricle is presented with the coronary sinus⁹ mapped. Second, a technique for parameterization of the left atrium was presented. For this, the orientation of the structure given by the image data is used to generate wrapping contour lines. These lines are unfolded into the 2D domain at uniform

⁸ One of the four chambers of the heart. It receives oxygenated blood from the lungs and passes the blood to the left ventricle.

⁹ The largest vein of the heart. Its function is to drain deoxygenated blood from the heart muscle into the right atrium.




Figure 20: Standardized map of the left atrium with a fixed zone layout. Image from [316] under Creative Commons license [70].

intervals, yielding a surface parameterization. A color map is used to encode data recorded with MRI on the map. Similarly, Nuñez-García et al. [230] studied the left and right atrium and provided a solution to parameterize these structures in a quasi-conformal manner. To achieve this, boundary conditions must be satisfied. Here, the edges of the structures are mapped onto a 2D disk and holes in the structures that indicate incoming or outgoing vessels are also mapped to predefined contours within the disk. To ensure quasi-conformity, two quadratic programming problems are solved that return the 3D vertex position in 2D to minimize the deviation of the angle. The resulting maps are used to encode wall thickness and delayed enhanced MRI data with a color map.

Paun et al. [242] propose a parameterization of the left and right ventricles. In the first step, the boundary surface mesh must be extracted. Then, two landmarks are manually placed on the ridge of the boundary. Finally, the ridge is mapped onto a 2D slice. By using Laplace's equation and the gradient of the Laplace field the 3D anatomy mesh is reconstructed in 2D. The result is a 2D representation with additional detailed anatomy information. Roney et al. [261] presented a parameterization method that maps the left or right atrium to a unit square. For this, as with the previous approach, the Laplace equation was solved by defining boundary conditions as well as manually placed landmarks. The parameterization can be used to map information with a color map and compare different data sets.

Williams et al, [316] propose a different unfolding technique to visualize the left atrium in an abstract plot. They unfold the atrium with all the structures to a predefined map illustrating the individual parts, see Figure 20. The mapping can be considered an extension of the BEP to the left atrium. It uses a 2D disk and fixed contours within the disk. Then, the left atrium with its incoming branches is identified and mapped onto the 2D slice. The simplified map can be used to illustrate data with color, such as local activation time or delayed enhanced MRI data. This standardized abstract representation is highly suitable for comparison tasks.

4.6.3.3 Stent Projections

The aortic valve separates the left ventricle from the trunk of the aorta, i.e., it lies between the heart and the main arterial system. Born et al. [33] introduced a map-based depiction to investigate side effects occurring after treatment of aortic valve stenosis using transcatheter aortic valve implantation (TAVI). TAVI is a minimally invasive treatment method in which the abnormal valve is displaced toward the aortic wall by an expanded stent¹⁰ containing a prosthetic aortic valve. After a TAVI procedure, paravalvular leakage may occur, which is the back-flow of blood between the stent and its surrounding structures into the left ventricle, resulting in volume overload and potential dilatation of this ventricle. Based on numerical simulations, different risk factors were identified for paravalvular leakage such as stent deformation, radial attachment forces, or calcifications. To analyze the influence of these risk factors Born et al. [33] presented a 2D visualization (stent maps) of the expanded stent. As input the stent mesh including its centerline as well as mesh-based representations of calcifications are needed. Moreover, six anatomical landmarks manually placed in CT images are used to add context information to the resulting stent maps. For generating the stent map, the stent geometry is projected to a curved cylinder around the stent leading to generalized cylindrical coordinates. Based on this cylinder representation the stent can be easily unrolled to 2D. On the stent map, a scalar field is color-coded. Moreover, anatomical landmarks are shown together with a contour rendering of the calcifications and the stent geometry. The stent map provides a comprehensive overview of possible risk factors and allows to contextualize relations between stent placement and thickness of aortic valve calcifications. Due to the generalized cylindrical coordinates, different data sets can be easily compared to identify patterns that are typical for paravalvular leakage. Furthermore, the stent map could be applied to other medical use cases such as treatment planning of other heart valve defects and cerebral aneurysm stenting.

¹⁰ A metallic or polymer mesh tube that is inserted into an artery to keep it open.

4.6.4 Arbitrary Vessel Trees

A 2D map-based representation of complex vascular trees with numerous branches supports diagnostic tasks such as the detection of narrowed, dilated, and thrombosed vessel areas as well as the detection of lung nodules.¹¹ Zhu et al. [334, 337] introduced a conformal parameterization to project trees of depth n to the 2D domain. As input, the centerline of the vessel tree is needed, which is divided into y-shaped parts consisting of one branch point and three endpoints. Between the branch point and the endpoints smooth curves are defined that are used as cut lines to unroll the 3D surface part. By stitching multiple unrolled parts together, the whole vessel tree is mapped. The resulting map provides an overview of the 3D domain where a scalar field can be color-coded on the map. Besides, vascular trees the authors discussed other applications, such as the mapping of bronchial trees. Moreover, they reused their conformal mapping to generate an area-preserving mapping [335, 336]. For this purpose, they computed a density map that encodes the area changes from 3D to 2D using conformal mapping. Based on this density map, they applied optimal mass transport to minimize the area changes. However, since in both approaches the y-shaped parts are mapped independently of each other, unexpected cuts may result.

In Chapter 5, a bijective mapping of 3D vascular trees to the 2D domain will be introduced [346]. The method automatically cuts and flattens vessel trees. No centerline or other inputs are necessary besides the vessel surface itself. First, vessel branches and endpoints are automatically detected. Subsequently, the cutting lines for projecting the mesh to the 2D domain are found automatically. Instead of a centerline, the graph topology of the cut is used to create an initial straightened layout in the 2D domain. For mapping the geometry, the area-preserving ARAP parameterization technique is used. As the boundary is not fixed, the final layout resembles an untangling of the input tree. The algorithm can be applied to any tree-like structure that can be represented by a surface mesh.

4.7 VESSEL NETWORK MAPS

The visual search for malformations in vascular networks can be understood as a combination of abstract network analysis tasks, like browsing and path following. The underlying idea of vessel network maps is to facilitate such overview tasks, for instance, feature identification and topology browsing. Similarly, they can be applied to analyze the topological context of features and navigate a complex interconnected domain. For tasks involving spatial memory and identification, 2D depictions of data have shown to be more effective than 3D

¹¹ Small growths in the lung that can be cancerous.




Figure 21: Untangling of the abdominal aorta centerline. (a) Surface rendering of the vessel tree. (b) Flattened centerline with vessel diameter mapping. (c) Bounding box representation. (d) Untangled bounding boxes. (e) Color coding of the width relative to the median width of the segment. Image from [318], © 2013 IEEE, reprinted with permission.

representations [63, 89, 154, 161, 300]. Therefore, vessel network maps transfer the spatial vascular topology into a 2D map-like visualization. Typically, node-link visualizations are used, as they are popular in visualizing networks and have been extensively studied [20, 82, 256, 311]. Vessel network maps are based on a graph data structure that describes the connectivity of the cardiovascular region. Usually, this is the vessel centerline derived from image data. Besides the topology, other attributes can be of interest, such as the vessel diameter.

4.7.1 Connectivity Maps of Large Arteries

Large vessels, like the aorta, femoral, or carotid arteries, are clearly visible in angiographic imaging modalities. They can be seen as a treelike system of tubes that quickly becomes complex if multiple branches are considered, resulting in intricate 3D structures. Won et al. [317–319] propose an uncluttered single-image visualization of the abdominal

aortic tree. They optimize the 3D geometry of a binary tree of depth n in a 2D layout to resolve overlaps. Their original method [317] built on Kanitsar et al. [144]. They use the stretched CPR mapping to initialize a 2D layout of the centerline, which may have overlaps of different branches, but already prevents overlaps of individual segments. Then, they relax the spatial configuration of segments but keep the branches connected. They prioritize large branches, such that they inhabit predictable positions and place small branches last. A radial sweep line algorithm is used to iteratively find the available region for the placement of each segment. Later adoptions optimized this procedure [318, 319]. The relaxation is performed using bounding boxes around each vessel segment, which are rotated to find an optimal configuration (see Figure 21). Each box is denoted by $v \in V$, where V is the index set of the vessel parts. The orientation of a bounding box is given by an angle Θ_v . An optimal configuration of bounding boxes minimizes their overlap given by $\Omega(\cdot)$ and, simultaneously, minimizes the deviation of the layout from the initial configuration, described by the energy $\Delta(\cdot)$. The idea is to preserve the relative location of segments and the natural shape of the vessel tree, i.e., an untangled layout is created. As the optimization criteria contradict each other, a score function is optimized, which is a linear combination of both:

$$\mathcal{E}(\Theta_1, \dots, \Theta_{|V|}) = \Omega(\Theta_1, \dots, \Theta_{|V|}) + \lambda \cdot \Delta(\Theta_1, \dots, \Theta_{|V|}). \quad (3)$$

To optimize this function, Won et al. [319] first used simulated annealing [153] and later proposed another approach that uses a technique inspired by the protein side-chain placement problem [184, 318]. Jeon et al. [141] also extended the procedure to parallel computing architecture. The resulting visualization uses a size and colormap encoding to highlight deviations in vessel width. The overview is thus helpful for identifying possible stenoses and aneurysms. It could, in theory, also be used to compare or navigate domains. As the initial layout is created by stretched CPR and only rotations of vessel segments occur, the relative arc length of the vessels is preserved but the technique is also view-dependent.

Borkin et al. [30] propose a diagram representation of coronary artery trees of depth n . They found that users are more efficient and accurate at identifying atherosclerotic lesions (possible stenoses), compared to using a 3D surface depiction. The layout they propose disconnects branches and straightens them parallel to each other, preserving segment length. They map the vessel circumference to the size of the visualized segments. As their centerlines are based on surface models they are also able to add a projection of surface properties into this area, for example, simulated WSS. The resulting field is encoded with a colormap (see Figure 22). They show how the map can be used for overview tasks and navigation by combining a view of a focus domain with a view of the full domain, where




Figure 22: Coronary artery tree visualized as a surface rendering (bottom) and as a map of segments (top). Line width encodes the vessel diameter, the colormap encodes a scalar field (shear stress). Image from [30], © 2011 IEEE, reprinted with permission.

the focus region is marked. While they limit the demonstration to coronary artery trees, other vessel trees could plausibly be visualized with the same approach.

Targeting a general solution, Marino et al. [205] propose a planarization of treelike 3D structures. They demonstrate the applicability of the approach to arterial trees of depth n . While they use a surface geometry as input, the creation of the map-like depiction effectively only requires centerline and width information, which is why we classify the output as a vessel network map. Depending on the 3D view direction, the centerline is projected onto the 2D view plane, which typically leads to overlaps. Separate from this first step, a radial planar layout of the centerline is constructed based on the embedding of phylogenetic trees [17]. Then, the original shape of the tree is recovered by iteratively converging to the angles that are measured in the projection from step one. Again, the final layout is an untangling of the input graph, aiming to achieve visual coherence with the 3D depiction. In the proposed visualization, they map two attributes to the layout:

vessel width and segment depth. The width is directly mapped to the width around the centerline geometry and depth cues can be added in form of a grayscale border. The authors also demonstrate how to map a surface into the resulting shape, however, all information visible in the final representation (width, z-axis depth) can be encoded on the centerline alone. Mapping the surface may be useful for future applications, although it should be noted that the surface area is highly distorted in the process, as the vessel width encodes the diameter, not the circumference.

Vessel network maps can also be used as part of multiple coordinated views. A possibility is to use the map to first locate points of interest and then navigate to these locations, e.g., by steering a camera in a 3D rendering linked to the map. In [Chapter 6](#), a framework for analyzing hemodynamics in carotid stenoses is proposed [342]. Therein, a network map of the vessel tree is integrated to find and navigate regions of interest. The input trees of the left and right carotids are straightened and the two sides are aligned such that their principal branches are in parallel. Side branches are connected at a fixed angle. Thus, the layout trades preservation of the original angles for a uniform and symmetric view of the vessel. Similar to other approaches that focus on stenosis analysis, the vessel diameter is mapped to the line width of the rendering. In the final application, the map is then used for overview purposes, comparing the left and right trees, and for quantifying stenosis length, which can be directly read on the map. Also, the spatial domain can be navigated with draggable markers on the branches that control probe positions in the 3D data.

Some applications require analysis of a highly branched structure, like the arterial tree of the liver investigated by Lichtenberg et al. [188]. One objective of this approach is to show the proximity of vessels and branch structures for intervention planning to minimize the risk of cutting through vital arteries. The technique visually separates individual subtrees in a straightened compact binary tree layout. A hierarchical or radial layout can be used. In both, overlaps are prevented and the subtrees are delineated. Then, they apply color coding on the subtrees for branch identification and compute a screen space parameterization for pattern mapping [190]. This parameterization enables an illustrative encoding as an overlay of the map. It shows the proximity of branches to selected structures, for example, vessels close to a tumor or the distance to a simulated surgical needle, which can be utilized for operation path planning. An application is shown in [Figure 23](#). Later, Lichtenberg and Lawonn [189] applied a similar straightening technique and radial embedding of subtrees to generate 2D overview plots. They use an HSV colormap defined in the 2D domain to identify subtrees in a 3D model. Both methods are based on topology analysis of a mesh centerline. The visualizations reduce complex structures to




Figure 23: A 3D view (top) and 2D map (bottom) of the arterial tree of a liver. Three possibilities to reach a large tumor structure with a surgical needle are shown. A distance field around the needle is calculated and color-coded on the map. Vessels with a minimal distance of less than 20 mm are at risk of being damaged. They are highlighted with a stippling pattern, which allows for determining an optimal path. Image from [188], © 2019 The Eurographics Association.

abstract map-like depictions. Original geometric features are lost in exchange for a visual focus on the branch structure.

4.7.2 Cerebral Connectivity Maps

The cerebral arterial system can be modeled as a graph with weighted edges. These are mostly directed edges, as blood flow is distributed into increasingly smaller arteries, ending in capillaries.¹² Compared to other areas of the cardiovascular system, the cerebral arterial network is a special case for three reasons. First, small disturbances in blood distribution can cause stroke and be life-threatening. Second, there are also regulating cyclical arteries, which are *undirected* graph edges. In our graph model, these hierarchical and non-hierarchical components form a mixed hierarchy. Third, the cerebral network topology is highly relevant to determining which parts of the brain are affected by a vessel pathology and how and if a malformation can be treated. The mechanical accessibility of a pathology is a key factor in clinical risk assessment. Neuroradiologists need to identify stenoses, aneurysms, or missing arteries and then plan a potential intervention, where the ideal path to reach the respective site with surgical instruments needs to be determined, starting from the trunk. The cerebral arterial network can be thought of as consisting of three components. There are four entry points where the two carotid and the two vertebral arteries bring


¹² Minuscule vessels through which oxygen and nutrients are exchanged with the surrounding tissue.

new blood into the system. They are connected in the circle of Willis, which regulates flow between them and acts as a layer of redundancy – a blockage in one artery can often be circumvented. Many anatomical variations of the circle of Willis exist, which necessitates individual assessment. Lastly, the cerebral arteries connect to the circle of Willis and distribute blood further throughout the brain.


Miao et al. [219] propose a standardized description of arteries in the circle of Willis in the *CoWRadar*. The *CoWRadar* is an abstract map-like depiction that allows visual quantification of segment width and connectivity information. They implemented an automatic pipeline to detect the circle of Willis in image data, create a centerline, derive the vessel graph, and label arteries. The network is fitted to a predetermined layout. The length of segments and vessel intensity in the image data are shown by arc glyphs at fixed locations. The connectivity is visualized with active/inactive line connections on the map. They show how the *CoWRadar* can be used to quickly identify the presence of vessels, understand the patient-specific network configuration, and quantify the blood perfusion of segments. An resulting depiction is shown in Figure 24.

Pandey et al. [239] developed *CerebroVis*, an abstract 2D network visualization of the cerebral arteries that aims to preserve spatial context. They report an increased accuracy in stenosis detection tasks in comparison to conventional 3D visualizations of the same network. The readability of network visualizations is dependent on multiple factors. Node-link depictions should minimize edge-crossings [256], path continuity facilitates visual browsing tasks [155, 311], and symmetry helps comparison [192]. The authors apply these principles to *CerebroVis*, resulting in an abstract fitting of the topology with spatial constraints, e.g., the left and right sides of the brain are separated, see Figure 25. Furthermore, they run undirected edges horizontally in the final layout, to differentiate them visually from the hierarchical components. This idea was originally proposed by Gutwenger et al. [119]. Width and color encoding are used to display vessel diameter and an arbitrary additional property, like blood diffusion. *CerebroVis* can be used to detect points of interest, place them in their topological context, and navigate a 3D rendering of the network by clicking on segments.

A further technique that should be mentioned is the animated un-tangling of the circle of Willis by Saalfeld et al. [268]. While technically not a 2D map visualization, the geometry is unfolded and acts as a map for annotations. The centerline approximates a planar layout and the depiction can be used from a single perspective, yet the surface geometry remains 3D. The approach is a proof-of-concept, as the unfolding and animation are created manually, which is unrealistic for time-sensitive clinical applications.



(a) Vessel graph with anatomical labels



(b) CoWRadar with labels

Figure 24: The CoWRadar, a standardized connectivity map for the circle of Willis. (a) Maximum intensity projection of the circle of Willis with a centerline and extracted labels. (b) The resulting CoWRadar. The arc glyphs encode the length of segments (outer arc) and the perfusion in the image data (inner arc). Image from [219], © 2016 The Eurographics Association and John Wiley & Sons Ltd.

4.7.3 Maps of Microvasculature Networks

The analysis of microvessels is crucial for research in tissue development, pathophysiology, and treatment [18, 92, 246]. Tissue degeneration, for example, initiated by diseases, results in subtle differences




Figure 25: Projection of a cerebral artery network (left) and an automatically generated network map (right). Line thickness encodes the vessel diameter and color encodes the vessel category. Image from [239], © 2020 IEEE, reprinted with permission.




Figure 26: Microvasculature map (top) and 3D vessel structure (bottom). Multiple steps are presented, in which a microvessel cluster is selected and explored. Image from [112], © 2019 IEEE, reprinted with permission.

between networks. Microvasculature consists of complex interconnected vessels across large volumes, which is difficult to visualize [111]. Tracing the topology of microvasculature results in a highly cyclical but sparse graph with a large number of redundant edges. Govyadinov et al. [112] propose a bi-modal visualization with a graph-based and a geometry-based view of the same microvasculature data. They extract the centerline of microscopy tissue slices with a GPU-based predictor-corrector algorithm [111], which results in a graph of the capillary branches. On this graph, they compute a spectra clustering over a user-selected property, like the microvessel length. Then, they show a highly intricate 3D representation of the segmented network, accompanied by a 2D map of clusters and connections between them. They use color to encode the cluster correspondence and arc glyphs to encode the inter- and intra-cluster relationships. The graph view can be used to navigate the volumetric view by brushing and linking. For example, selected regions in a cluster are highlighted in the 3D model.

An example of the application is shown in [Figure 26](#). They show how to use the representation to compare microvasculature networks and detect small changes caused by diseases like Alzheimer's.

4.8 RECOMMENDATIONS FOR VESSEL MAP CREATION

With the diverse range of existing vessel map techniques, it can be challenging to choose an appropriate method. We propose the following recipe to aid in this process. The outlined steps should be considered as overarching design considerations. They are not guidelines per se, as they are not founded on a well-enough base of perception studies on vessel maps alone. Large-scale quantitative evaluations of vessel maps are practically non-existent. Meaningful evaluation tasks that test the effectiveness of a vessel map require domain experts, often highly specialized physicians or medical researchers. Nonetheless, vessel maps have proven successful in terms of adaptation. They are a popular tool for abstracting information about the cardiovascular system, which is unsurprising considering the large body of map-like depictions used to illustrate cardiovascular anatomy in medical education. We base these recommendations on the domain requirements we distilled for vessel maps ([Section 4.1](#)) and recurring patterns found in the discussed literature. Also, the lack of quantitative studies can to some degree be substituted with observations from perception studies that involve similar map-like depictions from other contexts, like geospatial data [[154](#), [161](#)] or node-link visualizations [[82](#)].

4.8.1 Identifying Domain Tasks

The first and crucial step is to consider what the map is for. Which purpose does it serve? Along the lines of typical visualization design studies [[237](#), [277](#)], we recommend first clearly identifying and abstracting the user tasks that need to be facilitated. Typical motives, which are pursued by all the discussed vessel map approaches, are *overview* tasks such as identification of features, exploring distributions, or browsing the topology. The usual incentive is to provide a **2D** or abstracted view that is better suited for these tasks than overly detailed **3D** renderings of the data. Fitting a spatial domain in one picture may increase legibility, leading to higher efficiency and/or accuracy in these tasks. We do, however, encourage thinking beyond these aspects. Other tasks that could be facilitated may be *contextualization* tasks. Where exactly does a feature occur? What are its spatial relations to the cardiovascular structure? What is the surrounding anatomical context? What is the topological context in the vascular network? Also, some maps allow direct *quantification*. Can data values be quantified directly on the map? For example, can sizes be measured? The map coordinates may serve this purpose if the geometry mapping preserves the relevant

properties. There may also be *navigation* tasks. Could the map be used to navigate another domain? Does it need to represent all necessary information or could it serve as a secondary depiction, along the lines of focus and context? Examples, where this applies, are typically systems with multiple coordinated views. Keep in mind that maps can help navigation in spatial [342] and temporal [340] dimensions. Lastly, aiding *comparison* is also a common purpose for map-like depictions. May the map be a useful asset to compare multiple instances of a domain? These could be an individual before and after treatment or different patients up to cohorts examined in medical research.

4.8.2 *Choosing Data Type and Region Complexity*

Depending on the domain tasks, the appropriate data type and region of interest must be determined. Even for general mapping approaches, it is advisable to think about which specific tasks – they can be abstract – may be facilitated. The common data structures for vessel maps are described in [Section 4.2](#). Define, where specifically the information relevant for the tasks associated with the map is. For example, the data structure could be a vessel tree surface, but if our main interest for the map is to give a topology overview, it may be much more sensible to create a map based on data encoded on the centerline. Similar considerations should be made regarding the region's complexity. Knowing if a single segment, a tree, or a network is required to accommodate the tasks, enables choosing an appropriate map layout.

4.8.3 *Preservation versus Standardization*

The question of whether to preserve attributes or provide a more consistent, maybe even standardized map, is a core consideration that should be made before any algorithm is implemented. It is closely related to the question of how vigorously the present information should be simplified for the sake of uniformity and legibility. Preservation can refer to geometric proportions. For mesh parameterization, in particular, techniques are generally grouped by their preservation characteristics. A parameterization method may be area-preserving or angle-preserving (conformal) regarding the mapped mesh structure. But in vascular geometry, more general features may also be preserved, like the arc length of individual segments, the diameter of vessels, or the branching angles. Preservation can also refer to the topology of a vessel tree or network. Preserving the branching topology is required for most overview, navigation, and contextualization tasks.

The primary consideration for preserving attributes is that information crucial to a task is retained. For example, a task might be to assess the distribution of a certain attribute, like plaque occurrence, where the size is important to determine the severity. Sometimes, though,

preserving the actual data values might not be the ideal approach to facilitate a task. For instance, the variation in the vessel diameter may be important but too small to be visually discernible [342]. In very specific instances, it might make sense to explicitly exaggerate values to draw attention to points of interest. This possibility should be handled with care and only applied if the quantification of said values is done outside of the map. Another advantage of attribute preservation lies in the increased familiarity that the resulting depictions might have. Users may be familiar with the 3D shape of a particular vessel tree and quickly understand and correspond to a map that preserves, e.g., the overall layout.

The core benefit of uniform maps, on the other hand, is that they are more directly comparable. This applies to visual exploration but also algorithmic evaluation. The advantage particularly holds for approaches that attempt to define a standardized map, where map coordinates are fixed to specific cardiovascular substructures. The goal of standardization in vessel maps is often to develop quantifiable biomarkers [49, 54], where threshold values can support clinical decisions.

4.8.4 *Layout, Geometry, and Attribute Mapping*

Only after these considerations should the actual vessel map be designed. We propose to think of the process in terms of the taxonomy outlined in Section 4.3 and view the layout, attribute mapping, and geometry mapping as separate steps. If the region complexity is very low, e.g., for a surface patch or vessel segment without branches, the layout step can usually be solved implicitly. If the underlying data is only of the centerline graph type, the layout and geometry mapping are the same step. Mind that, depending on the input graph, the layout algorithm needs to build on either a hierarchical or mixed hierarchy structure. In most cases, general readability principles from network visualization apply, e.g., regarding edge-crossings [256], path continuity [155], and symmetry [192]. As a rule-of-thumb, untangling techniques lie on the preservation end of the abstraction-level scale, while fitting techniques are more applicable if standardization is required. Often intertwined with the layout step, the geometry of the data needs to be mapped. Here, similar rules apply: choose or develop a technique that fits the necessary preservation characteristics and shows the data required to solve the domain tasks. We also recommend paying attention to the algorithmic dependencies. Requiring landmarks or manual cuts may impose additional pre-processing steps, making the technique less viable for broader applications. Depending on the 3D view direction may defeat the advantage that the map gives a single-image overview. Still, this dependency may be useful in some scenarios. For example, a vessel tree may have a “default” viewing

direction, like facing the patient, which may be used to construct the map. Lastly, the data attributes need to be visualized on the map. This should be regarded as an individual step in the map design process, as it is distinct from the geometry mapping. An encoding shared with different (3D) representations of the data may increase familiarity and visual coherence. However, relying only on familiar encodings might mean that alternatives are overlooked. Common trends are to visualize vessel diameter with a scale/thickness encoding and to visualize scalar fields with a colormap. Categorical colors are sometimes used, for instance, to distinguish subtrees. The use of glyphs is recommended to encode multiple values or more complex information. Furthermore, we encourage thinking about possible combinations of attribute mappings for simultaneous use of multiple visual information channels.

4.9 OPEN CHALLENGES

The high number of recent works in the flow, wall, and network maps areas evinces that concurrent investigations into vessel maps are ongoing. While many techniques for mapping vascular structures, ranging in complexity from individual segments to highly branched microvasculature networks, have been investigated, some core challenges remain.

Flow volume maps. Map-like visualizations of blood flow, in particular, are underexplored. The core challenge lies in finding a suitable geometry mapping not only for a 2D embedded structure like the vessel wall but also for the internal 3D flow field. To foster an understanding of qualitative flow information in branched structures, abstractions that transfer flow features into a 2D graph layout appear to be a promising direction [296, 329].

Network maps. While several mapping techniques for treelike vascular structures have been proposed, only a few approaches exist that handle cyclical graphs and complex microvasculature [112, 239]. Current advances in 3D imaging using, e.g., micro-CT or light-sheet microscopy result in increasingly complex vascular data. All arteries eventually terminate in interconnected networks and any increase in imaging resolution poses a need for network-based algorithms to explore these data.

Temporally resolved data. A further insufficiently examined application area is temporally resolved data spaces. The cardiovascular system is, after all, a spatio-temporal domain that is highly dependent on the cardiac cycle. The possibility of using vessel maps for, e.g., temporal navigation tasks, has only been adopted in very few works [340, 217, 329]. Not only blood flow information is temporally resolved, but parts of the cardiovascular system, namely the heart and larger arteries, are also constantly in pulsating motion. How to reproduce

the motion on a map or in which cases this would be beneficial, has not been sufficiently investigated.

Data curation. Furthermore, we would like to draw attention to the general challenge of data curation [103]. In many areas of medical visualization, the applied techniques require data in very specific formats and of sufficient quality to produce adequate results. Vessel maps are no exception. To increase the robustness and comparability of techniques, the availability of standardized data sets must be encouraged.

Clinical adoption. Finally, clinical adoption requires either evidence-based quantifiable biomarkers or strong validations through effectiveness evaluations. Both need objective benchmarks to measure the potential benefits of a new method, which are currently rare.

4.10 CONCLUSION

This chapter presented a literature overview of map-like visualizations of the cardiovascular system. To give a structure to the diverse approaches for the creation of such vessel maps, it provides a classification centered around the four principal data types used to generate them: medical volume images, blood flow volume meshes, vessel surface meshes, and centerline graphs. Further, not only are the various algorithmic techniques to transfer these data types to a map layout discussed, but also the visualizations examined in terms of their applicability. For this purpose multiple aspects were filtered: their visual design properties (layout types, encodings), the domain tasks they facilitate, and the dependencies that must be met to use them. Based on the literature review, considerations were reviewed that are repeatedly made in the visual analysis of cardiovascular structures and condensed to recommendations for the design of vessel maps.

Within the general domain of map-like visualizations, vessel maps are schematizations of cardiovascular structures. They emphasize the display of data over anatomical accuracy. Hognräfer et al. [130] identified various techniques for performing schematizations. The majority of vessel maps naturally use either path-centered techniques, if the network structure is focused, or shape deforming techniques, if the aim is to display surface or volume data. Other types of techniques derived from the schematization of cartographic maps might be worth exploring, but many are not applicable. After all, vessel maps represent anatomical structures, which means data is typically not organized around, for example, borders. On a high level, vessel maps do, however, share their advantages and limitations with other map-like visualizations. Their benefits include that they are easier to decode than convoluted 3D renderings and features are filtered from irrelevant anatomical context information. Relevant parts of the data are emphasized, fostering an understanding of important relations.

A pitfall to avoid is that misinterpretations become possible due to geometric deformations and data abstractions. Therefore, the purpose of the map must be clear. Often, a certain task is well facilitated by the map but it also needs to be complemented by other techniques to provide an actual benefit. Despite persistent challenges, vessel maps prove to be a valuable component in the toolbox of cardiovascular structure analysis. Through domain simplification and standardization, they enable otherwise impractical or impossible visual and algorithmic assessments and are a significant asset for vascular disease monitoring, treatment planning, and research.

Part III

VISUAL EXPLORATION OF HEMODYNAMIC
SIMULATIONS IN STENOSES

5

Automatic Vessel Wall Maps

This chapter is partly based on:

Pepe Eulzer, Kevin Richter, Monique Meuschke, Anna Hundertmark, and Kai Lawonn. "Automatic Cutting and Flattening of Carotid Artery Geometries." In: *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, 2021.

As described in [Section 3.2](#), recent advances in simulated hemodynamics allow for the modeling of blood flow properties based on [CFD](#) [[44](#), [293](#)]. These simulations are performed on patient-specific arterial wall geometry. Typically, one flow inlet is defined, for instance, at a cross-section of the [CCA](#). Here, inflow is simulated following the pattern of a standardized human heart cycle. The flow exits the domain at one or more predefined outlets, which lie downstream from the inlet. Together with the wall, the inlet and outlets constitute boundary conditions that allow solving the internal flow field. This yields various parameters at arbitrarily high resolutions, some of which cannot be measured in-vivo. These simulation results can help understand the progression of cardiovascular disease and may even provide additional markers for medical treatment decisions in the future. A principal interest is the analysis of wall-related parameters, such as [NWS](#) or [WSS](#), as they show how the vessel wall is affected by the flow field.

The analysis of [CFD](#) data is a cumbersome process. Often, a large spatio-temporal data space needs to be examined, where comparisons need to be made regarding values over time and different parameterizations of the simulation. To provide better insights into this complex data, a variety of visualization techniques and frameworks have been proposed for medical flow visualization [[234](#)]. An overview of these works is provided in [Section 3.3](#). Following this line of research, this chapter introduces overview depictions for at-a-glance assessments of wall-related [CFD](#) results. Existing model-based techniques for the flattening of vessel surfaces either create patches and do not show a continuous domain [[9](#), [30](#), [110](#)], or do not preserve the surface area [[205](#)]. For example, Borkin et al. [[30](#)] proposed [2D](#) projections of cylindrical vessel segments. They map coronary artery trees to a diagram representation for overview assessment of simulated [WSS](#). However, the geometry is not preserved as they use a projection and vessel bifurcations cannot be shown. Marino et al. [[205](#)] introduced a mapping for arbitrary branch configurations, which shows the whole vessel and achieves bijectivity. They fit the surface into the fixed [2D](#) boundary shape with an adapted harmonic mapping. For the purposes outlined in the following, however, the method is not suited as the area is not preserved. Furthermore, the cuts used in many previous works optimize geodesic distance, which results in unintuitive cutting paths, as we show in [Section 5.2](#). Area-preserving maps of the carotids have also been proposed. Notable are, for instance, the works of Antiga et al. [[9](#)], Zhu et al. [[336](#)], and Choi et al. [[59](#)]. The

fundamental problem of these approaches is that they are specific to a limited segment of the carotid, whose geometry is well-defined by three connected cylinders: the CCA, ICA, and ECA. The extracted section around the carotid bifurcation is about 20–25 mm in length, which allows these techniques to assume the branches as straight and cut them with predefined planes along known landmarks. On average, the carotid covers more than 200 mm in length in adults [60]. If the focus region is too small, stenoses that occur anywhere but directly at the carotid bifurcation will be hidden. Furthermore, for accurate CFD simulations, larger segments with potentially more than two outlets are of interest.

The flattening approach presented in this chapter is a global surface parameterization that creates a single patch. The parameterization preserves the vertex and connectivity information, except for the cut, making it simple to map any data fields defined on vertices or faces of the mesh. For a realistic assessment of parameter distributions, we minimize area distortion while keeping the layout close to the original structure. The approach is also fully automatic and requires no manual cuts or predefined landmarks. In contrast to other methods, the single input is the 3D mesh of the vessel tree, we do not require a centerline. Further, the flattening results in intuitive 2D views of the unrolled inner vessel wall. We achieve this by introducing natural vessel cuts, which remain on one side of the cylindrical vessel shape. The proposed method supports vessel trees of various sizes with an arbitrary number of outlets, contrary to existing approaches to map the walls of the human carotid that only consider a limited region of the vessel with a specific geometry [55, 59].

Furthermore, we investigate a use case for the flattening. We integrated the method into a CFD pipeline that simulates blood flow in the carotids of real patients. We provide a graphical user interface (GUI) to generate flattened depictions of the extracted wall geometries, on which scalar fields can be mapped. The tool supports navigation of the 3D and 2D spaces as well as the temporal dimension through interpolated animation of time steps. The demonstration focuses on the carotid as an example of a highly relevant application area. However, we also show how our method can be generalized to any vessel tree.

5.1 OBJECTIVES

The objective is to create a flattening through surface parameterization. We assume the vessel tree is a triangulated surface with vertices \mathcal{V} and faces \mathcal{F} , where $(i, j, k) \in \mathcal{F}$ with $i, j, k \in \mathcal{V}$ means that the vertices i, j, k form a triangle. The 3D position of a vertex i is denoted by $\mathbf{v}_i \in \mathbb{R}^3$. We want to find a configuration of (u, v) parameter coordinates where each vertex $\mathbf{v}_i = (x_i, y_i, z_i)$ corresponds to exactly one point (u_i, v_i) . For an unambiguous representation of scalar field data on the flattened

surface, we aim for a bijective mapping between the 3D and parameter space. No overlaps or triangle flips should occur.

The goal is to facilitate use cases in CFD and medical research, where surface attributes on the vessel need to be analyzed. This analysis involves localization of certain attribute distributions to understand, e.g., where irregularities occur. Naturally, not only the location but also the size of these feature distributions need to be accurately portrayed for visual analysis. Therefore, the parameterization needs to preserve surface area if possible. Simultaneously, the surface layout should be intuitive, i.e., it should be easy to recognize anatomical features and mentally project them into the 3D space. All parts of the input surface should be visible, even if in specific cases some side branches might be less medically relevant. Mapping the total surface per default makes the technique more generally applicable. Also, for CFD research, it is important to examine and compare the behavior of parameters on the whole surface that defines the simulation domain.

We expect the input mesh to be triangulated and assume open boundaries at inlets and outlets, which is the typical morphology of vessel walls in CFD blood flow simulations. The rest of the mesh should be watertight and not contain self-intersections. The vertex indices of each triangle must be ordered consistently with normals pointing outwards. As we focus on vasculature, we assume a tree-like surface structure where sub-branches on average get thinner as compared to their root branch. Result data can be given as scalar/vector fields discretized on \mathcal{V} or \mathcal{F} . It should be kept in mind that vessels do not consist of perfect cylinders and small deformations are the norm. Also, stenoses are possible that can disturb the cylindrical shape.

5.2 CUTTING

To find a valid surface parameterization that can be rendered in a flat layout, the input mesh must be cut such that it becomes homeomorphic to a disk. We want to find *natural vessel cuts* that are necessary for an intuitive flattened layout. We define natural vessel cuts as follows:

1. The first cut starts at the inlet boundary of the vessel tree.
2. At each branch of the vessel tree, the cut splits in two child cuts, one for each branch.
3. At each outlet boundary exactly one cut terminates.
4. Cuts may only intersect at the point where they connect.
5. Wherever possible, cuts travel in the longitudinal direction of cylindrical vessel segments.

This results in a minimal number of connected cuts that slice open each vessel segment, creating a single surface patch. The definition is

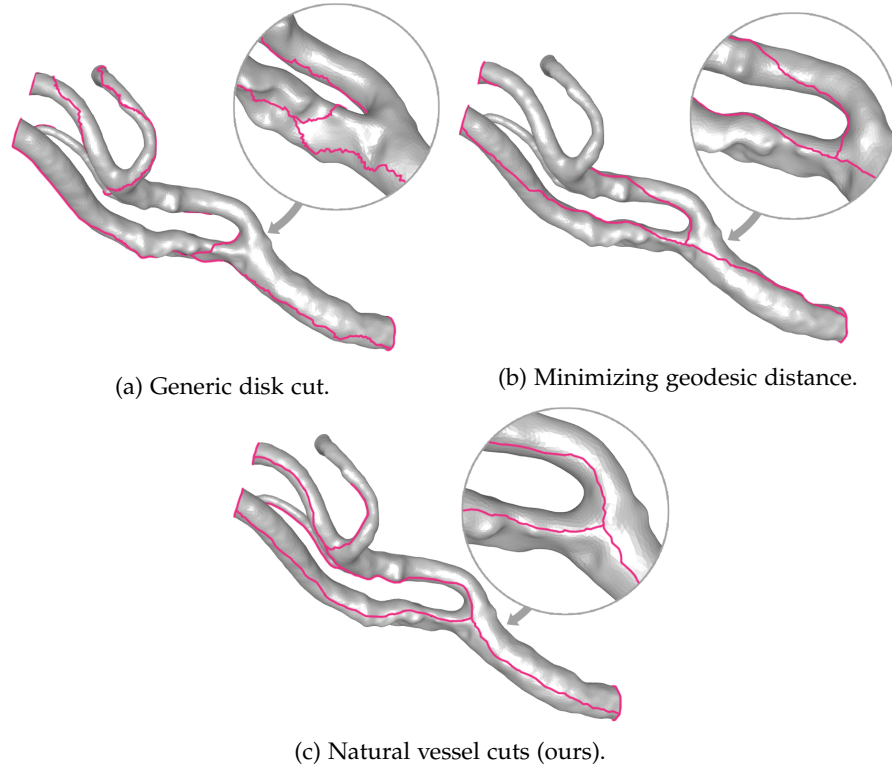


Figure 27: Different cutting approaches that all result in a mesh homeomorphic to a disk.

similar to the boundary creation as described by Marino et al. [205]. However, we additionally impose the last point as a crucial condition. Cuts should remain on one side of the cylinders and not wrap around them. By taking this requirement into account, we can generate cuts that enable intuitive mental mapping between 3D and 2D, as the surface will “unfold” in a consistent orientation. The advantage of these cuts becomes apparent when compared to the traditional approaches of either using a generic cutting algorithm [115] or minimizing geodesic distance (see Figure 27). We create natural vessel cuts by following the minimal principal curvature, which, on a cylindrical object, defines a tangent field along its length.

5.2.1 *Sorting boundaries*

To create the cuts we first find the set of existing boundary vertices $\mathcal{V}_B \subset \mathcal{V}$. We select an arbitrary vertex in \mathcal{V}_B and recursively search for one neighboring vertex that is also in \mathcal{V}_B until we found a loop. Then, the process is repeated with an unvisited boundary vertex until all boundary vertices are visited. This gives us one ordered boundary loop for the inlet and one for each outlet. To prepare for the flattening where these loops are effectively “unrolled”, we check that the loop orientation is consistent with the winding order of the mesh faces.




Figure 28: The minimal principal curvature is aligned with the vessel direction, serving as the basis to orient the natural vessel cuts.

This means a vertex candidate is only appended to a loop if it and the previous vertex would occur in the same order as they do in the face they share. They can only share one face as they are both on the boundary. This condition only needs to be tested for the second vertex of a new loop, as two neighbors to the first vertex are possible candidates. Ultimately, we sort the list of loops by the arc length of the loops. We expect the root of the vessel tree to lie at the thickest branch. Hence, the longest boundary loop is considered to be the inlet, all others are considered to be outlets. The advantage of this procedure is that no manual landmarks need to be set.

5.2.2 *Aligning curvature*

We compute the principal curvature directions at each vertex and extract the minimum principal curvature $\mathcal{C}_{\min} \subset \mathbb{R}^3$, which is a good geometric estimator of the longitudinal direction along a cylinder. The idea was also employed by Lichtenberg and Lawonn [189] to compute a continuously increasing scalar field across a vessel tree. We use the method by Panozzo et al. [241] to compute the principal curvature cross-field. The principal curvature is numerically solved on a best-fit quadric over the neighborhood of each vertex. This method is fast and robust. It allows the definition of a scale for the quadric based on the average edge length of the mesh edges e . As in the original proposal, we set the scale to a default of $5e$. Then, the quadric fitting will take into account a large portion of neighboring vertices, which effectively smoothens noise and parts of the surface that deviate from the ideal cylinder. As such, the curvature field will be aligned with the overall shape of the vessel tree. It is robust against irregularities in the vessel structure, small aneurysms, and deformations from stenoses.

A problem when using the principal curvature is that the sign of individual elements $\mathbf{c}_i \in \mathcal{C}_{\min}$ is ambiguous. At each vertex \mathbf{c}_i or $-\mathbf{c}_i$ can be computed. We need all directions in \mathcal{C}_{\min} to be aligned towards the root of the tree geometry. We resolve this ambiguity as proposed by Lichtenberg and Lawonn [189]. First, the geodesic distance $\mathcal{G} \subset \mathbb{R}$ is computed as a scalar field over \mathcal{V} , where each vertex is assigned the distance to the inlet boundary. We use the geodesics in heat method [69] as it is fast and sufficient for the purpose. As \mathcal{G} increases strictly monotonously over the vessel tree, we can use its gradient at vertex positions $\mathbf{g}_i \in \nabla \mathcal{G}$ to obtain the aligned curvature \mathcal{Z} :

$$\mathbf{z}_i = \begin{cases} \mathbf{c}_i, & \text{if } \langle \mathbf{c}_i, \mathbf{g}_i \rangle < 0 \\ -\mathbf{c}_i, & \text{otherwise.} \end{cases} \quad (4)$$

The resulting field of aligned minimal curvature is shown in [Figure 28](#).

5.2.3 Finding optimal paths

The mesh is now sliced open with natural vessel cuts that when traced from outlets to the inlet minimize deviation from the directions given in \mathcal{Z} . We will use cuts that run along existing mesh edges, i.e., no new topology will be introduced in the process. This makes it simple to map any data fields between the [3D](#) and [2D](#) representations, as a complete set of (u, v) coordinates will be defined.

Picking any two vertices on the mesh, an optimal path can be found between them by interpreting the mesh as a directed graph. In this graph, a unique node exists for each element in \mathcal{V} and each mesh edge is represented by two directed graph edges, one in each direction. By associating each directed edge with a travel cost, we can find the path between two selected nodes that minimizes the total cost. For simplicity, we use Dijkstra's algorithm [78] to find the path. Due to its logarithmic complexity, it also scales with large meshes. The cost function can be used to control the path. From a mesh vertex \mathbf{v}_i edges exist to any direct neighbor \mathbf{v}_j . When setting the cost to the Euclidean distance from \mathbf{v}_i to \mathbf{v}_j , the path – along existing edges – with minimal geodesic length is found:

$$E_g(\mathbf{v}_i, \mathbf{v}_j) = \|\mathbf{v}_j - \mathbf{v}_i\|. \quad (5)$$

In this case, the graph could be undirected. To minimize the deviation of the path from the aligned curvature \mathbf{z}_i at \mathbf{v}_i , however, we need directed edges with costs:

$$E_c(\mathbf{v}_i, \mathbf{v}_j) = 1 - \left\langle \frac{\mathbf{v}_j - \mathbf{v}_i}{\|\mathbf{v}_j - \mathbf{v}_i\|}, \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|} \right\rangle. \quad (6)$$




Figure 29: Robust cutting of a carotid bifurcation. (a) Schema of the path finding for cutting. Paths start at the outlets and connect to the inlet or the closest existing cut. (b) Resulting cut geometry applied to the original mesh edges.

Due to the normalization, the edge length is ignored. If edge lengths deviate strongly from their mean, it is possible to factor in higher costs for longer edges by interpolating the geodesic and curvature costs:

$$E_{cg}(\mathbf{v}_i, \mathbf{v}_j) = a \cdot E_c(\mathbf{v}_i, \mathbf{v}_j) + (1 - a) \cdot E_g(\mathbf{v}_i, \mathbf{v}_j). \quad (7)$$

This is usually not a problem for meshes used in CFD, which typically have a semiregular triangulation. Therefore, we found the most intuitive cuts ensue from using E_c alone (equivalent to $a = 1$) but reserve the option to use E_{cg} if required.

After creating the graph representation with edge costs, we compute one cut for each outlet. We resolve the ambiguity of where to split the path at branch locations by creating the cuts in reverse, i.e., we start one cut at each outlet, as illustrated in Figure 29. The first cutting path runs from the largest outlet to the inlet. This initial cut follows the vessel branches with the largest diameter and, therefore, the highest flow volume. On this principal branch, the cut will lie as straight on the vessel geometry as the triangulation allows (cf. Figure 30). Every vertex that lies on a cut is appended to a set $\mathcal{V}_C \subset \mathcal{V}$. For the following cutting paths, we iterate the remaining outlet boundaries from largest to smallest. In every iteration, we find the vertex in \mathcal{V}_C that is closest to the starting outlet. Choosing this vertex as the target ensures that we always hit an existing cut in the vicinity of a branch location and that cuts do not cross.

We provide two options to decide which point on an inlet/outlet boundary will be chosen to start or end the cut. (1) The points on a boundary loop can be connected to a common imaginary node with




Figure 30: Natural vessel cuts on different geometries. (a)-(b) Human carotid arteries. (c) Cerebral artery with an aneurysm. (d) Aorta with an aneurysm. (e) Large segment of a carotid.

zero edge costs. The imaginary node is not present on the mesh and is only used in the graph search. It has only outgoing edges at an outlet and only incoming edges at the inlet. When computing a cut, the imaginary nodes ensure that the most cost-efficient path is found. This requires no further input and is fully automatic. (2) If a cut should face the observer (which makes it easier to mentally reconstruct the unfolding process), a single 3D point can be given as a parameter, for example, the current viewpoint or a fixed location in front of the vessel. Then, we choose the point with the minimal Euclidean distance to this location from each boundary loop. Figure 30 shows the robustness of this approach with different vessel geometries. Finally, the cut is applied to the mesh by duplicating the vertices in \mathcal{V}_C and solving the topology in \mathcal{F} accordingly. If data fields are present on the affected vertices they are copied as well.

5.3 FLATTENING

The flattening should balance area preservation and recognizability of the anatomy. We use a boundary-free approach, which benefits both requirements. However, we fix the inlet boundary loop to an iso- u line

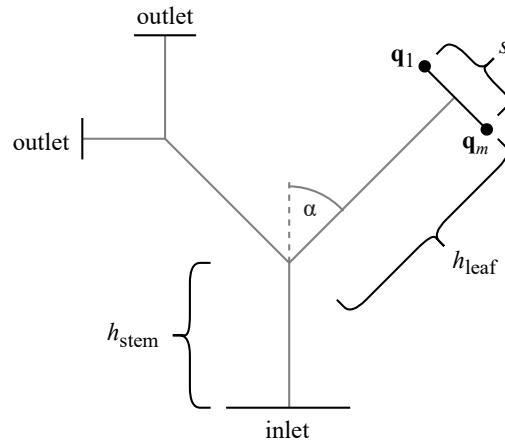


Figure 31: Possible layout initialization that fixes the (u, v) coordinates of inlet and outlet boundaries to resemble a tree. The angle is set to $\alpha = 45^\circ$ but can be arbitrarily chosen. All other values are read from the cuts.

centered on $(u, v) = (0, 0)$, which serves as a reference in parameter space above which the remaining coordinates can develop. Fixing the inlet this way increases the comparability of different flattenings. We use the [ARAP](#) method for mesh deformation introduced by Sorkine and Alexa [286], which was extended for single-patch surface parameterization by Liu et al. [196]. The [ARAP](#) method attempts to retain distances in triangles, preserving their shape up to a rigid transformation, which yields quasi-isometric results when applied to surface parameterization. A trade-off is achieved between area preservation and conformality. The drawback of this method is that it requires an initialization of (u, v) coordinates in parameter space. Typically, this initialization can be done by mapping the surface to a convex shape, such as a square or circle, for which fast and robust methods exist, e.g., harmonic parameterization. We found, however, that the tree-like structure of the mesh is not suited for this approach and extensive overlaps occur within the parameter domain. A better way is to provide a flat initialization that already resembles the expected target layout and then use [ARAP](#) deformation to restore the proportions as observed on the 3D surface.

Different ideas have been proposed for the layout of vessel maps, depending on the use case, as discussed in [Chapter 4](#). We first simplify the convoluted 3D vessel shape. We fix the inlet and outlets to predetermined (u, v) coordinates and move them to their desired locations, as shown in [Figure 31](#). Each inlet/outlet loop contains an ordered set of boundary vertices $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^3$. We re-order the loops such that they start at the cut and the last vertex \mathbf{p}_n is the copy created of \mathbf{p}_1 during the cutting. The loops are unrolled as straight lines, where

each boundary point $\mathbf{p}_i \in \mathbb{R}^3$ is associated with a point in parameter space $\mathbf{q}_i \in \mathbb{R}^2$. Given the arc length of one such loop

$$s = \sum_{i=2}^n \|\mathbf{p}_i - \mathbf{p}_{i-1}\|, \quad (8)$$

we create an intermediate initialization of the mapping that preserves the arc length:

$$\tilde{\mathbf{q}}_i = \begin{cases} \begin{pmatrix} -0.5s \\ h_{\text{leaf}} \end{pmatrix} & \text{if } i = 1 \\ \tilde{\mathbf{q}}_{i-1} + \begin{pmatrix} \|\mathbf{p}_i - \mathbf{p}_{i-1}\| \\ 0 \end{pmatrix} & \text{otherwise.} \end{cases} \quad (9)$$

The resulting line is centered around the v -axis in parameter space. The height h_{leaf} approximates the length of the vessel segment that is capped by the boundary loop (cf. [Figure 31](#)). The intermediate mapping is then recursively displaced by the transforms Φ_1, \dots, Φ_d , where d is the branch depth of an outlet:

$$\mathbf{q}_i = \Phi_1 \circ \dots \circ \Phi_d(\tilde{\mathbf{q}}_i), \quad \Phi(\tilde{\mathbf{q}}_i) = \mathbf{R}\tilde{\mathbf{q}}_i + \begin{pmatrix} 0 \\ h_{\text{stem}} \end{pmatrix}. \quad (10)$$

This is effectively a transformation hierarchy, where the displacement of a segment is passed on to all child segments. The 2×2 matrix \mathbf{R} describes a rotation around an angle α . The angle can be constant or vary depending on the geometry or desired layout. We automatically determine the sign of α by testing if the branch lies to the left or right of the stem it is attached to. A good approximation for the heights h_{stem} and h_{leaf} can be read from the scalar field \mathcal{G} by subtracting the value at the start from the value at the end of the corresponding vessel cut. Note that slight variations of the heights do not affect the result, as the (u, v) coordinates will be displaced in the [ARAP](#) stage anyway. The parameter coordinates of the [LSCM](#) step only serve as an initialization.

After the boundary loops are fixed to the (u, v) plane, the remaining surface coordinates need to be mapped to the parameter domain. It is not important which algorithm is used for this purpose, as long as it creates a bijective mapping that is adequate for use with the [ARAP](#) method. As originally proposed [[196](#)], we use [LSCM](#) [[185](#)] for this step. [LSCM](#) requires at least two fixed (u, v) coordinates but does not need a fixed boundary. We solve the [LSCM](#) with fixed inlet/outlets, which results in a flattening that already approximates the desired layout (see [Figure 32](#)). Finally, the (u, v) coordinates are optimized with the [ARAP](#) method, which displaces the [2D](#) mesh such that it retains the proportions of its [3D](#) counterpart wherever possible (a quasi-isometric parameterization). In this last step, only the inlet boundary remains fixed. Results of the flattening on different geometries are shown in [Figure 33](#).

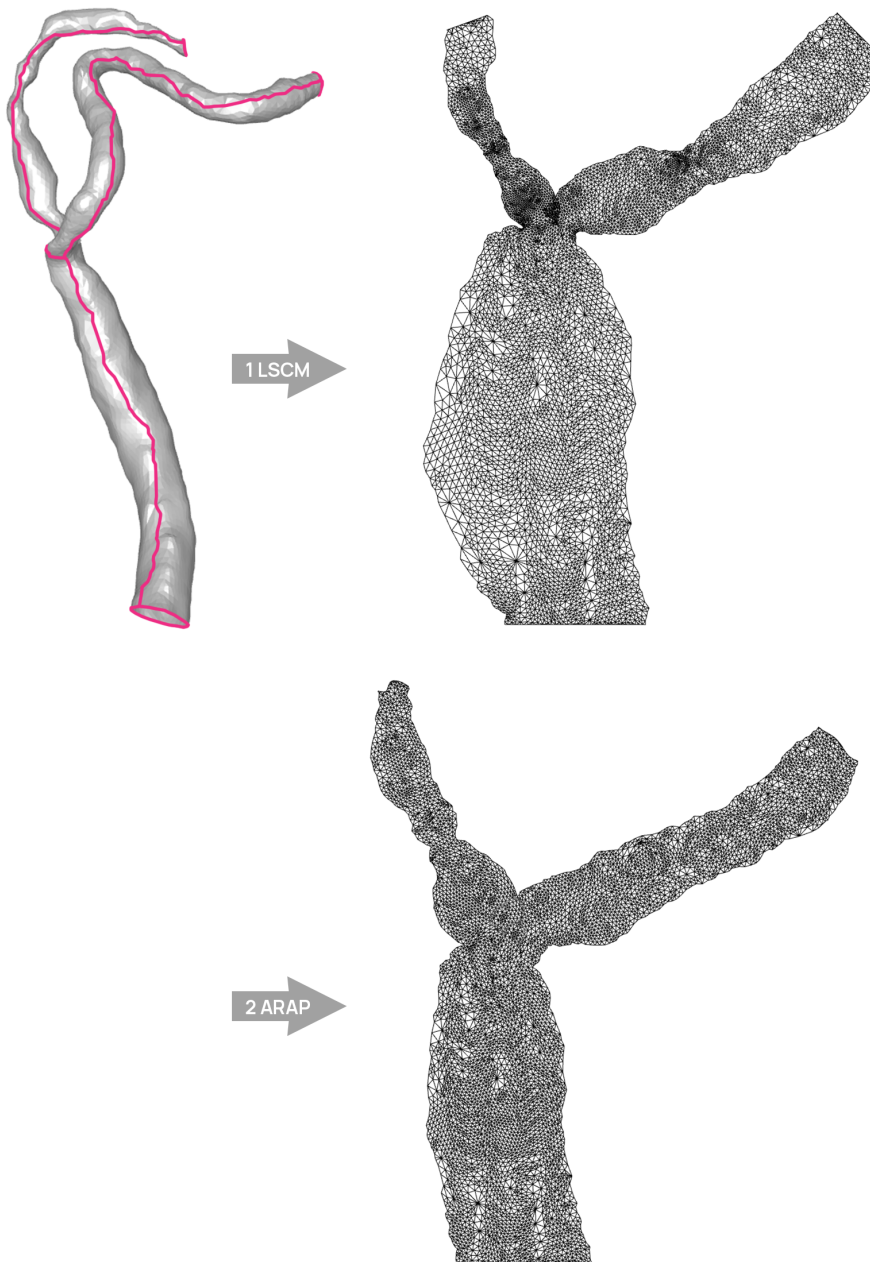


Figure 32: Stages of the flattening. The cut mesh is first flattened to a predetermined layout using [LSCM](#). Then, [ARAP](#) is applied to retain the original proportions.

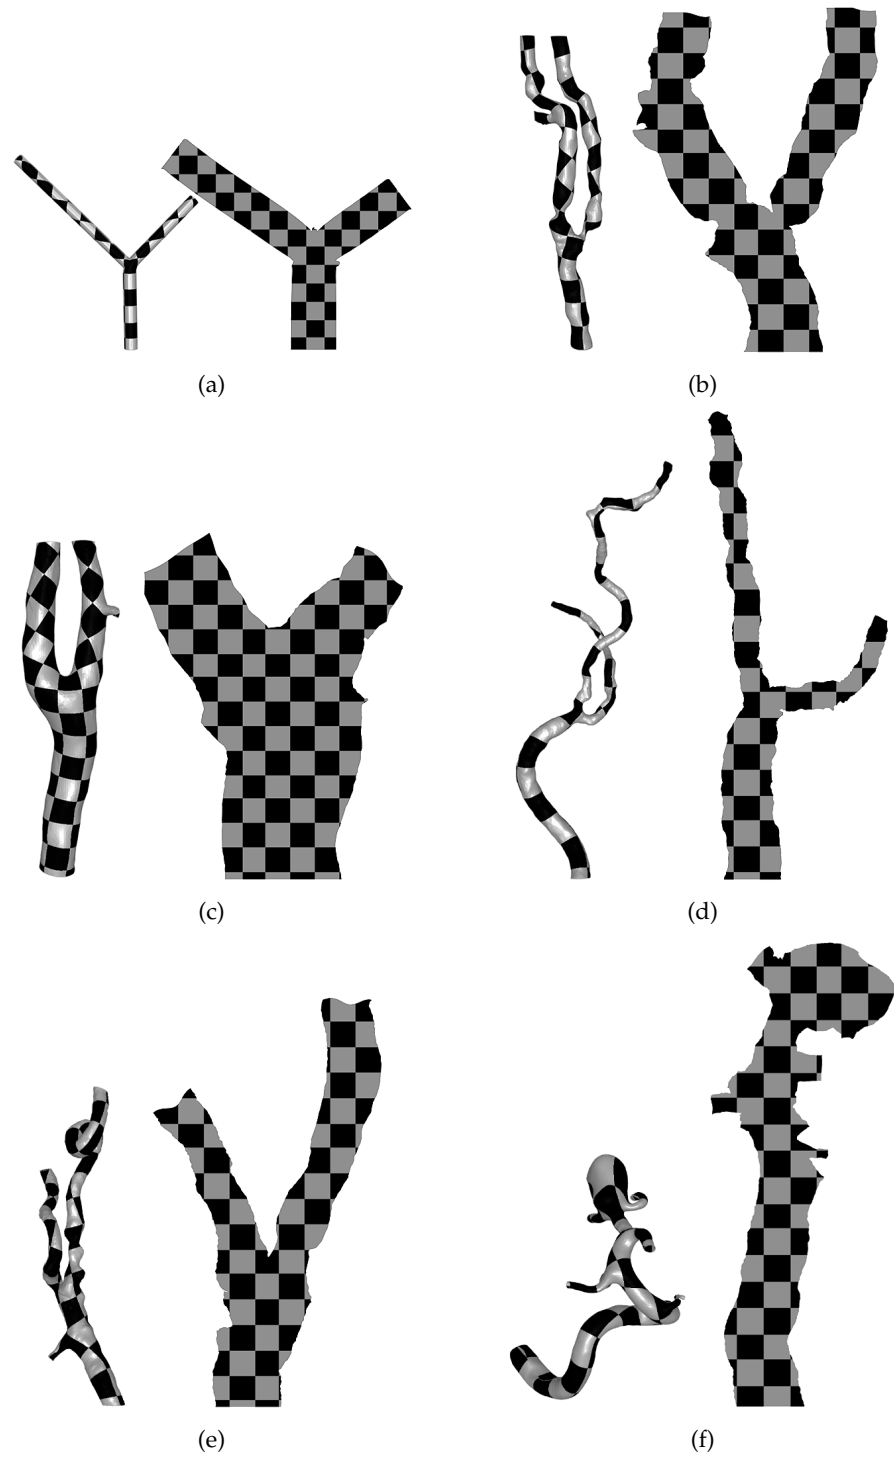


Figure 33: The flattening applied to different geometries. Overall, the area is visibly preserved. We keep the general layout of the vessel trees but complex shapes are straightened, yielding a better overview. (a) Phantom data set of a bifurcation. (b)-(e) Human carotid arteries. (f) Cerebral artery with an aneurysm.

5.4 APPLICATION

A possible application for the flattening is result analysis in CFD workflows. Therefore, we applied the flattening method to analyze simulations of blood flow in the carotid arteries. The overarching goal of this line of research is to better understand the pathological progression of vascular diseases, e.g., how the arterial wall is affected by blood flow in the vicinity of stenoses. These insights could reveal why stenoses often first develop in regions like the carotid bifurcation or siphon and lead to additional markers for clinical decision-making.

The input of the simulation is a model of the carotid inner and outer walls. The data is acquired by solving a fluid-structure interaction problem. The underlying simulation consists of the incompressible Navier-Stokes equations for a Newtonian fluid and its interaction with a linear elastic vessel wall modeled as a 3D geometry with thickness. The fluid equations are equipped with a velocity function at the inlet boundary and a “do-nothing” boundary condition for the normal stress at the outlets. The vessel wall is considered to behave as a linear elastic material with its outer wall moving freely. The continuity of fluid and solid velocities and the continuity of forces are maintained at the fluid-structure interface.

Artificial truncation of the vessel results in surfaces perpendicular to the blood flow which are kept stationary throughout the simulation. A density of $\rho_{blood} = 1000 \frac{\text{kg}}{\text{m}^3}$ [308] and constant blood viscosity of $\mu = 0.00345 \text{ Pa}\cdot\text{s}$ are chosen [210]. The vessel wall has a Young’s modulus between $E = 0.25\text{--}0.75 \text{ MPa}$, a Poisson’s ratio of $\nu = 0.17$, and a density of $\rho_{vessel} = 1100 \frac{\text{kg}}{\text{m}^3}$. The cyclic inflow profile of one heart cycle (0.9 s) has peak systolic and diastolic velocities of $0.59 \frac{\text{m}}{\text{s}}$ at 0.2 s and $0.29 \frac{\text{m}}{\text{s}}$ at 0.55 s, respectively. For insights into pathophysiological risk parameters we evaluate the WSS:

$$\text{WSS} := \sqrt{\sum_{j=1}^2 ((\mathbf{T} \cdot \mathbf{n}) \cdot \tau^{(j)})^2} \quad (11)$$

and the NWS on the fluid-structure interface:

$$\text{NWS} := (\mathbf{T} \cdot \mathbf{n}) \cdot \mathbf{n} \quad (12)$$

with the fluid Cauchy stress tensor \mathbf{T} , the unit outward normal \mathbf{n} , and the unit tangential vectors $\tau^{(j)}$. Parameters that are defined on the whole wall domain, such as displacement, stresses, and strains, can either be evaluated on the fluid-structure interface or the outer vessel wall since they stay approximately constant in radial directions. The CFD simulation was solved with the finite element method (FEM)-based partial differential equation toolbox COMSOL 6.0 [65]. With around 60k elements and an average element-to-volume ratio of 10^{-4} , the simulation over two cardiac cycles takes approximately 43 hours

on two Intel® Xeon® Gold 5115 CPUs @2.4 GHz and 256 GB DDR4 RAM.

The simulation results in temporally resolved scalar fields and a displacement vector defined on each vertex. For the output analysis, we use a discretization increment of 0.02 s over the span of 1.8 s. These are effectively 90 time steps that contain two heart cycles. Exploring this spatio-temporal domain is a non-trivial task, which is why we apply the map for the visual analysis of the result data. As in [Figure 34](#), we show the 3D wall geometry on which any computed scalar field can be color-mapped. The temporal dimension can be explored through animation or the selection of time points, where we smoothly interpolate the data between time steps. If wall displacement was computed it can be animated in the 3D view. We provide a GUI to apply the cutting and flattening to the observed model. The method can be seamlessly integrated into the analysis workflow. Cuts automatically face the viewer and the mesh is flattened on demand. The animated colormap is synchronized automatically between the 3D and 2D domains. We found the 2D view especially useful during temporal exploration (see [Figure 35](#)), as the whole domain is constantly visible. No interaction is necessary as compared to the 3D rendering, where only ever one side of the model can be shown and self-occlusions of geometry may occur.

5.4.1 Implementation

The cutting and flattening are implemented in Python, where time-critical tasks are executed on C++ libraries. The mesh processing largely relies on the implementations in libigl¹ [136]. We use a heap queue for Dijkstra’s algorithm to achieve logarithmic complexity. The user can interface with the cutting and flattening through a GUI that also allows the exploration of surface fields. All common mesh formats supported by libigl can be loaded. CFD data is optional and supplied through a table. The application integrates a PyQt² GUI with OpenGL³ 3D rendering. If CFD data is available, we use a texture to pass the surface deformation (three components) and the selected scalar field (one component) to the GPU. The texture contains one entry for each vertex and time step, allowing us to interpolate the temporal dimension at interactive speed.

5.5 RESULTS AND DISCUSSION

We tested our system with carotid vessel trees segmented from five patients. The models we used vary in length, morphology, degree

¹ <https://libigl.github.io/>

² <https://www.qt.io/>

³ <https://www.opengl.org/>




Figure 34: A surface parameter (*WSS*) is colormapped to a carotid artery. In a 3D projection (top) interaction is necessary to examine the whole surface. We propose an automatic method to cut and flatten vessel trees, which shows the full surface while retaining its proportions and layout (bottom).




Figure 35: The flattened geometry can be used to map parameters, like [NWS](#) (top) or [WSS](#) (bottom). This is particularly beneficial for analysis of the temporal evolution as the whole surface can be assessed at once.

of stenosis, and mesh resolution to cover a broad range of scenarios. Overall, we found the surface area to be well preserved in the [2D](#) mappings (see [Figure 33](#)), which is to be expected due to the [ARAP](#) optimization. This makes it possible to accurately assess the distribution of scalar fields on the surface. Additionally, the area preservation along with the natural vessel cuts lead to an undistorted visualization of the vessel circumference. This shows where narrow parts of the vessel are located. For instance, in [Figure 33b](#) the two stenoses close to the carotid bifurcation are visible in the flattened mesh. We also maintain the general shape of the vessel tree and the orientation of branch locations. This helps to mentally map the flattened surface to the original [3D](#) structure. Complex bends and twists are straightened by the layout constraints we impose (cf. [Figure 31](#)). This filters out unimportant geometrical information and provides a simpler view of the data.

5.5.1 Performance

All tests of the cutting and flattening were executed on an AMD Ryzen 7 5800X processor with 32 GB of RAM. See [Table 1](#) for the impact of different steps and model sizes on the overall performance. For typical use cases of meshes with 10k to 50k triangles, the cutting and surface parameterization executes in 1–12 s. Therefore, we found the flattening can be easily integrated into an analysis workflow. The user only needs to set the frontal direction before initiating the cutting, which determines the starting location of the cut on the mesh. The parameter is set automatically by positioning the camera, making this interaction efficient and intuitive. Currently, the slowest step is the application

Table 1: Timings in seconds of each cutting/flattening step for meshes of different sizes. $|\mathcal{F}|$: Number of triangles. \mathcal{Z} : Computing and aligning curvature. GC: Graph conversion. \mathcal{V}_C : Finding cutting paths with Dijkstra’s algorithm. AC: Applying the cuts to the mesh topology. L: Layout initialization. F: Flattening stages [LSCM](#) and [ARAP](#).

$ \mathcal{F} $	\mathcal{Z}	GC	\mathcal{V}_C	AC	L	F	Total
4k	0.1	0.1	0.1	0.1	< 0.1	< 0.1	0.5
13k	0.3	0.4	0.3	0.7	< 0.1	0.1	1.9
13k	0.3	0.4	0.4	0.4	< 0.1	0.1	1.7
19k	0.5	0.6	0.5	1.1	< 0.1	0.1	2.8
20k	0.4	0.7	0.6	1.7	0.1	0.2	3.6
24k	0.8	0.8	1.1	2.4	0.1	0.2	5.3
28k	0.6	0.9	1.2	2.9	0.1	0.3	5.9
44k	1.2	1.4	2.3	5.6	0.4	0.3	11.2
300k	16.2	9.1	38.1	109.0	1.0	2.3	175.8

of cuts to the mesh topology. This bottleneck is due to the indexed face-set data structure we use. Performance could be improved by integrating a mesh data structure that enables more efficient cutting.

5.5.2 Generalizability

The application demonstration shows the potential of the flattening for the analysis of carotid wall parameters. The presented algorithm is not specific to certain geometries but can be used for a variety of domains, as well as large and small sections of the carotids (see [Figure 33](#)). The inlet and outlets are automatically detected. Also, the branch layout does not affect the result, i.e., our method is robust against bends or twists of the vessel. The flattening with attribute mapping would also be equivalently applicable for measured instead of simulated wall properties, for instance, plaque occurrence or wall thickness. Furthermore, we demonstrate that the method can be applied to other vessel trees, not just the carotids. We show examples of cerebral and aortic vessel trees with aneurysms (cf. [Figure 30](#) and [Figure 33](#)), which pose similar problems in [CFD](#) and medical research. Last but not least, the cutting is independent of the flattening stage. As a result, natural vessel cuts could be used in other scenarios or with a completely different flattening approach.

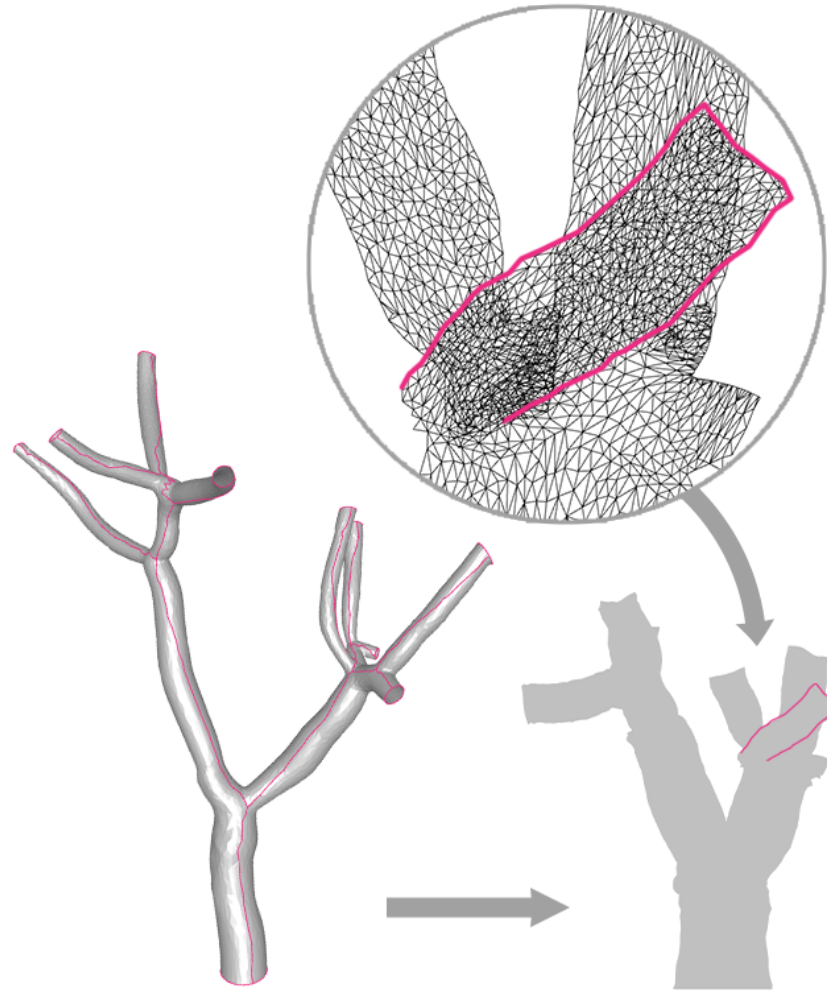


Figure 36: The flattening does not resolve self-intersections of the layout. This can be problematic for trees with many branches.

5.5.3 Limitations

[LSCM](#) and [ARAP](#) ensure that no triangle flips occur and the surface is parameterized as a single patch. However, self-intersections may still appear in the parameter domain. In some models, we could observe this happening close to branch locations. These overlaps could only be solved by introducing area deformation. Also, they only affect small portions of the parameter domain. Larger overlaps can ensue from more complex vessel trees with numerous branches and sub-branches. In [Figure 36](#) the flattening produces noticeable artifacts as the [ARAP](#) optimization relaxes the mesh such that it intersects itself. Our layout initialization (see [Figure 31](#)) does not account for possible overlaps of segments. Therefore, we would advise against using the flattening method for overly complex vessel trees. The cutting, however, is not affected by the number of branches and can be applied to any tree (see [Figure 37](#)).

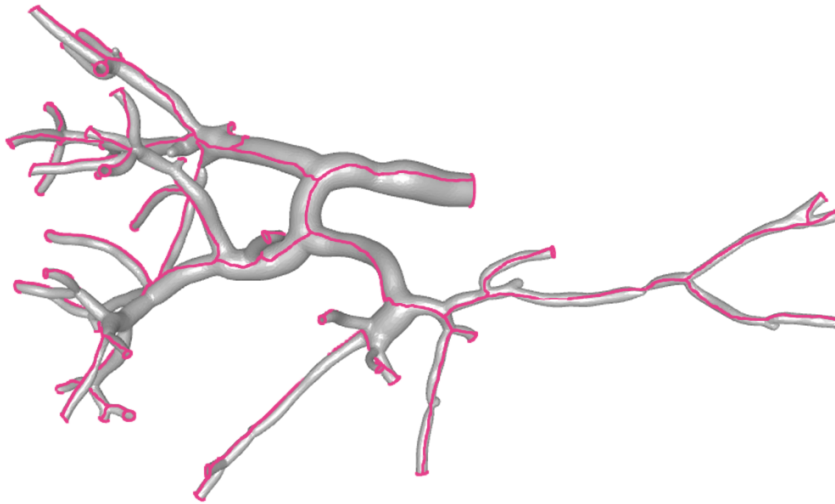


Figure 37: The cutting performs well even on complex structures like this arterial tree segmented from a liver.

5.6 CONCLUSION

In this chapter, a fully automatic technique for cutting and flattening vessel tree geometries was presented. Natural vessel cuts were introduced that remain on one side of cylindrical vessel segments and face the observer when possible. Furthermore, we demonstrated a use case for the analysis of fluid-wall interaction simulations in the human carotids. The application is not limited to the carotids but can be generalized to similar tree-like structures. In contrast to other approaches, the technique requires only the surface mesh as an input and creates a boundary-free parameterization that is quasi-isometric and retains the original vessel shape. The cutting is robust and will always return a valid disk cut given a tree-like geometry, where the cut connections will be at branch locations. The flattening always results in a single connected patch. We presented an automatic way to determine a layout for the flattened geometry that can be used for the carotids or vessel trees with similar complexity.

In the future, the layout initialization could be improved with additional constraints that prevent overlaps of segments. Further ideas include leveraging the flattened geometry for more advanced visual exploration and interaction techniques. Camera control, as well as brushing and linking of interesting regions, could be facilitated by the flattened geometry. Due to its visual simplicity, the 2D depiction could also be used to create comparative views of data set collections. Such overviews would be advantageous when comparing patient cohorts or CFD parameter sets.

6

Visualizing Blood Flow in Carotid Stenoses

This chapter is partly based on:

Pepe Eulzer, Monique Meuschke, Carsten M. Klingner, and Kai Lawonn. "Visualizing Carotid Blood Flow Simulations for Stroke Prevention." In: *Computer Graphics Forum* 40.3 (2021), pp. 435–446.

VISUALIZING BLOOD FLOW IN CAROTID STENOSES

Finding stenoses of the ICA and accurately determining their severity is a crucial clinical task in stroke prevention diagnostics, as reviewed in Chapter 2. A well-timed surgical intervention prevents the closure of the critical artery and drastically reduces the possibility of plaque particle separation (embolization). However, choosing the best moment for surgery and determining the optimal treatment strategy requires a thorough analysis of each case and typically involves multiple specialized experts. Today, two widely implemented diagnostic methods for carotid stenosis are CTA and Doppler ultrasonography. CTA is used to determine vessel *morphology*, i.e., the 3D shape. Ultrasonography reveals some *hemodynamic* properties, describing the blood flow. The parameters measured with each modality need to be mentally combined to derive a decision.

Recent years have seen advances in simulated hemodynamics, which derive properties of the flow field based on CFD [44, 271, 293]. This emerging domain could substitute or complement the information gained through sonographic assessments. CFD can generate high-resolution hemodynamic information based on patient-specific morphology and can also be applied to regions where ultrasound signals are distorted or occluded. The comprehensive work covering medical flow visualization shows that carefully designed exploration methods can support the transfer of simulated hemodynamics into medical research and practice [234], see also Section 3.3. We propose to extend these efforts to cover the domain of stenosis analysis in the carotids. In particular, better visualization tools would allow the integration of hemodynamic and morphological features, which are currently separately assessed. This chapter introduces concepts of how medical practitioners could navigate the spatio-temporal data space faster and more effectively by using sophisticated exploration methods. An application that integrates these methods was implemented, attempting to aid in the identification, comparison, and evaluation of stenoses. Ultimately, the goal is to facilitate clinical decision-making and, therefore, advance stroke prevention.

As this chapter is motivated by the conjecture that visualizing blood flow in the carotids could support real-world analysis tasks, we follow the recently proposed outline to a *data-first* design study [237]: We contribute a stakeholder analysis, concluding that medical practitioners could benefit from simulated flow in the human carotids. We then analyze and abstract the data and tasks involved in studying blood

flow properties w.r.t. stroke prevention. Based on our conclusions, we distill the most promising methods developed in medical flow visualization, which we adapt for the exploration of carotid hemodynamics and morphology. We implemented and tested these methods and discussed our approaches with an interdisciplinary group of physicians, simulation researchers, and flow visualization experts.

6.1 DOMAIN CHARACTERIZATION

The vascular models used in blood flow simulations are often segmented from already available volume images, which means that no additional diagnostic methods are necessary to acquire the data. In the case of carotid hemodynamics, CTA scans are routinely performed, vessel segmentation algorithms based on these scans exist [203], and a variety of CFD techniques were developed [198]. On this basis, we analyzed who would profit from this data, which types of tasks exist, and whether visualization methods needed to be adapted to better facilitate these use cases. We took this data-first perspective [237] to better understand the challenges and opportunities the data poses before applying it to a problem. Further, we intend to provide an initial frame to the domain of visualizing flow in the carotids and outline possible directions.

6.1.1 Data Abstraction

CFD simulations of blood flow span a wide range of methods producing different results. On an abstract layer, however, the types and semantics of the outputs are comparable. The simulation computes attribute fields over the domain, which can include scalars (e.g., pressure) and vectors (e.g., velocity). Each field is resolved in a spatial dimension, which can either refer to the volume inside of the vessel or the wall. Spatial data granularities contain the *region* of the vascular tree used in the simulation and the *spatial resolution*, describing the number and size of individual cells. If physical time is simulated, the fields are also temporally resolved. Temporal data granularities are the *aggregate*, for example, the number of cardiac cycles that are simulated in total, as well as the *temporal resolution*, which effectively describes the number of discrete time steps that are output.

In general, CFD allows full control over the spatial and temporal resolutions and the aggregate. The simulated region is only constrained by the lumen visibility in the underlying volume data. This creates opportunities for flow analysis beyond the capabilities of traditional diagnostic methods, as hemodynamic and morphological information are intrinsically combined in the simulation domain. In addition to values like velocity that are already used in clinical routine, simulations can compute supplementary parameters which could be linked to the

pathological progression of stenoses, such as wall shear stress [327]. It should be noted, however, that computational flow models necessarily rely on simplifications, which poses challenges. The role of CFD to derive hemodynamic factors is controversial due to the difficult validation of the generated data [323]. Results have been shown to vary [274], depending on the employed methods of segmentation, mesh creation, and flow simulation. Only gradually, increasingly sophisticated methods are able to produce robust calculations that are reasonably close to measured parameters [293].

6.1.2 Stakeholder Analysis

Currently, medical and simulation researchers are actively working with flow data generated by CFD in various domains, including carotid hemodynamics. Their tasks evolve around the validation and comparison of models, as well as hypothesis forming and testing. The second group of stakeholders who could profit from detailed carotid flow parameters are medical practitioners, foremost neurologists and neuro-radiologists. The additional information could help make therapeutic decisions, plan interventions, and evaluate treatment success.

We reached out to several independent stakeholders: a flow simulation researcher (13 years of experience) working with carotid data sets, a neurologist (15 years of practice), and a radiologist (30 years of practice). Both physicians are involved in the treatment of stroke patients. We discussed the prospects of simulated blood flow analysis for stroke prevention and the current state-of-the-art in each field.

Medical and simulation researchers use workflows that rely on standard depictions of the result data, e.g., color-coded vessel wall geometry. These depictions could presumably be improved to enable more direct comparisons and integrated evaluations of multiple parameters. However, when speaking with medical practitioners, we found another issue: Even though CFD use for carotid hemodynamics has been extensively studied [198], the transfer to medical practice is low. A possible reason is that such simulations miss validations and derived implications seem to have a small impact [323]. Another possibility might be that the generated data is difficult to grasp for clinicians and does not target their immediate necessities. Most likely, a combination of both factors is involved, which makes the applied use of computed carotid blood flow difficult. If the data is not broadly validated it cannot be used to derive implications and will not be integrated into clinical workflows. However, if it is not integrated, large-scale trials that could validate the data are less likely. We believe that exploration and visualization methods tailored to clinical needs could help bridge the gap between research and application. For this reason, this chapter focuses on possible clinical use cases. We conjecture that visualization methods developed to derive insights

from other types of hemodynamic simulations, like blood flow in aneurysms, can be selectively adapted to fit these scenarios.

6.1.3 Task Abstraction

Our goal is to identify which types of visual displays are required to map the information derived from CFD to the tasks involved in assessing carotid stenoses. We iteratively discussed possible approaches with the physicians and found that the typical tasks performed when analyzing stenoses condense to the following: First, possible stenoses need to be identified. This is a substantial task, as often patients exhibit more than one critical region with a buildup of atherosclerotic plaque. Second, morphological features are analyzed w.r.t. questions like: *What is the diameter inside and after the stenosis? Over which length is the vessel constricted?* If precise values are needed, they are usually measured based on CTA. We found that these first two aspects require large parts of the vessel to be visible, as the morphology of the whole artery is taken into account. Third, hemodynamic features are analyzed, typically based on ultrasonography. Physicians are mostly interested in the two time points of highest (systole) and lowest flow (diastole). *What is the peak velocity? Is turbulence visible?*

In addition to the primary tasks, further inquiries revealed processes that are performed in parallel and might not be obvious at first but are equivalently important. For one, the location of a stenosis determines possible treatment and intervention strategies. An intracranial stenosis needs to be handled differently than one at the bifurcation. This is a trivial factor, however, it needs consideration, as visualizations should not remove this context information. Further, it should be factored in that full assessments include the analysis of tissue around a stenosis. If plaque causes the vessel constriction, its type (e.g., soft or calcified) and configuration (e.g., even or fractured) can be used to judge the pathophysiological factors at play. This answers questions like: *How fast did the stenosis develop? How likely is the rupture of plaque?* Last but not least, the left and right-side carotid are often contrasted for a full picture to inquire, for example: *Does one side compensate for a stenosis in the other?* We summarize these findings in four types of tasks that a visualization needs to cover:

- T1** Gain an overview of the attribute space and find relevant points.
- T2** Probe selected spatio-temporal data points.
- T3** Compare observations regarding time and side (left vs. right).
- T4** Explore the context of global and local anatomical features.

We presume that CFD simulations of carotid blood flow could facilitate these tasks and enhance state-of-the-art diagnostics. Routine




Figure 38: The pipeline used in this work. The carotid lumen is reconstructed from CTA volumes. Based on the models, a simulation of the hemodynamics is performed, resulting in discrete flow fields for both sides over multiple time steps. The simulation results can then be explored with our framework. Navigation relies on the centerline of the segmented models and we incorporate the original data for contextualization and validation purposes.

CTAs could serve as a sensible basis for further data processing. If a sufficiently broad domain is segmented T_1 could be assisted. A 3D model reconstruction in combination with computed flow fields could further support advanced probing (T_2). Fine spatial resolutions can be achieved in the simulation and morphological attributes like the diameter are accessible. A temporal resolution is also available. For comparison (T_3), flow in the left and right carotid can be computed and the results can be displayed in a synchronized visualization. Lastly, the CTA and the simulation data complement each other. Therefore, we propose an integrated visualization, which includes the CTA volume and allows to evaluate near and far context (T_4).

6.2 PRE-PROCESSING

We obtained 13 CTA data sets of patients with differing degrees and locations of carotid stenoses. From these volumes, we segmented the lumen of the full CCA and ICA on both sides until the vessels split into smaller branches. We also include parts of the ECA close to the

bifurcation, which are required for a correct flow simulation. We used the level-set method from the [VMTK](#) [7] and subsequent manual corrections to perform the segmentations. Then, we computed the centerline of the extracted vessel tree based on the 3D Voronoi diagram of the geometry, as this method yields robust and accurate results [310]. Each centerline consists of around 3000 points, providing smooth and consistent lines. Furthermore, the vessel radius is computed as a convenient byproduct, since the procedure aims to maximize the enclosed spheres.

We used the freely available toolkit [OpenFOAM](#) [236] to generate volume meshes of the segmented carotid lumen and perform the [CFD](#) simulation. We applied a standard velocity profile at the inlet of the [CCA](#) which models a typical flow waveform of the cardiac cycle [5]. Solving the Navier-Stokes equations lead to patient-specific flow information, where blood is considered an incompressible Newtonian fluid. For each data set, we simulated three cardiac cycles (about 2.7 seconds) from which we saved 30 equidistantly spaced time steps. The full pipeline is illustrated in [Figure 38](#).

6.3 METHODS

We chose and adapted a variety of methods developed in flow visualization and implemented them in a framework that allows simultaneous exploration of the morphology and hemodynamics of both carotids. Over our implementation cycles, we converged on the layout shown in [Figure 39](#), combining four linked views. We present the 3D domain as either a surface rendering or a streamline visualization (A). The data space navigation is split into the temporal and spatial components of the [CFD](#) results. Orientation in the temporal dimension is facilitated by a line graph showing the volumetric flow rate in each vessel (B). This depiction is representative of the simulated cardiac cycles. A sought time step can be selected by dragging a slider across the graph. The spatial orientation is aided by a vessel map (C), which provides at-a-glance information on the vessel diameter. The location of the probes is determined by markers in this reduced depiction, which can be dragged over the domain and snap to the nearest vessel. The probe views (D) show cross-sections of the vessels at the selected spatio-temporal data points and enable a comparison of the two sides.

A simulation parameter can be selected and is mapped via a colormap shared in (A) and (D). Technically, the framework is capable of showing arbitrary scalar fields. However, in this study, we focused on velocity magnitude and pressure as these are the conventional attributes utilized in clinical routine. We use the [Viridis](#) colormap, as its colors are perceptually uniform and can also be perceived by most forms of color blindness [197].




Figure 39: Screenshot of the application. Four views are linked. (A) 3D depiction of the domain that either displays a surface model or streamlines. (B) Graph of the flow rate for *temporal navigation*. (C) Map of the carotids for *spatial navigation*. (D) Cross-sections of the vessels at the probe locations.

6.3.1 Domain Overview

We show an initial overview of the 3D geometry as a surface rendering to give a general impression of the morphology. The colormap encodes parameter values on the surface. This view can also be set to display a streamline rendering (see Figure 39A), showing the structure of the flow and the value distribution inside of the vessel for the selected parameter. We chose integral lines instead of glyphs or volume rendering, as they are best suited to reveal occurring turbulence. We found that physicians tend to evaluate singular time steps, e.g., at peak systole where the highest flow occurs. Therefore, we chose to display streamlines rather than other types of integral lines or animated particles because they display the state of the flow field at exactly one time step. This aids in the overview task (T₁) and minimizes necessary interaction.

The streamline seeds are placed in 100 random cells of each carotid volume mesh, resulting in a display that broadly covers the domain. We use the adaptive Runge-Kutta integration scheme [170, 266] of 5th order with a maximum propagation of 0.2 meters (about the length of the carotid). When rendering the streamlines we apply two techniques that have proven to be beneficial in similar scenarios [234]. First, we use front-face culling to add context information about the vessel wall without occluding the lines. Second, we render implied tubes instead of lines, which are shaded to enhance depth perception.




Figure 40: The map creation follows this blueprint. The branches are converted into lines, keeping the length l between the segments.

6.3.2 Vessel Map

In 3D depictions of vessel trees, self-occlusions are usually unavoidable. Considerable user interaction is required to assess all parts of the geometry. This makes the diameter of the vessel lumen difficult to compare across the full domain and means that developing stenoses might be overlooked. In the past years, dimensional reduction techniques were effective to overcome these types of issues [163] and many approaches to construct vessel maps were proposed [343], as reviewed in Chapter 4. In this sense, we construct a 2D map of the segmented vasculature to enable at-a-glance assessments of the diameter and increase the comparability of data sets based on this abstract representation.

In the map the CCA and ICA are combined to a fixed straight line, providing an overview of this clinically relevant segment. Neighboring vessels, like the ACE, are drawn splitting off at a constant angle, e.g., $\phi = 30^\circ$, to indicate branch locations and give the map a tree-like appearance. Computation of the map is performed as follows: We treat each branch of the vessel centerline as a curve described by the 3D points $\mathbf{p}_1, \dots, \mathbf{p}_n$ connected in ascending order. This curve is straightened to a line made up of the corresponding 2D points $\mathbf{q}_1, \dots, \mathbf{q}_n$. The main branch (CCA and ICA) is laid out first and side branches are recursively connected at their relative positions. Given the recursion depth $d \in \mathbb{N}_0$ and a start point \mathbf{q}_1 , which is either at $\mathbf{0}$ for $d = 0$ or the branch location otherwise, the mapping for one curve is defined by:

$$\mathbf{q}_i = \mathbf{q}_{i-1} + \begin{pmatrix} \cos(\phi d) \\ \sin(\phi d) \end{pmatrix} \cdot \|\mathbf{p}_i - \mathbf{p}_{i-1}\| \quad \forall i \in [2, n] \quad (13)$$




Figure 41: Map of the carotids used for overview and navigation. The line thickness encodes the vessel radius. Zooming in shows that the stenosis is about 20 millimeters long.

This procedure preserves the geodesic distance between points, showing the accurate length of each vessel segment in the map. A schematic of this approach is shown in Figure 40.

As mentioned earlier, we already retrieved the radius information of the inscribed sphere at each sample location of the centerline. We now map this information to the flattened representation by scaling the thickness of the rendered line. We first used a linear scaling, however, we found that differences in the line thickness did not become obvious. We overcame this problem by using a non-linear scaling which shrinks “small” radii and enlarges “large” ones. This emphasizes narrow regions and thus highlights possible stenoses. Small and large radii are distinguished with a threshold value t . We use $t = 3$ millimeters, which is the average radius of the carotid in a healthy adult [60]. Each radius r_i is then scaled by

$$s(r_i) = t^{1-a} r_i^a, \quad (14)$$

where $a \in \mathbb{R} \geq 1$ determines the “linearity” of the scaling. A value of $a = 1$ produces a linear scaling. We found stenoses to be suitably emphasized at around $a = 2$.

The vessel map is constructed for each of the two carotids. Then, the maps are displayed in a combined depiction, where we flip the left-side

carotid vertically and align the vessel trees in a stacked format (see [Figure 41](#)). We register them horizontally at the bifurcation, which is the defining landmark. Our plot allows zooming and panning and we label the horizontal axis with millimeters. This way, the flat depiction can be used to measure the length of a stenosis in the ICA. The CCA, ICA, and ECA are automatically labeled.

6.3.3 *Flow Rate Graph*

To give the user a sense of the temporal dimension, we plot the inflow volume of each side over time, resulting in a view showing superposed line graphs (see [Figure 39B](#)). This is a beneficial depiction for two reasons: First, temporal patterns regarding the vascular system are commonly based on the cardiac cycle. The flow rate graphs provide an overview of the simulated cardiac cycles and, therefore, tie in with the background knowledge of the users. Second, the volumetric flow rate is a vital global parameter, which essentially determines the quantity of oxygen a vessel can transport. As the flow rate cannot be easily measured in-vivo, it is typically not used in practice. In a simulation, however, the in- and outflow volumes are always known or can be derived. Therefore, they should be visually integrated to provide an impression of the transported blood volume.

As two graphs are used, the flow rate of the left and right carotid can be directly compared. Further, the display is leveraged to show the flow rate inside specific vessel branches. A second line is rendered per side in a dashed style, illustrating the flow in one additional branch. This can be used to compare how much of the total flow is received by one branch, e.g., the ICA. Displaying the flow rate in more than one additional branch per side would clutter the superposed line graphs. Therefore, we select this branch based on where the spatial probe is located. As the simulation assumes incompressibility of the fluid, the flow rate is equal at every location in one vessel branch.

6.3.4 *Navigation and Probing*

We aim to visually guide users to derive relevant insights quickly, thereby minimizing the necessary interaction. Simultaneously, we want to make the spatio-temporal data space (T_2) accessible, allowing all parameters to be thoroughly assessed. In our use case, the temporal dimension appears to be utilized less during exploration. It is often set at the beginning of the analysis process and is only infrequently adjusted. Therefore, we approach this optimization problem by partitioning the controls for data exploration into temporal and spatial navigation.

We use the flow rate graph as a tool for temporal navigation. It reflects the cardiac cycle, providing an intuitive sense of time. Points

of interest, typically the systolic and diastolic time steps, can be easily discerned. We integrate a slider that can be dragged over the plot, which selects the nearest time step in the data series and updates all views accordingly.

Probes are placed via the vessel map, as it presents the full spatial domain in one depiction and is designed to highlight relevant regions, where stenoses could be located. We use a triangular marker to indicate the probe position, as it is conveniently selectable and does not hide information about the vessel thickness. This omits the need for ambiguous placement interactions in the 3D domain. There exists a variety of options for how a spatial probe could be implemented. In the discussions with the domain experts, we concluded that vessel cross-sections, i.e., plane cuts, would be most useful. They show the vessel shape for morphological assessments and the parameter distribution (e.g., the velocity profile) for hemodynamic analysis. Also, they resolve the full scale of the simulation and cover the whole domain, which complements the streamlines that can only show a subset of values. The optimal plane cut is usually orthogonal to the vessel direction. This is the most accurate representation of the diameter, which would otherwise be elongated in at least one axis. Therefore, we compute the direction of the cutting plane as follows. The plane is positioned on the selected 3D point \mathbf{p}_i and oriented according to the normal $\hat{\mathbf{n}}$ defined by:

$$\mathbf{n} = \frac{1}{5} \sum_{j=i+1}^{i+5} \mathbf{p}_j - \frac{1}{5} \sum_{k=i-5}^{i-1} \mathbf{p}_k, \quad \hat{\mathbf{n}} = \frac{\mathbf{n}}{\|\mathbf{n}\|} \quad (15)$$

By averaging the predecessors and successors, noise on the centerline is reduced, which would otherwise lead to jittering when the plane is dragged. The method requires $i \in [6, n - 6]$, which means the plane cannot be placed at the far edges of the centerline. However, regions proximal to the inlets and outlets should not be included in the analysis anyway, as they are likely to contain simulation errors. By relying on the vessel map for positioning and by automatically orienting the cutting plane, only minimal interaction is required to place the plane and none to orient it.

The resulting cut could be shown in the 3D view. However, such a display of the vessel cross-section would suffer from perspective distortion and the viewer's direction would need to be changed depending on the probe position, to see it fully. Therefore, we implemented a separate 2D view of the cross-section, which is aligned on the vessel center (see Figure 39D). We overlay the probe view with a millimeter grid to enable measurements and size estimations. Simultaneously, we indicate the probe location in the 3D view with the same 20×20 millimeter grid. Furthermore, the computed diameter is displayed below the probe view, as it is an important quantifiable parameter. It is also hinted at as a circle to allow validation of the numerical output




Figure 42: Dragging the probe reveals the hemodynamics of stenoses. Usually, the velocity profile should be homogeneous with slower flow in proximity to vessel walls (1). Velocity increases in stenoses (2), followed by visible turbulence (3 and 4).

in regions where the vessel is more elliptical. For direct comparisons (T3) we display two probe views, one for the left and right carotid. We differentiate between the two sides in all views using labels and unified colors. To facilitate a fluent evaluation of the parameter domain along the vessel centerline, the probe can be interactively dragged over the vessel map. The probe view then shows an animation of how the diameter changes and helps to analyze how turbulences behave (see Figure 42).

6.3.5 Contextualization

We determined that, for a full assessment, surrounding structures need to be evaluated (T4). To give a sense of the global context, we embed the original CTA data on-demand into the 3D scene using volume rendering. The definition of transfer functions for volume rendering is a demanding task. As the depiction is only intended to show additional context information, we simplify this for the user. We use preset transfer functions (intensity to color, intensity to opacity, and gradient to opacity), which we derived from functions that have been found suitable in coronary artery visualization [106].

Some control, however, is given to the user, as the region of interest can vary depending on the task, e.g., whether vessels or tissue should be shown. Also, the value ranges of CTAs are not universal due to differences in the contrast agent diffusion. Two sliders are implemented for this purpose: The first shifts the window of the visible value range and the second defines the overall opacity. This way, the user can quickly adjust the density region shown and determine how dominant the volume rendering should be. The final scene is rendered using volume ray casting with composite blending. This method allows a




Figure 43: Hemodynamics visualization with contextual volume rendering enabled. (a) The context rendering shows that the stenosis is located in the upper neck area. (b) Adjusting the thresholding and opacity reveals that the internal jugular vein is in front of this stenosis, potentially complicating surgery.

hybrid depiction of the polygonal models (segmented carotids and streamlines) within the CTA volume. Figure 43 illustrates how this could facilitate far context analysis.

The contextualization is also extended to the 2D probe view, to show the near context of tissue and possible plaque surrounding the vessel cross-section. We sample the CTA volume with the probe plane. The resulting texture is optionally displayed in the background of the probe view. It is rendered with shades of gray similar to the standard CTA representations. We map the Hounsfield Unit range $[-800, 1200]$ linearly to the gray levels $[0, 255]$.

The primary purpose of the contextual CTA slice is to enable the analysis of plaque on the vessel walls. However, it is also useful to validate the segmentation, as demonstrated in Figure 44. The user can adjust the opacity of the model cross-section to reveal all of the




Figure 44: The probe can also sample the CTA volume. This allows to (a) assess plaque and (b) validate the segmentation at crucial locations.




Figure 45: The probe follows the vessel's direction when dragged. This gives an accurate depiction of the lumen diameter, which is essential when assessing the stenosis at (c). Probing the CTA volume unveils the cause of this stenosis – a buildup of plaque on the vessel wall.

original data if desired. With this representation, the CTA volume can be explored similarly to the spatial simulation domain. Dragging the probe across the vessel map reveals the varying shape of the carotid, occurrences of plaque, and changes in the contrast agent diffusion. Aligning the view plane orientation to the vessel direction allows a consecutive evaluation of the vessel following the orthogonal cross-section (see Figure 45). This gives a more accurate view of the vessel morphology as compared to using a fixed slice orientation, which is the typical exploration method in practice. Furthermore, no deformation of the volume is introduced, as is the case when applying planar reformation methods [314].

6.4 EVALUATION

We conducted qualitative reviews with seven domain specialists (2 female, 5 male; 26–61 years old), including the three experts we spoke to during the stakeholder analysis. To evaluate if the methods we adapted aid in the defined tasks, we first individually interviewed two radiologists, one neurologist, and one neuroradiologist who routinely treat stroke patients. We reviewed the concept of flow simulation based on [CFD](#) and explained the pipeline used to generate our data sets. Then, we introduced them to our framework and demonstrated how it can be used to simultaneously explore hemodynamic and morphological information. We explained each view and the ideas behind it. Then, we consecutively presented three data sets of differing types (full stenosis, developing stenosis, and no stenosis). For each set, we asked the participants to interact with the framework themselves and explore the data as if they were to perform a diagnosis. We inquired what they could tell us about the data sets and which information they could gain from the visualizations. We encouraged them to think aloud during this process. Afterward, participants filled out a questionnaire. We associated each view with a series of statements relating to the tasks the view was designed for. Participants were to rate how much they agreed with the statements on a five-point Likert scale (—, —, ○, +, ++). We also showed two pictures of the vessel map, one with linearly scaled diameter mapping and one based on our non-linear scaling method. Without knowing the scaling parameters, participants were to choose on which map they could better identify possible stenoses.

While physicians can evaluate how useful the proposed methods would be in practice, they are less aware of which exploration and visualization approaches exist to display [CFD](#) data. To determine whether we chose and transferred fitting flow visualization techniques and which other options might exist, we extended the evaluation. We interviewed two researchers working with [CFD](#) simulations of the carotids and a flow visualization expert with a research background in aneurysm hemodynamics. We followed a similar evaluation procedure, however, we did not ask the experts to perform mock-up diagnoses. Instead, we demonstrated the steps clinicians took when exploring the data sets, summarizing what was important during clinical assessments. In the questionnaire, we omitted the parts relevant only to physicians, e.g., regarding the analysis of plaque in embedded [CTA](#) data. We discussed whether other visualization methods were known and could be applied.

6.5 RESULTS AND DISCUSSION

All participants quickly understood the linked views and were able to explore the data without noticeable difficulties. When the clinicians examined the example data sets, they mostly defaulted to a single systolic time step and analyzed the domain at this point of peak flow. They immediately identified the stenoses (**T1**) and predominantly used the framework to probe the surrounding velocity magnitudes and vessel diameters. For this task, the vessel map was extensively used as a navigational tool. All experts preferred the non-linear over the linear diameter scaling. They did not see an issue with over- or underestimation of the diameter, as the map was mainly used for localization of potential stenoses, which were then assessed in detail using the probe view. The radiologists pointed out that dragging the plane probe along the carotid shows cross-sections well suited for probing (**T2**), as they could instantly evaluate the lumen diameter, shape, and condition of the wall over a vessel segment. In some cases, participants placed a probe at an equivalent point on the opposite carotid to compare the diameters (**T3**). All physicians extensively used the contextual CTA slice around the probe (**T4**) and appreciated the integration of hemodynamic information on top of the CTA. However, participants also mentioned that they felt uncertain about the correctness of the computed values. Parts of the carotid branches in the cranium form a vascular short circuit between the left and right sides. Due to this connection, a severe stenosis of one side often leads to a reversed flow in the following section of the ICA. As flow is typically simulated separately for each carotid, this phenomenon is not represented. Therefore, we argue that hemodynamic simulations of carotid stenosis could be improved by enlarging the computational domain and taking further cranial branches into account.

The interviewed simulation and visualization experts generally found the selected methods suitably chosen and adapted. They had only minor recommendations and improvement suggestions. One CFD expert proposed to include quantified information on stenoses, e.g., the minimum diameter at a predicted stenosis. This could be included in the vessel map view for a faster initial comparison of multiple stenoses. The flow visualization expert suggested extracting plaque geometry in addition to the vessel lumen and integrating it into an overview depiction (the 3D or the map view). Calcified plaque could be automatically segmented from CTA and used to indicate regions where pathological processes occur.

The questionnaire results can be seen in [Figure 46](#). They reflect the generally positive impressions of all participants. Still, some aspects should be discussed. The volume rendering was well received for the localization of stenoses, but was less used to explore surrounding structures, as participants found the noise distracting. Filtering could

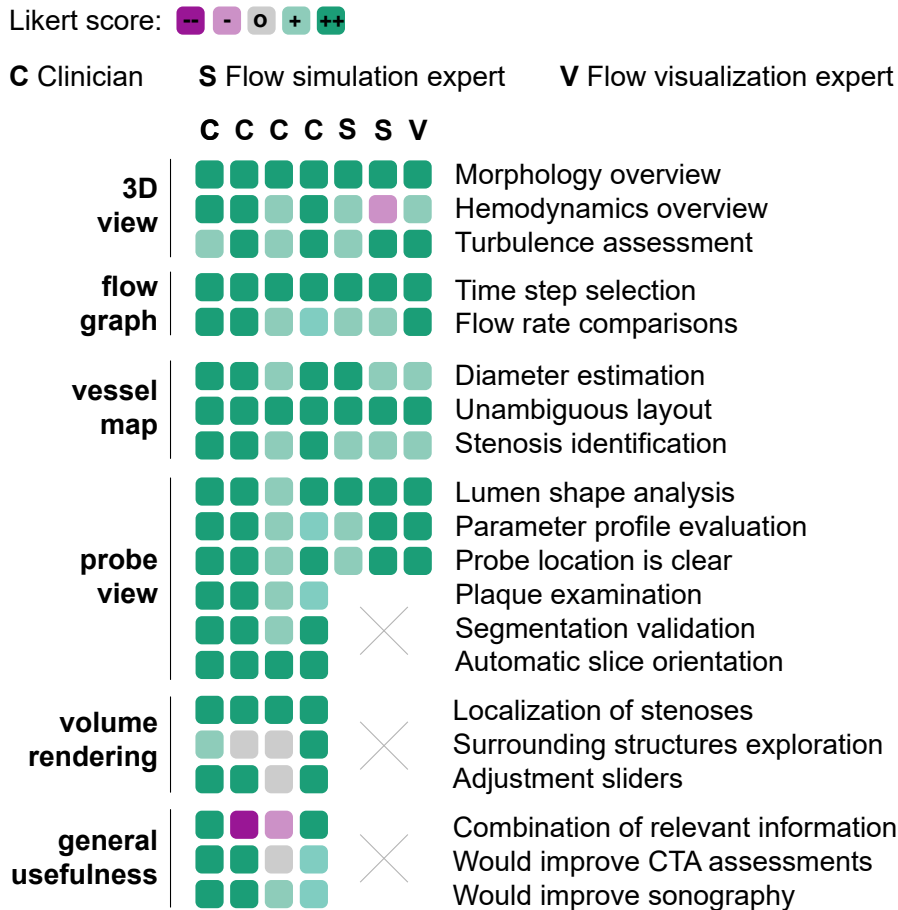


Figure 46: Results of the questionnaire with color-encoded Likert scores. Each box shows the answer of one participant.

be applied to remedy this issue, however, this would not be a trivial task, as context information should not be occluded. Different types of structures (e.g., tissue, vessels, or bones) can be of interest, which makes the definition of automatic filters difficult.

One assertion in the survey led to disagreements: *The framework appropriately combines hemodynamic (as in Doppler ultrasonography) and morphological information (as in CTA)*. Two participants argued that the hemodynamic data shown is simulated not measured, i.e., the framework does not combine information but rather shows additional data. One could suggest that we should have phrased the statement differently, to better reflect our actual workflow. Then, we could argue that the validity will increase with better simulations available. As this would likely not affect how the data should be visualized, we could dismiss the issue as outside of our research domain. However, we want to draw attention to this point, as it is, again, indicative of an important challenge: Many clinicians see the validity of the simulation results as a barrier to the practical adoption of CFD methods. We revisited this discussion with researchers from both sides (medical and

CFD) and singled out a core problem: advances in the CFD domain are rarely brought to the attention of physicians, and the requirements of clinicians are often unknown to researchers in the field of CFD. This bottlenecks clinical integration of computed flow data, which would be necessary for large trials that could support data validation. We believe that visualizations of simulated data could be key to solving this issue, as they can provide clinicians with appropriate views on the computation results, which has a high potential to improve interdisciplinary communication. Furthermore, our discussions led to two options for how visualizations could directly address the uncertainty around simulated flow data: A possible approach would be to register measured data from CTA and ultrasonography in a combined visual display. This is a challenging objective due to the low resolution and high noise associated with ultrasound assessments. Also, ultrasonography is usually not recorded in a way that CTA or MRA are but used mainly for live exploration. Another option is to use sonographic measurements to initialize boundary conditions for the simulation in areas where ultrasonography can be performed. Ideally, this would combine the advantages of ultrasonography and CFD.

6.6 CONCLUSION AND FUTURE WORK

In this chapter, we investigated how medical flow visualization can be applied to enhance the diagnostic evaluation of carotid stenoses. We developed and tested an application with four linked views, which enables physicians to explore the vessel morphology in combination with simulated hemodynamics. Our findings indicate that many concepts from flow visualization research can be transferred to match specific tasks in this scenario. We found geometry abstractions to support overview tasks and the identification of points of interest. We stress that this appeared particularly useful for spatial navigation, as demonstrated by the vessel map we employ to place data probes. Further, we would like to emphasize the importance of contextualization in medical visualization, which is a key feature of this work. Context information can not only be used to localize stenoses but also to explore important additional features (such as plaque), and perform on-the-fly validations. We also realized that some caveats should be considered when developing visualizations with respect to stenosis analysis. For clinical use cases, the local view of flow patterns and parameters in the vicinity of stenoses should be complemented by a more global perspective. To gain a full picture, it is crucial that the general distribution of blood flow and the occurrences of multiple stenoses can be explored. Furthermore, we found the temporal dimension to be of less importance than we first anticipated. From our evaluations and discussions, we assume that a depiction focused on the systolic and diastolic values could be more useful in practice

than the integration of a high number of time steps. Lastly, it should be noted that CFD can generate an extensive amount of data. The exploration of this data in application scenarios, such as medical diagnostics, can be cumbersome and prevent the integration into routine. It is therefore essential that the core tasks of users are distilled first, followed by tailoring methods to match these tasks exactly, aiming to minimize necessary interaction. In this study, developing the right exploration techniques played a similarly important role as choosing and adapting visualization methods.

In the future, a central objective should be to increase the validity and, therefore, the applicability of hemodynamics computed with CFD. We mentioned two options for how visualization could advance these efforts. First, different modalities (e.g., CTA and ultrasonography) could be visually integrated. Then, data visualizations could, for example, illustrate which parts of the data were measured and which were simulated. Second, a larger vascular domain could be segmented to compute the correct flow properties of arterial short circuits. This would pose the question of how such an extended domain could be depicted without cluttering the view. Furthermore, we prioritized parameters that are widely used, such as flow velocity and vessel diameter. However, CFD is often employed to generate complementary attribute fields, which could also play a role in pathogenesis. To foster a better understanding of such additional factors, future work should aid in the investigation of occurring patterns. This would contribute to understanding the initiation and development of stenoses. Cohort visualizations could facilitate this endeavor by enabling comparisons across patient collectives. Potentially, this would allow further insights to be derived which could then be applied to supplement diagnostic assessments.

Part IV

IDENTIFYING AND SOLVING CLINICAL
TRANSFER CHALLENGES

7

A Fully Integrated Pipeline Architecture

This chapter is partly based on:

Pepe Eulzer, Fabienne von Deylen, Ralf Wickenhöfer, Carsten M. Klingner, and Kai Lawonn. "A Fully Integrated Pipeline for Visual Carotid Morphology Analysis."

In: *Computer Graphics Forum* 42.3 (2023), pp. 25–37.

Pepe Eulzer, Kevin Richter, Anna Hundertmark, Monique Meuschke, Ralph Wickenhöfer, Carsten Klingner, and Kai Lawonn. "Visualizing Carotid Stenoses for Stroke Treatment and Prevention."

In: *EuroVis – Dirk Bartz Prize*. The Eurographics Association, 2023.

Model-based vessel visualization provides a wide range of specialized tools for aiding clinical analysis tasks [251, 252]. It also opens up the possibility of automatic feature classification [176]. The drawback is that advanced model-based vessel visualizations tend to have a high adoption cost, as they typically depend on multiple fragmented tools for data curation and processing. The same holds for the approaches described in Chapter 5 and Chapter 6. A comprehensive visualization of the stenosis shape, anatomical context, and flow information is only possible with sophisticated prior data processing. Keeping this processing opaque to the end user may not only hinder efficient adoption but also negatively impact user trust. With highly processed data, how can users assure the correctness of what they see? In this chapter, we approach this challenge by proposing a fully integrated pipeline for segmentation, pre-processing, and visual analysis of the carotids. As its fundamental principle, the pipeline interface enables interactive and retrospective user control over all data processing stages. Through the pipeline design, we target increasing user trust by making the underlying data validatable on the fly, decreasing adoption costs by minimizing external dependencies, and optimizing scalability by streamlining the data processing.

7.1 INTRODUCTION

There already exist a number of data processing frameworks designed for researchers working with vascular models. The VMTK [7, 135] 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 [151] and ImageJ [276]. Due to their high versatility, they are broadly used. However, due to their steep learning curve and vast range of functions, in terms of usability for specialized (clinical) tasks they cannot compete with individualized tools specifically designed for these purposes. 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 [167], image reformations [221, 263], map-like depictions [343], and visualizations of hemodynamics (blood flow) [234]. These individualized tools, however, are difficult to implement in practice due to their typically high number of exter-

nal dependencies and complex pre-processing requirements. 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. We identify four major bottlenecks for the practical transfer of methods that we aim to address with our approach: (1) the users' trust in what is visualized, (2) the efficiency of using a visualization with new data, (3) the compatibility of the used data formats, and (4) the scale at which such systems are validated.

To analyze carotid stenosis, essential objectives are the proper assessment of the stenosis location, degree, composition, and shape, i.e., the vessel morphology. In this chapter, we introduce a fully integrated pipeline to efficiently extract and visualize 3D vascular models from 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 47. It aims to target the identified transfer 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. We elaborate on the design of the pipeline for the analysis of stenoses of the internal carotids. The concept could be transferred to other domains like cranial circulation [219], aneurysm risk-assessment [213, 214], liver surgery planning [189], or the analysis of heart valves [340, 341]. In summary, this chapter contributes:

- 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 the approach, the pipeline has been used to generate 152 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 [48, 347].




Figure 47: The integrated pipeline is split into three pre-processing modules (the crop module, the segmentation module, and the centerline module) 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.

7.2 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 [253] – 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 not lucrative to sustainably determine the benefits of new vessel visualizations. From the literature on vessel visualization systems and 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 possess a high potential value for diagnosis and treatment. So why are they only sparsely adopted? While various probable reasons exist, in discussions we kept iterating on a particular aspect: commonly, in prototypical development, the focus is not 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 [103].

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 data sets, however, may lead to unwanted side effects, such as algorithms and visualization techniques that are overfitted, i.e., 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.

7.3 SYSTEM DESIGN

This chapter introduces an 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, as outlined in Chapter 2. Of particular importance is the stenosis degree $S_{NASCET} \in [0\%, 100\%]$, based on the NASCET criteria [88]. It requires the minimal inner diameter inside a stenosis d_{min} and the normal poststenotic vessel diameter d_{normal} , see also Equation 1. In practice, the exact positions for measuring d_{min} and d_{normal} are not known and are subject to the estimate of the 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




Figure 48: 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.

CTA image data in standardized Digital Imaging and Communications in Medicine (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.

7.4 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 47, 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 48. 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 [47]. The 3D visualizations are based on the Python VTK wrapper and 2D graphs are created with PyQtGraph¹ to enable smooth real-time interaction.

7.4.1 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 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 48 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

¹ <https://www.pyqtgraph.org/>

interest is also indicated in a volume rendering (see [Figure 48](#) 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 [182]. 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.

7.4.2 Convolutional Neural Network

Deep neural networks, CNNs in particular, have rapidly evolved and been widely used in medical image analysis [195]. Inspired by the biological visual system [191], 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 [262], which is a CNN variant 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 [62, 81, 284]. To segment the carotids in CTA volumes, we chose a state-of-the-art 3D U-Net implemented in PyTorch from MONAI [47]. 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 pre-processing 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 [291] 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 also contains the plaque label whose shape and proportions vary extensively between datasets and is thus hard to train for.




Figure 49: 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.

7.4.3 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 49). 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 [114] to the segmentation label map. It is smoothed with a windowed sinc filter [297] (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




Figure 50: 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.

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.

7.4.4 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. [6]. 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 50](#), such that the user can quickly visually inspect the results.

7.4.5 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 51](#) right). The most important information is the vessel diameter, for 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 51D](#)). Hovering over a graph highlights the respective vessel segment in the 3D model to intuitively connect the two visualizations. This interaction is shown in [Figure 52](#).

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 51](#) 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 [50].

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 $\frac{1}{2}$ 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

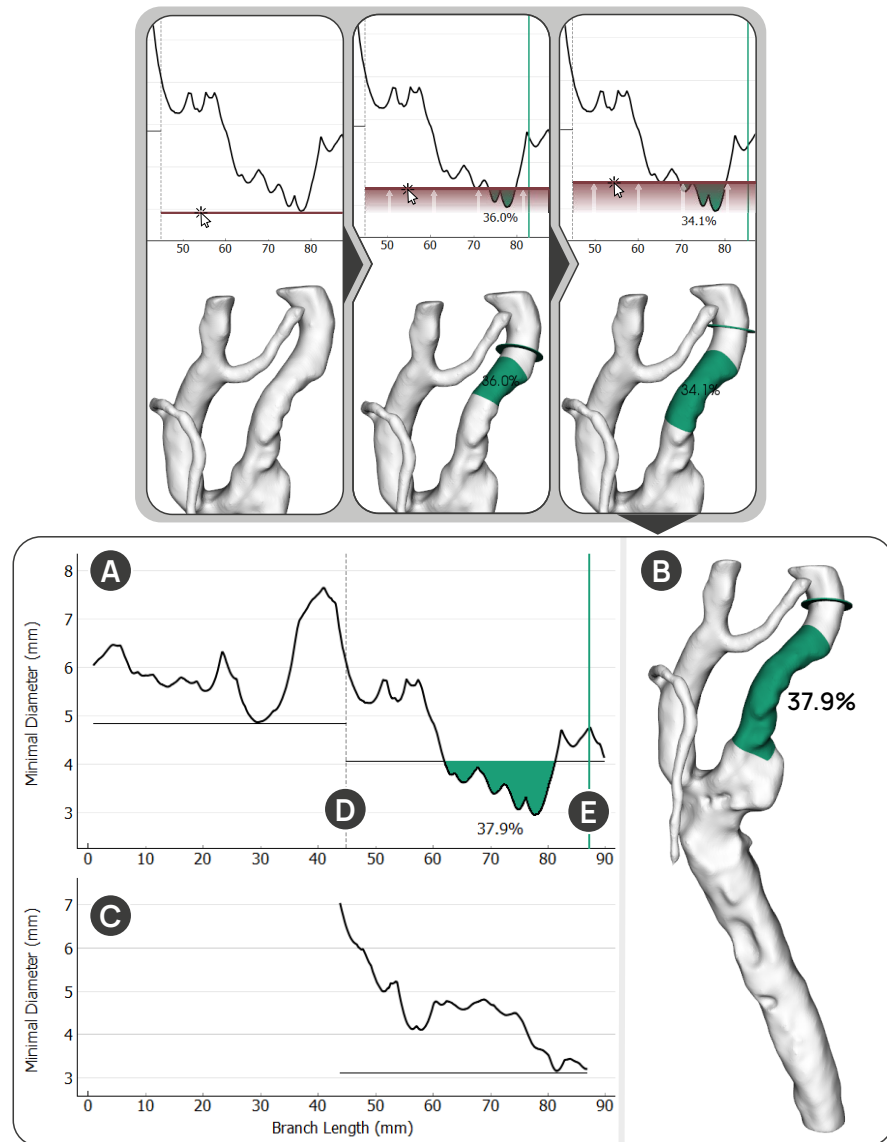


Figure 51: The classifier module is composed of a diameter graph for each branch modeled by a centerline and a 3D view of the vessel lumen. As shown on the top, dragging the branch-specific threshold, indicated by a horizontal line, segments the corresponding stenosis candidate in the surface model. 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.

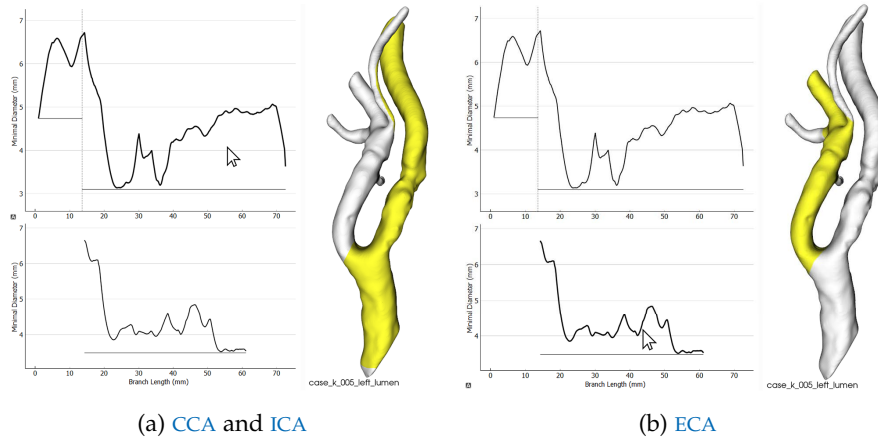


Figure 52: Hovering over the plots indicates the corresponding branch in the vessel model.

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} . Results of the automatic stenosis classification with different degrees of severity can be seen in Figure 53.

7.5 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 62 patients, resulting in 110 carotid bifurcation models², one each for the left and right sides. Arteries with 100% stenosis were omitted, as fully occluded arteries are not visible in the CTA data. The database is available online [347]. 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 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.

² The database was later extended and currently contains 152 carotid models extracted from 79 patients.




Figure 53: Exemplary results of the stenosis classifier with increasing stenosis degree from (a) to (g). 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).

7.5.1 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. 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 pre-processed, 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 54, regarding the workflow and modules, which the participants were to score on a five-point Likert scale (–, –, o, +, ++). 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.

7.5.2 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 or stenting. 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)

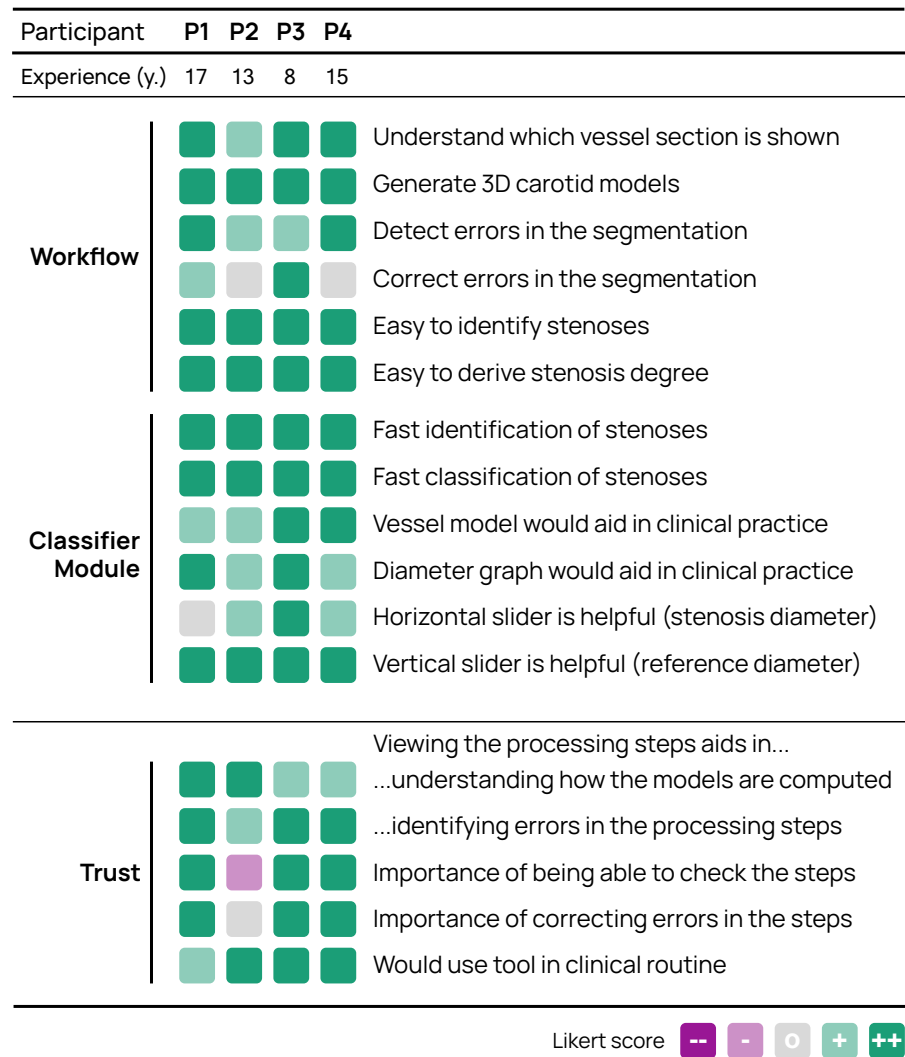


Figure 54: Results of the questionnaire with color-encoded Likert scores. Each box represents the answer of one physician.

All physicians stated they would find the tool highly useful to support the screening of the carotids. P3 and P4 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 54](#). 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.

“Seeing the dips in the graph, I can very quickly highlight the stenosis. That it instantly gives me the stenosis degree is very impressive.” (P3)

*“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.” (P4)*

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)

7.6 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 53](#), 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.

In this chapter, 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 only needed 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 62 different patients [347]. 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 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 the next step, we intend to deploy the pipeline in clinical research, gathering further feedback and testing the applicability in the real world.

7.6.1 *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 [189], 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.

7.6.2 *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 ultrasonography. We plan to evaluate the automatic stenosis classification in a clinical comparative study to investigate the differences between these methods.

7.7 CONCLUSION

This chapter presented a system for the visual analysis of the carotid morphology to improve cardiovascular disease screening and treatment. The methods are implemented 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.

8

Similarity-based Hemodynamics Analysis

This chapter is partly based on:

Pepe Eulzer, Kevin Richter, Anna Hundertmark, Ralph Wickenhöfer, Carsten M. Klingner, and Kai Lawonn. “Instantaneous Visual Analysis of Blood Flow in Stenoses Using Morphological Similarity.” In: *Computer Graphics Forum* 43.3 (2024), in print.

Carotid stenosis is commonly diagnosed using CTA and Doppler ultrasonography, as reviewed in Section 2.2. CTA provides a static volume image of the head and neck region with highlighted arteries. Doppler ultrasonography is used to get a dynamic view of the blood flow and to measure the flow velocity. The blood flow velocity increases with a smaller vessel diameter, and as such it is used as a primary clinical marker for the severeness of the stenosis. A disadvantage of these methods is that the data from the CTA and sonography need to be mentally combined. Two separate procedures, involving different experts and equipment need to be performed. Doppler sonography, in particular, requires a highly trained medical professional to be carried out correctly. The necessary personnel and devices are not available in many smaller clinics, which means measuring flow can be inaccessible.

In the last decades, a large body of research has been dedicated to simulating arterial blood flow [27, 44, 198, 293]. These simulations are numerical predictions commonly using CFD to derive flow properties based on the vessel wall geometry. This geometry can be extracted from static imaging like CTA, theoretically omitting the need for additional diagnostic procedures. Hemodynamic simulation can also yield a much higher resolution than sonography and it can be applied to regions where ultrasound signals are occluded. It can further be used to derive many additional parameters beyond flow velocity, where some are relevant biomarkers for predicting the progression of atherosclerotic diseases [124, 229]. Studies based on hemodynamic simulation could show that carotid bifurcation plaque formation is influenced by WSS, dynamic pressure, strain rate, and the total pressure gradient [186]. However, these factors are not taken into account in clinical decision-making, as performing hemodynamic simulations on site remains practically impossible.

While hemodynamic simulations have the potential to significantly impact patient care and treatment outcomes, they are difficult to integrate into clinical workflows, due to the often time-consuming, resource-intensive, and complex nature of the simulations. The simulations used in this work, for instance, required an average computation time of 4 hours and 50 minutes *per artery*, excluding pre- and post-processing steps like geometry extraction. Hemodynamic simulations are often developed by specialized research groups and require significant expertise, as well as specific software and hardware to perform.

This chapter introduces a method to circumvent performing hemodynamic simulation for each new case individually. An overview




Figure 55: Overview of the data processing and visualization steps. The vessel geometry and features are extracted from an imported DICOM series. Using similarity criteria, the new model is matched against a flow model database. Then, we use comparative visualization techniques to enable the exploration of flow parameters in the best-matching vessels.

of the approach is shown in Figure 55. We created a database of high-resolution flow models for a wide range of carotid bifurcation geometries. Then, we developed a framework that, given a new carotid wall geometry, finds the candidates with the highest similarity from the simulation database. We enable exploring these candidates and visually comparing their geometries, simulated flow, and wall-related parameters. We investigated how similarity between vascular models can be established and whether similarity-based analysis is a viable approach to circumvent complex individual flow simulations. We

enabled exploring multiple candidates and visually comparing their geometries, simulated flow, and wall stress. In summary, the contributions of this chapter are:

- Extending the open-source database of carotid bifurcation models created in [Chapter 7](#) with high-resolution blood flow simulation data. At the time of writing, the database encompasses 152 individual bifurcation models extracted from real patients [347].
- Developing difference metrics to create an alignment space and derive model similarity.
- Implementing and testing an interactive framework that, given a new vessel wall model, can be used to instantly find and visualize the most fitting flow simulations.

The developed software is open source and implemented as an extension module of the CarotidAnalyzer framework [48], which was introduced in [Chapter 7](#). The carotid flow database is available online [347]. It contains fluid data (velocity, pressure) and surface data (WSS) for the two clinically relevant time steps of peak systole and end-diastole.

8.1 TASK ANALYSIS

For the development of our framework, we collaborated with three independent specialists: a neurologist leading a stroke unit (P1, 17 years of practice), a radiologist (P2, 32 years of practice), and a hemodynamic simulation expert (H1, 15 years of experience). At the beginning of the project, we conducted multiple focus group interviews to extract the clinical tasks and requirements a similarity-based blood flow exploration framework would need to address. Both physicians emphasized the potential impact of such a system:

“If I could get flow information immediately from the CT, that would be transformative. The blood flow is really ultimately what I’m interested in. If I can get a reasonable prediction of the flow velocity inside a stenosis, for example, that would be highly valuable. A good estimate would already be better than nothing, we do not work with highly exact values anyway. Even when we measure the flow with Doppler, we have a high observer variability.” (P1)

“I have worked with CFD simulations before. I think the potential of CFD in medicine is really restricted by the fact that doing these simulations is so laborious. Building a database of many pre-simulated cases could be a breakthrough idea. I suspect for regions where we often find very similar anatomical variations, like the carotid bifurcation, this could work. It would be important to find out how many cases we need in the database.” (P2)

They stressed the significance of good automatic matching of the models and efficient exploration of similar cases. *“No physician would go through the database manually. Ideally, the best matches are shown automatically and I can get a peak of what they look like before I read out any specific values (P2).”* When asked about which flow parameters should be derived from the simulations, the physicians argued that to complement current clinical decision-making the maximal flow velocity should be the primary focus. It should be possible to probe the maximal velocity in different branches, e.g., inside a stenosis, but also in the CCA, ECA, and after a stenosis. P1 added that the occurrence of poststenotic turbulence and resulting reverse flow can sometimes be observed in sonographic screenings and that such flow patterns can provide cues about the severity of a stenosis and the probability of plaque separation. Further, P2 noted that studying WSS distributions is increasingly important for understanding stenosis progression. Regarding the temporal flow dimension, the physicians argued they are essentially only interested in the peak systolic flow, i.e., the point in the cardiac cycle with the maximal arterial flow. In some cases, the systolic flow velocity is also compared against the diastolic (lowest) flow.

Last but not least, we must stress that the applicability of the proposed framework depends highly on the efficiency of the model extraction. To compute any similarity metric between a new case and a pre-computed CFD simulation, a geometric vessel model must be available. We quickly determined that using CTA imaging is the most sensible starting point, as CTAs are acquired routinely in stroke diagnostics, in both preventive and acute cases. CTA is a comparably fast procedure that images the arterial lumen with good contrast, shows plaques and is simultaneously used to rule out internal bleeding. Therefore, to make the system usable in practice, we must integrate an efficient way to generate a geometric model from a CTA volume that we can then compare against the flow database. In summary, the following tasks should be fulfilled:

- T1** Extract a carotid bifurcation vessel wall model from CTA.
- T2** Find matching models with pre-computed flow fields.
- T3** Visually compare the new model against the selected models.
- T4** Explore relevant flow-derived parameters in the selected models: velocity, reverse flow, WSS.
- T5** Probe any subregion of the flow fields and extract the maximal local flow velocity.

Visually comparing multiple vessels with hemodynamic information has been investigated before, however, only for a limited number of instances. For instance, techniques for integrated comparison of two

data sets with internal flow were proposed by Meuschke et al. [351]. They provide several exploration tools for comparing two data sets, e.g., before and after treatment. To the best of our knowledge, there are no frameworks geared for comparative visual exploration of larger sets of hemodynamic simulations, such as patient cohorts. The most promising techniques that could be used to compare multiple hemodynamic simulation instances fall in the category of vessel maps, see [Chapter 4](#). Vessel maps have been, for example, applied to track plaque development in the carotids over multiple sonography scanning sessions [58, 59].

8.2 STREAMLINE CLUSTERING

A popular method to reveal the structure of internal flow fields is the use of integral lines, such as streamlines. To minimize visual clutter that can lead to occlusion and to highlight flow patterns, various streamline clustering techniques were proposed and many of them specifically target medical applications. The motivation behind clustering in the context of streamlines is usually to find similar patterns and reduce the number of lines. A clustering method needs a metric that measures how similar two lines are and an approach to group lines based on their similarity [325]. Corouge et al. [68] applied clustering methods on fiber bundles. As a similarity measure, they proposed three different metrics: the closest point distance, the mean distance of closest distances, and the Hausdorff metric. The clustering itself is based on a threshold, whenever the distance between two fibers is less than the threshold, they are grouped in a cluster. McLoughlin et al. [208] provided a similarity metric that was based on geometrical properties, e.g., curvature and torsion. Afterward, agglomerative hierarchical clustering (AHC) with average linking was used to group the streamlines. Clustering of blood flow inside an aneurysm was first applied by Oeltze et al. [233]. They slightly changed the distance measure by Corouge et al. [68] and compared various clustering methods. To account for temporally resolved flow fields, Meuschke et al. [351, 218] developed a physical-based similarity metric. Then, they also apply AHC to group the lines. To further identify blood flow patterns in an aneurysm independent of their individual shape, Meuschke et al. [213] clustered streamlines based on the method of Oeltze et al. [233]. This method identifies classical flow patterns, enabling comparability between different aneurysm data sets.

8.3 METHODS

This work is a first attempt at comparison-based blood flow analysis. The primary objective is to test the viability of this approach. For this, we require an initial database of hemodynamic simulations, which

allows us to find geometrically proximate matches to a new carotid bifurcation wall model. When choosing the size of the database, we must consider the number of instances required to test the viability of the system, but also the feasibility of performing a numerical simulation for each instance. Studies that employ hemodynamic simulations in patient-specific vascular anatomy typically only focus on very few subjects. The processes of image data acquisition, model extraction, pre-processing, and numerical simulation incur substantial time and computational costs. In their systematic review, Lopes et al. [198] identified 49 studies on the simulation of blood flow in patient-specific geometries of the carotid artery bifurcation. Out of these studies, the majority are based on 1–5 bifurcation geometries. Recently, some larger studies with up to 50 carotid bifurcations were conducted [16, 178], but these remain exceptions. While statistical evidence is lacking regarding the number of cases sufficient to cover the anatomical variations of carotid stenosis, we assume that we require more cases than even these larger studies. Consulting our collaborating physicians, we determined that it would be sensible to include at least 100 variations in the database to be able to reasonably test the approach. We also need to consider additional overhead for testing, i.e., we need models with simulated flow that are not part of the test database. Consequently, we have factored in a 20% overhead and extracted 120 individual carotid bifurcation models from CTA volumes to construct the numerical flow database.¹ From the radiological report of each case, we have prior knowledge regarding the severity of the ICA constriction. We chose cases with ICA stenosis close to the bifurcation and aimed to cover the full spectrum of this type of stenosis. Therefore, we systematically selected an approximately equal number of cases with no/mild (0–49%), moderate (50–69%), and severe (70–99%) stenosis. Generally, we extracted both the left and right carotid in each scan, however, we excluded arteries with full occlusion (100% stenosis), as these vessels are not discernible in the CTA. Each extraction covers 80 mm in height with the bifurcation approximately at the center. This section contains the carotid bifurcation and any nearby stenoses. We extract the vessel lumen, therefore, the surfaces model the inner vessel wall. The wall thickness is not reliably detectable in CTA. The data was sourced from two different hospitals. In Section 8.4.2 we test the accuracy of this initial system in predicting flow parameters and in Section 8.4.3 we investigate its clinical applicability.

¹ The current database is larger with 152 models at the time of writing. This is because the database was later extended with new cases. The evaluations in the following sections are based on the initial 120 models.




Figure 56: Capping of the branch geometry to create inflow and outflow boundaries for simulation.

8.3.1 Creating a Database of High-Resolution Carotid Flow Models

We performed flow simulations on the segmented vessels, including the full extracted bifurcation domain. Only the tips of the vessel branches were trimmed to define cross-sectional areas for the flow inlets and outlets, an example is shown in Figure 56. Each model has one flow inlet at the CCA and one outlet at the ICA. The ECA splits into an arbitrary number of sub-branches, for each of which we modelled an outlet. We setup the boundary conditions such that both primary branches may contain an arbitrary number of subbranches, which is necessary to correctly model the flow behaviour of the ECA, which tends to quickly split into multiple smaller arteries. All cuts are made to preserve the branching morphology as exactly as possible and furthest away from the bifurcation point. That way, the impact of the numerical boundary conditions is minimized in the crucial areas of the simulation, i.e., the bifurcation areas and narrowed sections of possible stenosis in the ICA.

The blood flow in the carotid artery tree was simulated by solving the underlying incompressible Navier-Stokes equations for Newtonian fluids. This was achieved in the FEM software COMSOL 6.0 [65]. The fluid-flow problem was solved on a rigid domain with constant fluid density $\rho_{blood} = 1000 \frac{\text{kg}}{\text{m}^3}$ [308]. We set the viscosity to $\mu = 0.00345 \text{ Pa} \cdot \text{s}$, which is typically used for blood as a Newtonian fluid as it is the lower viscosity limit of the shear-thinning models μ_∞ [210]. At the inflow, we prescribe a cyclic velocity function that mimics a cardiac cycle. We use an instantaneous inflow velocity boundary condition corresponding to a typical CCA temporal flow rate profile, see Richter et al. [352]. The mean flow is $432 \frac{\text{ml}}{\text{min}}$. This results in an average inflow velocity of $0.22 \frac{\text{m}}{\text{s}}$, which corresponds to the normal flow velocity measured in the CCA of adults [181]. The no-slip boundary conditions are prescribed on the vessel walls. The outflows are modeled with a boundary condition that simulates the resistance of the down-stream vessel tree. At each outflow boundary, we implement a varying pressure condition with a resistance approach to split volume flow through

the ICA and ECA with a 70% : 30% ratio. As long as a stenosis does not fully block the blood flow, this split is a good approximation [148, 313]. This is because a developing stenosis is first compensated through dynamic cerebral autoregulation, which is an important physiological mechanism that maintains constant cerebral blood flow [240]. The result is an increase in flow velocity inside the stenosed region. As Vignon-Clementel et al. [306], we specify the total vascular resistance $R_{tot} = P_{mean}/Q_{mean}$ with the mean pressure obtained as a combination of systolic (120 mmHg) and diastolic (80 mmHg) pressures. At a normal resting heart rate, systole is typically shorter than diastole. We use $P_{mean} = \frac{1}{3} \cdot 120 \text{ mmHg} + \frac{2}{3} \cdot 80 \text{ mmHg}$, which resembles the ratio between the systolic and diastolic phases we use at the inflow. The mean flow rate is $Q_{mean} \approx 7.2 \frac{\text{ml}}{\text{s}}$. According to the law of adding resistances in parallel, we get resistances $R_{int} = R_{tot}/0.7$ and $R_{ext} = R_{tot}/0.3$ for the ICA and ECA, respectively. Further branching of the ECA distributes the flow volumes according to their cross-sectional outflow areas A_{ext_i} for $i = 1, \dots, k$ in relation to the total external outflow area $A_{ext} = \sum_i A_{ext_i}$. After an initial cardiac cycle, the second cycle is used to evaluate the flow field parameters to mitigate initialization errors. We derive velocity, pressure, and the risk parameter WSS. Our volume meshes vary from $9 \cdot 10^5$ to $45 \cdot 10^5$ elements with an average around $20 \cdot 10^5$. The meshes are mostly composed of tetrahedral elements and include two boundary layers of prism elements to adequately resolve the velocity gradients near the wall. We apply linear finite elements representing the piece-wise linear numerical solution for both the velocity and pressure field. For equal polynomial order for velocity and pressure representations, the FEM is stabilized by streamline and crosswind diffusion. We use an adaptive time-stepping scheme with a maximum time step constraint of 0.01 s.

The mean simulation time for one artery, excluding pre- and post-processing, is 4 hours 50 minutes on an Intel® Xeon® Gold 5122 CPU at 3.60 GHz with 4 cores. The total pure computation time for all arteries was 30 days and 15 hours. While the time requirements for simulation could be reduced by faster hardware, the pre- and post-processing are also time-intensive. Each data set required an expert additionally 1–3 hours to prepare, setup the simulation, and export the results.

The export of a reliable wall stress parameter is dependent on the magnitude and the orientation of fluid stresses at the vessel wall. For our purpose, we focus on evaluating both, WSS amplitude as well as WSS alignment with the mean flow direction represented by the centerline. In the latter approach, the centerline's tangent is projected onto the vessel surface to achieve a consistently aligned tangential vector field in each vertex. For the details on longitudinal WSS evaluation and its advantage over the conventional amplitude-based indicators, we refer to Richter et al. [352]. Negative values of the longitudinal WSS




Figure 57: Visualizing the negative range of longitudinal **WSS** reveals reverse flow zones. These regions are highly indicative of turbulent vortices, a gradation criterion for stenoses. The first model shows signs of reverse flow and closer exploration indeed reveals the vortex flow. The second model is a counterexample of a stenosis of similar degree but without turbulence.

are indicative of regions with reverse flow, which often occurs after a stenosis due to turbulent vortices, see [Figure 57](#).

8.3.2 *Extracting Carotid Bifurcation Geometries*

While the **CFD** simulation can be externally computed to build the flow database, any new bifurcation geometry must be extracted from **CTA** imaging within the clinical workflow (**T₁**). We address **T₁** by building our flow visualization methods as an extension module to the integrated pipeline developed in [Chapter 7](#) [48]. The pipeline already provides a user interface for importing **CTA** volumes in stan-

standardized [DICOM](#) format, efficiently labeling the carotid bifurcation lumen, and extracting surface models. The pipeline can further be used to compute vessel tree centerlines using the method of Antiga et al. [6]. These centerlines also provide diameter information at every point inside the model. Using this information, the stenosis degree can be derived, following the [NASCET](#) guidelines [88]. In summary, we use the integrated pipeline to efficiently extract carotid bifurcation surface geometries, centerlines, and the inner vessel diameter, and derive the stenosis degree. During testing, we found that after a short introduction users can perform the segmentation and pre-processing steps for a new dataset within 2–10 minutes. The extraction is executed in a single application, it is fully guided by a user interface, and it requires no knowledge of the underlying algorithms. The extracted models are seamlessly imported into our extension module, where they can be instantaneously compared against the flow simulation database.

8.3.3 *Aligning the Carotid Data*

To visualize different carotid data comparably and measure their similarity, we need to align them, as the individual models are arbitrarily positioned. For consistent alignment, we consider the centerlines of different parts of the carotid artery: the [CCA](#), the [ICA](#), and the [ECA](#). For a detailed description refer to [Section 2.1](#) and [Figure 1](#). We denote the point of the carotid bifurcation as \mathbf{p}_B and define it as the point where the [CCA](#) centerline splits into the two carotid subbranches. The centerlines we receive from the extraction pipeline are already labeled accordingly. To align different carotid datasets, we must ensure that the corresponding parts are aligned, i.e., the [ICA](#) of the first dataset should not be aligned with the [ECA](#) of the second dataset, even if they are more similar in shape. The main feature point \mathbf{p}_B should be placed at the origin. Therefore, each carotid is translated so that the corresponding point $\mathbf{p}_B = (0, 0, 0)$. Next, we register the parts of the carotid artery. We have to find optimal rotations

$$\begin{aligned}
 R_{\alpha,\beta,\gamma} &= R_x(\alpha), R_y(\beta), R_z(\gamma) \\
 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{pmatrix} \cdot \begin{pmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{pmatrix} \\
 &\quad \cdot \begin{pmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (16)
 \end{aligned}$$

such that the parts are optimally aligned. For simplicity, we omit the angles: $R := R_{\alpha,\beta,\gamma}$. Since the morphology of the parts can be very dif-




Figure 58: Range of the weighting function applied to the carotid surfaces.

ferent, we create an artificial carotid artery and align the other carotid arteries accordingly. The artificial CCA consist of the points $(0, 0, -\lambda)$, the ICA of the points $(\lambda, 0, \lambda)$, and the ECA of the points $(-\lambda, 0, \lambda)$ with $\lambda \in [0, 0.5, 1, 1.5, \dots, 20]$. Decreasing the step size has no visual effect on the result. The set of points are denoted with $\mathbf{A}^C, \mathbf{A}^I, \mathbf{A}^E$. When aligning the parts (CCA, ICA, ECA) with their corresponding set of points $\mathbf{P}_i^C, \mathbf{P}_i^I, \mathbf{P}_i^E$ of an arbitrary carotid i to the artificial ones, we have to find an optimal rotation such that:

$$\arg \min_R \max \{d_H(\mathbf{A}^C, R \cdot \mathbf{P}_i^C), d_H(\mathbf{A}^I, R \cdot \mathbf{P}_i^I), d_H(\mathbf{A}^E, R \cdot \mathbf{P}_i^E)\}, \quad (17)$$

where $d_H(\cdot, \cdot)$ denotes the Hausdorff metric. We solve Equation 17 with the approach by Ugray et al. [302] provided in the Matlab framework. We solve this optimization problem for every carotid and apply the rotation yielding aligned data sets. Unfortunately, if we stop here, some carotid models are not well aligned. This is due to the high variance in the morphology at the end of the centerlines. Therefore, we weigh the points so that it becomes more important to align the points near the origin \mathbf{p}_B . We define a weighting function and multiply the points of the carotid data sets and the artificial carotid points by the weighting. The weight is defined based on the distance to the origin:

$$W(\mathbf{p}) = 5 \cdot \exp(-\|\mathbf{p}\|^2/10). \quad (18)$$

The rationale behind this weighting is that we emphasize the region around the critical landmark \mathbf{p}_B . The weighting drops from 5 to 1 when the distance exceeds 4 cm, see Figure 58. The results of the alignment are shown in Figure 59.

8.3.4 Visualizing the Alignment Space Neighborhood

To facilitate T2 we need to query the flow database for matching cases. First, we derive the geometric similarity S_g between the models. We define S_g as the remaining Hausdorff distance after registration, i.e., it is a byproduct of the branch alignment. We found that for finding matching cases, using this similarity of the vessel tree shape alone




Figure 59: Carotid geometries automatically aligned with our method. The ICA branches to the right, the ECA to the left, and the bifurcation point is centered.

does not yield ideal results, as it does not incorporate information on the vessel thickness. Therefore, we use the layout-independent information we gain from the model extraction pipeline as additional similarity measures, namely, the point-wise difference of the inner vessel diameter S_d and the difference of the stenosis degree S_s . As a standard clinical measure, the stenosis degree is an established comparison metric already. However, S_s alone can be misleading, as different stenosis shapes and locations can produce the same degree. Therefore, we further use the variation in vessel diameter of the CCA and ICA branches. We align the diameter profile of any two subjects at the BP and compute their point-wise distance. As the extracted branches differ in arc length, the longer profile is clamped to match the shorter one when computing S_d . When experimenting with test cases manually, we found that good search results ensue from merging the three metrics. S_g measures similarity regarding the branching geometry and layout. S_d measures similarity regarding the vessel diameter and *where* a stenosis is located, i.e., how far it is from the bifurcation. S_s measures similarity regarding the severeness of the constriction. To make the values comparable, we normalize each metric linearly to $\bar{S}_g, \bar{S}_d, \bar{S}_s \in [0, 1]$. For each metric, the maximum 1 is a theoretical identical match and the minimum 0 is assigned to the database model with lowest fit. Given a new input geometry, $\bar{S}_g, \bar{S}_d, \bar{S}_s$ are relative measures of how well a model in the database fulfills




Figure 60: The user interface of the visualization framework. (A) Toolbar with parameter selection and color map. (B) 3D view of the extracted vessel surface geometry. (C) 3D views of the selected comparison cases from the flow database. (D) Panel to control which comparison metrics are queried. (E) Bar charts showing how similar a match is regarding each metric. (F) Vessel surface maps giving an overview of the flow parameters.

each metric. The normalized metrics can be arbitrarily weighted to derive a total distance between two data sets. As a starting point for our evaluations, we use an evenly weighted linear combination $\bar{S}_{gds} = \frac{1}{3}(\bar{S}_g + \bar{S}_d + \bar{S}_s)$, which showed good preliminary results. A first measure of the performance of this matching is provided in [Section 8.4.2](#). The system is designed such that other metrics could also be added.

The candidates with highest similarity \bar{S}_{gds} are provided for visual comparison. We found that only showing a single match may not be the best approach, as the similarity search is not a trivial task. Different candidates may be relevant, depending on what the user is looking for. Therefore, we explicitly visualize the subset of the flow database closest to the new model regarding \bar{S}_{gds} . We call this subset the alignment space neighborhood. As we do not know how many samples are required, the user can choose the neighborhood size arbitrarily. To show how similar a model is regarding each metric, we visualize a bar chart for each candidate, where each metric \bar{S}_g , \bar{S}_d , \bar{S}_s is shown by one bar. We use these bar charts, as they are instantly readable even at a small size. The axis values are omitted, as the numerical value of each bar is meaningless and only the relative length of each bar is relevant. We sort the candidates automatically and display them below a view of the newly extracted 3D carotid model, see [Figure 60E](#). The candidates are shown with decreasing

fit from left to right. We implement an option to toggle individual metrics on or off, which re-sorts the candidate list. This allows, for example, to run the following queries: *Is there a model with the same stenosis degree? Give me models that are very similar in shape. Give me a model with a similar diameter profile.*

We quickly noticed that when displaying the alignment space neighborhood, a visual representation of each candidate's geometry is highly beneficial for navigation. Providing visual cues about the geometry and flow parameters increases recognizability and accelerates finding the ideal match. An overview of the flow parameters also allows the efficient identification of different types of models. For example, the user might want to get a best and worst-case comparison.

When visualizing ten or more 3D vascular models, scalability becomes an issue. Visually assessing multiple complex 3D renderings is problematic, as small depictions become hard to read, 3D interaction would be required to view the full surface, and self-occlusion of the branches obstructs the view. To address these shortcomings of 3D renderings, vessel maps have been shown to be an effective tool. They can provide an instant overview of surface fields, require no interaction, and allow the comparison of many instances since they remain readable even at small scales [343]. Therefore, we show flattened maps of the vessel wall surfaces in the candidates view. We computed a surface map of every instance in the flow database and visualize these maps below the corresponding bar chart, see Figure 6oF. We use the method developed in Chapter 5 to compute the maps, which is a global surface parameterization technique aimed at creating a single patch for vessel trees. The method automatically cuts and flattens tree-like surface meshes with an arbitrary number of branches. It uses vessel cuts that are positioned on one side of the cylindrical vessel branches, enabling intuitive 2D views of the unrolled vessel wall. The approach preserves vertex and connectivity information, facilitating mapping of data fields defined on mesh vertices. It minimizes area distortion while maintaining proximity to the original branch layout. The area optimization is also useful to identify stenosis candidates. We use the resulting maps to show the flow-derived surface parameters. These are the WSS and the reverse flow field that we compute from the stress vector and the vessel directionality, as described in Section 8.3.1. We pre-render the maps as 1000×1000 pixel textures, which means they can be quickly loaded and displayed. The textures encode a scalar value of each surface field. During the rendering, we apply a colormap to the surface parameter. The active surface parameter, the colormap, and the visualized value range are chosen by the user. Any perceptually uniform colorCET colormap [162] can be applied. The user can zoom into the candidate list, to display fewer candidates and read details on the maps, and zoom out to show more candidates. How many candidates are sensible depends on the screen size. With the

maps, specific queries can be made at a glance. For example, during testing we encountered the following queries: *Is there a similar model that shows signs of turbulence? Give me a similar model with high WSS at the bifurcation. How fast is the flow in a slightly worse stenosis?*

8.3.5 Comparative Visualization of Multiple Vascular Flow Models

To enable **T3**, we need to facilitate a comparison of the new case with any selected matching models from the flow database. Visual comparison can be implemented by a superimposition of instances, juxtaposition, interchange, or explicit encoding [152]. In our case, the complex vascular geometry excludes superimposition as an option, as the view would just become cluttered. We do not know how many comparisons are required and the user should be able to inspect each geometry individually and simultaneously. Therefore, we use a juxtaposition to enable a comparison of the 3D geometries and flow fields. By clicking on a vessel map or the bar chart above, a view of the corresponding 3D model is toggled on or off. This way, an arbitrary number of models can be displayed, which are automatically laid out in a grid beside the inspected case, see Figure 60C. A trackball-style interaction can be used to interact with the vessels. An apparent disadvantage of these side-by-side views is that the 3D interaction becomes cumbersome for multiple models. For a more fluent interaction, we integrated a *synchronized* juxtaposition that can be enabled on demand. In the synchronized view, we transform the displayed models into the alignment space where the CCA, ICA, and ECA branches are registered. Then, the camera is synchronized between all views, i.e., all views are simultaneously rotated, zoomed, or translated.

The flow information is mapped on the model to address **T4**. The same colormap that is set for the vessel maps is shared across all depicted surface models. To analyze the systolic or diastolic velocity field, we need to provide a view inside the vessel. If instead of a surface field a velocity field is chosen, we render streamlines of the flow and contextualize them with the backside of the vessel wall. The absolute velocity is color mapped on the lines. We chose this streamline encoding, because primarily the single time step of peak systolic velocity is relevant to the users and the streamlines reveal features of interest, such as poststenotic vortices. For the integration, we use the 5th order Runge-Kutta method [170, 266] with a maximum propagation of 0.2 m, which is about the length of one carotid bifurcation model. The streamline seeds are placed in 300 random cells of each carotid volume mesh to distribute them across the mesh. The streamline geometry is pre-computed and cached for each case. Therefore, they can be immediately shown once a case is chosen for display or the time steps are switched.




Figure 61: The adaptive streamline clustering ensures that features are not lost to occlusion when the 3D rendering is small. (a) The inner streamlines are hidden, giving a false impression of the velocity. (b) Clustering reduces the number of active lines while retaining flow features.

8.3.5.1 Adaptive Streamline Clustering

If multiple flow models are displayed their viewports become increasingly smaller, resulting in streamlines overlapping each other. At some point, this can obstruct the view of the velocity information at the inner layers of the flow field. Further, the rendering of many high-resolution lines can lead to low framerates, depending on the rendering hardware. Therefore, we introduce an *adaptive* streamline clustering that groups similar lines with decreasing viewport size. To cluster the streamlines of the blood flow, we need to determine the similarity between different streamlines. To do this, we compute a distance matrix based on the same Hausdorff metric as utilized in Section 8.3.3, followed by an AHC. We show 300, 200, 100, 30, or 10 clusters for up to 1, 4, 9, 16, and > 16 juxtaposed views, respectively. We found these thresholds to provide good results, where the rendered clusters cover the domain and display relevant flow features but do not excessively occlude each other. Using more thresholds for even finer granularity is possible, but seemed redundant in our testing. Like the streamlines, the clusters are pre-computed and can be instantly applied to every model in the database. Through adaptive clustering, we ensure that important features are not lost if the views become smaller. We show this effect in Figure 61. Using clustering is more feature-sensitive than using fewer lines, where some regions would be undersampled. The clustering ensures that lines are chosen which are representative of a collection of stream traces, preserving larger flow features.

8.3.5.2 *Probing*

Lastly, **T5** requires functionality to probe the flow field at user-defined locations and extract the maximal local flow velocity. We implemented a type of flow lens, which can be activated when the flow velocity is displayed. When the flow lens is enabled, we shoot a ray into the scene of the active viewport, originating from the cursor position on the screen. We determine the first hit position of the viewport's vascular model or, if no hit is found, the closest surface point to the ray. Around this point, we create a sphere with a user-adjustable radius. The sphere is used as an implicit function to filter the streamlines and we crop the line geometry outside the sphere. We filter the actual 3D geometry, not the image space, which allows us to accurately determine the highest flow velocity inside this region. The extracted value is displayed next to the probed region, see [Figure 62](#). The probe can be continuously moved by dragging the cursor, allowing smooth exploration of the flow field and determining the maximal flow velocity in all required regions, e.g., inside the stenosis, but also in the [CCA](#) or behind the stenosis.

8.4 EVALUATION

We evaluated the numerical soundness of the comparison-based flow analysis to get a baseline of how well flow parameters can be predicted with the system. We also tested the interactive framework with five physicians working in stroke care.

8.4.1 *Performance Evaluation*

To test how the interactivity of the system scales with the database size, we measured the performance of the operations that are dependent on the number of underlying flow models. These are (1) the computation of similarity metrics when loading a new geometry and (2) the sorting of the candidates list when changing the weighting of the metrics. The metric computation has linear time complexity, as it compares the input to each database entry once. The candidate sorting computes the weighting and then sorts a list with one floating point entry per data set. Both operations should not run into performance issues. In [Table 2](#), we provide benchmarks for the two operations on a system with an Intel® Core i9-9900K CPU @3.6 GHz and 32 GB of DDR4 RAM. For the database sizes > 120 , we have added additional models by duplicating the existing ones.

Technically, the initial loading of the database entries into memory is also dependent on the database size, but this is only noticeable once when starting the tool. Note that the actual flow data, mesh geometry, and field data, are only loaded for a specific model when




Figure 62: A probe lens can be interactively dragged over the vessel to focus on the flow in a specific region. The maximal flow velocity within the selection is extracted automatically.

the candidate is selected. All remaining operations (displaying maps, selecting candidates, exploring flow features, changing the active scalar field, changing the colormap) are not dependent on the database size.

8.4.2 Numerical Evaluation

We want to quantify the current error of the prediction when compared to performing a case-specific CFD simulation. For this purpose, we applied a standard 80/20 split on the database, leading to 96 remaining models in the prediction database and 24 test instances. The test cases were randomly selected. For each test instance, we measured the deviation between the simulation ground truth and the system's prediction. For the prediction, we use the database geometry with the highest similarity metric \bar{S}_{gds} , which for simplicity we will call the *best match*. Furthermore, to assess the robustness of the similarity metrics, we compare \bar{S}_{gds} against choosing the best match based on

Table 2: Results of performance testing of operations that are dependent on the size of the database.

Database Size	Computing Metrics	Sorting Candidates
120 flow models	0.01195 s	0.00025 s
240 flow models	0.02335 s	0.00043 s
1200 flow models	0.11433 s	0.00157 s
12000 flow models	1.26789 s	0.00956 s

the smallest Hausdorff distance, which is a widely applied standard for matching geometric objects [15, 267]. In summary, for each of the test instances, we find the best matching model in the database, once using our combined similarity metric and once using the Hausdorff distance. For each match, we compute how the flow properties differ between the ground truth and the prediction. For our comparisons, we focus on flow parameters relevant to clinical contexts as well as flow field analysis. We determined the prediction error regarding velocity, vorticity, *WSS*, and pressure. The systolic peak velocity is highly relevant as it is used for stenosis grading. The occurrence of poststenotic turbulent or vortex-like flow is also used as a secondary marker in clinical stenosis grading [76]. We measure the vorticity as a property of the flow structure. Note, however, that vorticity has limited expressiveness, as it is not only influenced by vortices but also the velocity and shear flow near vessel walls. *WSS* and pressure measure the forces acting on the vessel walls and have been linked to plaque formation in the carotids [186].

Comparing the flow properties of a best match to the ground truth is not a straightforward task, since the wall geometry and layout of the test instance can differ from the match, even after aligning the models. A match may be very similar in terms of the diameter and stenosis degree, leading to a high similarity score, but have branches of a different length and/or layout. This means a direct comparison of the flow fields is impractical. To enable a sensible comparison, we use predetermined samples of the flow field at the clinically most significant and uniformly recognizable landmarks: inside the *CCA* close to the bifurcation, inside the *ICA* stenosis, and shortly after the stenosis. These sample positions are indicated in Figure 62, positions (b), (c), (d). They are the relevant locations examined in carotid ultrasound diagnostics. The hemodynamic parameters at the sample location are averaged within a 5 mm radius to prevent noise from affecting the measurement. We also measure the peak velocity, i.e., the maximum value within each sample region. For every test instance, we sample the three landmarks at the peak systolic and end-diastolic time points, as these are the two time points used in clinical assessments. By computing the absolute deviation for each sample location and time point

between the case-specific simulation and the prediction, we determine the prediction error.

8.4.3 Structured Interviews

To validate if the tasks **T1–T5** can be performed, we conducted qualitative interviews with five physicians (P1–P5; two female, three male; ages 30–63). They have 17, 32, 13, 9, and 15 years of working experience in clinics for neurology or radiology. P1 and P2 have also taken part in earlier discussions and the task analysis.² Also, they acquired the **CTA** data and validated the lumen/plaque segmentations we used to create the flow database. None of the participants had used the developed software framework before.

We conducted individual in-person interviews lasting about 60 minutes, where we first acquainted the participants with the concept of the similarity-based hemodynamics analysis. Then, we demonstrated how a model can be extracted from **CTA** using the integrated pipeline and how the wall geometry can be viewed in our extension module. We explained and showed all exploration features, e.g., how the neighborhood can be filtered, how candidates can be selected, and how the flow parameters can be analyzed. After the introduction, the participants were asked to analyze three new **CTAs**, which were not included in the flow database. One represented a healthy patient, one contained a mild, and one a severe stenosis. We asked the participants to extract one geometry from each **CTA** volume (**T1**), find one or multiple matching flow models (**T2**), and compare them to the extracted geometry (**T3**). We further asked which flow parameters they could gather from their selection (**T4**, **T5**) and if these confirmed their expectations. During these case studies, we employed a think-aloud protocol. We followed the case studies with a questionnaire, where we asked questions regarding the fulfillment of each task and the comprehensibility and usefulness of each visual encoding (vessel maps, bar charts, color maps, streamlines, probe). The participants rated each question on a five-point Likert scale (–, –, ◦, +, ++) and we noted down verbal comments regarding each item.

8.5 RESULTS AND DISCUSSION

The results of the numerical evaluation are summarized in [Table 3](#). For the majority of the samples, the predictions are highly accurate. For the peak velocity, we achieve a median error of $0.06 \frac{\text{m}}{\text{s}}$ using the \bar{S}_{gds} matching metric. For the average velocity, the median error is $0.04 \frac{\text{m}}{\text{s}}$. Our mean errors are $0.12 \frac{\text{m}}{\text{s}}$ (peak velocity) and $0.08 \frac{\text{m}}{\text{s}}$ (average velocity). The magnitude of these velocity errors can be better interpreted by looking at the clinical measurement variability. The mean

² Their identifiers, P1 and P2, are consistent with [Section 8.1](#).

Table 3: Results of the numerical evaluation. The deviation from the ground truth is measured individually per sample – for each position and point in time. The median, mean, and maximum of these errors are shown. Hausdorff: The prediction was selected based on the Hausdorff distance between aligned models. \bar{S}_{gds} : The prediction was selected based on the proposed combined similarity metric. For each measured parameter, the observed value range in the ground truth samples is given as a reference for the scale of the absolute errors.

		Peak Velocity (m/s)	Average Velocity (m/s)	Average Vorticity (1/s)	Average WSS (Pa)	Average Pressure (Pa)
Absolute Error	Median (Hausdorff)	0.18	0.09	0.18	2.38	220
	Median (\bar{S}_{gds})	0.06	0.04	0.1	1.08	121
	Mean (Hausdorff)	0.34	0.15	0.44	5.22	498
	Mean (\bar{S}_{gds})	0.12	0.08	0.22	2.43	337
	Maximum (Hausdorff)	3.75	1.59	5.99	74	4487
	Maximum (\bar{S}_{gds})	0.82	0.79	1.93	21	3605
Ground Truth Minimum		0.1	0.05	0.03	0.3	10177
Ground Truth Maximum		4.5	1.95	6.49	80.3	32463

interobserver variabilities for ultrasound assessments of the CCA and ICA peak velocities were observed to range from -0.08 to $0.14 \frac{m}{s}$ [209]. The limits of agreement for individual patients were found to be -0.68 and $0.85 \frac{m}{s}$ [209]. Our limits of agreement for individual samples are -0.72 and $0.82 \frac{m}{s}$ for the peak velocity and -0.79 to $0.28 \frac{m}{s}$ for the average velocity. This means, for the clinically relevant flow velocity, our prediction errors are within the currently tolerated limits. The system showed to be able to find similar cases and predict flow velocities with errors comparable to the disagreement between physicians when performing velocity measurements.

Generally, the median error of all assessed parameters is low with regard to the typically observed value range. The mean error, however, in all cases is higher than the median error, indicating a skewed distribution due to some predictions with disproportionately larger errors. This assumption can be verified by looking at the boxplots of the error distributions, see Figure 63. For all flow parameters, at least half of the samples only show small errors close to 0. However, there are outliers with a larger error which skew the distribution. 25% of the peak velocity samples and 16% of the average velocity samples show an error larger than the observed interobserver variability of $0.14 \frac{m}{s}$. These cases seem to be insufficiently covered by the pre-computed flow fields, leading to the conclusion that the database will need to be extended.

Lastly, we observe that the combined similarity metric \bar{S}_{gds} is consistently better at choosing a prediction model than using the Hausdorff




Figure 63: Boxplots of the error between the ground truth simulation and prediction. H: The prediction was selected based on the Hausdorff distance. \bar{S}_{gds} : The prediction was selected based on our combined similarity metric.

metric (H). The median error is significantly lower for all tested parameters and the spread of the error distribution is smaller. These results indicate that the relevant flow parameters are more dependent on the [CCA/ICA](#) vessel diameter, stenosis location, and stenosis degree than on the shape or layout of the vessel branches. A core observation of this study is that the clinically relevant blood flow parameters can be predicted by a vessel model with a similar stenosis (regarding location, length, and diameter), even if the branch layout is different and the number of modelled [ECA](#) branches varies. Our proposed algorithm to build the alignment space appears to accurately match cases with similar flow in the relevant areas, i.e., inside the [ICA](#) and potential stenoses, even if no model exists that has a globally equivalent surface geometry. This observation may aid future research regarding comparison-based flow analysis, as it shows that the database does not need to cover all potential geometric variations of the target anatomy but can be constrained to focus on variations of the crucial features, such as the stenosis position and shape.

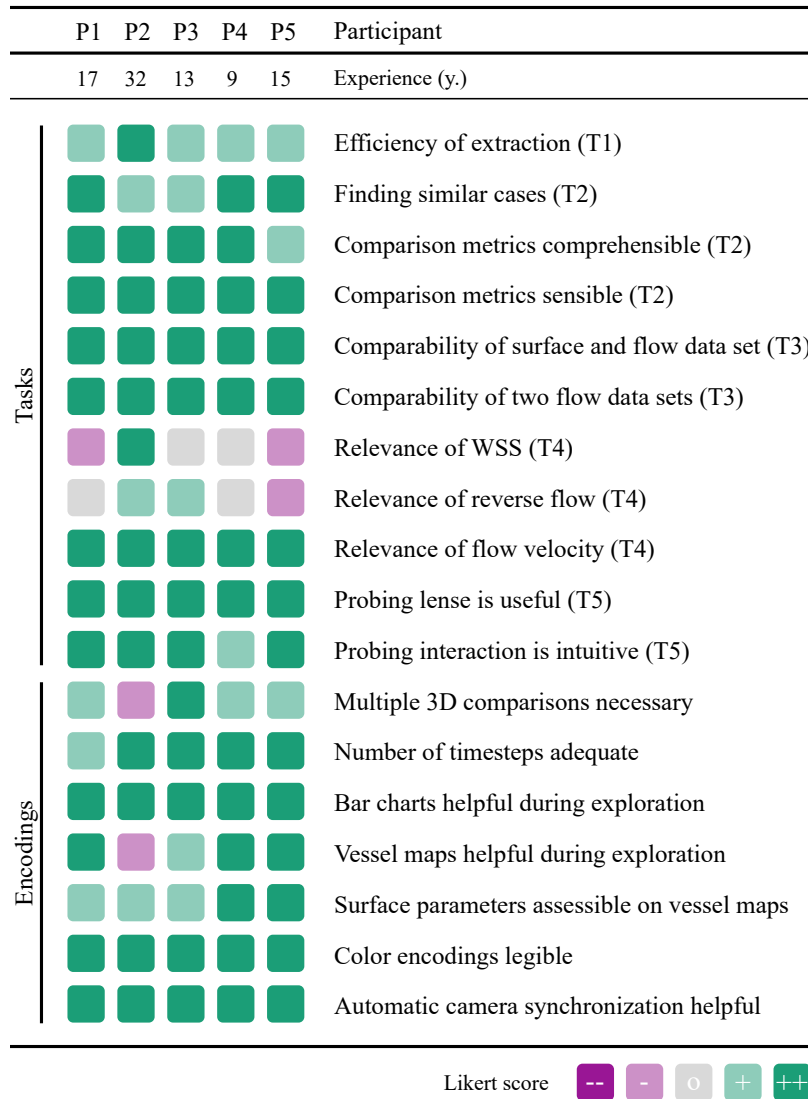


Figure 64: Likert scores assigned in the questionnaire. Each box represents the answer of one physician.

8.5.1 Interview Results

The Likert score results of the questionnaire are shown in Figure 64. We report on frequent comments and feedback from the interviews. The overall impressions of the participants were highly positive. The processing pipeline was well received and the model extraction (T1) was deemed efficient enough to be applicable.

“I highly appreciate that the software is usable as is. It is very well integrated. I can import a DICOM image and very quickly get an impression of the internal flow field.” (P3)

The participants were generally able to find similar enough comparison cases (T2).

“Considering you only have about 100 models, I am surprised how well the pairing works. I think for some we could get even better results, but that is just a question of extending the database.”
(P4)

The comparison metrics were understood and all experts noted they found the metrics sensible. P1 emphasized, *“Especially the diameter makes total sense. I think the exact progression of the branches is not so important, so it’s nice that I can turn that filter [shape similarity] off.”* We observed that typically one or two and rarely more than four comparison cases were used simultaneously. Two participants noted they would only want to see the closest case and that *“the less interaction is required the better”* (P2). However, opinions on this matter diverged. P5 noted, *“It highly depends on the variance. If I can get one really good match it is sufficient. But otherwise, I definitely need three or more similar cases for comparison.”* Also, the option of comparing a slightly worse stenosis to visualize the possible progression was mentioned.

Regarding the clinical relevance of the shown parameters, we gathered the impression that as long as no clinical guidelines exist, using WSS and reverse flow is only of supplemental importance. For instance, P5 said, *“It is nice to have but cannot be the basis for a treatment decision.”* All experts agreed that the peak flow velocity is currently of the highest clinical relevance. P2 argued, *“I would still definitely keep the other parameters, especially WSS is increasing in importance. Having it available is never a drawback.”* When asked about the temporal resolution, every interviewed expert underlined that they are almost exclusively interested in the peak systolic flow and would not see the upside of including more time steps.

Furthermore, all participants stated that the bar charts used to depict the model similarity are highly useful to judge the appropriateness of a case. The vessel maps were also positively received as a way to get an overview of the displayed candidates.

“I used the maps extensively to choose a model that reflects my presumptions about the flow. I want to see them and make the decision which model should be compared.” (P5)

“The maps give a good impression of the overall form already.”
(P4)

When assessing the 3D models (T3, T4), the participants mentioned that the streamlines were helpful to judge if turbulence occurred. All five physicians remarked that the ability to synchronize the model orientations is highly beneficial. For example, P3 said, *“The linked perspective is extremely helpful. I first struggled to align all the cases manually, it is so much faster with the linking. The branches are really nicely aligned.”* They also stated that they found the probing interaction (T5) very useful and intuitive.

“The probe is very helpful, also that I can change its size. It’s like an enhanced ultrasound, I can peak into every branch and get a quick reading of the velocity, this is exactly what I need. Dragging the region is very smooth and the visual presentation is elegant.” (P2)

After the interview, three participants expressed immediate interest to deploy the tool and proposed to perform a validation study to compare the results against the measured flow. Notably, P4 described a further application scenario that we did not consider at the beginning:

“Every day, I have to decide if we transfer stroke patients from smaller clinics to the university hospital. They cannot do a Doppler ultrasound there, I have to judge the stenoses on the CTA alone. This is also a question of resource management, we cannot transfer everyone and very often we transfer unnecessarily, most patients do not need an endarterectomy [surgical plaque removal]. If I could get a good prediction of the flow from the CT images, that would vastly improve these decisions. I would use this on a daily basis.” (P4)

8.5.2 Limitations

In many of the test instances of the numerical evaluation, we observed that the approach produces suitable results. If vessel wall models with high similarity in the crucial areas are available in the database, the framework accurately predicts the flow properties of a case-specific CFD simulation at the clinically relevant landmarks. However, we also observed test instances with a higher prediction error. These instances are visible mainly as outliers in the distribution of the sample errors. Unusually large errors occurred if no model with a high similarity metric was available. Keeping in mind that the database we used during the testing only involved 96 models, we still consider the results of the numerical evaluation a positive outcome. We demonstrate that with the existing pre-simulated cases, we can already predict the clinically relevant flow parameters in a large portion of carotid stenosis cases and achieve velocity errors similar to the interobserver variability of physicians. However, for a usable application, the current database needs to be extended. This extension should not only increase the database size but also improve the data variation. When selecting CTA scans, we only differentiated them by stenosis degree. To fully sample the shape space of ICA stenosis, other geometric factors should be considered. This includes the stenosis position, length, and also local properties such as the regularity of the internal wall structure. We also noticed that since the velocity does not increase linearly with the stenosis degree, we require further artery models with severe stenoses to achieve more precise results in these edge cases. One option

to improve the database size and variation would be to introduce artificial modulations to the data. By varying the position and degree of stenoses, their shape space could be systematically sampled. With artificial geometries, however, it could be difficult to reproduce exactly how stenoses develop. In general, more testing will be required to form a better understanding of the relation between the database size and quality of results.

We tested one set of similarity metrics to relate vascular geometries and CFD results. While the developed combined metric \bar{S}_{gds} showed better matching results than the Hausdorff distance, we did not test other weightings or other measures. It would be conceivable to use more refined approaches. For instance, when comparing the vessel diameters, \bar{S}_d is based on a fixed alignment, which could benefit from techniques like the Fréchet distance that would optimize the alignment of the diameter profiles. When the database of flow models grows and geometric features become more nuanced, the discriminative power of the metrics may also be insufficient. Overcoming this issue would require an adaptation of the metrics that takes the actual surface geometry more into account.

CFD necessarily needs to employ simplifications. Standard boundary conditions model realistic but averaged parameters that might not apply to every patient. Some CFD studies use patient-specific boundary conditions, such as inflow velocities that were recorded using Doppler sonography. Using such boundary conditions requires a measure first, which defeats a substantial advantage of computed hemodynamics – dynamic flow data can already be generated from static image data, without additional procedures. Such simulated flow provides an estimate of important hemodynamic parameters from a single CTA scan but necessarily requires the use of standardized boundary conditions. Many of the resulting parameters, e.g., the peak velocity and the occurrence of poststenotic turbulence, are key factors for stenosis evaluation but cannot be directly derived from CTA. An estimate of the blood flow can already provide a substantial benefit, especially where measurements have not been or cannot be conducted, for example, if a stroke patient was delivered to a hospital where Doppler ultrasonography is not performed. Ultimately, the accuracy of the simulations determines their viability. During the interviews, we received the feedback that the validation of flow-derived parameters is yet to be improved. First studies have shown that velocities measured with Doppler sonography can be reproduced with CFD [116, 287]. However, larger studies with more subjects are pending. While many studies exist that link WSS to plaque progression [124], further long-time studies are necessary to make flow-derived parameters part of clinical guidelines. The feasibility of extensive clinical CFD studies depends highly on the applicability of hemodynamic simulation in real-world contexts. The method described in this work addresses the

latter, specifically the challenging time and resource-intensive nature of these simulations. Our approach should be seen as a complement to research targeting the validity of hemodynamic simulation.

Lastly, we need to note that the extraction efficiency (T_1) is key for applicability. Currently, the extraction speed highly depends on the segmentation step. If the labels predicted by the neural network are good, this step is fully automatic. However, we still see cases where manual corrections to the segmentations are required, which can slow the whole process down. Fortunately, segmentation networks have shown constant improvements in the past decade, minimizing necessary manual corrections.

8.5.3 *Opportunities and Future Work*

Two interviewed physicians independently mentioned that the flow parameters of matching models could be interpolated depending on their similarity distance. This could provide more accurate values and projecting them on the new geometry would make them directly readable on the extracted vessel. However, we note that this approach might also be riskier than the purely comparison-based analysis, as the interpolation could be misleading if no good match is available. Similarly, first studies have shown that hemodynamic surface parameters like *WSS* can be predicted using *CNNs* [292]. The disadvantage compared to our approach is that the prediction model is obscure to the observer, i.e., it is not clear where the information comes from. In our case, the user can immediately see if no reasonably similar vessel geometry is found, instead of getting a potentially misleading prediction.

In the future, we intend to extend both, the flow database and the simulation domain. The latter was proposed by one study participant, who noted that extending the region to include intracranial vessels would be highly beneficial. Measuring intracranial flow velocities with Doppler ultrasonography is extremely challenging due to the distortion of the ultrasound signal by the skull. Yet, it is still a widely adopted procedure, as non-invasive alternatives are missing. Likewise, we plan to apply the similarity-based flow analysis to aneurysms, for many of which only static imaging is available. Similarity-based analysis could also help to find similar cases with other information than flow. For instance, from an aneurysm database, geometrically close cases could be queried for which a rupture status is known. This would allow to derive a geometry-based rupture probability.

8.6 CONCLUSION

We presented the first approach for similarity-based visual analysis of hemodynamics. While we focus our implementation and testing

on blood flow in carotid bifurcations, the concept could equally be transferred to other transport problems, like intracranial circulation, blood flow in the aorta, or hemodynamics of aneurysms. We propose a method to form a vascular alignment space that registers vessel branches in a common coordinate system. Using clinically motivated similarity metrics, the distances between a new vessel model and models in a database can be derived. For the models in the database, we pre-computed flow fields, relevant flow-derived parameters, and velocity streamlines. Given a new carotid bifurcation, the most similar geometries with computed flow can be queried from the database and visualized, providing an instantaneous approximation of the blood flow. We provide an overview of fitting candidates using map-like visualization and facilitate simultaneous exploration and comparison of flow in selected models. The framework is built as an extension to the carotid analysis pipeline described in [Chapter 7](#), which provides an efficient toolset for the extraction of carotid bifurcation models from CTA images [48]. It uses the vessel wall map developed in [Chapter 5](#) to create scalable overview visualizations of a matching subset of the model database. We show that for models in the database that are geometrically close, clinically relevant flow parameters, like the peak systolic velocity, are highly similar. The user study confirmed that the similarity-based flow analysis could benefit critical clinical processes, which demand fast access to relevant parameters such as flow velocity. In conclusion, we investigated a possibility to combat one of the strongest weaknesses of CFD in medicine – due to the time-consuming and resource-intensive nature of hemodynamic simulations, they are difficult to directly apply in clinical workflows. Our hope is that the insights gathered in this work will support efforts to make hemodynamic simulations more accessible and useful for clinical decision-making.

Part V

CONCLUSION

9

Conclusion

CONCLUSION

In recent years, data and simulation-based methods have gained increasing popularity as an approach to personalize and optimize medical diagnostics and treatment [27]. For the prevention of ischemic strokes, information gained through typical imaging techniques like CTA, MRA, and ultrasonography could be supplemented by hemodynamic simulations based on patient-specific arterial wall models [186, 198, 352]. To explore and analyze the multidimensional hemodynamic data gained with these simulations, visualization methods are indispensable [234].

Arterial stenoses are a frequent cause of ischemic stroke – the most common type of stroke [80] – yet the visualization of blood flow in developing stenoses has only been sparsely covered. This thesis focuses on two major areas. First, blood vessel and flow visualization methods are adapted and refined for the exploration of carotid stenoses, filling a gap in the literature. The carotid arteries serve as a basis for the contributions of this thesis, as they are a primary factor for stenosis-induced strokes [76]. The second focus of this work lies in the integration of distinct sophisticated processing and analysis stages, with the objective to increase the real-world applicability of the developed methods and to address persisting clinical transfer challenges.

9.1 SUMMARY

This thesis is composed of four parts. After an introduction to the topic in Part i, Part ii covers the relevant medical and computational background. CFD simulation methods, processing steps, as well as the state-of-the-art in vessel and blood flow visualization are reviewed. The concept of vessel maps is introduced, including a comprehensive survey of methods to dimensionally reduce 3D vascular structures to map-like depictions. Vessel maps emphasize data display over anatomical accuracy and utilize techniques like path-centered or shape-deforming methods. A taxonomy of vessel maps is proposed that categorizes these visualizations based on four types of underlying data: medical volume images, blood flow volume meshes, vessel surface meshes, and centerline graphs. The techniques to create vessel maps and their applicability are discussed, considering visual design properties, domain tasks, and dependencies. Vessel maps are valuable for cardiovascular analysis, aiding disease monitoring, treatment planning, and research.

Part iii focuses on developing techniques for the visual exploration and analysis of vascular stenoses. An automatic technique for cutting and flattening vessel tree geometries is introduced, featuring natural cuts that maintain visibility and positioning for analysis. The method is demonstrated using simulations of fluid-wall interactions in human carotid arteries, but its applicability extends to similar tree-like structures. Unlike other approaches, the technique operates solely on the surface mesh input, without additional dependencies, and produces a boundary-free quasi-isometric parameterization that preserves the original vessel shape. The cutting process is robust, ensuring valid disk cuts with connections at branch locations, and the flattening results in a single connected patch. The simplicity of the 2D representation can aid in creating comparative views of datasets for tasks like studying patient groups or analyzing CFD parameter variations.

After focusing on the vessel wall features, techniques are developed utilizing visualizations of the simulated internal flow field to improve the diagnostic assessment of carotid stenoses. An application with interconnected views is created, allowing physicians to examine vessel morphology alongside simulated hemodynamics. The study shows that concepts from flow visualization research can be effectively applied to specific tasks in this context. Of particular significance are appropriate exploration techniques alongside visualization methods. A vessel map encoding the diameter profile of the two carotids proved to be highly useful for gaining an overview and identifying key points. The flattened vessel tree aids in spatial navigation and data probing. Embedding contextualization in the visualizations is vital, as it can assist in localizing stenoses, exploring additional features like plaque, and performing on-the-fly validations of the underlying models.

Some further considerations emerged when developing these tools for stenosis analysis. Clinical use cases require both, a combination of local flow pattern views near stenoses and a broader global perspective. The temporal dimension was found to be less significant than anticipated, suggesting that focusing on systolic and diastolic values might be more practical than including numerous time steps. Additionally, dealing with the substantial data generated by CFD can be challenging for medical diagnostics. Prioritizing core user tasks and tailoring methods to minimize interaction is essential.

Part iv of this work approaches challenges in the applicability and transferability of methods, especially considering clinical use cases. Four major bottlenecks in the transfer of advanced vessel visualizations are identified: user trust in the visualizations, the efficiency of the data processing, the compatibility of the underlying tools and file formats, as well as the validation scale applied to new methods. An integrated pipeline is built to approach these issues. The pipeline combines all preprocessing and visual analysis stages for the example of carotid stenosis grading. Advanced techniques, such as neural

network-based segmentation and interactive visual analysis methods, are combined in a single framework. A significant takeaway of the ensuing study is that actively visualizing the processing steps and enabling retrospective control allowed users to better understand the different processing stages and assure the correctness of what they see. In some of the cases, the stenosis degrees determined by the physicians testing the framework differed considerably from the degree given in the patient's radiology report. Remarkably, when asked which value they thought to be more accurate, all clinicians claimed they would rather trust the application than the report, as they found the numerical derivation comprehensible and were able to assure the correctness of the geometry.

The integration of advanced data processing and visualization is challenging. The developed system required considerable collaborative and time-intensive efforts to connect previously separate solutions for, e.g., model extraction and automatic stenosis grading. This integration, however, has the potential to make the adoption of new methods that enhance medical imaging analysis feasible. A key consideration is determining where users can interrupt and adjust automatic processing. Balancing result accuracy with user time and effort is crucial, particularly when dealing with complex algorithms in medical applications. The results of this study suggest that if extensive processing is essential for new medical visualization tools, such as image segmentation, any required user interaction must be fast enough to not disrupt clinical workflows significantly. With the increasing capabilities of machine learning models like CNNs, there is an opportunity to replace time-intensive tasks in this process. The developed framework was applied to reconstruct more than 150 geometric models from CTA scans, which were also made openly accessible [347].

Lastly, a novel approach to similarity-based visual analysis of hemodynamics is introduced, which can circumvent complex individual flow simulations. The application extends the integrated pipeline, focusing on blood flow in carotid bifurcations. The concept itself, however, would be applicable to other transport problems like intracranial circulation, aortic blood flow, or aneurysm hemodynamics. The method involves creating a vascular alignment space that registers vessel branches in a common coordinate system. Similarity metrics are used to compute distances between new vessel models and a database of models. The database contains pre-computed flow fields, flow-derived parameters, and velocity streamlines. Given a new carotid bifurcation, the most similar geometries with computed flow can be retrieved from the database and visualized, offering a quick approximation of the blood flow properties. Since the framework extends the integrated carotid analysis pipeline, efficient tools for carotid bifurcation model extraction from CTA images are already provided. A study was performed to assess the prediction errors of this system regarding

the clinically relevant flow parameters. The results show that if geometrically close models exist in the database, the relevant flow parameters, such as peak systolic velocity, can be predicted with the necessary accuracy. With 120 pre-computed flow models, the measured prediction errors are already in the same range as the interobserver variability that exists in carotid Doppler ultrasonography. It also showed that a good approximation of the relevant blood flow parameters is not dependent on a globally similar surface geometry, but features such as the stenosis degree and vessel diameter are more important for similar flow behavior. This observation implies that for similarity-based analysis to succeed, the database does not need to cover all anatomic variations (branch layouts, vessel lengths) and should rather focus on variations of the features that will dictate the relevant flow parameters (vessel diameter, stenosis degree, stenosis length). In a user study involving five radiologists and neurologists, the potential of similarity-based flow analysis to enhance clinical processes requiring quick access to these flow parameters was validated. While this novel approach is inevitably prototypical, it addresses a significant limitation of hemodynamic simulations in medicine – their time and resource intensity, which hinder direct application in clinical workflows. The approach offers a way to overcome this limitation, with the potential to advance personalized medicine and make hemodynamic simulations more accessible and valuable for clinical decision-making.

9.2 FUTURE WORK

Validation of computational hemodynamics. The validation of simulated flow and flow-derived parameters in carotid stenoses is yet to be improved. First studies have shown that velocities measured with Doppler ultrasonography can be reproduced with CFD [116, 287]. We also successfully compared the flow between simulation and an in-vitro model of the carotid bifurcation [353]. However, since these studies are limited to a handful of subjects, larger studies and generalizable results are pending. Therefore, in the future, a key goal should be to enhance the validity and applicability of hemodynamics computed through CFD. To enable such validations, the integration of different imaging modalities, such as CTA and ultrasonography, could be achieved through visual means. Visualizations of registered data spaces could make measured and simulated parameters directly comparable. Furthermore, this work has shown that the extraction and flow simulation of 100 and more patient-specific carotid bifurcation geometries is feasible, which means that large-scale evaluations of measured versus simulated flow properties are within reach.

Multi-factor analysis. In this work, the clinically most relevant properties of stenoses are targeted, which are primarily the stenosis degree and the systolic peak velocity. The visualization of hemodynamic scalar

data, like the *WSS*, is also considered in some contexts. In many concurrent studies from the field of *CFD*, the interplay of different parameters, such as high/low *WSS* or the *OSI* are analyzed for their potential as disease progression markers [124, 186, 229]. Future research could further investigate these supplementary factors to uncover patterns and insights into the initiation and progression of stenoses. In this context, visualization techniques could greatly facilitate understanding the interplay of multiple factors in stenosis and atherosclerotic plaque development.

Intracranial vessels. This work examined the carotid bifurcation as it is a typical predilection site for stenosis development. Extending the analyzed domain to include intracranial vessels (or the vertebral arteries), could be valuable and has been mentioned by collaborating physicians multiple times. Detecting stenoses in larger vessel trees is considerably more challenging and can benefit from automated processing and vessel visualization techniques. Also, expanding the segmented vascular domain to accurately compute flow properties of arterial short circuits could be explored. In many anatomical variations of the circle of Willis, the two internal carotids are directly connected, which means one side can compensate for a stenosis in the other. Hence, the specific layout of the arterial subtree following the *ICA* is highly relevant for clinical decision-making. Furthermore, measuring intracranial flow velocities via Doppler ultrasonography is challenging because of skull-induced signal distortion, but remains a common procedure due to the lack of non-invasive alternatives. High-resolution computed hemodynamics could be substantially beneficial in these scenarios. Presenting such an extended domain without visual clutter is a considerable challenge.

Pipeline extensions. The integrated pipeline proposed in Chapter 7, in conjunction with the carotid bifurcation model database [347], serves as a platform for efficient extension development. Based on the modular pipeline architecture, additional processing or visualization stages can be seamlessly incorporated. Possible extensions encompass modules designed for epidemiological cohort studies or the integration of multimodal data – such as *MRA* or Doppler ultrasonography. Enhancements to current processes are also feasible, such as enabling the segmentation prediction network to continuously and autonomously learn from user-corrected masks.

Extend comparison-based analysis. During interviews with physicians, the idea of interpolating flow parameters based on the similarity distance for matching models was mentioned. If we interpret a model as a vector in the similarity space, an interpolation of the nearest cases seems to be sensible to “sample” the unknown data point. This approach could potentially yield more accurate values, and projecting the results onto the new geometry could enhance readability. However,

there is a risk of misinterpretation if good matches are unavailable, i.e., if the sampling density in the similarity space neighborhood of the input model is insufficient.

Similarity-based flow analysis could also be applied to other structures, such as aneurysms. In many clinical cases, only static imaging is available, where flow parameters could enhance the evaluation. In the long term, similarity-based analysis could also aid in finding similar cases with more comprehensive information than just flow. For example, it could be possible to query geometrically similar aneurysm cases from a database with recorded parameters beyond simulated flow, such as the rupture status of similar aneurysms.

Part VI

APPENDIX

BIBLIOGRAPHY

- [1] Stephan Achenbach, Werner Moshage, Dieter Ropers, and Kurt Bachmann. "Curved Multiplanar Reconstructions for the Evaluation of Contrast-Enhanced Electron Beam CT of the Coronary Arteries." In: *American Journal of Roentgenology* 170.4 (1998), pp. 895–899. DOI: [10.2214/ajr.170.4.9530029](https://doi.org/10.2214/ajr.170.4.9530029).
- [2] Klaus Kaae Andersen, Tom Skyhøj Olsen, Christian Dehrendorff, and Lars Peter Kammergaard. "Hemorrhagic and Ischemic Strokes Compared: Stroke Severity, Mortality, and Risk Factors." In: *Stroke* 40.6 (June 2009), pp. 2068–2072. ISSN: 1524-4628. DOI: [10.1161/strokeaha.108.540112](https://doi.org/10.1161/strokeaha.108.540112).
- [3] Magnus Andersson, Jonas Lantz, Tino Ebbers, and Matts Karlsson. "Quantitative assessment of turbulence and flow eccentricity in an aortic coarctation: impact of virtual interventions." In: *Cardiovascular Engineering and Technology* 6.3 (2015), pp. 281–293. DOI: [10.1007/s13239-015-0218-x](https://doi.org/10.1007/s13239-015-0218-x).
- [4] Paolo Angelelli and Helwig Hauser. "Straightening tubular flow for side-by-side visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2063–2070. DOI: [10.1109/TVCG.2011.235](https://doi.org/10.1109/TVCG.2011.235).
- [5] Luca Antiga. "Patient-specific Modeling of Geometry and Blood Flow in Large Arteries." PhD thesis. Politecnico di Milano, 2002.
- [6] Luca Antiga, Bogdan Ene-Iordache, and Andrea Remuzzi. "Computational geometry for patient-specific reconstruction and meshing of blood vessels from MR and CT angiography." In: *IEEE Transactions on Medical Imaging* 22.5 (2003), pp. 674–684. DOI: [10.1109/tmi.2003.812261](https://doi.org/10.1109/tmi.2003.812261).
- [7] Luca Antiga, Marina Piccinelli, Lorenzo Botti, Bogdan Ene-Iordache, Andrea Remuzzi, and David A. Steinman. "An image-based modeling framework for patient-specific computational hemodynamics." In: *Medical & Biological Engineering & Computing* 46.11 (2008), p. 1097. DOI: [10.1007/s11517-008-0420-1](https://doi.org/10.1007/s11517-008-0420-1).
- [8] Luca Antiga and David A. Steinman. *Automated Parameterization and Patching of Bifurcating Vessels*. Tech. rep. EPFL Infoscience, 2003.
- [9] Luca Antiga and David A. Steinman. "Robust and objective decomposition and mapping of bifurcating vessels." In: *IEEE Transactions on Medical Imaging* 23.6 (2004), pp. 704–713. DOI: [10.1109/TMI.2004.826946](https://doi.org/10.1109/TMI.2004.826946).

- [10] Sunil Appanaboyina, Fernando Mut, Rainald Löhner, Christopher M. Putman, and Juan R. Cebal. "Simulation of intracranial aneurysm stenting: techniques and challenges." In: *Computer Methods in Applied Mechanics and Engineering* 198.45-46 (2009), pp. 3567–3582. DOI: [10.1016/j.cma.2009.01.017](https://doi.org/10.1016/j.cma.2009.01.017).
- [11] Amirhossein Arzani, Petter Dyverfeldt, Tino Ebbers, and Shawn C. Shadden. "In vivo validation of numerical prediction for turbulence intensity in an aortic coarctation." In: *Annals of Biomedical Engineering* 40.4 (2012), pp. 860–870.
- [12] Oscar Kin-Chung Au, Chiew-Lan Tai, Hung-Kuo Chu, Daniel Cohen-Or, and Tong-Yee Lee. "Skeleton extraction by mesh contraction." In: *ACM Transactions on Graphics* 27.3 (2008), pp. 1–10. DOI: [10.1145/1360612.1360643](https://doi.org/10.1145/1360612.1360643).
- [13] Thomas Auzinger, Gabriel Mistelbauer, Ivan Baclija, Rüdiger Schernthaner, Arnold Köchl, Michael Wimmer, M. Eduard Gröller, and Stefan Bruckner. "Vessel visualization using curved surface reformation." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2858–2867. DOI: [10.1109/TVCG.2013.215](https://doi.org/10.1109/TVCG.2013.215).
- [14] Ali K Al-Awami, Johanna Beyer, Hendrik Strobel, Narayanan Kasthuri, Jeff W Lichtman, Hanspeter Pfister, and Markus Hadwiger. "NeuroLines: a subway map metaphor for visualizing nanoscale neuronal connectivity." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2369–2378. DOI: [10.1109/TVCG.2014.2346312](https://doi.org/10.1109/TVCG.2014.2346312).
- [15] Orhun Utku Aydin, Abdel Aziz Taha, Adam Hilbert, Ahmed A. Khalil, Ivana Galinovic, Jochen B. Fiebach, Dietmar Frey, and Vince Istvan Madai. "On the usage of average Hausdorff distance for segmentation performance assessment: hidden error when used for ranking." In: *European Radiology Experimental* 5.1 (2021). DOI: [10.1186/s41747-020-00200-2](https://doi.org/10.1186/s41747-020-00200-2).
- [16] Dara Azar, William M. Torres, Lindsey A. Davis, Taylor Shaw, John F. Eberth, Vijaya B. Kolachalama, Susan M. Lessner, and Tarek Shazly. "Geometric determinants of local hemodynamics in severe carotid artery stenosis." In: *Computers in Biology and Medicine* 114 (2019), p. 103436. DOI: [10.1016/j.compbiomed.2019.103436](https://doi.org/10.1016/j.compbiomed.2019.103436).
- [17] Christian Bachmaier, Ulrik Brandes, and Barbara Schlieper. "Drawing Phylogenetic Trees." In: *Algorithms and Computation*. Springer Berlin Heidelberg, 2005, pp. 1110–1121. DOI: [10.1007/11602613_110](https://doi.org/10.1007/11602613_110).
- [18] Stavros J. Baloyannis, Ioannis Mavroudis, Demetrios Mitilneos, Ioannis S. Baloyannis, and Vassiliki G. Costa. "The Hypothalamus in Alzheimer's Disease." In: *American Journal of*

- Alzheimer's Disease & Other Dementias* 30.5 (2014), pp. 478–487. DOI: [10.1177/1533317514556876](https://doi.org/10.1177/1533317514556876).
- [19] Alister J. Bates, Denis J. Doorly, Raul Cetto, Hadrien Calmet, A. M. Gambaruto, N. S. Tolley, Guillaume Houzeaux, and R. C. Schroter. “Dynamics of airflow in a short inhalation.” In: *Journal of the Royal Society Interface* 12.102 (2015), p. 20140880. DOI: [10.1098/rsif.2014.0880](https://doi.org/10.1098/rsif.2014.0880).
- [20] Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall PTR, 1998.
- [21] Christian Bauer, Thomas Pock, Erich Sorantin, Horst Bischof, and Reinhard Beichel. “Segmentation of interwoven 3D tubular tree structures utilizing shape priors and graph cuts.” In: *Medical Image Analysis* 14.2 (2010), pp. 172–184. DOI: [10.1016/j.media.2009.11.003](https://doi.org/10.1016/j.media.2009.11.003).
- [22] Benjamin Behrendt, David Pleuss-Engelhardt, Matthias Gutberlet, and Bernhard Preim. “2.5D Geometric Mapping of Aortic Blood Flow Data for Cohort Visualization.” In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2021, pp. 91–100. DOI: [10.2312/vcbm.20211348](https://doi.org/10.2312/vcbm.20211348).
- [23] Philipp Berg, Christoph Roloff, Oliver Beuing, Samuel Voß, Shin-Ichiro Sugiyama, Nicolas Aristokleous, Andreas S Anayiotos, Neil Ashton, Alistair Revell, and Neil W Bressloff. “The computational fluid dynamics rupture challenge 2013—phase II: variability of hemodynamic simulations in two intracranial aneurysms.” In: *Journal of Biomechanical Engineering* 137.12 (2015).
- [24] Hiram G. Bezerra, Marco A. Costa, Giulio Guagliumi, Andrew M. Rollins, and Daniel I. Simon. “Intracoronary Optical Coherence Tomography: A Comprehensive Review.” In: *JACC: Cardiovascular Interventions* 2.11 (2009), pp. 1035–1046. DOI: [10.1016/j.jcin.2009.06.019](https://doi.org/10.1016/j.jcin.2009.06.019).
- [25] I. Bitter, A. E. Kaufman, and M. Sato. “Penalized-distance volumetric skeleton algorithm.” In: *IEEE Transactions on Visualization and Computer Graphics* 7.3 (2001), pp. 195–206. DOI: [10.1109/2945.942688](https://doi.org/10.1109/2945.942688).
- [26] I. Bitter, M. Sato, M. Bender, K. T. McDonnell, A. E. Kaufman, and Ming Wan. “CEASAR: a smooth, accurate and robust centerline extraction algorithm.” In: *Proceedings Visualization (VIS)*. 2000. DOI: [10.1109/visual.2000.885675](https://doi.org/10.1109/visual.2000.885675).
- [27] Danny Bluestein. “Utilizing Computational Fluid Dynamics in Cardiovascular Engineering and Medicine – What You Need to Know. Its Translation to the Clinic/Bedside.” In: *Artificial Organs* 41.2 (2017), pp. 117–121. DOI: [10.1111/aor.12914](https://doi.org/10.1111/aor.12914).

- [28] Harry Blum et al. *A transformation for extracting new descriptors of shape*. MIT Press, 1967.
- [29] Georges-Pierre Bonneau, Hans-Christian Hege, Chris R. Johnson, Manuel M. Oliveira, Kristin Potter, Penny Rheingans, and Thomas Schultz. "Overview and state-of-the-art of uncertainty visualization." In: *Scientific Visualization*. 2014, pp. 3–27. DOI: [10.1007/978-1-4471-6497-5_1](https://doi.org/10.1007/978-1-4471-6497-5_1).
- [30] Michelle Borkin, Krzysztof Gajos, Amanda Peters, Dimitrios Mitsouras, Simone Melchionna, Frank Rybicki, Charles Feldman, and Hanspeter Pfister. "Evaluation of artery visualizations for heart disease diagnosis." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2479–2488. DOI: [10.1109/TVCG.2011.192](https://doi.org/10.1109/TVCG.2011.192).
- [31] Silvia Born, Michael Markl, Matthias Gutberlet, and Gerik Scheuermann. "Illustrative visualization of cardiac and aortic blood flow from 4D MRI data." In: *IEEE Pacific Visualization Symposium*. 2013, pp. 129–136. DOI: [10.1109/pacificvis.2013.6596137](https://doi.org/10.1109/pacificvis.2013.6596137).
- [32] Silvia Born, Matthias Pfeifle, Michael Markl, Matthias Gutberlet, and Gerik Scheuermann. "Visual analysis of cardiac 4D MRI blood flow using line predicates." In: *IEEE Transactions on Visualization and Computer Graphics* 19.6 (2012), pp. 900–912. DOI: [10.1109/tvcg.2012.318](https://doi.org/10.1109/tvcg.2012.318).
- [33] Silvia Born, Simon H. Sundermann, Christoph Russ, Raoul Hopf, Carlos E. Ruiz, Volkmar Falk, and Michael Gessat. "Stent Maps – Comparative Visualization for the Prediction of Adverse Events of Transcatheter Aortic Valve Implantations." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2704–2713. DOI: [10.1109/tvcg.2014.2346459](https://doi.org/10.1109/tvcg.2014.2346459).
- [34] Ferhat Bozkurt, Cemal Köse, and Ahmet Sari. "An inverse approach for automatic segmentation of carotid and vertebral arteries in CTA." In: *Expert Systems with Applications* 93 (2018), pp. 358–375. DOI: [10.1016/j.eswa.2017.10.041](https://doi.org/10.1016/j.eswa.2017.10.041).
- [35] Ken Brodlie, Rodolfo Allendes Osorio, and Adriano Lopes. "A review of uncertainty in data visualization." In: *Expanding the Frontiers of Visual Analytics and Visualization*. 2012, pp. 81–109. DOI: [10.1007/978-1-4471-2804-5_6](https://doi.org/10.1007/978-1-4471-2804-5_6).
- [36] Arjan J. M. Broos, Niels H. L. C. de Hoon, Patrick J. H. de Koning, Rob J. van der Geest, Anna Vilanova, and Andrei C. Jalba. "A Framework for Fast Initial Exploration of PC-MRI Cardiac Flow." In: *Visual Computing for Biology and Medicine*. 2016, pp. 69–78.

- [37] Thomas G. Brott, Jonathan L. Halperin, Suhny Abbara, et al. "2011 ASA/ACCF/AHA/AANN/AANS/ACR/ASNR/CNS/SAIP/SCAI/SIR/SNIS/SVM/SVS Guideline on the Management of Patients With Extracranial Carotid and Vertebral Artery Disease: Executive Summary." In: *Journal of the American College of Cardiology* 57.8 (2011), pp. 1002–1044. DOI: [10.1016/j.jacc.2010.11.005](https://doi.org/10.1016/j.jacc.2010.11.005).
- [38] Stefan Bruckner and M. Eduard Gröller. "Instant Volume Visualization using Maximum Intensity Difference Accumulation." In: *Computer Graphics Forum* 28.3 (2009), pp. 775–782. DOI: [10.1111/j.1467-8659.2009.01474.x](https://doi.org/10.1111/j.1467-8659.2009.01474.x).
- [39] Katja Bühler, Petr Felkel, and Alexandra La Cruz. "Geometric Methods for Vessel Visualization and Quantification – A Survey." In: *Geometric Modeling for Scientific Visualization*. Ed. by Guido Brunnett, Bernd Hamann, Heinrich Müller, and Lars Linsen. Berlin, Heidelberg: Springer, 2004, pp. 399–419. DOI: [10.1007/978-3-662-07443-5_24](https://doi.org/10.1007/978-3-662-07443-5_24).
- [40] Michael H Buonocore. "Visualizing blood flow patterns using streamlines, arrows, and particle paths." In: *Magnetic Resonance in Medicine* 40.2 (1998), pp. 210–226. DOI: [10.1002/mrm.1910400207](https://doi.org/10.1002/mrm.1910400207).
- [41] Mariana Bustamante, Sven Petersson, Jonatan Eriksson, Urban Alehagen, Petter Dyverfeldt, Carl-Johan Carlhäll, and Tino Ebbers. "Atlas-based analysis of 4D flow CMR: Automated vessel segmentation and flow quantification." In: *Journal of Cardiovascular Magnetic Resonance* 17.1 (2015). DOI: [10.1186/s12968-015-0190-5](https://doi.org/10.1186/s12968-015-0190-5).
- [42] Greg Byrne and Juan R. Cebral. "Vortex dynamics in cerebral aneurysms." In: *ArXiv e-prints* (2013).
- [43] Greg Byrne, Fernando Mut, and Juan R. Cebral. "Quantifying the large-scale hemodynamics of intracranial aneurysms." In: *American Journal of Neuroradiology* 35.2 (2014), pp. 333–338. DOI: [10.3174/ajnr.A3678](https://doi.org/10.3174/ajnr.A3678).
- [44] Andres D. Caballero and Santiago Laín. "A review on computational fluid dynamics modelling in human thoracic aorta." In: *Cardiovascular Engineering and Technology* 4.2 (2013), pp. 103–130. DOI: [10.1007/s13239-013-0146-6](https://doi.org/10.1007/s13239-013-0146-6).
- [45] Wenli Cai. "3D planar reformation of vascular central axis surface with biconvex slab." In: *Computerized Medical Imaging and Graphics* 31.7 (2007), pp. 570–576. DOI: [10.1016/j.compmedimag.2007.06.007](https://doi.org/10.1016/j.compmedimag.2007.06.007).

- [46] Paul S. Calhoun, Brian S. Kuszyk, David G. Heath, Jennifer C. Carley, and Elliot K. Fishman. "Three-dimensional volume rendering of spiral CT data: Theory and method." In: *Radiographics* 19.3 (1999), pp. 745–764. DOI: [10.1148/radiographics.19.3.g99ma14745](https://doi.org/10.1148/radiographics.19.3.g99ma14745).
- [47] M. Jorge Cardoso, Wenqi Li, Richard Brown, et al. "MONAI: An open-source framework for deep learning in healthcare." In: (2022). DOI: <https://doi.org/10.48550/arXiv.2211.02701>.
- [48] Pepe Eulzer and Kai Lawonn. *Carotid Analyzer*. Friedrich Schiller University Jena. 2023. URL: <https://github.com/PepeEulzer/CarotidAnalyzer>.
- [49] Paula Casademunt, Oscar Camara, Bart Bijmens, Èric Lluch, and Hernan G. Morales. "Valve Flattening with Functional Biomarkers for the Assessment of Mitral Valve Repair." In: *Lecture Notes in Computer Science*. Springer International Publishing, 2022, pp. 169–178. DOI: [10.1007/978-3-030-93722-5_19](https://doi.org/10.1007/978-3-030-93722-5_19).
- [50] Cynthia A. Brewer and Mark Harrower. *ColorBrewer*. Pennsylvania State University. 2022. URL: <https://colorbrewer2.org/>.
- [51] Juan R. Cebral, Bongjae Chung, Fernando Mut, Fred van Nijmegen, and Danny Ruijters. "Comparison of cerebral aneurysm flow fields obtained from CFD and DSA." In: *Proceedings of Summer Biomechanics, Bioengineering and Biotransport Conference*. 2016.
- [52] Juan R. Cebral, Mariano Vazquez, Daniel M Sforza, Guillaume Houzeaux, Satoshi Tateshima, Esteban Scrivano, Carlos Bleise, Pedro Lylyk, and Christopher M. Putman. "Analysis of hemodynamics and wall mechanics at sites of cerebral aneurysm rupture." In: *Journal of NeuroInterventional Surgery* 7.7 (2015), pp. 530–536.
- [53] Jieyu Cheng, Eranga Ukwatta, Shadi Shavakh, Tommy WS Chow, Grace Parraga, J David Spence, and Bernard Chiu. "Sensitive three-dimensional ultrasound assessment of carotid atherosclerosis by weighted average of local vessel wall and plaque thickness change." In: *Medical Physics* 44.10 (2017), pp. 5280–5292. DOI: doi.org/10.1002/mp.12507.
- [54] Bernard Chiu, Weifu Chen, and Jieyu Cheng. "Concise biomarker for spatial-temporal change in three-dimensional ultrasound measurement of carotid vessel wall and plaque thickness based on a graph-based random walk framework: Towards sensitive evaluation of response to therapy." In: *Computers in Biology and Medicine* 79 (2016), pp. 149–162. DOI: [10.1016/j.combiomed.2016.10.015](https://doi.org/10.1016/j.combiomed.2016.10.015).

- [55] Bernard Chiu, Micaela Egger, David J Spence, Grace Parraga, and Aaron Fenster. "Area-preserving flattening maps of 3D ultrasound carotid arteries images." In: *Medical Image Analysis* 12.6 (2008), pp. 676–688. DOI: [10.1016/j.media.2008.04.002](https://doi.org/10.1016/j.media.2008.04.002).
- [56] Bernard Chiu, Bing Li, and Tommy WS Chow. "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." In: *Physics in Medicine & Biology* 58.17 (2013), pp. 5959–5982. DOI: [10.1088/0031-9155/58/17/5959](https://doi.org/10.1088/0031-9155/58/17/5959).
- [57] Bernard Chiu, Eranga Ukwatta, Shadi Shavakh, and Aaron Fenster. "Quantification and visualization of carotid segmentation accuracy and precision using a 2D standardized carotid map." In: *Physics in Medicine and Biology* 58.11 (2013), pp. 3671–3703. DOI: [10.1088/0031-9155/58/11/3671](https://doi.org/10.1088/0031-9155/58/11/3671).
- [58] Gary P. T. Choi, Yimin Chen, Lok Ming Lui, and Bernard Chiu. "Conformal mapping of carotid vessel wall and plaque thickness measured from 3D ultrasound images." In: *Medical & Biological Engineering & Computing* 55.12 (2017), pp. 2183–2195. DOI: [10.1007/s11517-017-1656-4](https://doi.org/10.1007/s11517-017-1656-4).
- [59] Gary P. T. Choi, Bernard Chiu, and Chris H. Rycroft. "Area-preserving mapping of 3D carotid ultrasound images using density-equalizing reference map." In: *IEEE Transactions on Biomedical Engineering* 67.9 (2020), pp. 2507–2517. DOI: [10.1109/TBME.2019.2963783](https://doi.org/10.1109/TBME.2019.2963783).
- [60] Farooq A. Choudhry, John T. Grantham, Ansaar T. Rai, and Jeffery P. Hogg. "Vascular geometry of the extracranial carotid arteries: an analysis of length, diameter, and tortuosity." In: *Journal of Neurointerventional Surgery* 8.5 (2016), pp. 536–540. DOI: [10.1136/neurintsurg-2015-011671](https://doi.org/10.1136/neurintsurg-2015-011671).
- [61] Bongjae Chung and Juan Raul Cebal. "CFD for evaluation and treatment planning of aneurysms: review of proposed clinical uses and their challenges." In: *Annals of Biomedical Engineering* 43.1 (2015), pp. 122–138.
- [62] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. "3D U-Net: learning dense volumetric segmentation from sparse annotation." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2016, pp. 424–432.
- [63] Andy Cockburn and Bruce McKenzie. "Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments." In: *CHI Conference on Human Factors in Computing Systems*. ACM, 2002, pp. 203–210. DOI: [10.1145/503376.503413](https://doi.org/10.1145/503376.503413).

- [64] Mark S. Cohen. "Real-Time Functional Magnetic Resonance Imaging." In: *Methods* 25.2 (2001), pp. 201–220. DOI: [10.1006/meth.2001.1235](https://doi.org/10.1006/meth.2001.1235).
- [65] COMSOL Multiphysics. 2023. URL: <https://doc.comsol.com/6.0/docserver/>.
- [66] Michele Conti, Chris Long, Michele Marconi, Raffaella Berchiolli, Yuri Bazilevs, and Alessandro Reali. "Carotid artery hemodynamics before and after stenting: A patient specific CFD study." In: *Computers & Fluids* 141 (2016), pp. 62–74. DOI: [10.1016/j.compfluid.2016.04.006](https://doi.org/10.1016/j.compfluid.2016.04.006).
- [67] Steven A Coons. *Surfaces for computer-aided design of space forms*. Tech. rep. MIT, 1967.
- [68] Isabelle Corouge, Sylvain Gouttard, and Guido Gerig. "Towards a shape model of white matter fiber bundles using diffusion tensor MRI." In: *IEEE International Symposium on Biomedical Imaging*. Vol. 1. 2004, pp. 344–347. DOI: [10.1109/ISBI.2004.1398545](https://doi.org/10.1109/ISBI.2004.1398545).
- [69] Keenan Crane, Clarisse Weischedel, and Max Wardetzky. "Geodesics in heat: a new approach to computing distance based on heat flow." In: *ACM Transactions on Graphics* 32.5 (2013), pp. 1–11. DOI: [10.1145/2516971.2516977](https://doi.org/10.1145/2516971.2516977).
- [70] Creative Commons. *Attribution 4.0 International License*. URL: <http://creativecommons.org/licenses/by/4.0/>.
- [71] Alejandro C. Crespo, Jose M. Dominguez, Anxo Barreiro, Moncho Gómez-Gesteira, and Benedict D. Rogers. "GPUs, a new tool of acceleration in CFD: efficiency and reliability on smoothed particle hydrodynamics methods." In: *PloS one* 6.6 (2011), e20685. DOI: [10.1371/journal.pone.0020685](https://doi.org/10.1371/journal.pone.0020685).
- [72] Olivier Cuisenaire, Sunny Virmani, Mark E. Olszewski, and Roberto Ardon. "Fully automated segmentation of carotid and vertebral arteries from contrast enhanced CTA." In: *SPIE Proceedings*. Ed. by Joseph M. Reinhardt and Josien P. W. Pluim. SPIE, 2008. DOI: [10.1117/12.770481](https://doi.org/10.1117/12.770481).
- [73] Tim Culver, John Keyser, and Dinesh Manocha. "Exact computation of the medial axis of a polyhedron." In: *Computer Aided Geometric Design* 21.1 (2004), pp. 65–98. DOI: [10.1016/j.cagd.2003.07.008](https://doi.org/10.1016/j.cagd.2003.07.008).
- [74] Anna Christina Daub. "Numerical Haemodynamics in the Human Heart." PhD thesis. Karlsruher Institut für Technologie, 2018.
- [75] Ivana Despotović, Bart Goossens, and Wilfried Philips. "MRI Segmentation of the Human Brain: Challenges, Methods, and Applications." In: *Computational and Mathematical Methods in Medicine* 2015 (2015), pp. 1–23. DOI: [10.1155/2015/450341](https://doi.org/10.1155/2015/450341).

- [76] Deutsche Gesellschaft für Gefäßchirurgie und Gefäßmedizin (DGG), Deutsche Gesellschaft für Neurologie (DGN), Deutsche Schlaganfall-Gesellschaft (DSG), et al. *S3-Leitlinie zur Diagnostik, Therapie und Nachsorge der extracraniellen Carotisstenose, Version 2.1*. 2020. URL: <https://register.awmf.org/de/leitlinien/detail/004-028>.
- [77] Stefan Diepenbrock, Sven Hermann, Michael Schäfers, Michael Kuhlmann, and Klaus Hinrichs. “Comparative visualization of tracer uptake in in vivo small animal PET/CT imaging of the carotid arteries.” In: *Computer Graphics Forum* 32.3pt2 (2013), pp. 241–250. DOI: [10.1111/cgf.12111](https://doi.org/10.1111/cgf.12111).
- [78] Edsger W. Dijkstra. “A note on two problems in connexion with graphs.” In: *Numerische Mathematik* 1.1 (1959), pp. 269–271. DOI: [10.1007/BF01386390](https://doi.org/10.1007/BF01386390).
- [79] Jingliang Dong, Kelvin K. L. Wong, and Jiyuan Tu. “Hemodynamics analysis of patient-specific carotid bifurcation: A CFD model of downstream peripheral vascular impedance.” In: *International Journal for Numerical Methods in Biomedical Engineering* 29.4 (2013), pp. 476–491. DOI: [10.1002/cnm.2529](https://doi.org/10.1002/cnm.2529).
- [80] Geoffrey A. Donnan, Marc Fisher, Malcom Macleod, and Stephen M. Davis. “Stroke.” In: *The Lancet* 371.9624 (2008), pp. 1612–1623. DOI: [10.1016/S0140-6736\(08\)60694-7](https://doi.org/10.1016/S0140-6736(08)60694-7).
- [81] Getao Du, Xu Cao, Jimin Liang, Xueli Chen, and Yonghua Zhan. “Medical Image Segmentation based on U-Net: A Review.” In: *Journal of Imaging Science and Technology* 64.2 (2020), 20508-1–20508-12. DOI: [10.2352/j.imagingsci.technol.2020.64.2.020508](https://doi.org/10.2352/j.imagingsci.technol.2020.64.2.020508).
- [82] C. Dunne, S. I. Ross, B. Shneiderman, and M. Martino. “Readability metric feedback for aiding node-link visualization designers.” In: *IBM Journal of Research and Development* 59.2/3 (2015), 14:1–14:16. DOI: [10.1147/jrd.2015.2411412](https://doi.org/10.1147/jrd.2015.2411412).
- [83] Rickard Englund, Timo Ropinski, and Ingrid Hotz. “Coherence Maps for Blood Flow Exploration.” In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2016, pp. 79–88. DOI: [10.2312/vcbm.20161274](https://doi.org/10.2312/vcbm.20161274).
- [84] S. Esneault, C. Lafon, and J.-L. Dillenseger. “Liver Vessels Segmentation Using a Hybrid Geometrical Moments/Graph Cuts Method.” In: *IEEE Transactions on Biomedical Engineering* 57.2 (2010), pp. 276–283. DOI: [10.1109/tbme.2009.2032161](https://doi.org/10.1109/tbme.2009.2032161).
- [85] Maarten H. Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. “Depth-Dependent Halos: Illustrative Rendering of Dense Line Data.” In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1299–306. DOI: [10.1109/TVCG.2009.138](https://doi.org/10.1109/TVCG.2009.138).

- [86] Giacomo Feliciani, Wouter V. Potters, Pim van Ooij, Joppe J. Schneiders, Aart J. Nederveen, Ed van Bavel, Charles B. Majoie, and Henk A. Marquering. "Multiscale 3-D + t intracranial aneurysmal flow vortex detection." In: *IEEE Transactions on Biomedical Engineering* 62.5 (2015), pp. 1355–1362. DOI: [10.1109/TBME.2014.2387874](https://doi.org/10.1109/TBME.2014.2387874).
- [87] S. Ferchichi and Shengrui Wang. "A Clustering-based Algorithm for Extracting the Centerlines of 2D and 3D Objects." In: *18th International Conference on Pattern Recognition*. 2006. DOI: [10.1109/icpr.2006.44](https://doi.org/10.1109/icpr.2006.44).
- [88] Gary G. Ferguson et al. "The North American Symptomatic Carotid Endarterectomy Trial." In: *Stroke* 30.9 (1999), pp. 1751–1758. DOI: [10.1161/01.str.30.9.1751](https://doi.org/10.1161/01.str.30.9.1751).
- [89] A. Forsberg, Jian Chen, and D. Laidlaw. "Comparing 3D Vector Field Visualization Methods: A User Study." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1219–1226. DOI: [10.1109/tvcg.2009.126](https://doi.org/10.1109/tvcg.2009.126).
- [90] Alejandro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever. "Multiscale vessel enhancement filtering." In: *Medical Image Computing and Computer-Assisted Intervention*. Springer Berlin Heidelberg, 1998, pp. 130–137. DOI: [10.1007/bfb0056195](https://doi.org/10.1007/bfb0056195).
- [91] M. M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. R. Rudnicka, C. G. Owen, and S. A. Barman. "Blood vessel segmentation methodologies in retinal images – A survey." In: *Computer Methods and Programs in Biomedicine* 108.1 (2012), pp. 407–433. DOI: [10.1016/j.cmpb.2012.03.009](https://doi.org/10.1016/j.cmpb.2012.03.009).
- [92] Dai Fukumura and Rakesh K. Jain. "Tumor microvasculature and microenvironment: Targets for anti-angiogenesis and normalization." In: *Microvascular Research* 74.2-3 (2007), pp. 72–84. DOI: [10.1016/j.mvr.2007.05.003](https://doi.org/10.1016/j.mvr.2007.05.003).
- [93] AM Gambaruto and AJ João. "Flow structures in cerebral aneurysms." In: *Computers & Fluids* 65 (2012), pp. 56–65. DOI: [10.1016/j.compfluid.2012.02.020](https://doi.org/10.1016/j.compfluid.2012.02.020).
- [94] Guilherme JM Garcia, Neil Bailie, Dário A Martins, and Julia S Kimbell. "Atrophic rhinitis: a CFD study of air conditioning in the nasal cavity." In: *Journal of Applied Physiology* 103.3 (2007), pp. 1082–1092. DOI: [10.1152/jappphysiol.01118.2006](https://doi.org/10.1152/jappphysiol.01118.2006).
- [95] Hector M. García-García, Bill D. Gogas, Patrick W. Serruys, and Nico Bruining. "IVUS-based imaging modalities for tissue characterization: similarities and differences." In: *The International Journal of Cardiovascular Imaging* 27.2 (2011), pp. 215–224. DOI: [10.1007/s10554-010-9789-7](https://doi.org/10.1007/s10554-010-9789-7).

- [96] Rocco Gasteiger, Mathias Neugebauer, Oliver Beuing, and Bernhard Preim. "The FLOWLENS: A focus-and-context visualization approach for exploration of blood flow in cerebral aneurysms." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2183–2192. DOI: [10.1109/TVCG.2011.243](https://doi.org/10.1109/TVCG.2011.243).
- [97] Rocco Gasteiger, Mathias Neugebauer, Christoph Kubisch, and Bernhard Preim. "Adapted surface visualization of cerebral aneurysms with embedded blood flow information." In: *Visual Computing for Biology and Medicine*. Ed. by Dirk Bartz, Charl Botha, Joachim Hornegger, Raghu Machiraju, Alexander Wiebel, and Bernhard Preim. 2010, pp. 25–32. DOI: [10.2312/VCBM/VCBM10/025-032](https://doi.org/10.2312/VCBM/VCBM10/025-032).
- [98] GBD 2015 Mortality and Causes of Death Collaborators. "Global, regional, and national life expectancy, all-cause mortality, and cause-specific mortality for 249 causes of death, 1980–2015: a systematic analysis for the Global Burden of Disease Study 2015." In: *The Lancet* 388.10053 (2016), pp. 1459–1544. ISSN: 0140-6736. DOI: [10.1016/S0140-6736\(16\)31012-1](https://doi.org/10.1016/S0140-6736(16)31012-1).
- [99] GBD 2016 Stroke Collaborators. "Global, regional, and national burden of stroke, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016." In: *The Lancet Neurology* 18.5 (2019), pp. 439–458. ISSN: 1474-4422. DOI: [10.1016/S1474-4422\(19\)30034-1](https://doi.org/10.1016/S1474-4422(19)30034-1).
- [100] GBD 2019 Stroke Collaborators. "Global, regional, and national burden of stroke and its risk factors, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019." In: *The Lancet Neurology* 20.10 (2021), pp. 795–820. DOI: [10.1016/s1474-4422\(21\)00252-0](https://doi.org/10.1016/s1474-4422(21)00252-0).
- [101] Yaorong Ge, David R. Stelts, and David J. Vining. "3D skeleton for virtual colonoscopy." In: *Lecture Notes in Computer Science*. 1996, pp. 449–454. DOI: [10.1007/bfb0046985](https://doi.org/10.1007/bfb0046985).
- [102] P. Giblin and B. B. Kimia. "A formal classification of 3D medial axis points and their local geometry." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26.2 (2004), pp. 238–251. DOI: [10.1109/tpami.2004.1262192](https://doi.org/10.1109/tpami.2004.1262192).
- [103] Christina Gillmann, Noeska N. Smit, M. Eduard Gröller, Bernhard Preim, Anna Vilanova, and Thomas Wischgoll. "Ten Open Challenges in Medical Visualization." In: *IEEE Computer Graphics and Applications* 41.5 (2021), pp. 7–15. DOI: [10.1109/mcg.2021.3094858](https://doi.org/10.1109/mcg.2021.3094858).

- [104] Christina Gillmann, Thomas Wischgoll, and Hans Hagen. "Visual Exploration in Surgery Monitoring for Coronary Vessels." In: *IEEE VIS Practitioner Event: Visualization in Practice*. 2015. URL: <https://corescholar.libraries.wright.edu/cse/491/>.
- [105] Sylvia Glaßer, Kai Lawonn, Thomas Hoffmann, Martin Skalej, and Bernhard Preim. "Combined visualization of wall thickness and wall shear stress for the evaluation of aneurysms." In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2506–2515. DOI: [10.1109/TVCG.2014.2346406](https://doi.org/10.1109/TVCG.2014.2346406).
- [106] Sylvia Glasser, Steffen Oeltze, Anja Hennemuth, Christoph Kubisch, A Mahnken, Skadi Wilhelmsen, and Bernhard Preim. "Automatic transfer function specification for visual emphasis of coronary artery plaque." In: *Computer Graphics Forum* 29.1 (2010), pp. 191–201. DOI: [10.1111/j.1467-8659.2009.01590.x](https://doi.org/10.1111/j.1467-8659.2009.01590.x).
- [107] Philip B Gorelick. "The global burden of stroke: persistent and disabling." In: *The Lancet Neurology* 18.5 (2019), pp. 417–418. DOI: [10.1016/S1474-4422\(19\)30030-4](https://doi.org/10.1016/S1474-4422(19)30030-4).
- [108] Akshat Gotra, Lojan Sivakumaran, Gabriel Chartrand, Kim-Nhien Vu, Franck Vandenbroucke-Menu, Claude Kauffmann, Samuel Kadoury, Benoît Gallix, Jacques A. de Guise, and An Tang. "Liver segmentation: indications, techniques and future directions." In: *Insights into Imaging* 8.4 (2017), pp. 377–392. DOI: [10.1007/s13244-017-0558-1](https://doi.org/10.1007/s13244-017-0558-1).
- [109] Leonid Goubergrits, Eugenie Riesenkauff, Pavlo Yevtushenko, Jens Schaller, Ulrich Kertzscher, Anja Hennemuth, Felix Berger, Stephan Schubert, and Titus Kuehne. "MRI-based computational fluid dynamics for diagnosis and treatment prediction: Clinical validation study in patients with coarctation of aorta." In: *Journal of Magnetic Resonance Imaging* 41.4 (2015), pp. 909–916.
- [110] Leonid Goubergrits, Jens Schaller, Ulrich Kertzscher, Nils van den Bruck, Kai Pöthkow, Christoph Petz, Hans-Christian Hege, and Andreas Spuler. "Statistical wall shear stress maps of ruptured and unruptured middle cerebral artery aneurysms." In: *Journal of the Royal Society Interface* 9.69 (2012), pp. 677–688. DOI: [10.1098/rsif.2011.0490](https://doi.org/10.1098/rsif.2011.0490).
- [111] Pavel A. Govyadinov, Tasha Womack, Jason L. Eriksen, Guoning Chen, and David Mayerich. "Robust Tracing and Visualization of Heterogeneous Microvascular Networks." In: *IEEE Transactions on Visualization and Computer Graphics* 25.4 (2019), pp. 1760–1773. DOI: [10.1109/tvcg.2018.2818701](https://doi.org/10.1109/tvcg.2018.2818701).

- [112] Pavel A. Govyadinov, Tasha Womack, Jason L. Eriksen, David Mayerich, and Guoning Chen. “Graph-Assisted Visualization of Microvascular Networks.” In: *IEEE Visualization Conference Short Papers*. 2019, pp. 6–10. DOI: [10.1109/visual.2019.8933682](https://doi.org/10.1109/visual.2019.8933682).
- [113] Henry Gray. *Anatomy of the Human Body*. Ed. by Warren H. Lewis. Lea & Febiger, 1918.
- [114] Roman Grothausmann. “Providing values of adjacent voxel with vtkDiscreteMarchingCubes.” In: *The VTK Journal* (2016). DOI: [10.54294/2aeqx3](https://doi.org/10.54294/2aeqx3).
- [115] Xianfeng Gu, Steven J. Gortler, and Hugues Hoppe. “Geometry Images.” In: *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*. 2002, pp. 355–361. DOI: [10.1145/566570.566589](https://doi.org/10.1145/566570.566589).
- [116] Bruno Guerciotti and Christian Vergara. “Computational Comparison Between Newtonian and Non-Newtonian Blood Rheologies in Stenotic Vessels.” In: *Biomedical Technology*. Springer International Publishing, 2017, pp. 169–183. DOI: [10.1007/978-3-319-59548-1_10](https://doi.org/10.1007/978-3-319-59548-1_10).
- [117] Tobias Günther, Christian Rössl, and Holger Theisel. “Opacity optimization for 3D line fields.” In: *ACM Transactions on Graphics* 32.4 (2013), 120:1–120:8. DOI: [10.1145/2461912.2461930](https://doi.org/10.1145/2461912.2461930).
- [118] Rajiv Gupta, Arnold C. Cheung, Soenke H. Bartling, Jennifer Lisauskas, Michael Grasruck, Christianne Leidecker, Bernhard Schmidt, Thomas Flohr, and Thomas J. Brady. “Flat-Panel Volume CT: Fundamental Principles, Technology, and Applications.” In: *RadioGraphics* 28.7 (2008), pp. 2009–2022. DOI: [10.1148/rg.287085004](https://doi.org/10.1148/rg.287085004).
- [119] Carsten Gutwenger, Michael Jünger, Karsten Klein, Joachim Kupke, Sebastian Leipert, and Petra Mutzel. “A new approach for visualizing UML class diagrams.” In: *Proceedings of the ACM Symposium on Software Visualization*. 2003, pp. 179–188. DOI: [10.1145/774833.774859](https://doi.org/10.1145/774833.774859).
- [120] Alison Halliday, Michael Harrison, Elizabeth Hayter, Xiangling Kong, Averil Mansfield, Joanna Marro, Hongchao Pan, Richard Peto, John Potter, Kazem Rahimi, et al. “10-year stroke prevention after successful carotid endarterectomy for asymptomatic stenosis (ACST-1): a multicentre randomised trial.” In: *The Lancet* 376.9746 (2010), pp. 1074–1084. DOI: [10.1016/S0140-6736\(10\)61197-X](https://doi.org/10.1016/S0140-6736(10)61197-X).
- [121] Ethan A Halm, Stanley Tuhim, Jason J Wang, Caron Rockman, Thomas S Riles, and Mark R Chassin. “Risk factors for perioperative death and stroke after carotid endarterectomy: results of the New York carotid artery surgery study.” In: *Stroke* 40.1 (2009), pp. 221–229. DOI: [10.1161/STROKEAHA.108.524785](https://doi.org/10.1161/STROKEAHA.108.524785).

- [122] K. Hameeteman, M. Freiman, M. A. Zuluaga, et al. "Carotid Lumen Segmentation and Stenosis Grading Challenge." In: *The MIDAS Journal* (2010). DOI: [10.54294/fl5463](https://doi.org/10.54294/fl5463).
- [123] Graeme J Hankey. "Stroke." In: *The Lancet* 389.10069 (2017), pp. 641–654. DOI: [10.1016/s0140-6736\(16\)30962-x](https://doi.org/10.1016/s0140-6736(16)30962-x).
- [124] Eline M. J. Hartman, Giuseppe De Nisco, Frank J. H. Gijzen, Suze-Anne Korteland, Anton F. W. van der Steen, Joost Daemen, and Jolanda J. Wentzel. "The definition of low wall shear stress and its effect on plaque progression estimation in human coronary arteries." In: *Scientific Reports* 11.1 (2021). DOI: [10.1038/s41598-021-01232-3](https://doi.org/10.1038/s41598-021-01232-3).
- [125] P. Hastreiter, C. Rezk-Salama, B. Tomandl, K. E. W. Eberhardt, and T. Ertl. "Fast analysis of intracranial aneurysms based on interactive direct volume rendering and CTA." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 1998, pp. 660–669. DOI: [10.1007/BFb0056252](https://doi.org/10.1007/BFb0056252).
- [126] Sha He, Ruping Dai, Bin Lu, Cheng Cao, Hua Bai, and Baolian Jing. "Medial Axis Reformation: A New Visualization Method for CT Angiography." In: *Academic Radiology* 8.8 (2001), pp. 726–733. DOI: [10.1016/s1076-6332\(03\)80579-2](https://doi.org/10.1016/s1076-6332(03)80579-2).
- [127] Anja Hennemuth, Ola Friman, Christian Schumann, Jelena Bock, Johann Drexl, Markus Hüllebrand, Michael Markl, and Heinz-Otto Peitgen. "Fast interactive exploration of 4D MRI flow data." In: *Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling*. Vol. 7964. 2011, 79640E. DOI: [10.1117/12.878202](https://doi.org/10.1117/12.878202).
- [128] Fernando V. Higuera, Natascha Sauber, Bernd Tomandl, Christopher Nimsky, Guenther Greiner, and Peter Hastreiter. "Automatic adjustment of bidimensional transfer functions for direct volume visualization of intracranial aneurysms." In: *Medical Imaging 2004: Visualization, Image-Guided Procedures, and Display*. Vol. 5367. 2004, pp. 275–284. DOI: [10.1117/12.535534](https://doi.org/10.1117/12.535534).
- [129] Fernando V. Higuera, Natascha Sauber, Bernd Tomandl, Christopher Nimsky, Günther Greiner, and Peter Hastreiter. "Enhanced 3D-visualization of intracranial aneurysms involving the skull base." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 2003, pp. 256–263. DOI: [10.1007/978-3-540-39903-2_32](https://doi.org/10.1007/978-3-540-39903-2_32).
- [130] Marius Hografer, Magnus Heitzler, and Hans-Jörg Schulz. "The State of the Art in Map-Like Visualization." In: *Computer Graphics Forum* 39.3 (2020), pp. 647–674. DOI: [10.1111/cgf.14031](https://doi.org/10.1111/cgf.14031).

- [131] Yiemeng Hoi, Bruce A Wasserman, Edward G Lakatta, and David A. Steinman. "Carotid bifurcation hemodynamics in older adults: effect of measured versus assumed flow waveform." In: *Journal of Biomechanical Engineering* 132.7 (2010). DOI: [10.1115/1.4001265](https://doi.org/10.1115/1.4001265).
- [132] Niels H. L. C. de Hoon, Roy F. P. van Pelt, Andrei C. Jalba, and Anna Vilanova. "4D MRI flow coupled to physics-based fluid simulation for blood-flow visualization." In: *Computer Graphics Forum* 33.3 (2014), pp. 121–130. DOI: [10.1111/cgf.12368](https://doi.org/10.1111/cgf.12368).
- [133] Peter R. Hoskins, Kevin Martin, and Abigail Thrush, eds. *Diagnostic Ultrasound: Physics and Equipment*. 3rd ed. CRC Press, 2019. DOI: [10.1201/9781138893603](https://doi.org/10.1201/9781138893603).
- [134] Petra Isenberg, Florian Heimerl, Steffen Koch, Tobias Isenberg, Panpan Xu, Charles D. Stolper, Michael Sedlmair, Jian Chen, Torsten Moller, and John Stasko. "Vispubdata.org: A Metadata Collection About IEEE Visualization (VIS) Publications." In: *IEEE Transactions on Visualization and Computer Graphics* 23.9 (2017), pp. 2199–2206. DOI: [10.1109/tvcg.2016.2615308](https://doi.org/10.1109/tvcg.2016.2615308).
- [135] Richard Izzo, David Steinman, Simone Manini, and Luca Antiga. "The vascular modeling toolkit: a Python library for the analysis of tubular structures in medical images." In: *Journal of Open Source Software* 3.25 (2018), p. 745. DOI: [10.21105/joss.00745](https://doi.org/10.21105/joss.00745).
- [136] Alec Jacobson, Daniele Panozzo, et al. *libigl: A simple C++ geometry processing library*. 2019. URL: <https://libigl.github.io/>.
- [137] G. Janiga, P. Berg, S. Sugiyama, K. Kono, and D. A. Steinman. "The computational fluid dynamics rupture challenge 2013—phase I: prediction of rupture status in intracranial aneurysms." In: *American Journal of Neuroradiology* 36.3 (2015), pp. 530–536.
- [138] Jochen Jankowai, Rickard Englund, Timo Ropinski, and Ingrid Hotz. "Interactive 4D MRI blood flow exploration and analysis using line predicates." In: *Proceedings of SIGRAD*. 2016, pp. 35–42.
- [139] Muhammad Moazzam Jawaaid, Sanam Narejo, Imran Ali Qureshi, and Nasrullah Pirzada. "A Review of the State-of-the-Art Methods for Non-calcified Plaque Detection in Cardiac CT Angiography." In: *International Journal of Computer Theory and Engineering* 10.3 (2018), pp. 84–92. DOI: [10.7763/ijcte.2018.v10.1204](https://doi.org/10.7763/ijcte.2018.v10.1204).
- [140] W. D. Jeans. "The development and use of digital subtraction angiography." In: *The British Journal of Radiology* 63.747 (1990), pp. 161–168. DOI: [10.1259/0007-1285-63-747-161](https://doi.org/10.1259/0007-1285-63-747-161).

- [141] Yongkweon Jeon, Joong-Ho Won, and Sungroh Yoon. "Massively Parallel Energy Space Exploration for Uncluttered Visualization of Vascular Structures." In: *IEEE Transactions on Biomedical Engineering* 60.1 (2013), pp. 240–244. DOI: [10.1109/tbme.2012.2214386](https://doi.org/10.1109/tbme.2012.2214386).
- [142] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller. "CPR – curved planar reformation." In: *IEEE Visualization*. 2002, pp. 37–44. DOI: [10.1109/VISUAL.2002.1183754](https://doi.org/10.1109/VISUAL.2002.1183754).
- [143] A. Kanitsar, D. Fleischmann, R. Wegenkittl, D. Sandner, P. Felkel, and M. E. Gröller. "Computed tomography angiography: a case study of peripheral vessel investigation." In: *Proceedings Visualization, 2001. VIS '01*. IEEE, 2001. DOI: [10.1109/visual.2001.964555](https://doi.org/10.1109/visual.2001.964555).
- [144] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. E. Gröller. "Advanced curved planar reformation: flattening of vascular structures." In: *IEEE Conference on Visualization*. 2003. DOI: [10.1109/visual.2003.1250353](https://doi.org/10.1109/visual.2003.1250353).
- [145] Armin Kanitsar, Dominik Fleischmann, Rainer Wegenkittl, and M. Eduard Gröller. "Diagnostic Relevant Visualization of Vascular Structures." In: *Scientific Visualization: The Visual Extraction of Knowledge from Data. Mathematics and Visualization*. Ed. by G. P. Bonneau, T. Ertl, and G. M. Nielson. Berlin, Heidelberg: Springer, 2006, pp. 207–228. DOI: [10.1007/3-540-30790-7_13](https://doi.org/10.1007/3-540-30790-7_13).
- [146] Mathias Kanzler, Florian Ferstl, and Rüdiger Westermann. "Line density control in screen-space via balanced line hierarchies." In: *Computers & Graphics* 61 (2016), pp. 29–39. DOI: [10.1016/j.cag.2016.08.001](https://doi.org/10.1016/j.cag.2016.08.001).
- [147] Rashed Karim et al. "Surface flattening of the human left atrium and proof-of-concept clinical applications." In: *Computerized Medical Imaging and Graphics* 38.4 (2014), pp. 251–266. DOI: [10.1016/j.compmedimag.2014.01.004](https://doi.org/10.1016/j.compmedimag.2014.01.004).
- [148] Piotr Kaszczewski, Michał Elwertowski, Jerzy Leszczyński, Tomasz Ostrowski, Joanna Kaszczewska, Tomasz Brzeziński, Daniel Jarosz, Siavash Świeczkowski-Feiz, and Zbigniew Gałazka. "Volumetric Flow Assessment in Extracranial Arteries in Patients with 70–99% Internal Carotid Artery Stenosis." In: *Diagnostics* 12.9 (2022). DOI: [10.3390/diagnostics12092216](https://doi.org/10.3390/diagnostics12092216).
- [149] Lilli Kaufhold, Andreas Harloff, Christian Schumann, Axel J. Krafft, Juergen Hennig, and Anja Hennemuth. "Image-based assessment of uncertainty in quantification of carotid lumen." In: *Journal of Medical Imaging* 5.3 (2018), p. 034003. DOI: [10.1117/1.jmi.5.3.034003](https://doi.org/10.1117/1.jmi.5.3.034003).

- [150] Marta Kersten-Oertel, Sean Jy-Shyang Chen, and D. Louis Collins. "An evaluation of depth enhancing perceptual cues for vascular volume visualization in neurosurgery." In: *IEEE Transactions on Visualization and Computer Graphics* 20.3 (2013), pp. 391–403. DOI: [10.1109/tvcg.2013.240](https://doi.org/10.1109/tvcg.2013.240).
- [151] Ron Kikinis, Steve D. Pieper, and Kirby G. Vosburgh. "3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support." In: *Intraoperative Imaging and Image-Guided Therapy*. 2013, pp. 277–289. DOI: [10.1007/978-1-4614-7657-3_19](https://doi.org/10.1007/978-1-4614-7657-3_19).
- [152] Kyungyoon Kim, John V. Carlis, and Daniel F. Keefe. "Comparison techniques utilized in spatial 3D and 4D data visualizations: A survey and future directions." In: *Computers & Graphics* 67 (2017), pp. 138–147.
- [153] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. "Optimization by Simulated Annealing." In: *Science* 220.4598 (1983), pp. 671–680. DOI: [10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671).
- [154] Andreas Kjellin, Lars Winkler Pettersson, Stefan Seipel, and Mats Lind. "Evaluating 2D and 3D visualizations of spatiotemporal information." In: *ACM Transactions on Applied Perception* 7.3 (2010), 19:1–19:23. DOI: [10.1145/1773965.1773970](https://doi.org/10.1145/1773965.1773970).
- [155] Stephen G. Kobourov, Tamara Mchedlidze, and Laura Vonessen. "Gestalt Principles in Graph Drawing." In: *Graph Drawing and Network Visualization*. Ed. by Emilio Di Giacomo and Anna Lubiw. 2015, pp. 558–560. DOI: [10.1007/978-3-319-27261-0_50](https://doi.org/10.1007/978-3-319-27261-0_50).
- [156] Benjamin Köhler, Silvia Born, Roy F. P. van Pelt, Anja Hennemuth, Uta Preim, and Bernhard Preim. "A Survey of Cardiac 4D PC-MRI Data Processing." In: *Computer Graphics Forum* 36.6 (2016), pp. 5–35. DOI: [10.1111/cgf.12803](https://doi.org/10.1111/cgf.12803).
- [157] Benjamin Köhler, Rocco Gasteiger, Uta Preim, Holger Theisel, Matthias Gutberlet, and Bernhard Preim. "Semi-automatic vortex extraction in 4D PC-MRI cardiac blood flow data using line predicates." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2773–2782. DOI: [10.1109/TVCG.2013.189](https://doi.org/10.1109/TVCG.2013.189).
- [158] Benjamin Köhler, Monique Meuschke, Uta Preim, Katharina Fischbach, Matthias Gutberlet, and Bernhard Preim. "2D plot visualization of aortic vortex flow in cardiac 4D PC-MRI data." In: *Bildverarbeitung für die Medizin 2015*. Ed. by Heinz Handels, Thomas Martin Deserno, Hans-Peter Meinzer, and Thomas Tolxdorff. Springer Berlin Heidelberg, 2015, pp. 257–262. DOI: [10.1007/978-3-662-46224-9_45](https://doi.org/10.1007/978-3-662-46224-9_45).

- [159] Benjamin Köhler, Uta Preim, Matthias Grothoff, Matthias Guterlet, Katharina Fischbach, and Bernhard Preim. "Motion-aware stroke volume quantification in 4D PC-MRI data of the human aorta." In: *International Journal of Computer Assisted Radiology and Surgery* 11.2 (2015), pp. 169–179. DOI: [10.1007/s11548-015-1256-4](https://doi.org/10.1007/s11548-015-1256-4).
- [160] Benjamin Köhler, Uta Preim, Matthias Grothoff, Matthias Guterlet, Katharina Fischbach, and Bernhard Preim. "Robust cardiac function assessment in 4D PC-MRI data of the aorta and pulmonary artery." In: *Computer Graphics Forum* 35.1 (2016), pp. 32–43.
- [161] E. L. Koua, A. Maceachren, and M.-J. Kraak. "Evaluating the usability of visualization methods in an exploratory geovisualization environment." In: *International Journal of Geographical Information Science* 20.4 (2006), pp. 425–448. DOI: [10.1080/13658810600607550](https://doi.org/10.1080/13658810600607550).
- [162] Peter Kovesi. *Good Colour Maps: How to Design Them*. 2015. DOI: [10.48550/ARXIV.1509.03700](https://doi.org/10.48550/ARXIV.1509.03700). URL: <https://colorcet.com>.
- [163] Julian Kreiser, Monique Meuschke, Gabriel Mistelbauer, Bernhard Preim, and Timo Ropinski. "A survey of flattening-based medical visualization techniques." In: *Computer Graphics Forum* 37.3 (2018), pp. 597–624. DOI: [10.1111/cgf.13445](https://doi.org/10.1111/cgf.13445).
- [164] Jan Kretschmer, Bernhard Preim, and Marc Stamminger. "Bilateral Depth Filtering for Enhanced Vessel Reformation." In: *EuroVis – Short Papers*. Ed. by N. Elmqvist, M. Hlawitschka, and J. Kennedy. The Eurographics Association, 2014. DOI: [10.2312/EUROVISSHORT.20141150](https://doi.org/10.2312/EUROVISSHORT.20141150).
- [165] Harinarayan Krishnan, Christoph Garth, Jens Guhring, Mehmet Akif Gulsun, Andreas Greiser, and Kenneth I Joy. "Analysis of time-dependent flow-sensitive PC-MRI data." In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2011), pp. 966–977. DOI: [10.1109/tvcg.2011.80](https://doi.org/10.1109/tvcg.2011.80).
- [166] Karl Krissian, Gregoire Malandain, and Nicholas Ayache. "Directional anisotropic diffusion applied to segmentation of vessels in 3D images." In: *Scale-Space Theory in Computer Vision*. Springer Berlin Heidelberg, 1997, pp. 345–348. DOI: [10.1007/3-540-63167-4_68](https://doi.org/10.1007/3-540-63167-4_68).
- [167] Christoph Kubisch, Sylvia Glaßer, Mathias Neugebauer, and Bernhard Preim. "Vessel visualization with volume rendering." In: *Visualization in Medicine and Life Sciences II*. 2012, pp. 109–132. DOI: [10.1007/978-3-642-21608-4_7](https://doi.org/10.1007/978-3-642-21608-4_7).

- [168] C. Kuehnel, A. Hennemuth, T. Boskamp, S. Oeltze, S. Bock, S. Krass, B. Preim, and H.-O. Peitgen. "New software assistants for cardiovascular diagnosis." In: *INFORMATIK 2006 – Informatik für Menschen, Band 1*. Ed. by Christian Hochberger and Rüdiger Liskowsky. Bonn: Gesellschaft für Informatik e.V., 2006, pp. 491–498.
- [169] Alexander Kuhn, Dirk J. Lehmann, Rocco Gasteiger, Mathias Neugebauer, Bernhard Preim, and Holger Theisel. "A Clustering-based Visualization Technique to Emphasize Meaningful Regions of Vector Fields." In: *Vision, Modeling and Visualization*. 2011, pp. 191–198.
- [170] Wilhelm Kutta. "Beitrag zur näherungsweise Integration totaler Differentialgleichungen." In: *Z. Math. Phys.* 46 (1901), pp. 435–453.
- [171] Jonas Lamy, Odyssee Merveille, Bertrand Kerautret, Nicolas Passat, and Antoine Vacavant. "Vesselness Filters: A Survey with Benchmarks Applied to Liver Imaging." In: *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021. DOI: [10.1109/icpr48806.2021.9412362](https://doi.org/10.1109/icpr48806.2021.9412362).
- [172] Jonas Lantz. "On Aortic Blood Flow Simulations: Scale-Resolved Image-Based CFD." PhD thesis. Linköping University, 2013. ISBN: 978-91-7519-720-3.
- [173] I. Larrabide, M. L. Aguilar, H. G. Morales, A. J. Geers, Z. Kulcsár, D. Rufenacht, and A. F. Frangi. "Intra-aneurysmal pressure and flow changes induced by flow diverters: relation to aneurysm size and shape." In: *American Journal of Neuroradiology* 34.4 (2013), pp. 816–822.
- [174] Kai Lawonn, Rocco Gasteiger, and Bernhard Preim. "Adaptive surface visualization of vessels with animated blood flow." In: *Computer Graphics Forum* 33.8 (2014), pp. 16–27. DOI: [10.1111/cgf.12355](https://doi.org/10.1111/cgf.12355).
- [175] Kai Lawonn, Sylvia Glaßer, Anna Vilanova, Bernhard Preim, and Tobias Isenberg. "Occlusion-free blood flow animation with wall thickness visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2015), pp. 728–737. DOI: [10.1109/TVCG.2015.2467961](https://doi.org/10.1109/TVCG.2015.2467961).
- [176] Kai Lawonn, Monique Meuschke, Ralph Wickenhöfer, Bernhard Preim, and Klaus Hildebrandt. "A Geometric Optimization Approach for the Detection and Segmentation of Multiple Aneurysms." In: *Computer Graphics Forum* 38.3 (2019), pp. 413–425. DOI: [10.1111/cgf.13699](https://doi.org/10.1111/cgf.13699).

- [177] Noah Lee and Matthias Rasch. "Tangential curved planar reformation for topological and orientation invariant visualization of vascular trees." In: *International Conference of the IEEE Engineering in Medicine and Biology Society*. 2006. DOI: [10.1109/iembs.2006.259518](https://doi.org/10.1109/iembs.2006.259518).
- [178] Sang Hyuk Lee, Kap-Soo Han, Nahmkeon Hur, Young I. Cho, and Seul-Ki Jeong. "The Effect of Patient-Specific non-Newtonian Blood Viscosity on Arterial Hemodynamics Predictions." In: *Journal of Mechanics in Medicine and Biology* 19.08 (2019), p. 1940054. DOI: [10.1142/S0219519419400542](https://doi.org/10.1142/S0219519419400542).
- [179] Sang-Wook Lee, Luca Antiga, and David A. Steinman. "Correlations among indicators of disturbed flow at the normal carotid bifurcation." In: *Journal of Biomechanical Engineering* 131.6 (2009). DOI: [10.1115/1.3127252](https://doi.org/10.1115/1.3127252).
- [180] T. C. Lee, R. L. Kashyap, and C. N. Chu. "Building Skeleton Models via 3-D Medial Surface Axis Thinning Algorithms." In: *CVGIP: Graphical Models and Image Processing* 56.6 (1994), pp. 462–478. DOI: [10.1006/cgip.1994.1042](https://doi.org/10.1006/cgip.1994.1042).
- [181] Whal Lee. "General principles of carotid Doppler ultrasonography." In: *Ultrasonography* 33.1 (2013), pp. 11–17. DOI: [10.14366/usg.13018](https://doi.org/10.14366/usg.13018).
- [182] T. M. Lehmann, C. Gonner, and K. Spitzer. "Survey: interpolation methods in medical image processing." In: *IEEE Transactions on Medical Imaging* 18.11 (1999), pp. 1049–1075. DOI: [10.1109/42.816070](https://doi.org/10.1109/42.816070).
- [183] David Lesage, Elsa D. Angelini, Isabelle Bloch, and Gareth Funka-Lea. "A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes." In: *Medical Image Analysis* 13.6 (2009), pp. 819–845. DOI: [10.1016/j.media.2009.07.011](https://doi.org/10.1016/j.media.2009.07.011).
- [184] Michael Levitt and Arieh Warshel. "Computer simulation of protein folding." In: *Nature* 253.5494 (1975), pp. 694–698. DOI: [10.1038/253694a0](https://doi.org/10.1038/253694a0).
- [185] Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. "Least squares conformal maps for automatic texture atlas generation." In: *ACM Transactions on Graphics* 21.3 (2002), pp. 362–371. DOI: [10.1145/566654.566590](https://doi.org/10.1145/566654.566590).
- [186] Cong-Hui Li, Bu-Lang Gao, Ji-Wei Wang, Jian-Feng Liu, Hui Li, and Song-Tao Yang. "Hemodynamic Factors Affecting Carotid Sinus Atherosclerotic Stenosis." In: *World Neurosurgery* 121 (2019), e262–e276. DOI: [10.1016/j.wneu.2018.09.091](https://doi.org/10.1016/j.wneu.2018.09.091).

- [187] EY Liang, Michael Chan, JH Hsiang, SB Walkden, WS Poon, WW Lam, and C Metreweli. "Detection and assessment of intracranial aneurysms: value of CT angiography with shaded-surface display." In: *American Journal of Roentgenology* 165.6 (1995), pp. 1497–1502. DOI: [10.2214/ajr.165.6.7484596](https://doi.org/10.2214/ajr.165.6.7484596).
- [188] Nils Lichtenberg, Bastian Kraye, Christian Hansen, Stefan Müller, and Kai Lawonn. "Distance Field Visualization and 2D Abstraction of Vessel Tree Structures with on-the-fly Parameterization." In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2019, pp. 265–278. DOI: [10.2312/VCBM.20191251](https://doi.org/10.2312/VCBM.20191251).
- [189] Nils Lichtenberg and Kai Lawonn. "Parameterization, Feature Extraction and Binary Encoding for the Visualization of Tree-Like Structures." In: *Computer Graphics Forum* 39.1 (2020), pp. 497–510. DOI: [10.1111/cgf.13888](https://doi.org/10.1111/cgf.13888).
- [190] Nils Lichtenberg, Noeska Smit, Christian Hansen, and Kai Lawonn. "Real-time field aligned stripe patterns." In: *Computers & Graphics* 74 (2018), pp. 137–149. DOI: [10.1016/j.cag.2018.04.008](https://doi.org/10.1016/j.cag.2018.04.008).
- [191] Grace W. Lindsay. "Convolutional neural networks as a model of the visual system: Past, present, and future." In: *Journal of Cognitive Neuroscience* 33.10 (2021), pp. 2017–2031.
- [192] R. J. Lipton, S. C. North, and J. S. Sandberg. "A method for drawing graphs." In: *Proceedings of the First Annual Symposium on Computational Geometry*. ACM Press, 1985. DOI: [10.1145/323233.323254](https://doi.org/10.1145/323233.323254).
- [193] Vladimir D. Liseikin. *Grid Generation Methods*. 3rd ed. Springer, 2017. DOI: [10.1007/978-3-319-57846-0](https://doi.org/10.1007/978-3-319-57846-0).
- [194] Paweł Liskowski and Krzysztof Krawiec. "Segmenting Retinal Blood Vessels With Deep Neural Networks." In: *IEEE Transactions on Medical Imaging* 35.11 (2016), pp. 2369–2380. DOI: [10.1109/tmi.2016.2546227](https://doi.org/10.1109/tmi.2016.2546227).
- [195] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen AWM Van Der Laak, Bram Van Ginneken, and Clara I. Sánchez. "A survey on deep learning in medical image analysis." In: *Medical Image Analysis* 42 (2017), pp. 60–88.
- [196] Ligang Liu, Lei Zhang, Yin Xu, Craig Gotsman, and Steven J. Gortler. "A local/global approach to mesh parameterization." In: *Computer Graphics Forum* 27.5 (2008), pp. 1495–1504. DOI: [10.1111/j.1467-8659.2008.01290.x](https://doi.org/10.1111/j.1467-8659.2008.01290.x).

- [197] Yang Liu and Jeffrey Heer. "Somewhere over the rainbow: An empirical assessment of quantitative colormaps." In: *CHI Conference on Human Factors in Computing Systems*. 2018. DOI: [10.1145/3173574.3174172](https://doi.org/10.1145/3173574.3174172).
- [198] D. Lopes, H. Puga, J. Teixeira, and R. Lima. "Blood flow simulations in patient-specific geometries of the carotid artery: A systematic review." In: *Journal of Biomechanics* 111 (2020), p. 110019. ISSN: 0021-9290. DOI: [10.1016/j.jbiomech.2020.110019](https://doi.org/10.1016/j.jbiomech.2020.110019).
- [199] L. M. Lorigo, O. D. Faugeras, W. E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin. "CURVES: Curve evolution for vessel segmentation." In: *Medical Image Analysis* 5.3 (2001), pp. 195–206. DOI: [10.1016/s1361-8415\(01\)00040-8](https://doi.org/10.1016/s1361-8415(01)00040-8).
- [200] Ramon Luengo-Fernandez, Nicola L.M. Paul, Alastair M. Gray, Sarah T. Pendlebury, Linda M. Bull, Sarah J.V. Welch, Fiona C. Cuthbertson, and Peter M. Rothwell. "Population-Based Study of Disability and Institutionalization After Transient Ischemic Attack and Stroke: 10-Year Results of the Oxford Vascular Study." In: *Stroke* 44.10 (2013), pp. 2854–2861. DOI: [10.1161/strokeaha.113.001584](https://doi.org/10.1161/strokeaha.113.001584).
- [201] Xinrong Lv, Xinbo Gao, and Hua Zou. "Interactive curved planar reformation based on snake model." In: *Computerized Medical Imaging and Graphics* 32.8 (2008), pp. 662–669. DOI: [10.1016/j.compmedimag.2008.08.002](https://doi.org/10.1016/j.compmedimag.2008.08.002).
- [202] YingLiang Ma, Rashed Karim, R. James Housden, Geert Gijssbers, Roland Bullens, Christopher Aldo Rinaldi, Reza Razavi, Tobias Schaeffter, and Kawal S. Rhode. "Cardiac Unfold: A Novel Technique for Image-Guided Cardiac Catheterization Procedures." In: *Information Processing in Computer-Assisted Interventions*. Springer Berlin Heidelberg, 2012, pp. 104–114. DOI: [10.1007/978-3-642-30618-1_11](https://doi.org/10.1007/978-3-642-30618-1_11).
- [203] Rashindra Manniesing, Michiel Schaap, Sietske Rozie, Reinhard Hameeteman, Danijela Vukadinovic, Aad van der Lugt, and Wiro Niessen. "Robust CTA lumen segmentation of the atherosclerotic carotid artery bifurcation in a large patient population." In: *Medical Image Analysis* 14.6 (2010), pp. 759–769. ISSN: 1361-8415. DOI: [10.1016/j.media.2010.05.001](https://doi.org/10.1016/j.media.2010.05.001).
- [204] Rashindra Manniesing, Max A. Viergever, and Wiro J. Niessen. "Vessel enhancing diffusion: A scale space representation of vessel structures." In: *Medical Image Analysis* 10.6 (2006), pp. 815–825. DOI: [10.1016/j.media.2006.06.003](https://doi.org/10.1016/j.media.2006.06.003).
- [205] Joseph Marino and Arie Kaufman. "Planar visualization of treelike structures." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 906–915. DOI: [10.1109/TVCG.2015.2467413](https://doi.org/10.1109/TVCG.2015.2467413).

- [206] M. Markl, S. Brendecke, J. Simon, A. Frydrychowicz, and A. Harloff. "Coregistration of wall shear stress and plaque distribution within the thoracic aorta of acute stroke patients." In: *Proceedings of the International Society for Magnetic Resonance in Medicine*. Vol. 18. 2010, p. 63.
- [207] Henry C. McGill, C. Alex McMahan, and Samuel S. Gidding. "Preventing Heart Disease in the 21st Century." In: *Circulation* 117.9 (2008), pp. 1216–1227. DOI: [10.1161/circulationaha.107.717033](https://doi.org/10.1161/circulationaha.107.717033).
- [208] Tony McLoughlin, Mark W. Jones, Robert S. Laramée, Rami Malki, Ian Masters, and Charles D. Hansen. "Similarity Measures for Enhancing Interactive Streamline Seeding." In: *IEEE Transactions on Visualization and Computer Graphics* 19.8 (2013), pp. 1342–1353. DOI: [10.1109/TVCG.2012.150](https://doi.org/10.1109/TVCG.2012.150).
- [209] Gillian E. Mead, Stephanie C. Lewis, and Joanna M. Wardlaw. "Variability in Doppler ultrasound influences referral of patients for carotid surgery." In: *European Journal of Ultrasound* 12.2 (2000), pp. 137–143. DOI: [10.1016/s0929-8266\(00\)00111-7](https://doi.org/10.1016/s0929-8266(00)00111-7).
- [210] Jessica Benitez Mendieta, Davide Fontanarosa, Jiaqiu Wang, Phani Kumari Paritala, Tim McGahan, Thomas Lloyd, and Zhiyong Li. "The importance of blood rheology in patient-specific computational fluid dynamics simulation of stenotic carotid arteries." In: *Biomechanics and Modeling in Mechanobiology* 19.5 (2020), pp. 1477–1490. DOI: [10.1007/s10237-019-01282-7](https://doi.org/10.1007/s10237-019-01282-7).
- [211] Shanthi Mendis, Pekka Puska, Bo Norrving, et al. *Global Atlas on Cardiovascular Disease Prevention and Control*. Geneva: World Health Organization, 2011.
- [212] M. Meuschke, B. Köhler, U. Preim, B. Preim, and K. Lawonn. "Semi-automatic Vortex Flow Classification in 4D PC-MRI Data of the Aorta." In: *Computer Graphics Forum* 35.3 (2016), pp. 351–360. DOI: [10.1111/cgf.12911](https://doi.org/10.1111/cgf.12911).
- [213] Monique Meuschke, Tobias Günther, Philipp Berg, Ralph Wickenhofer, Bernhard Preim, and Kai Lawonn. "Visual Analysis of Aneurysm Data using Statistical Graphics." In: *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2019), pp. 997–1007. DOI: [10.1109/tvcg.2018.2864509](https://doi.org/10.1109/tvcg.2018.2864509).
- [214] Monique Meuschke, Bernhard Preim, and Kai Lawonn. "Aneulysis – A system for the visual analysis of aneurysm data." In: *Computers & Graphics* 98 (2021), pp. 197–209. DOI: [10.1016/j.cag.2021.06.001](https://doi.org/10.1016/j.cag.2021.06.001).

- [215] Monique Meuschke, Samuel Voß, Oliver Beuing, Bernhard Preim, and Kai Lawonn. “Combined visualization of vessel deformation and hemodynamics in cerebral aneurysms.” In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2017), pp. 761–770. DOI: [10.1109/TVCG.2016.2598795](https://doi.org/10.1109/TVCG.2016.2598795).
- [216] Monique Meuschke, Samuel Voß, Oliver Beuing, Bernhard Preim, and Kai Lawonn. “Glyph-based comparative stress tensor visualization in cerebral aneurysms.” In: *Computer Graphics Forum* 36.3 (2017), pp. 99–108. DOI: [10.1111/cgf.13171](https://doi.org/10.1111/cgf.13171).
- [217] Monique Meuschke, Samuel Voß, Franziska Gaidzik, Bernhard Preim, and Kai Lawonn. “Skyscraper visualization of multiple time-dependent scalar fields on surfaces.” In: *Computers & Graphics* 99 (2021), pp. 22–42. DOI: [10.1016/j.cag.2021.05.005](https://doi.org/10.1016/j.cag.2021.05.005).
- [218] Monique Meuschke, Samuel Voß, Bernhard Preim, and Kai Lawonn. “Exploration of blood flow patterns in cerebral aneurysms during the cardiac cycle.” In: *Computers & Graphics* 72 (2018), pp. 12–25. DOI: [10.1016/j.cag.2018.01.012](https://doi.org/10.1016/j.cag.2018.01.012).
- [219] H. Miao, G. Mistelbauer, C. Našel, and M. E. Gröller. “Visual Quantification of the Circle of Willis: An Automated Identification and Standardized Representation.” In: *Computer Graphics Forum* 36.6 (2016), pp. 393–404. DOI: [10.1111/cgf.12988](https://doi.org/10.1111/cgf.12988).
- [220] Gabriel Mistelbauer. “Smart Interactive Vessel Visualization in Radiology.” PhD thesis. Institute of Computer Graphics and Algorithms TU Wien, 2013.
- [221] Gabriel Mistelbauer, Anca Morar, Andrej Varchola, Rüdiger Scherthaner, Ivan Baclija, Arnold Köchl, Armin Kanitsar, Stefan Bruckner, and M. Eduard Gröller. “Vessel visualization using curvicircular feature aggregation.” In: *Computer Graphics Forum* 32.3pt2 (2013), pp. 231–240. DOI: [10.1111/cgf.12110](https://doi.org/10.1111/cgf.12110).
- [222] Gabriel Mistelbauer, Andrej Varchola, Hamed Bouzari, Juraj Starinsky, Arnold Köchl, Rüdiger Scherthaner, Dominik Fleischmann, M. Eduard Gröller, and Milos Sramek. “Centerline reformations of complex vascular structures.” In: *IEEE Pacific Visualization Symposium*. 2012, pp. 233–240. DOI: [10.1109/PacificVis.2012.6183596](https://doi.org/10.1109/PacificVis.2012.6183596).
- [223] Sara Moccia, Elena De Momi, Sara El Hadji, and Leonardo S. Mattos. “Blood vessel segmentation algorithms – Review of methods, datasets and evaluation metrics.” In: *Computer Methods and Programs in Biomedicine* 158 (2018), pp. 71–91. DOI: [10.1016/j.cmpb.2018.02.001](https://doi.org/10.1016/j.cmpb.2018.02.001).

- [224] Pim Moeskops, Jelmer M. Wolterink, Bas H. M. van der Velden, Kenneth G. A. Gilhuijs, Tim Leiner, Max A. Viergever, and Ivana Išgum. “Deep Learning for Multi-task Medical Image Segmentation in Multiple Modalities.” In: *Medical Image Computing and Computer-Assisted Intervention*. 2016, pp. 478–486. DOI: [10.1007/978-3-319-46723-8_55](https://doi.org/10.1007/978-3-319-46723-8_55).
- [225] Patrick Mullen, Yiying Tong, Pierre Alliez, and Mathieu Desbrun. “Spectral Conformal Parameterization.” In: *Computer Graphics Forum* 27.5 (2008), pp. 1487–1494. DOI: [10.1111/j.1467-8659.2008.01289.x](https://doi.org/10.1111/j.1467-8659.2008.01289.x).
- [226] Sandy Napel, Michael P Marks, Geoffrey D. Rubin, Michael D. Dake, Charles H. McDonnell, Samuel M. Song, Dieter R. Enzmann, and R. B. Jeffrey Jr. “CT angiography with spiral CT and maximum intensity projection.” In: *Radiology* 185.2 (1992), pp. 607–610. DOI: [10.1148/radiology.185.2.1410382](https://doi.org/10.1148/radiology.185.2.1410382).
- [227] Mathias Neugebauer, Rocco Gasteiger, Oliver Beuing, Volker Diehl, Martin Skalej, and Bernhard Preim. “Map displays for the analysis of scalar data on cerebral aneurysm surfaces.” In: *Computer Graphics Forum* 28.3 (2009), pp. 895–902. DOI: [10.1111/j.1467-8659.2009.01459.x](https://doi.org/10.1111/j.1467-8659.2009.01459.x).
- [228] Mathias Neugebauer, Rocco Gasteiger, Volker Diehl, Oliver Beuing, and Bernhard Preim. “Automatic generation of context visualizations for cerebral aneurysms from MRA datasets.” In: *International Journal of Computer Assisted Radiology and Surgery* 4.Supplement 1 (2009), pp. 112–113.
- [229] Giuseppe De Nisco, Ayla Hoogendoorn, Claudio Chiastra, Diego Gallo, Annette M. Kok, Umberto Morbiducci, and Jolanda J. Wentzel. “The impact of helical flow on coronary atherosclerotic plaque development.” In: *Atherosclerosis* 300 (2020), pp. 39–46. DOI: [10.1016/j.atherosclerosis.2020.01.027](https://doi.org/10.1016/j.atherosclerosis.2020.01.027).
- [230] Marta Nuñez-Garcia, Gabriel Bernardino, Ruben Doste, Jichao Zhao, Oscar Camara, and Constantine Butakoff. “Standard Quasi-Conformal Flattening of the Right and Left Atria.” In: *Functional Imaging and Modeling of the Heart*. Ed. by Yves Coudière, Valéry Ozenne, Edward Vigmond, and Nejib Zemzemi. Springer International Publishing, 2019, pp. 85–93. DOI: [10.1007/978-3-030-21949-9_10](https://doi.org/10.1007/978-3-030-21949-9_10).
- [231] Martin J. O’Donnell, Siu Lim Chin, Sumathy Rangarajan, et al. “Global and regional effects of potentially modifiable risk factors associated with acute stroke in 32 countries (INTERSTROKE): a case-control study.” In: *The Lancet* 388.10046 (2016), pp. 761–775. DOI: [10.1016/s0140-6736\(16\)30506-2](https://doi.org/10.1016/s0140-6736(16)30506-2).

- [232] S. Oeltze, A. Kuß, F. Grothues, A. Hennemuth, and B. Preim. "Integrated Visualization of Morphologic and Perfusion Data for the Analysis of Coronary Artery Disease." In: *Proceedings of the Eighth Joint Eurographics / IEEE VGTC Conference on Visualization*. Eurographics Association, 2006, pp. 131–138.
- [233] Steffen Oeltze, Dirk J. Lehmann, Alexander Kuhn, Gábor Janiga, Holger Theisel, and Bernhard Preim. "Blood flow clustering and applications in virtual stenting of intracranial aneurysms." In: *IEEE Transactions on Visualization and Computer Graphics* 20.5 (2014), pp. 686–701. DOI: [10.1109/TVCG.2013.2297914](https://doi.org/10.1109/TVCG.2013.2297914).
- [234] Steffen Oeltze-Jafra, Monique Meuschke, Mathias Neugebauer, Sylvia Saalfeld, Kai Lawonn, Gabor Janiga, Hans-Christian Hege, Stefan Zachow, and Bernhard Preim. "Generation and visual exploration of medical flow data: Survey, research trends and future challenges." In: *Computer Graphics Forum* 38.1 (2019), pp. 87–125. DOI: [10.1111/cgf.13394](https://doi.org/10.1111/cgf.13394).
- [235] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial tessellations: concepts and applications of Voronoi diagrams*. 2nd ed. John Wiley & Sons, 2000.
- [236] OpenCFD Ltd. *OpenFOAM user guide*. Version v1906. 2019. URL: <https://www.openfoam.com/>.
- [237] Michael Oppermann and Tamara Munzner. "Data-First Visualization Design Studies." In: *IEEE Workshop on Evaluation and Beyond-Methodological Approaches to Visualization (BELIV)*. 2020, pp. 74–80. DOI: [10.1109/BELIV51497.2020.00016](https://doi.org/10.1109/BELIV51497.2020.00016).
- [238] David S. Paik, Christopher F. Beaulieu, R. Brooke Jeffrey, Geoffrey D. Rubin, and Sandy Napel. "Automated flight path planning for virtual endoscopy." In: *Medical Physics* 25.5 (1998), pp. 629–637. DOI: [10.1118/1.598244](https://doi.org/10.1118/1.598244).
- [239] Aditeya Pandey, Harsh Shukla, Geoffrey S. Young, Lei Qin, Amir A Zamani, Liangge Hsu, Raymond Huang, Cody Dunne, and Michelle A. Borkin. "Cerebrovis: designing an abstract yet spatially contextualized cerebral artery network visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2020), pp. 938–948. DOI: [10.1109/TVCG.2019.2934402](https://doi.org/10.1109/TVCG.2019.2934402).
- [240] Ronney B. Panerai. "Cerebral Autoregulation: From Models to Clinical Applications." In: *Cardiovascular Engineering* 8.1 (2007), pp. 42–59. DOI: [10.1007/s10558-007-9044-6](https://doi.org/10.1007/s10558-007-9044-6).
- [241] Daniele Panozzo, Enrico Puppo, and Luigi Rocca. "Efficient multi-scale curvature and crease estimation." In: *Proceedings of Computer Graphics, Computer Vision and Mathematics* 1.6 (2010).

- [242] Bruno Paun, Bart Bijmens, Tinen Iles, Paul A. Iaizzo, and Constantine Butakoff. "Patient independent representation of the detailed cardiac ventricular anatomy." In: *Medical Image Analysis* 35 (2017), pp. 270–287. DOI: [10.1016/j.media.2016.07.006](https://doi.org/10.1016/j.media.2016.07.006).
- [243] Theo Pavlidis. "A thinning algorithm for discrete binary images." In: *Computer Graphics and Image Processing* 13.2 (1980), pp. 142–157. DOI: [10.1016/s0146-664x\(80\)80037-2](https://doi.org/10.1016/s0146-664x(80)80037-2).
- [244] Roy F. P. van Pelt, Andrea Fuster, Geert G. H. Claassen, and Anna Vilanova. "Characterization of blood-flow patterns from phase-contrast MRI velocity fields." In: *EuroVis – Short Papers*. 2014. DOI: [10.2312/eurovisshort.20141158](https://doi.org/10.2312/eurovisshort.20141158).
- [245] Karl Perktold and Gerhard Rappitsch. "Computer simulation of local blood flow and vessel mechanics in a compliant carotid artery bifurcation model." In: *Journal of Biomechanics* 28.7 (1995), pp. 845–856. DOI: [10.1016/0021-9290\(95\)95273-8](https://doi.org/10.1016/0021-9290(95)95273-8).
- [246] Lynn S. Perlmutter and Helena Chang Chui. "Microangiopathy, the vascular basement membrane and Alzheimer's disease: a review." In: *Brain Research Bulletin* 24.5 (1990), pp. 677–686. DOI: [10.1016/0361-9230\(90\)90007-m](https://doi.org/10.1016/0361-9230(90)90007-m).
- [247] S. Pieper, M. Halle, and R. Kikinis. "3D Slicer." In: *2nd IEEE International Symposium on Biomedical Imaging: Macro to Nano*. 2005. DOI: [10.1109/isbi.2004.1398617](https://doi.org/10.1109/isbi.2004.1398617).
- [248] S. I. S. Pinto, E. Doutel, J. B. L. M. Campos, and J. M. Miranda. "Blood analog fluid flow in vessels with stenosis: Development of an OpenFOAM code to simulate pulsatile flow and elasticity of the fluid." In: *APCBEE Procedia* 7 (2013), pp. 73–79. DOI: [10.1016/j.apcbee.2013.08.015](https://doi.org/10.1016/j.apcbee.2013.08.015).
- [249] Senol Piskin, Onur Dur, and Kerem Pekkan. "Blood flow in realistic neonatal aorta using open source software OpenFOAM®." In: *8th International OpenFOAM® Workshop* (2013).
- [250] Andrzej Polanczyk, Michal Podgorski, Tomasz Wozniak, Ludomir Stefanczyk, and Michal Strzelecki. "Computational fluid dynamics as an engineering tool for the reconstruction of hemodynamics after carotid artery stenosis operation: a case study." In: *Medicina* 54.3 (2018), p. 42. DOI: [10.3390/medicina54030042](https://doi.org/10.3390/medicina54030042).
- [251] Bernhard Preim and Steffen Oeltze. "3D Visualization of Vasculature: An Overview." In: *Visualization in Medicine and Life Sciences. Mathematics and Visualization*. Ed. by L. Linsen, H. Hagen, and B. Hamann. Springer Berlin Heidelberg, 2008, pp. 39–59. DOI: [10.1007/978-3-540-72630-2_3](https://doi.org/10.1007/978-3-540-72630-2_3).
- [252] Bernhard Preim, Renata Raidou, Noeska Smit, and Kai Lawonn. "Advanced vessel visualization." In: *Visualization, Visual Analytics and Virtual Reality in Medicine*. Elsevier, 2023, pp. 27–57. DOI: [10.1016/b978-0-12-822962-0.00010-9](https://doi.org/10.1016/b978-0-12-822962-0.00010-9).

- [253] Bernhard Preim, Timo Ropinski, and Petra Isenberg. "A Critical Analysis of the Evaluation Practice in Medical Visualization." In: *Eurographics Workshop on Visual Computing for Biology and Medicine* (2018), pp. 45–56. DOI: [10.2312/VCBM.20181228](https://doi.org/10.2312/VCBM.20181228).
- [254] Pavle Prentašić, Morgan Heisler, Zaid Mammo, Sieun Lee, Andrew Merkur, Eduardo Navajas, Mirza Faisal Beg, Marinko Šarunic, and Sven Loncaric. "Segmentation of the foveal microvasculature using deep learning networks." In: *Journal of Biomedical Optics* 21.7 (2016), p. 075008. DOI: [10.1117/1.jbo.21.7.075008](https://doi.org/10.1117/1.jbo.21.7.075008).
- [255] Daniel J. Price. "Smoothed particle hydrodynamics and magnetohydrodynamics." In: *Journal of Computational Physics* 231.3 (2012), pp. 759–794. DOI: [10.1016/j.jcp.2010.12.011](https://doi.org/10.1016/j.jcp.2010.12.011).
- [256] Helen Purchase. "Which aesthetic has the greatest effect on human understanding?" In: *Graph Drawing*. 1997, pp. 248–261. DOI: [10.1007/3-540-63938-1_67](https://doi.org/10.1007/3-540-63938-1_67).
- [257] S. Rajsic, H. Gothe, H. H. Borba, G. Sroczynski, J. Vujicic, T. Toell, and U. Siebert. "Economic burden of stroke: a systematic review on post-stroke care." In: *The European Journal of Health Economics* 20.1 (2019), pp. 107–134. DOI: [10.1007/s10198-018-0984-0](https://doi.org/10.1007/s10198-018-0984-0).
- [258] A. Rerkasem, S. Orrapin, D. P. J. Howard, and K. Rerkasem. "Carotid endarterectomy for symptomatic carotid stenosis." In: *Cochrane Database of Systematic Reviews* 9 (2020). DOI: [10.1002/14651858.CD001081.pub4](https://doi.org/10.1002/14651858.CD001081.pub4).
- [259] Amarporn Rerkasem, Saritphat Orrapin, Dominic PJ Howard, and Kittipan Rerkasem. "Carotid endarterectomy for symptomatic carotid stenosis." In: *Cochrane Database of Systematic Reviews* 2020.9 (2020). DOI: [10.1002/14651858.cd001081.pub4](https://doi.org/10.1002/14651858.cd001081.pub4).
- [260] E. Robertson, V. Choudhury, S. Bhushan, and D. K. Walters. "Validation of OpenFOAM numerical methods and turbulence models for incompressible bluff body flows." In: *Computers & Fluids* 123 (2015), pp. 122–145. DOI: [10.1016/j.compfluid.2015.09.010](https://doi.org/10.1016/j.compfluid.2015.09.010).
- [261] Caroline H. Roney, Ali Pashaei, Marianna Meo, Rémi Dubois, Patrick M. Boyle, Natalia A. Trayanova, Hubert Cochet, Steven A. Niederer, and Edward J. Vigmond. "Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes." In: *Medical Image Analysis* 55 (2019), pp. 65–75. DOI: [10.1016/j.media.2019.04.004](https://doi.org/10.1016/j.media.2019.04.004).
- [262] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer. 2015, pp. 234–241.

- [263] Justus E. Roos, Dominik Fleischmann, Arnold Koechl, Tejas Rakshe, Matus Straka, Alessandro Napoli, Armin Kanitsar, Milos Sramek, and M. Eduard Gröller. "Multipath Curved Planar Reformation of the Peripheral Arterial Tree in CT Angiography." In: *Radiology* 244.1 (2007), pp. 281–290. DOI: [10.1148/radiol.2441060976](https://doi.org/10.1148/radiol.2441060976).
- [264] Timo Ropinski, Sven Hermann, Rainer Reich, Michael Schafers, and Klaus Hinrichs. "Multimodal Vessel Visualization of Mouse Aorta PET/CT Scans." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1515–1522. DOI: [10.1109/tvcg.2009.169](https://doi.org/10.1109/tvcg.2009.169).
- [265] Russell Ross. "The pathogenesis of atherosclerosis: a perspective for the 1990s." In: *Nature* 362.6423 (1993), pp. 801–809. DOI: [10.1038/362801a0](https://doi.org/10.1038/362801a0).
- [266] Carl Runge. "Über die numerische Auflösung von Differentialgleichungen." In: *Mathematische Annalen* 46.2 (1895), pp. 167–178. DOI: [10.1007/BF01446807](https://doi.org/10.1007/BF01446807).
- [267] Jegoon Ryu and Sei-ichiro Kamata. "An efficient computational algorithm for Hausdorff distance based on points-ruling-out and systematic random sampling." In: *Pattern Recognition* 114 (June 2021), p. 107857. DOI: [10.1016/j.patcog.2021.107857](https://doi.org/10.1016/j.patcog.2021.107857).
- [268] Patrick Saalfeld, Sylvia Glaßer, Oliver Beuing, and Bernhard Preim. "The FAUST framework: Free-form annotations on unfolding vascular structures for treatment planning." In: *Computers & Graphics* 65 (2017), pp. 12–21. DOI: [10.1016/j.cag.2017.03.003](https://doi.org/10.1016/j.cag.2017.03.003).
- [269] Luca Saba and Giorgio Mallarini. "MDCTA of carotid plaque degree of stenosis: Evaluation of interobserver agreement." In: *American Journal of Roentgenology* 190.1 (2008), W41–W46. DOI: [10.2214/ajr.07.2604](https://doi.org/10.2214/ajr.07.2604).
- [270] Punam K. Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja. "A survey on skeletonization algorithms and their applications." In: *Pattern Recognition Letters* 76 (2016), pp. 3–12. DOI: [10.1016/j.patrec.2015.04.006](https://doi.org/10.1016/j.patrec.2015.04.006).
- [271] Khalid M. Saqr, Sherif Rashad, Simon Tupin, Kuniyasu Nizuma, Tamer Hassan, Teiji Tominaga, and Makoto Ohta. "What does computational fluid dynamics tell us about intracranial aneurysms? A meta-analysis and critical review." In: *Journal of Cerebral Blood Flow & Metabolism* 40.5 (2020), pp. 1021–1039. DOI: [10.1177/0271678X19854640](https://doi.org/10.1177/0271678X19854640).
- [272] Laurent Saroul, Sebastian Gerlach, and Roger D. Herch. "Exploring Curved Anatomic Structures with Surface Sections." In: *IEEE Visualization*. 2003, pp. 27–34. DOI: [10.1109/visual.2003.1250351](https://doi.org/10.1109/visual.2003.1250351).

- [273] M. Sato, I. Bitter, M. A. Bender, A. E. Kaufman, and M. Nakajima. "TEASAR: tree-structure extraction algorithm for accurate and robust skeletons." In: *Proceedings the Eighth Pacific Conference on Computer Graphics and Applications*. IEEE, 2000. DOI: [10.1109/pccga.2000.883951](https://doi.org/10.1109/pccga.2000.883951).
- [274] Alexander Scheid-Rehder, Kai Lawonn, and Monique Meuschke. "Robustness evaluation of CFD simulations to mesh deformation." In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2019, pp. 189–199. DOI: [10.2312/vcbm.20191244](https://doi.org/10.2312/vcbm.20191244).
- [275] Hermann Schlichting and Klaus Gersten. *Boundary-Layer Theory*. 9th ed. Springer, 2016. DOI: [10.1007/978-3-662-52919-5](https://doi.org/10.1007/978-3-662-52919-5).
- [276] Caroline A Schneider, Wayne S Rasband, and Kevin W Eliceiri. "NIH Image to ImageJ: 25 years of image analysis." In: *Nature Methods* 9.7 (2012), pp. 671–675. DOI: [10.1038/nmeth.2089](https://doi.org/10.1038/nmeth.2089).
- [277] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. "Design Study Methodology: Reflections from the Trenches and the Stacks." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2431–2440. DOI: [10.1109/tvcg.2012.213](https://doi.org/10.1109/tvcg.2012.213).
- [278] Robert Seifert, Aaron Scherzinger, Friedemann Kiefer, Sven Hermann, Xiaoyi Jiang, and Michael A. Schäfers. "Statistical Permutation-based Artery Mapping (SPAM): a novel approach to evaluate imaging signals in the vessel wall." In: *BMC Medical Imaging* 17.1 (2017). DOI: [10.1186/s12880-017-0207-7](https://doi.org/10.1186/s12880-017-0207-7).
- [279] J. Serena, P. Irimia, S. Calleja, M. Blanco, J. Vivancos, and Ó. Ayo-Martín. "Ultrasound measurement of carotid stenosis: Recommendations from the Spanish Society of Neurosonology." In: *Neurología (English Edition)* 28.7 (2013), pp. 435–442. DOI: [10.1016/j.nrleng.2013.09.004](https://doi.org/10.1016/j.nrleng.2013.09.004).
- [280] M. Safdari Shadloo, G. Oger, and David Le Touzé. "Smoothed particle hydrodynamics method for fluid flows, towards industrial applications: Motivations, current state, and challenges." In: *Computers & Fluids* 136 (2016), pp. 11–34. DOI: [10.1016/j.compfluid.2016.05.029](https://doi.org/10.1016/j.compfluid.2016.05.029).
- [281] Ali Shahrokni, Hamid Soltanian-Zadeh, and Reza A. Zoroofi. "Fast skeletonization algorithm for 3D elongated objects." In: *Proc. SPIE 4322, Medical Imaging 2001: Image Processing*. Ed. by Milan Sonka and Kenneth M. Hanson. 2001. DOI: [10.1117/12.431102](https://doi.org/10.1117/12.431102).
- [282] Ali Sheharyar, Teodora Chitiboi, Eric Keller, Ozair Rahman, Susanne Schnell, Michael Markl, Othmane Bouhali, and Lars Linsen. "Spatio-temporal Visualization of Regional Myocardial Velocities." In: *Eurographics Workshop on Visual Computing*

- for Biology and Medicine*. 2016, pp. 89–98. DOI: [10.2312/VCBM.20161275](https://doi.org/10.2312/VCBM.20161275).
- [283] Dominik Sibbing, Hans-Christian Ebke, Kai Ingo Esser, and Leif Kobbelt. "Topology Aware Quad Dominant Meshing for Vascular Structures." In: *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2012, pp. 147–158. DOI: [10.1007/978-3-642-33463-4_15](https://doi.org/10.1007/978-3-642-33463-4_15).
- [284] Nahian Siddique, Sidike Paheding, Colin P. Elkin, and Vijay Devabhaktuni. "U-Net and Its Variants for Medical Image Segmentation: A Review of Theory and Applications." In: *IEEE Access* 9 (2021), pp. 82031–82057. DOI: [10.1109/access.2021.3086020](https://doi.org/10.1109/access.2021.3086020).
- [285] J. V. B. Soares, J. J. G. Leandro, R. M. Cesar, H. F. Jelinek, and M. J. Cree. "Retinal vessel segmentation using the 2-D Gabor wavelet and supervised classification." In: *IEEE Transactions on Medical Imaging* 25.9 (2006), pp. 1214–1222. DOI: [10.1109/tmi.2006.879967](https://doi.org/10.1109/tmi.2006.879967).
- [286] Olga Sorkine and Marc Alexa. "As-rigid-as-possible surface modeling." In: *Symposium on Geometry Processing*. Vol. 4. 2007, pp. 109–116.
- [287] Luísa C. Sousa, Catarina F. Castro, Carlos C. António, Fernando Sousa, Rosa Santos, Pedro Castro, and Elsa Azevedo. "Computational simulation of carotid stenosis and flow dynamics based on patient ultrasound data – A new tool for risk assessment and surgical planning." In: *Advances in Medical Sciences* 61.1 (2016), pp. 32–39. DOI: [10.1016/j.advms.2015.07.009](https://doi.org/10.1016/j.advms.2015.07.009).
- [288] Zoran Stankovic, Bradley D. Allen, Julio Garcia, Kelly B. Jarvis, and Michael Markl. "4D flow imaging with MRI." In: *Cardiovascular Diagnosis and Therapy* 4.2 (2014). DOI: [10.3978/j.issn.2223-3652.2014.01.02](https://doi.org/10.3978/j.issn.2223-3652.2014.01.02).
- [289] Ludwig Steffgen. "Klinikleitfaden Sonographie Common Trunk." In: Elsevier, 2011. Chap. Grundlagen der Ultraschalldiagnostik.
- [290] M. Straka, M. Červeňanský, A. La Cruz, A. Köchl, M. Šrámek, M. E. Gröller, and D. Fleischmann. "The VesselGlyph: Focus & Context Visualization in CT-Angiography." In: *IEEE Visualization*. 2004, pp. 385–392. DOI: [10.1109/visual.2004.104](https://doi.org/10.1109/visual.2004.104).
- [291] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. "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*. Ed. by M. Jorge Cardoso et al. Springer, 2017, pp. 240–248. DOI: [10.1007/978-3-319-67558-9_28](https://doi.org/10.1007/978-3-319-67558-9_28).

- [292] Zhiyu Sun, Ethan Rooke, Jerome Charton, Yusen He, Jia Lu, and Stephen Baek. “ZerNet: Convolutional Neural Networks on Arbitrary Surfaces Via Zernike Local Tangent Space Estimation.” In: *Computer Graphics Forum* 39.6 (2020), pp. 204–216. DOI: [10.1111/cgf.14012](https://doi.org/10.1111/cgf.14012).
- [293] Jeremy Szajer and Kevin Ho-Shon. “A comparison of 4D flow MRI-derived wall shear stress with computational fluid dynamics methods for intracranial aneurysms and carotid bifurcations – A review.” In: *Magnetic Resonance Imaging* 48 (2018), pp. 62–69. ISSN: 0730-725X. DOI: [10.1016/j.mri.2017.12.005](https://doi.org/10.1016/j.mri.2017.12.005).
- [294] Andrea Tagliasacchi. *Skeletal Representations and Applications*. Tech. rep. School of Computing Science, Simon Fraser University, 2014. URL: <https://arxiv.org/pdf/1301.6809.pdf>.
- [295] Hui Tang, Theo van Walsum, Reinhard Hameeteman, Rahil Shahzad, Lucas J. van Vliet, and Wiro J. Niessen. “Lumen segmentation and stenosis quantification of atherosclerotic carotid arteries in CTA utilizing a centerline intensity prior.” In: *Medical Physics* 40.5 (2013), p. 051721. DOI: [10.1118/1.4802751](https://doi.org/10.1118/1.4802751).
- [296] Jun Tao, Xiaoke Huang, Feng Qiu, Chaoli Wang, Jingfeng Jiang, Ching-Kuang Shene, Ye Zhao, and Daphne Yu. “VesselMap: A web interface to explore multivariate vascular data.” In: *Computers & Graphics* 59 (2016), pp. 79–92. DOI: [10.1016/j.cag.2016.05.024](https://doi.org/10.1016/j.cag.2016.05.024).
- [297] Gabriel Taubin, Tong Zhang, and Gene Golub. “Optimal surface smoothing as filter design.” In: *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1996, pp. 283–292. DOI: [10.1007/bfb0015544](https://doi.org/10.1007/bfb0015544).
- [298] Maurice Termeer, Javier Oliván Bescós, Marcel Breeuwer, Anna Vilanova, Frans Gerritsen, and M. Eduard Gröller. “CoViCAD: Comprehensive Visualization of Coronary Artery Disease.” In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1632–1639. DOI: [10.1109/tvcg.2007.70550](https://doi.org/10.1109/tvcg.2007.70550).
- [299] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. “Interactive Lenses for Visualization: An Extended Survey.” In: *Computer Graphics Forum* 36.6 (2016), pp. 173–200. DOI: [10.1111/cgf.12871](https://doi.org/10.1111/cgf.12871).
- [300] Melanie Tory, David Sprague, Fuqu Wu, Wing Yan So, and Tamara Munzner. “Spatialization Design: Comparing Points and Landscapes.” In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1262–1269. DOI: [10.1109/tvcg.2007.70596](https://doi.org/10.1109/tvcg.2007.70596).

- [301] Connie W. Tsao, Aaron W. Aday, Zaid I. Almarzooq, et al. "Heart Disease and Stroke Statistics—2023 Update: A Report From the American Heart Association." In: *Circulation* 147.8 (2023), e93–e621. DOI: [10.1161/cir.0000000000001123](https://doi.org/10.1161/cir.0000000000001123).
- [302] Zsolt Ugray, Leon Lasdon, John Plummer, Fred Glover, James Kelly, and Rafael Martí. "Scatter Search and Local NLP Solvers: A Multistart Framework for Global Optimization." In: *INFORMS Journal on Computing* 19.3 (2007), pp. 328–340. DOI: [10.1287/ijoc.1060.0175](https://doi.org/10.1287/ijoc.1060.0175).
- [303] R. F. P. Van Pelt, S. S. A. M. Jacobs, Bart M. ter Haar Romeny, and A. Vilanova. "Visualization of 4D blood-flow fields by spatiotemporal hierarchical clustering." In: *Computer Graphics Forum* 31.3pt2 (2012), pp. 1065–1074. DOI: [10.1111/j.1467-8659.2012.03099.x](https://doi.org/10.1111/j.1467-8659.2012.03099.x).
- [304] Roy F. P. Van Pelt, Javier Oliván Bescos, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, Bart ter Haar Romenij, and Anna Vilanova. "Exploration of 4D MRI blood flow using stylistic visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1339–1347. DOI: [10.1109/TVCG.2010.153](https://doi.org/10.1109/TVCG.2010.153).
- [305] Roy F. P. Van Pelt, Rocco Gasteiger, Kai Lawonn, Monique Meuschke, and Bernhard Preim. "Comparative blood flow visualization for cerebral aneurysm treatment assessment." In: *Computer Graphics Forum* 33.3 (2014), pp. 131–140. DOI: [10.1111/cgf.12369](https://doi.org/10.1111/cgf.12369).
- [306] Irene E. Vignon-Clementel, C. Alberto Figueroa, Kenneth E. Jansen, and Charles A. Taylor. "Outflow boundary conditions for three-dimensional finite element modeling of blood flow and pressure in arteries." In: *Computer Methods In Applied Mechanics And Engineering* 195.29–32 (2006), pp. 3776–3796. DOI: [10.1016/j.cma.2005.04.014](https://doi.org/10.1016/j.cma.2005.04.014).
- [307] Anna Vilanova, Bernhard Preim, Roy F. P. van Pelt, Rocco Gasteiger, Mathias Neugebauer, and Thomas Wischgoll. "Visual exploration of simulated and measured blood flow." In: *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization*. Ed. by Charles D. Hansen, Min Chen, Christopher R. Johnson, Arie E. Kaufman, and Hans Hagen. Springer London, 2014, pp. 305–324. DOI: [10.1007/978-1-4471-6497-5_25](https://doi.org/10.1007/978-1-4471-6497-5_25).
- [308] Dominic J. Vitello, Richard M. Ripper, Michael R. Fettiplace, Guy L. Weinberg, and Joseph M. Vitello. "Blood Density Is Nearly Equal to Water Density: A Validation Study of the Gravimetric Method of Measuring Intraoperative Blood Loss." In: *Journal of Veterinary Medicine* 2015 (2015), pp. 1–4. DOI: [10.1155/2015/152730](https://doi.org/10.1155/2015/152730).

- [309] Yu-Shuen Wang and Tong-Yee Lee. "Curve-Skeleton Extraction Using Iterative Least Squares Optimization." In: *IEEE Transactions on Visualization and Computer Graphics* 14.4 (2008), pp. 926–936. DOI: [10.1109/tvcg.2008.38](https://doi.org/10.1109/tvcg.2008.38).
- [310] Zhimin Wang, Yanling Chi, Weimin Huang, Sudhakar K. Venkatesh, Qi Tian, Thiha Oo, Jiayin Zhou, Wei Xiong, and Jimin Liu. "Comparisons of centerline extraction methods for liver blood vessels in ImageJ and 3D slicer." In: *APSIPA ASC* (2010), pp. 276–279.
- [311] Colin Ware, Helen Purchase, Linda Colpoys, and Matthew McGill. "Cognitive Measurements of Graph Aesthetics." In: *Information Visualization* 1.2 (2002), pp. 103–110. DOI: [10.1057/palgrave.ivs.9500013](https://doi.org/10.1057/palgrave.ivs.9500013).
- [312] E. M. W. Weber. *Schemata der Leitungsbahnen des Menschen*. Berlin, Heidelberg: Springer, 1978.
- [313] Claus Weigand. "Strömungsanalysen in der Karotisbifurkation." PhD thesis. Technische Universität München, 2000.
- [314] Stefan Wesarg, M. Fawad Khan, and Evelyn A. Firle. "Localizing calcifications in cardiac CT data sets using a new vessel segmentation approach." In: *Journal of Digital Imaging* 19.3 (2006), pp. 249–257. DOI: [10.1007/s10278-006-9947-6](https://doi.org/10.1007/s10278-006-9947-6).
- [315] Lars Wigström, Tino Ebbers, Anna Fyrenius, Matts Karlsson, Jan Engvall, Bengt Wranne, and Ann F. Bolger. "Particle trace visualization of intracardiac flow using time-resolved 3D phase contrast MRI." In: *Magnetic Resonance in Medicine* 41.4 (1999), pp. 793–799. DOI: [10.1002/\(sici\)1522-2594\(199904\)41:4<793::aid-mrm19>3.0.co;2-2](https://doi.org/10.1002/(sici)1522-2594(199904)41:4<793::aid-mrm19>3.0.co;2-2).
- [316] Steven E. Williams, Catalina Tobon-Gomez, Maria A. Zuluaga, Henry Chubb, Constantine Butakoff, Rashed Karim, Elena Ahmed, Oscar Camara, and Kawal S. Rhode. "Standardized unfold mapping: a technique to permit left atrial regional data display and analysis." In: *Journal of Interventional Cardiac Electrophysiology* 50.1 (2017), pp. 125–131. DOI: [10.1007/s10840-017-0281-3](https://doi.org/10.1007/s10840-017-0281-3).
- [317] Joong Ho Won, Geoffrey D. Rubin, and Sandy Napel. "Flattening the Abdominal Aortic Tree for Effective Visualization." In: *2006 International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2006. DOI: [10.1109/iembs.2006.259584](https://doi.org/10.1109/iembs.2006.259584).
- [318] Joong-Ho Won, Yongkweon Jeon, J. K. Rosenberg, Sungroh Yoon, G. D. Rubin, and S. Napel. "Uncluttered Single-Image Visualization of Vascular Structures Using GPU and Integer Programming." In: *IEEE Transactions on Visualization and Computer Graphics* 19.1 (2013), pp. 81–93. DOI: [10.1109/tvcg.2012.25](https://doi.org/10.1109/tvcg.2012.25).

- [319] Joong-Ho Won, Jarrett Rosenberg, Geoffrey D. Rubin, and Sandy Napel. "Uncluttered single-image visualization of the abdominal aortic vessel tree: Method and evaluation." In: *Medical Physics* 36.11 (2009), pp. 5245–5260. DOI: [10.1118/1.3243866](https://doi.org/10.1118/1.3243866).
- [320] Kelvin K. L. Wong, Richard M. Kelso, Stephen G. Worthley, Prashanthan Sanders, Jagannath Mazumdar, and Derek Abbott. "Cardiac flow analysis applied to phase contrast magnetic resonance imaging of the heart." In: *Annals of Biomedical Engineering* 37.8 (2009), pp. 1495–1515. DOI: [10.1007/s10439-009-9709-y](https://doi.org/10.1007/s10439-009-9709-y).
- [321] Aaron Wu, Ziyue Xu, Mingchen Gao, Mario Buty, and Daniel J. Mollura. "Deep vessel tracking: A generalized probabilistic approach via deep learning." In: *IEEE 13th International Symposium on Biomedical Imaging*. 2016. DOI: [10.1109/isbi.2016.7493520](https://doi.org/10.1109/isbi.2016.7493520).
- [322] Jinxiang Xi, Xiuhua Si, Yue Zhou, JongWon Kim, and Ariel Berlinski. "Growth of nasal and laryngeal airways in children: Implications in breathing and inhaled aerosol dynamics." In: *Respiratory Care* 59.2 (2014), pp. 263–273.
- [323] J. Xiang, V. M. Tutino, K. V. Snyder, and H. Meng. "CFD: computational fluid dynamics or confounding factor dissemination? The role of hemodynamics in intracranial aneurysm rupture risk assessment." In: *American Journal of Neuroradiology* 35.10 (2014), pp. 1849–1857. DOI: [10.3174/ajnr.A3710](https://doi.org/10.3174/ajnr.A3710).
- [324] Pengcheng Xu, Xin Liu, Heye Zhang, Dhanjoo Ghista, Dong Zhang, Changzheng Shi, and Wenhua Huang. "Assessment of boundary conditions for CFD simulation in human carotid artery." In: *Biomechanics and Modeling in Mechanobiology* 17.6 (2018), pp. 1581–1597. DOI: [10.1007/s10237-018-1045-4](https://doi.org/10.1007/s10237-018-1045-4).
- [325] Rui Xu and D. Wunsch. "Survey of clustering algorithms." In: *IEEE Transactions on Neural Networks* 16.3 (2005), pp. 645–678. DOI: [10.1109/TNN.2005.845141](https://doi.org/10.1109/TNN.2005.845141).
- [326] Jiayong Yan, Binsheng Zhao, Liang Wang, Andrew Zelenetz, and Lawrence H. Schwartz. "Marker-controlled watershed for lymphoma segmentation in sequential CT images." In: *Medical Physics* 33.7 (2006), pp. 2452–2460. DOI: [10.1118/1.2207133](https://doi.org/10.1118/1.2207133).
- [327] H. F. Younis, Ma R. Kaazempur-Mofrad, R. C. Chan, A. G. Isasi, D. P. Hinton, A. H. Chau, L. A. Kim, and R. D. Kamm. "Hemodynamics and wall mechanics in human carotid bifurcation and its consequences for atherogenesis: investigation of inter-individual variation." In: *Biomechanics and Modeling in Mechanobiology* 3.1 (2004), pp. 17–32. DOI: [10.1007/s10237-004-0046-7](https://doi.org/10.1007/s10237-004-0046-7).

- [328] Mohamad Shukri Zakaria, Farzad Ismail, Masaaki Tamagawa, Ahmad Fazli Abdul Azi, Surjatin Wiriadidjaya, Adi Azrif Basri, and Kamarul Arifin Ahmad. "Computational fluid dynamics study of blood flow in aorta using OpenFOAM." In: *Journal of Advanced Research in Fluid Mechanics and Thermal Sciences* 43 (2018), pp. 81–89.
- [329] Jun Zhang, Jun Tao, Jian-Xun Wang, and Chaoli Wang. "SurfRiver: Flattening Stream Surfaces for Comparative Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 27.6 (2021), pp. 2783–2795. DOI: [10.1109/tvcg.2021.3074585](https://doi.org/10.1109/tvcg.2021.3074585).
- [330] Fengjun Zhao, Yanrong Chen, Yuqing Hou, and Xiaowei He. "Segmentation of blood vessels using rule-based and machine-learning-based methods: a review." In: *Multimedia Systems* 25.2 (2017), pp. 109–118. DOI: [10.1007/s00530-017-0580-7](https://doi.org/10.1007/s00530-017-0580-7).
- [331] Yuan Zhao, J. David Spence, and Bernard Chiu. "Three-dimensional ultrasound assessment of effects of therapies on carotid atherosclerosis using vessel wall thickness maps." In: *Ultrasound in Medicine & Biology* 47.9 (2021), pp. 2502–2513. DOI: [10.1016/j.ultrasmedbio.2021.04.015](https://doi.org/10.1016/j.ultrasmedbio.2021.04.015).
- [332] Tianshu Zhou, Tao Tan, Xiaoyan Pan, Hui Tang, and Jingsong Li. "Fully automatic deep learning trained on limited data for carotid artery segmentation from large image volumes." In: *Quantitative Imaging in Medicine and Surgery* 11.1 (2021), pp. 67–83. DOI: [10.21037/qims-20-286](https://doi.org/10.21037/qims-20-286).
- [333] Y. Zhou and A.W. Toga. "Efficient skeletonization of volumetric objects." In: *IEEE Transactions on Visualization and Computer Graphics* 5.3 (1999), pp. 196–209. DOI: [10.1109/2945.795212](https://doi.org/10.1109/2945.795212).
- [334] Lei Zhu, Steven Haker, and Allen R. Tannenbaum. "Conformal flattening maps for the visualization of vessels." In: *Medical Imaging 2002: Visualization, Image-Guided Procedures, and Display*. Ed. by Seong K. Mun. SPIE, 2002. DOI: [10.1117/12.466985](https://doi.org/10.1117/12.466985).
- [335] Lei Zhu, Steven Haker, and Allen R. Tannenbaum. "Area-Preserving Mappings for the Visualization of Medical Structures." In: *Medical Image Computing and Computer-Assisted Intervention*. Ed. by Randy E. Ellis and Terry M. Peters. Berlin, Heidelberg, 2003, pp. 277–284. ISBN: 978-3-540-39903-2. DOI: [10.1007/978-3-540-39903-2_35](https://doi.org/10.1007/978-3-540-39903-2_35).
- [336] Lei Zhu, Steven Haker, and Allen R. Tannenbaum. "Flattening maps for the visualization of multibranching vessels." In: *IEEE Transactions on Medical Imaging* 24.2 (2005), pp. 191–198. DOI: [10.1109/TMI.2004.839368](https://doi.org/10.1109/TMI.2004.839368).

- [337] Lei Zhu, Steven Haker, Allen R. Tannenbaum, S. Bouix, and K. Siddiqi. "Angle-preserving mappings for the visualization of multi-branched vessels." In: *IEEE International Conference on Image Processing*. 2002. DOI: [10.1109/icip.2002.1040108](https://doi.org/10.1109/icip.2002.1040108).

LIST OF PUBLICATIONS

- [338] **Pepe Eulzer**, Sabine Bauer, Francis Kilian, and Kai Lawonn. “Visualization of Human Spine Biomechanics for Spinal Surgery.” In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2021), pp. 700–710. DOI: [10.1109/tvcg.2020.3030388](https://doi.org/10.1109/tvcg.2020.3030388).
- [339] **Pepe Eulzer**, Fabienne von Deylen, Ralf Wickenhöfer, Carsten M. Klingner, and Kai Lawonn. “A Fully Integrated Pipeline for Visual Carotid Morphology Analysis.” In: *Computer Graphics Forum* 42.3 (2023), pp. 25–37. DOI: [10.1111/cgf.14808](https://doi.org/10.1111/cgf.14808).
- [340] **Pepe Eulzer**, Sandy Engelhardt, Nils Lichtenberg, Raffaele De Simone, and Kai Lawonn. “Temporal Views of Flattened Mitral Valve Geometries.” In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2019), pp. 971–980. DOI: [10.1109/TVCG.2019.2934337](https://doi.org/10.1109/TVCG.2019.2934337).
- [341] **Pepe Eulzer**, Nils Lichtenberg, Rawa Arif, Andreas Brcic, Matthias Karck, Kai Lawonn, Raffaele De Simone, and Sandy Engelhardt. “Mitral Valve Quantification at a Glance.” In: *Bildverarbeitung für die Medizin*. Springer Fachmedien Wiesbaden, 2019, pp. 296–301. DOI: [10.1007/978-3-658-25326-4_66](https://doi.org/10.1007/978-3-658-25326-4_66).
- [342] **Pepe Eulzer**, Monique Meuschke, Carsten M. Klingner, and Kai Lawonn. “Visualizing Carotid Blood Flow Simulations for Stroke Prevention.” In: *Computer Graphics Forum* 40.3 (2021), pp. 435–446. DOI: [10.1111/cgf.14319](https://doi.org/10.1111/cgf.14319).
- [343] **Pepe Eulzer**, Monique Meuschke, Gabriel Mistelbauer, and Kai Lawonn. “Vessel Maps: A Survey of Map-Like Visualizations of the Cardiovascular System.” In: *Computer Graphics Forum* 41.3 (2022), pp. 645–673. DOI: [10.1111/cgf.14576](https://doi.org/10.1111/cgf.14576).
- [344] **Pepe Eulzer**, Kevin Richter, Anna Hundertmark, Monique Meuschke, Ralph Wickenhöfer, Carsten Klingner, and Kai Lawonn. “Visualizing Carotid Stenoses for Stroke Treatment and Prevention.” In: *EuroVis 2023 – Dirk Bartz Prize*. The Eurographics Association, 2023. DOI: [10.2312/EVM.20231086](https://doi.org/10.2312/EVM.20231086).
- [345] **Pepe Eulzer**, Kevin Richter, Anna Hundertmark, Ralph Wickenhöfer, Carsten M. Klingner, and Kai Lawonn. “Instantaneous Visual Analysis of Blood Flow in Stenoses Using Morphological Similarity.” In: *Computer Graphics Forum* 43.3 (2024), in print. DOI: [10.1111/cgf.15081](https://doi.org/10.1111/cgf.15081).

- [346] **Pepe Eulzer**, Kevin Richter, Monique Meuschke, Anna Hundertmark, and Kai Lawonn. "Automatic Cutting and Flattening of Carotid Artery Geometries." In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2021, pp. 79–89. DOI: [10.2312/vcbm.20211347](https://doi.org/10.2312/vcbm.20211347).
- [347] **Pepe Eulzer**, Kevin Richter, Tristan Probst, Anna Hundertmark, and Kai Lawonn. *A Dataset of Reconstructed Carotid Bifurcation Lumen and Plaque Models with Centerline Tree and Simulated Hemodynamics*. 2024. DOI: [10.5281/ZENODO.7634643](https://doi.org/10.5281/ZENODO.7634643).
- [348] **Pepe Eulzer**, Robert Rockenfeller, and Kai Lawonn. "HAExplorer: Understanding Interdependent Biomechanical Motions with Interactive Helical Axes." In: *CHI Conference on Human Factors in Computing Systems*. ACM, 2022. DOI: [10.1145/3491102.3501841](https://doi.org/10.1145/3491102.3501841).
- [349] Kai Lawonn, Monique Meuschke, **Pepe Eulzer**, Matthias Mitterreiter, Joachim Giesen, and Tobias Günther. "GRay: Ray Casting for Visualization and Interactive Data Exploration of Gaussian Mixture Models." In: *IEEE Transactions on Visualization and Computer Graphics* 29.1 (2023), pp. 526–536. DOI: [10.1109/tvcg.2022.3209374](https://doi.org/10.1109/tvcg.2022.3209374).
- [350] Nils Lichtenberg, **Pepe Eulzer**, Gabriele Romano, Andreas Brcic, Matthias Karck, Kai Lawonn, Raffaele de Simone, and Sandy Engelhardt. "Mitral Valve Flattening and Parameter Mapping for Patient-Specific Valve Diagnosis." In: *International Journal of Computer Assisted Radiology and Surgery* 15 (2020), pp. 617–627. DOI: [10.1007/s11548-019-02114-w](https://doi.org/10.1007/s11548-019-02114-w).
- [351] Monique Meuschke, Samuel Voß, **Pepe Eulzer**, Gabor Janiga, Christoph Arens, Ralph Wickenhöfer, Bernhard Preim, and Kai Lawonn. "COMFIS – Comparative Visualization of Simulated Medical Flow Data." In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. 2022, pp. 29–40. DOI: [10.2312/vcbm.20221185](https://doi.org/10.2312/vcbm.20221185).
- [352] Kevin Richter, Tristan Probst, Anna Hundertmark, **Pepe Eulzer**, and Kai Lawonn. "Longitudinal Wall Shear Stress Evaluation using Centerline Projection Approach in the Numerical Simulations of the Patient-based Carotid Artery." In: *Computer Methods in Biomechanics and Biomedical Engineering* 27.3 (2023), pp. 347–364. DOI: [10.1080/10255842.2023.2185478](https://doi.org/10.1080/10255842.2023.2185478).
- [353] Ashkan Shiravand, Kevin Richter, Pia Willmann, **Pepe Eulzer**, Kai Lawonn, Anna Hundertmark, and Giorgio Cattaneo. "Fabrication, Characterization and Numerical Validation of a Novel Thin-Wall Hydrogel Vessel Model for Cardiovascular Research Based on a Patient-Specific Stenotic Carotid Artery Bifurcation." In: *Scientific Reports* 14.1 (2024), p. 16301. DOI: [10.1038/s41598-024-66777-5](https://doi.org/10.1038/s41598-024-66777-5).

EHRENWÖRTLICHE ERKLÄRUNG

Hiermit erkläre ich ehrenwörtlich,

- dass mir die am heutigen Tage geltende Promotionsordnung der Fakultät für Mathematik und Informatik an der Friedrich-Schiller-Universität Jena bekannt ist,
- dass ich die Dissertation selbstständig angefertigt habe, keine Textabschnitte oder Ergebnisse eines Dritten oder eigenen Prüfungsarbeiten ohne Kennzeichnung übernommen habe und alle von mir benutzten Hilfsmittel, persönliche Mitteilungen und Quellen in meiner Arbeit angegeben habe,
- dass die Auswahl und Auswertung sämtlichen Materials eigenständig geschah,
- dass ich die Hilfe einer Promotionsberaterin/eines Promotionsberaters nicht in Anspruch genommen habe und dass Dritte weder unmittelbar noch mittelbar geldwerte Leistungen von mir für Arbeiten erhalten haben, die im Zusammenhang mit dem Inhalt der vorgelegten Dissertation stehen,
- dass ich die vorliegende Arbeit noch nicht als Prüfungsarbeit für eine staatliche oder andere wissenschaftliche Prüfung eingereicht habe,
- dass ich keine identische, eine in wesentlichen Teilen ähnliche bzw. eine andere Abhandlung bei einer anderen Hochschule als Dissertation eingereicht habe.

Jena, 05. August 2024

Pepe Eulzer