

Real-time Illustrative Visualization of Cardiovascular Hemodynamics



Real-time Illustrative Visualization of Cardiovascular Hemodynamics

Colophon


The thesis cover was designed by the author (Roy van Pelt).

The comic was drawn by Koert Stavenuiter (www.koertstavenuiter.nl).
The comic was digitally inked by the author.

This contents were typeset by the author using $\text{\LaTeX}2_{\epsilon}$.
The main body of the text was set using a 10-point Latin Modern Sans Serif font.

Travel grants were awarded by the Austrian Science Fund (FWF) and the Vereniging voor Biofysica en Biomedische Technologie (VvB-BMT).

Financial support for the publication of this thesis was kindly provided by the Advanced School for Computing and Imaging (ASCI) and Eindhoven University of Technology.



This work was carried out in the ASCI graduate school.
ASCI dissertation series number 240.

Printed by Off Page, Amsterdam, the Netherlands

A catalogue record is available from Eindhoven University of Technology library:
ISBN-13: 978-90-386-3153-0

© 2012 Roy F.P. van Pelt, Tilburg, The Netherlands, unless stated otherwise on chapter front pages. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the copyright owner.

Real-time Illustrative Visualization of Cardiovascular Hemodynamics

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen op
woensdag 13 juni 2012 om 14.00 uur

door

Roy Franciscus Petrus van Pelt

geboren te Tilburg

Dit proefschrift is goedgekeurd door de promotoren:


prof.dr.ir. B.M. ter Haar Romeny

en

prof.dr.ir. M. Breeuwer

Copromotor:

dr. A. Vilanova



Contents

Colophon	ii
Contents	vii
Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 Preface	2
1.2 Context of this research	4
1.3 Outline	6
2 Clinical background	9
2.1 Introduction	10
2.1.1 A short history	10
2.1.2 Motivation	12
2.2 Blood-flow acquisition	14
2.2.1 Modeling blood-flow mechanics	14
2.2.2 Measuring blood-flow mechanics	15
2.3 Blood-flow analysis	17
2.3.1 Quantitative analysis	17
2.3.2 Qualitative analysis	18
2.4 Cardiovascular disease - adult conditions	19
2.4.1 Epidemiology	19
2.4.2 Pathology example	21
2.5 Cardiovascular disease - congenital conditions	24
2.5.1 Epidemiology	24
2.5.2 Pathology example	25
3 Data acquisition using magnetic resonance	29
3.1 Magnetic resonance imaging	30
3.1.1 Spin physics	30
3.1.2 Volumetric imaging sequences	31
3.1.3 Blood flow imaging	33
3.2 Phase-contrast acquisition	34
3.2.1 Velocity encoding	34
3.2.2 Image formation	38
3.2.3 Data acquisition	40
3.2.4 Data characteristics	41
3.2.5 Acquisition accuracy and artifacts	42

4	Technical background	47
4.1	Introduction	48
4.2	Illustrative visualization	52
4.2.1	Motivation	52
4.2.2	Techniques - point and line	54
4.2.3	Techniques - shading and style	57
4.3	Flow visualization	58
4.3.1	Motivation	58
4.3.2	Texture-based graphics	60
4.3.3	Feature-based graphics	61
4.3.4	Partition-based graphics	64
4.3.5	Geometry-based graphics	64
4.4	Interactive visualization	67
4.4.1	Motivation	67
4.4.2	Consumer graphics hardware	68
5	Blood-flow abstraction	69
5.1	Temporal projections	70
5.1.1	Motivation	70
5.1.2	Temporal maximum intensity projection	70
5.1.3	Temporal mean orientation projection	71
5.2	Segmentation	72
5.2.1	Motivation	72
5.2.2	Features	73
5.2.3	Active Surface Model	75
5.2.4	Validation	77
5.2.5	Discussion and conclusions	78
5.3	Clustering	79
5.3.1	Motivation	79
5.3.2	Spatiotemporal dissimilarity measures	80
5.3.3	Coarse hierarchical clustering	81
5.3.4	Visualization	83
5.3.5	Validation	88
5.3.6	Discussion and conclusions	89
6	Blood-flow inspection: the thoracic arteries	91
6.1	Motivation	92
6.2	Probing	92
6.2.1	Appearance	92
6.2.2	Positioning	92
6.3	Seeding	93
6.3.1	Seeding strategies	93
6.3.2	Seeding distributions	93
6.4	Visualization	94
6.4.1	Planar reformat	94

6.4.2	Integral lines	96
6.4.3	Arrows	97
6.4.4	Anatomical context	100
6.5	Evaluation	101
6.6	Discussion and conclusions	104
7	Blood-flow inspection: the cardiovascular system	107
7.1	Motivation	108
7.2	Probing	108
7.2.1	Appearance	108
7.2.2	Positioning	109
7.2.3	Fitting	110
7.2.4	Refinement	111
7.3	Seeding	112
7.3.1	Seeding strategies	112
7.3.2	Seeding distributions	112
7.4	Visualization	113
7.4.1	Particles	113
7.4.2	Integral lines	115
7.4.3	Integral surfaces	117
7.4.4	Anatomical context	119
7.5	Evaluation	122
7.6	Discussion and conclusions	125
8	Blood-flow inspection: an ultrasound perspective	129
8.1	Motivation	130
8.2	Probing	131
8.2.1	Appearance	131
8.2.2	Interaction	131
8.3	Visualization	134
8.3.1	Mimicking ultrasound renditions	134
8.3.2	Enhancing ultrasound renditions	136
8.4	Evaluation	140
8.5	Discussion and conclusions	143
9	Discussion and conclusion	147
9.1	Contributions	148
9.2	Discussion	149
9.3	Future research	153
9.4	Conclusion	155
A	The cardiovascular system	157
B	Appendix to Chapter 6	161
C	Appendix to Chapter 7	173

References	187
List of figures	203
List of tables	205
List of abbreviations	205
List of symbols	209
Curriculum vitae	217
List of publications	219
PhD portfolio	221

Summary

Real-time Illustrative Visualization of Cardiovascular Hemodynamics

Healthcare institutions generate vast amounts of clinical imaging data. Because of the advances in acquisition techniques, contemporary imaging data can be multi-modal, multi-dimensional and multi-valued by nature. In particular, modern magnetic resonance imaging (MRI) techniques enable acquisition of multiple image series that supply anatomical and functional information. In this thesis, we concentrate on visual analysis of MRI-acquired blood-flow information in the heart and the thoracic arteries.

In addition to anatomical information, MRI enables non-invasive acquisition of time-resolved blood-flow velocity data that capture the intricate cardiovascular hemodynamics. These quantitative velocity data describe the blood flow by means of volumetric velocity fields during a heart beat. This is often referred to as four-dimensional blood-flow data, based on the three spatial dimensions plus the time.

For this relatively new MRI acquisition technique, physicians are rather unsure what to expect from the data. Nevertheless, there are clear indications that the data contain valuable information. Quantitative and qualitative analyses of these data should provide insight into the blood-flow dynamics, improving the understanding of the cardiovascular system and its pathologies. This improved understanding conceivably leads to better diagnosis and prognosis of cardiovascular diseases, and may facilitate risk assessment, as well as evaluation of treatment and follow-up studies.

With qualitative analyses, physicians aim for newfound insight into the intricate blood-flow dynamics, and therefore there are no a priori questions to be answered, or tasks to be performed. The visual analysis should enable exploration of the complex high-dimensional data. However, exploration through the typical series-by-series and slice-by-slice inspection requires a full mental reconstruction of the unsteady blood-flow velocity data, as well as the cardiovascular morphology. This is a tedious and highly challenging task, even for skilled physicians. Therefore, we aim to alleviate this task by means of comprehensive exploratory and interactive visualization techniques. These techniques incorporate domain knowledge, and provide a more abstract representation of the data that can be steered interactively by the physicians.

Prior to visual analysis, sensible abstraction of the high-dimensional data is generally required. We have investigated various approaches to simplify the abundance of information contained in the acquired blood-flow data. On the one hand, we present a segmentation of the luminal geometry, using both direction and speed of the blood-flow velocities. We show that the inclusion of directional information leads to more accurate segmentation results. On the other hand, we abstract the time-resolved blood-flow data using spatiotemporal hierarchical clustering. The resulting cluster tree allows for intuitive level-of-detail selection, using a single user-defined parameter. For sparser detail levels, we use the cluster results in various visualization techniques, providing an abstract overview of the blood-flow data.

To facilitate interactive exploration of the four-dimensional blood-flow data, we introduce different probing tools, enabling local analysis of the hemodynamics. The probes enclose a region-of-interest, and serve a basis for various visualizations.

The first probing approach focusses on the thoracic arteries, using an automated technique to select vessel cross-sections, perpendicular to the centerline of the vessel. With these cross-sections as a basis, we introduce novel geometry-based blood-flow visualization approaches, such as exploded planar reformats, and flow-rate arrow-trails. In addition, we present improvements on established flow visualization techniques, such as dynamic pathline seeding, and animated pathline highlights. All blood-flow visualizations are combined with an illustrative context to communicate the anatomy.

The second probing technique enables exploration throughout the cardiovascular system. To this end, we introduce a virtual probe that resides in the blood-flow field. The virtual probe can be translocated by means of elementary two-dimensional interactions, enabling exploration. Based on the location of the virtual probe, we introduce novel visualization techniques, such as comic-inspired particles, illustrative pathlines, and nested pathtubes. Furthermore, we have investigated approaches to communicate the anatomical context, using volume projections and volume clipping.

The results of both probing approaches were evaluated with domain experts, measuring the value of the visualizations, the interaction approaches, and the involved user parameters. The evaluation questionnaires were carried out with several physicians, who are actively involved with advancements in MRI blood-flow acquisition, and have in-depth knowledge of diagnosis and treatment of cardiovascular diseases. The feedback obtained from these evaluation studies have yielded valuable insights concerning the presented visualization and interaction techniques.

Furthermore, we have extended the use of the virtual probe, visualizing the four-dimensional MRI blood flow in a similar way as used with color Doppler ultrasound imaging. Ultrasound is an established technique for blood-flow measurements, and the typical red-blue visualizations are familiar to the physicians. We introduce a compound view with different visualizations, inspired by ultrasound imaging, while exploiting the merits of the volumetric MRI blood-flow velocity data.

All presented visualization techniques perform in real-time, enabling interactive exploration of the four-dimensional blood-flow data. The renditions update instantaneously when moving the probe, or when the user parameterizes the visualization. Furthermore, real-time interaction with the virtual camera facilitates the visual inspection, providing different viewpoints and enhancing perception of depth in the animated volumetric representations of the blood-flow. To achieve this performance, we have employed modern consumer graphics hardware for our visualizations, enabling parallel processing of the graphics and associated algorithms.

Based on the evaluation studies with the involved physicians, we believe that real-time exploration of time-resolved volumetric blood-flow data, by means of illustrative visualizations, facilitates qualitative analysis of the hemodynamic behavior. We were able to present exemplary pathological cases. Time will reveal what new insights can be obtained by means of exploratory qualitative analyses.

Samenvatting

Gezondheidsinstellingen genereren grote hoeveelheden klinische beeldgegevens. Door de ontwikkelingen in beeldvormingstechnieken zijn moderne beelddata vaak multi-modaal, multi-dimensionaal en bevatten ze meerdere waarden per beeldpunt. Moderne magnetic resonance imaging (MRI) technieken maken het mogelijk om beeldreeksen te acquireren. Deze reeksen kunnen zowel anatomische als functionele informatie bevatten. In dit proefschrift concentreren we ons op de visuele analyse van bloedstromingen in het hart en de thoracale slagaders, op basis van MRI beelden.

Naast anatomische beelden kan MRI ook op een niet-invasieve wijze tijdsgebonden snelheidsgegevens van de bloedstroming meten. Deze gegevens leggen de complexe dynamiek van de cardiovasculaire haemodynamica vast. Deze kwantitatieve gegevens beschrijven de bloedstroming door middel van een aantal volumetrische snelheidsvelden gedurende de hartslag. Deze beelddata worden vaak vier-dimensionaal genoemd, gebaseerd op de drie spatiële dimensies plus de tijd.

Artsen weten aanvankelijk niet precies wat zij kunnen verwachten van deze data, die zijn verkregen via nieuwe beeldvormingstechnieken. Er zijn echter aanwijzingen dat de gegevens belangrijke informatie bevatten. Kwalitatieve and kwantitatieve analyse van de gegevens moeten daarom inzicht bieden in het dynamische gedrag van de bloedstroming. Dit kan leiden tot beter begrip van het cardiovasculaire systeem en zijn pathologieën en kan het bepalen van risico's ondersteunen. Tevens ondersteunt een beter begrip de evaluatie van behandelingen, alsmede opvolgstudies na behandeling.

Door middel van kwalitatieve analyse zoeken artsen naar nieuwe inzichten in de complexe haemodynamica. Er zijn in dit geval aanvankelijk geen duidelijke taken die uitgevoerd worden of vragen die beantwoord dienen te worden. De visuele analyse moet het daarom mogelijk maken om de complexe multi-dimensionale data te verkennen. Het is gebruikelijk om de gegevens te verkennen door middel van inspectie van de individuele plakken per reeks. Dit vereist een volledige mentale reconstructie van zowel het tijdsgebonden stromingsveld, alsmede de cardiovasculaire morfologie. Dit is een moeizame en uiterst uitdagende taak, zelfs voor geschoolde artsen. Wij willen deze taak vereenvoudigen door middel van uitgebreide interactieve visualisatietechnieken. Deze technieken maken gebruik van kennis uit het medische vakgebied en leveren abstracte en aanpasbare representaties van de bloedstroming.

Voorafgaand aan de visuele analyse dienen de multi-dimensionale gegevens op een gepaste manier te worden geabstraheerd. We hebben daarom onderzocht hoe de grote hoeveelheid informatie in de metingen kan worden gereduceerd. Ten eerste hebben we gekeken naar segmentatie van de stromingssnelheden, gebruikmakend van zowel richting als snelheid in het veld. We tonen aan dat het gebruik van richtingsinformatie leidt tot een nauwkeuriger segmentatie. Daarnaast hebben we de tijdsgebonden bloedstromingsvelden vereenvoudigd door middel van spatiotemporele hiërarchische clustering. De cluster-boom maakt het mogelijk om op intuïtieve wijze een detail-niveau te kiezen, op basis van één gebruikersparameter. Op abstractere niveaus gebruiken we de clusters in verschillende visualisaties en bieden zo een overzicht van de bloedstroming.

Om het vier-dimensionale bloedstromingsveld interactief te kunnen verkennen introduceren we verschillende virtuele probes. Deze probes maken het mogelijk om de haemodynamica lokaal te inspecteren. De verschillende probes behelzen een lokaal interessegebied in de data en dienen als basis voor verschillende visualisaties.

De eerste techniek richt zich op de thoracale slagaders en plaatst automatisch een probe op de dwarsdoorsnede van het bloedvat, loodrecht op de middellijn van het vat. We gebruiken deze dwarsdoorsnede als basis voor innovatieve visualisatietechnieken, zoals 'exploded planar reformats' (uitgewaaierde vlakken) en 'flow-rate arrow-trails' (een spoor van stroomsnelheidspijlen). Daarnaast hebben we verbeteringen geïntroduceerd voor bestaande technieken, zoals het dynamisch initiëren van pathlines en geanimeerde markeringen op deze deeltjesbanen. Alle visualisaties van de bloedstroming zijn gecombineerd met een illustratieve anatomische context.

De tweede techniek met een virtuele probe maakt het mogelijk om het cardiovasculaire systeem in zijn geheel te verkennen. De probe bevindt zich daarom in het stromingsveld en kan voor verkenning worden verplaatst door middel van eenvoudige twee-dimensionale interacties. Op basis van de locatie van de virtuele probe hebben we verschillende innovatieve visualisatietechnieken geïntroduceerd. We presenteren particles (deeltjes) geïnspireerd op stripboeken, alsmede illustratieve pathlines en geneste pathtubes (buisvormige oppervlakken, op basis van deeltjesbanen).

De technieken op basis van de virtuele probes zijn geëvalueerd met experts in het veld, waarbij we de waarde van de visualisaties, de interacties en de verschillende parameters hebben gemeten. Door middel van vragenlijsten hebben we informatie verkregen van verschillende artsen die betrokken zijn bij de ontwikkelingen van MRI metingen van de bloedstroming. Zij hebben tevens een grondige kennis op het gebied van diagnose en behandeling van cardiovasculaire ziekten.

Bovendien hebben we de virtuele probe gebruikt om de vier-dimensionale bloedstromingsvelden weer te geven als beelden van Doppler echografie. Echografie is een gevestigde beeldvormingstechniek voor bloedstromingen en de rood-blauw visualisaties van echografie zijn vertrouwd voor veel artsen. We introduceren een samengesteld overzicht met verschillende visualisaties, geïnspireerd op echografie, waarbij we de voordelen van de volumetrische MRI gegevens benutten.

Alle gepresenteerde technieken zijn real-time, waardoor interactieve verkenning van de vier-dimensionale data mogelijk is. De visualisaties verversen onmiddellijk wanneer de probe wordt verplaatst of parameters worden aangepast. De virtuele camera kan interactief worden bediend. Hierdoor zijn de animaties vanuit meerdere standpunten te bekijken, en kan de ruimtelijke diepte worden ingeschat. Om deze snelheden te bereiken, hebben we gebruik gemaakt van moderne grafische hardware, waardoor parallelle verwerking van de visualisaties en algoritmen mogelijk is.

Op basis van de evaluatie studies met de betrokken artsen zijn we van mening dat interactieve verkenning van tijdsgebonden volumetrische bloedstromingsvelden, door middel van illustratieve visualisaties, de kwalitatieve analyse van de haemodynamica bevordert. We hebben verschillende voorbeelden laten zien van pathologische bloedstromingsvelden. De tijd moet echter uitwijzen welke inzichten kunnen worden verkregen met behulp van interactieve kwalitatieve analyse.

"Science and the human heart, there is no limit"

U2

Introduction

1

1.1 Preface

In the wide variety of specialized research disciplines, theses describe the continuous advancements and novelties in their respective fields. The term *thesis* literally means a proposition to be proven, typically offered by a candidate for scholastic honors (source: Merriam-Webster). What better way to start such a thesis than with a ventured statement, open for debate, and ready to be substantiated?

“A thesis in visualization entails gobbledygook for perception”

We start the underpinning of our assertion, by considering the individual keywords. Since this thesis ventures into the realms of *visualization*, this is the first term that will be addressed. In brief, visualization is the process of forming a mental image. The human visual system is able to receive large quantities of information instantly, when presented in the right form. Visualization as a research discipline uses computer systems to communicate a great variety of data, aiming for visual cognition.

The next keyword to be addressed is the quintessentially English term *gobbledygook*: Gob·ble·dy·gook *noun* \ˈgä-b əl-dē-gük\ - wordy and generally unintelligible jargon (source: Merriam Webster). Every research field inevitably communicates with a certain jargon, to which the visualization community is no exception. Although complicated for people inexperienced in the domains, technical disciplines successfully employ the intricate terminology to unambiguously communicate, and accurately and concisely formulate, their specific treatises.

This thesis is concerned with medical visualization, aiming to facilitate the qualitative analysis of clinical imaging data. Consequently, specific wording from the medical domain is required. Since measured data generally needs processing prior to visualization, terminology from the field of computer vision is introduced as well. The gobbledygook that results from this intertwined nomenclature easily leads to ambiguous terms, and it was our challenge to avoid incomprehension between the fields of expertise. Ambiguous terms will be clarified in separate frames throughout this thesis, while some terminology inevitably remains targeted at a specific readership.

Paradoxically, the involved use of terminology in the visualization community is intended to achieve simplicity. Through simple yet comprehensive visual representations of complex data, we aim for cognition in the observer. Hence we need to understand how an observer perceives the provided information, which brings us to the last keyword of our statement, namely *perception*.

Perception is the process of attaining understanding or insight through recognition and interpretation of sensory stimuli. In particular, we rely on visual perception, which is a research field on its own. The total understanding of the information contained in a given data set through visual analysis is the ultimate goal: the holy grail of visualization research.

Considering our initial statement, it follows that indeed visualization research relies on a plurality of terms, both in the technical domain, as well as in the domain of the target users. This jargon effectively underpins the scientific communication, striving for visual representations that are simple enough to comprehend, and complete enough to communicate the desired information. In other words, we need the gobbledygook to achieve the sought-after perception.

The achievement of perception is, however, nontrivial. Although reception and interpretation of sensory information is acute and intuitive, there is a caveat that impedes the noble quest for cognition. Perception is based chiefly on memory, reducing the acquired information to empirical factual knowledge. Fortunately, experiences and factual knowledge are partially universal, largely imparted during childhood, and dependent on the social and cultural upbringing.

A suitable example is found in comics. Knowingly or unknowingly, we all learned to read comics, perceiving time spatially while advancing from frame to frame. We intrinsically know how to read a comic, based on previous experience and empirical knowledge. To understand the value of this visual medium, we look at a definition: Comic *noun* \kom-ik \ - juxtaposed pictorial and other images in deliberate sequence, intended to convey information and/or to produce an aesthetic response in the viewer (source: Understanding Comics [158]).

From this definition, it becomes apparent that comics strive for two goals. On the one hand, comics need to appeal to the reader, and hence aesthetics play a role. On the other hand, comics need to convey information, telling the story as intended by the author. Effective communication of information is an important common goal between comics and visualization research.

The process of creating comics is associated with a high degree of artistic freedom. The author decides how the story evolves, and how it is communicated to the reader. Visualization of a given data set does not allow this level of artistic freedom. Nevertheless, comics artists possess a variety of interesting techniques that effectively convey the intended message. These techniques could be employed to communicate information contained in a data set, by embedding these techniques into specialized visualizations. One particular abstraction technique often used in comics is cartooning. Simplistic representations capture the necessary information. This process is not so much about eliminating details, but to focus on parts of the scene or data [158].

Different media have different benefits for communicating visual information. Instead of comic books, visualization research generally employs computer systems to effectuate the necessary imagery. It employs various techniques that originate from the area of computer graphics. This enables direct user interaction with the visualization, and fast qualitative analysis of large quantities of data.

1.2 Context of this research

The visualization work presented in this thesis investigates ways to facilitate the visual analysis of medical data. Healthcare institutions generate vast amounts of clinical imaging data, which in itself are becoming more complex by nature. Advances in acquisition techniques provide multi-modal, multi-dimensional and multi-valued data sets. These developments are furthermore fostered by clinical research. In this thesis, the term clinical research refers to an early research phase, comprising fundamental studies prior to clinical trials. This research aims to improve acquisition methods and investigates open clinical questions. Rapid visual analysis of the increasingly large and complex medical imaging data is therefore of importance.

Our research concentrates on data that capture the cardiovascular hemodynamics. Typically, the blood flow is expressed by velocity fields. In particular, we rely on volumetric and time-resolved velocity data, acquired by magnetic resonance imaging. These non-invasively measured velocity data describe the direction and speed of the bloodstream during a full heart beat. The velocity data are often combined with data that represent the cardiovascular morphology. Consequently, the blood-flow acquisitions provide exemplary high-dimensional and multi-valued data. An overview of the cardiovascular system and the blood-flow circulation is provided in appendix A.

For relatively new acquisition techniques, as is the case for the time-resolved volumetric blood-flow data, physicians are rather unsure what to expect from their data. Although there are clear indications that the data contain valuable information that is worthwhile to investigate, it is often unclear to the expert what exactly he seeks. The expert strives for newfound insight into the data, advancing the state of the art in his or her domain [245]. This may potentially improve future diagnosis and prognosis. Furthermore, the acquired understanding may facilitate risk assessment, and pre- and post-operative evaluations.

This situation is a prime example where perception of the given data is a key issue. To gain insight into the large and intricate data, visualization and interaction techniques are required that enable *exploration* of these data. The challenge for these exploratory techniques is to communicate the appropriate information, while there is no specific question to be answered, or particular task to be performed. This especially holds true for the large volumetric and time-resolved blood-flow data, which are virtually impossible to mentally grasp to its full extent.

Specific exploratory visualization techniques are therefore proposed, typically as part of prototype software. These prototypes comprise an extensive assortment of tools to inspect the data, turning into application-specific versatile visualization platforms. Generally, these prototypes are quite overwhelming in terms of user parameters. The software should be operated by researchers, who aim for a better understanding of their data. These new insights can lead to more specific *task-driven* tools, enabling a fixed workflow for studies in the clinical research phase. After that, there is a lengthy process towards *clinically* viable software, including validation.

Exploratory prototypes constitute the early beginning towards clinically valuable methods. For the exploration of large and complex data, interactivity is of vital importance. This requires fast processing, which can for instance be found in modern consumer graphics hardware. This hardware enables parallel processing of graphics and associated algorithms in real-time. These capabilities can likewise enable an interactive visual analysis of the blood-flow data, including real-time interaction with a virtual camera, and parametrization of the visualizations.

Under the assumption that an interactive processing is feasible with modern computer hardware, then the question remains how to effectively communicate the information contained within such large data? Therefore, we go back to the preface, where the notion of perception was discussed. An example of effective visual communication was found in cartooning, frequently employed by comic books. Cartooning enables simplistic representations of complex entities, while retaining the necessary information. Such techniques can potentially provide comprehensive and abstract renditions to convey the large and complex blood-flow dynamics.

Cartooning and similar abstraction techniques have been utilized in medical visualizations for centuries. Early drawings of medical conditions and interventions, such as the ones by Leonardo da Vinci in the late 15th century, clearly provided abstractions of reality, to induce cognition in the viewer. Also nowadays, simplified representations are still actively used to convey medical information, as for instance evidenced by modern-day atlases of the human physiology.

From the foregoing, the question arises what level of abstraction is most effective? Photorealism often provides excessive details, and needs simplification. Too much abstraction results in loss of important data characteristics in the visual representation. Continual simplification leads to the ultimate achievable abstraction, namely words and numbers. These representations are intensively used in quantitative analyses, providing figures of various physical entities.

However, linguistic means communicate the essence of the data indirectly. Especially for spatial patterns, words and numbers require much more time to interpret, in contrast to genuine visual representations, which are perceived instantaneously. Choosing the right level of abstraction is an important ongoing challenge in visualization research. The aim is to effectively convey the necessary information, without compromising essential characteristics of the data.

Considering the aforementioned, we presume that perception of the large and intricate blood-flow velocity data can be achieved by interactive visual exploration. We argue that abstract representations, for instance through cartooning, facilitate qualitative analyses of the hemodynamics. Hence, our research is underpinned by this postulation:

Interactive and real-time visual exploration of volumetric and time-resolved cardiovascular blood-flow velocity data, communicated through simplified representations, facilitates qualitative analyses for clinical research, and potentially effectuates better understanding of the hemodynamics in both healthy and pathological conditions.

1.3 Outline

In this thesis, we investigate a variety of visualization and interaction techniques to analyze measured time-resolved blood-flow velocity data. This thesis primarily comprises a technical treatise of the visualization approaches, and the associated interaction and performance aspects. Besides that, considerable attention is given to the medical applications, and the impact of the presented techniques for clinical research in cardiovascular hemodynamics.

In **chapter 2**, we provide the clinical background. After a short exposition on the role of blood in the history of medicine, we provide the overall rationale behind the analysis of cardiovascular hemodynamics, underpinned by the necessary statistics. Different ways to obtain blood-flow velocity information are addressed, followed by a description of common analysis techniques to gain insight into the hemodynamics. At the end of this chapter, two example cases underpin the need for the qualitative analysis of cardiovascular hemodynamics, both in congenital and adult pathologies.

Chapter 3 elaborates on the acquisition of blood-flow velocities, by magnetic resonance imaging. A short introduction into magnetic resonance imaging is followed by a more detailed treatment of velocity measurements, and the characteristics of the measured data.

The technical background of our approaches is then provided in **chapter 4**. First, an introduction contemplates on the visualization research field, then we discuss related work, which we have divided into two areas of interest. On the one hand, we motivate and discuss illustrative visualization techniques. On the other hand, we elaborate on the established field of flow visualization.

In order to attain an effective visualization of the measured blood-flow fields, adequate abstractions are required as a basis for different visualization methods. **Chapter 5** first discusses two temporal projection techniques, which have been utilized for various purposes throughout our research. Subsequently, we elaborate on two complementary approaches that abstract the information contained within a blood-flow measurement. We propose a novel segmentation technique, which enables extraction of the luminal geometry. We also introduce a novel spatiotemporal hierarchical clustering of the blood-flow velocities, in combination with an accompanying visualization approach. The clustering enables abstraction through a user-defined level-of-detail selection.

In **chapter 6**, we introduce a visualization and interaction approach, designated to applications concerned with the thoracic arteries. We present a probing approach to locally inspect blood-flow dynamics, based on a variety of novel or enhanced geometry-based flow visualization techniques. Furthermore, we discuss the evaluation of the various techniques, measuring the effectiveness of the visualizations with domain experts.

An extension of the techniques of chapter 6 to the cardiovascular system as a whole is presented in **chapter 7**. We introduce a volumetric probe, operated by elementary two-dimensional interactions. This probe enables real-time exploration of the entire blood-flow velocity field in a piecewise fashion. Novel visualization techniques, such as illustrative particles and surfaces, are presented to communicate the hemodynamics. The chapter concludes with an evaluation study that assesses the effectiveness of the various visualization and interaction techniques.

Although our visualization approaches rely on velocity data acquired by magnetic resonance imaging, color Doppler ultrasound remains the standard for clinical blood-flow velocity imaging. The field of Doppler ultrasound imaging relies on different visualization techniques, typically driven by red and blue color coding of the blood flow. In **chapter 8**, we present a framework that mimics Doppler ultrasound renditions, based on velocity data measured by magnetic resonance imaging. Two-dimensional renditions are combined and linked to volumetric visualizations, including a novel velocity field raycasting.

Chapter 9 provides a general discussion of the presented research, assessing our postulation. Furthermore, indications for future work are discussed.

"All this talk of blood and slaying has put me off my tea."

Cheshire Cat

Clinical background



2.1 Introduction

2.1.1 A short history

From early times, when medical practice was ruled by diviners and witches, blood has been regarded the liquid of life, nourishing the body [182]. With the advancement of civilization, the medical system slowly developed, combining religious rites with empirical treatments.

The extensive work 'Corpus', assigned to Hippocrates (ca. 460-370 BC), revolutionized medical practice. In his work, health and illness were described in terms of four key fluids, known as the four humors: blood, choler, phlegm, and black bile. These humors were associated with the human physique, providing a flexible explanatory framework, underpinned by archetypal contrasts (i.e., dry or wet, warm or cold). For instance, tenuity was caused by excessive choler, while overweight was associated with phlegm. Also body coloration was explained by the four humors: blood colors red, choler yellow, while phlegm led to paleness, and black bile implied darkness.

Galen of Pergamon (ca. 129-200 AD) substantially extended Hippocrates' theory of the humors, relating mental state to the four key fluid. As depicted in figure 2.1a, blood, choler, phlegm, and black bile were associated with a sanguine (sanguin), choleric (coleric), phlegmatic (flegmat), and melancholic (melanc) state-of-mind respectively. The prevailing conception was that illness resulted from humoral imbalance, and hence that health could be regained by balancing of these humors.

Throughout the Middle Ages, medical practice remained a wait-and-watch profession. The physician's demeanor was deemed distinguished, trustworthy, and sober; soothing circumstances instead of curing disease. One notable intervention persisted far into the Renaissance: bloodletting [82]. Besides a broad range of quackery, the practice to withdraw considerable quantities of blood was propelled by the authority of Galen, who believed that blood was produced by the liver, and subsequently consumed by the body for nourishment. A surplus of blood was drained to balance the humors, supposedly improving the patient's health.

In the late Renaissance, human dissection became accepted in the western world, instigated by the work of Vesalius (1514-1564). Human dissection provided a valuable new source of information, causing physicians to rethink the human body. However, direct access to the anatomy did not instantly lead to a profound understanding of human physiology, as reasonings were often compromised by traditional views. Only in 1602, William Harvey (1578-1657) largely unraveled the cardiovascular circulation and properties of blood, opposing longstanding dogmas. The full circulatory pathway was discovered years later, when microscopy revealed the capillary veins.

In the ages to follow the Renaissance, anatomical research evolved, and experimental investigation became commonplace. Experiments on blood transfusion commenced, against all ethical discretions at the time [230]. Besides the pulse, which was already known in Galen's age, clinical diagnosis had not included noteworthy measurements.




Figure 2.1 (a) Historical artwork by Thurn-Heisser (1574) of a half male and half female figure, illustrating the four humors. Illness was described in terms of these key fluids. Recuperation involved balancing the four humors through medicine or surgery. (b) Stephen Hales was the first person to measure blood pressure. He cannulated the carotid artery of a horse, observing the blood level.

In 1711, Stephen Hales (1677-1761) was the first to carry out a blood-pressure measurement, though crude and invasive [168]. As depicted in figure 2.1b, he cannulated the carotid artery of a horse with a 2.5 cm wide brass pipe, and observed the blood level in an attached 2.5 meter long glass pipe. Furthermore, he investigated the cardiac output of the horse's heart, by injecting bee's wax. In his work [81], he motivates these experiments:

“And as the healthy state of an animal principally consists, in the maintaining of a due equilibrium between solids and fluids; it has, ever since the important discovery of the circulation of blood, been looked upon as a matter well worth the enquiring into, to find the force and velocity with which these fluids are impelled; as a likely means to give a considerable insight into the animal œconomy.”

Ever since, many apparatus were invented to measure blood-flow characteristics. In 1834, the first non-invasive manometer was introduced, followed by the Sphygmograph in 1855, plotting the amplitude of the pulse [168]. These developments led to the modern arm cuff, which was first introduced in 1896. In the meantime, Poiseuille formulated the first physical laws to determine fluid flow in a tube, effectuating the foundation for contemporary computer simulations.

2.1.2 Motivation

After ages of crude measurements and in retrospect erroneous views, one opinion remains: blood is an influential factor in the cardiovascular system [194]. Ever since Hales' experiments, research has gained a much better understanding of the circulatory system. However, to a certain extent his motivation still holds today, as many aspects of the hemodynamics, and the interaction with surrounding biology, are still not fully understood. In this thesis, we focus on hemodynamics of the cardiovascular system and related diseases. In particular, we investigate the blood-flow in the heart and thoracic arteries, while we briefly touch upon a cerebrovascular application.

Cardiovascular disease (CVD) comprises a group of conditions with a prevalence of over thirty percent of the American population [5], and is the main cause of disease burden in Europe [3]. CVD is currently the leading cause of death worldwide (Fig. 2.2). Diagnosis and prognosis of CVD are largely based on acquired morphology, often extended by functional information from time-resolved scans. In this process, the blood-flow dynamics are largely neglected, while they are of vital importance for a healthy cardiovascular system. Understanding the hemodynamics potentially harbors a considerable value for future diagnosis and risk assessment.

A large body of clinical research indicates that atypical flow behavior directly relates to medical conditions [15, 91, 152, 235, 242, 262]. It is shown that anomalous blood-flow influences the morphology of surrounding tissue, for instance leading to vascular remodeling, and hence directly affects pathogenesis. This becomes apparent when examining the innermost layer of the vessel wall, the tunica intima. This layer interacts with the bloodstream to regulate coagulation [170]. Abnormalities in this process potentially cause atherosclerotic plaques, which may lead to blood clots, and hence myocardial infarction or stroke.

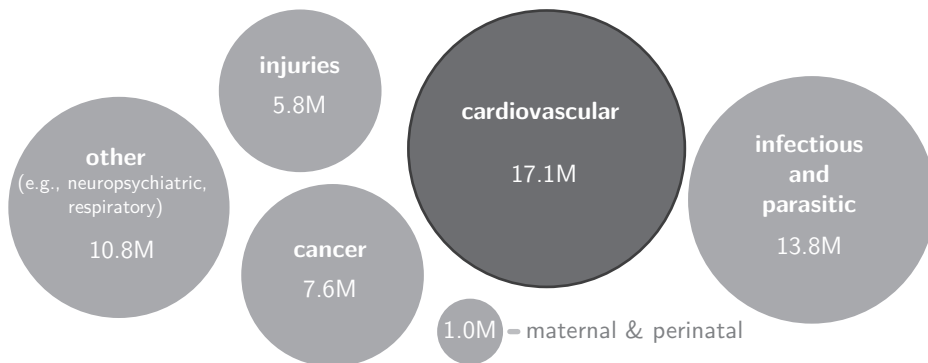


Figure 2.2 A taxonomy of the global disease categories, and the associated mortality worldwide in the year 2004. The statistics are based on indicators from World Health Organization (WHO) [266] (table A5). The size of the circles is scaled with the number of deaths. Cardiovascular disease is currently the leading cause of death worldwide, as emphasized by the darker circle.

For many cardiovascular applications, a profound understanding of blood-flow behavior can be a decisive factor for the diagnosis and prognosis of the disease. Figure 2.3 shows a global taxonomy of cardiovascular diseases, emphasizing the areas where hemodynamic information could be of considerable importance. Especially for conditions with malformations of the cardiovascular morphology, understanding the blood flow behavior is of interest. Malformation may affect the efficiency of the recirculation and cause anomalous vortices and blood-flow jets, which in turn influence the surrounding biology.

To gain understanding of the hemodynamics, an appropriate imaging modality is required. Contemporary modalities enable in-vivo and non-invasive acquisition of blood-flow characteristics, typically measuring velocity information. The benefits and drawbacks of the various imaging approaches will be elaborated in section 2.2.

While proper measurements provide an invaluable source of information, insight can only be obtained through an extensive exploration and analysis of the acquired data. The aim is to discover appropriate quantitative indicators and qualitative patterns that evidence pathology, or prognose future disorders. Clinical research currently strives to reveal the relations between the hemodynamics and a broad range of cardiovascular diseases, and hence the analysis of blood-flow information is part of the daily research practice. Both quantitative and qualitative analysis will be described in more detail in section 2.3.

A better understanding of blood-flow behavior in relation to different diseases, does not imply a direct yield for clinical practice. Conclusive clinical trials are necessary to substantiate the initial findings. This may eventually lead to clinical protocols, enabling a routine use of the acquired knowledge.

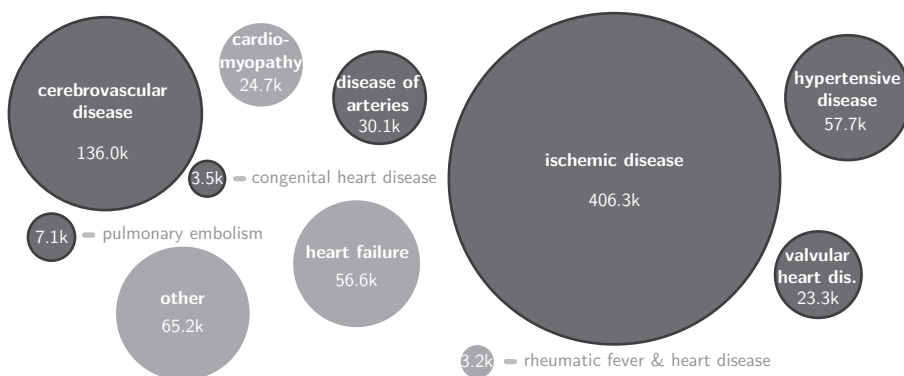


Figure 2.3 Cardiovascular disease mortality in the United States for the year 2007, accounting for 33.6% of all deaths. The statistics are based on indicators from the American Heart Association (AHA) [195], and the classification is based on the WHO International Classification of Diseases (ICD). For the majority of the cardiovascular diseases, blood-flow analysis potentially provides valuable diagnostic and prognostic information. The groups with the most potential are indicated by the darker bubbles. The other groups can nevertheless also benefit from blood-flow information.

2.2 Blood-flow acquisition

2.2.1 Modeling blood-flow mechanics

As described in subsection 2.1.1, many aspects of the blood-flow behavior, and of associated physical properties, were revealed by experimental investigations. The behavior of blood flow can also be addressed from a theoretical perspective, formulating the physical characteristics.

In physics, the field of fluid mechanics is concerned with the characteristics of fluids, comprising liquids and gases [94, 127]. A fluid is defined as a substance that deforms continuously under the action of a shearing force of any magnitude [166]. In particular, fluids at rest are studied in the branch of fluid statics, while fluids in motion are investigated in the area of fluid kinematics. Furthermore, the area of fluid dynamics considers the effect of forces on fluid motion, which is of particular importance for the understanding of blood-flow dynamics.

Fluids can be studied at the most elementary level of molecular interactions. Although this is addressed in kinetic theory, generally the macroscopic behavior of a fluid is of interest. To that end, the molecular level is neglected, and a continuous distribution of the matter is considered. This continuum is valid whenever the spacing between molecules is considerably smaller than the size of the flow system.

A fluid flow field is determined by a velocity vector, and two thermodynamic properties (e.g., pressure, density, temperature) at each position in the field. A flow field is therefore defined by a system of independent equations, which relies on the laws of conservation (i.e., conservation of mass, momentum, and energy). These systems become mathematically complex, and are therefore often solved numerically. The discipline of computational fluid dynamics (CFD) is dedicated to the numerical analysis of problems in fluid mechanics.

The unsteady blood flow is incompressible, meaning that the material density remains constant for isothermal pressure changes. Furthermore, blood is a non-Newtonian fluid, meaning that the relation between shear stress and strain is non-linear [127]. Its viscosity decreases with a higher rate of shear stress, which is referred to as a shear-thinning fluid. However, in veins with a diameter larger than a few millimeters, the decrease in viscosity is marginal, and blood behaves approximately as a Newtonian fluid. The motion of a viscous Newtonian fluid is described by the Navier-Stokes equations. These equations state that changes in momentum of fluid particles only depend on the external pressure and internal viscous forces acting on the fluid. The Navier-Stokes equations have been studied intensively, and various numerical approximations have been proposed, depending on the application (e.g., Treuille et al. [228]).

Whenever a viscous fluid flows in layers, or laminae, the behavior is termed laminar. In contrast, fluid flow is considered turbulent when the velocity components have random fluctuations imposed on their mean values [94]. The transition from

laminar to turbulent flow occurs at a fixed Reynolds number of approximately 2000, describing the dimensionless ratio of the inertial forces to the viscous forces [127]. Although chaotic behavior does occur in the cardiovascular blood flow, the behavior is predominantly laminar. Laminar flow occurs in different patterns; e.g., helical flow.

In the past few decades, CFD studies have performed a wide range of simulations to analyze the cardiovascular blood-flow dynamics. Some research focussed on the aortic blood flow [205], while others inspected the cardiac hemodynamics [52, 165, 221, 249]. Moreover, patient-specific simulations have been carried out to investigate pathological flow. For instance, the blood flow in an abdominal aortic aneurysm [21], an aortic bypass [226], and a univentricular heart [176] have been investigated. Additionally, the effect of implants on the hemodynamics has been studied, such as mechanical heart valves [213]. Besides modeling of the cardiac blood flow, also cerebrovascular blood flow is investigated using CFD simulations [32].

These models provide valuable information about the hemodynamics. However, as the term 'model' suggests, simulations remain a simplification of the actual flow field. Despite the thorough understanding of the physical characteristics, models rely on a range of assumptions that affect the simulation results. Many of these assumptions are applied consciously, for instance to confine the computational cost. Other assumptions are implicitly imposed on the model, and the effects are more difficult to control. For instance, the boundary conditions for patient-specific simulations are determined by the anatomy, and hence the segmentation quality plays a vital role.

2.2.2 Measuring blood-flow mechanics

Instead of modeling blood-flow dynamics, imaging techniques allow to measure the flow field. Various characteristics of the blood flow can be measured, while most often velocities are acquired. In current clinical practice, color Doppler ultrasound (US) is the reference standard for the imaging of blood-flow dynamics [85]. Ultrasound is cost-effective, and enables the acquisition of tissue structures and blood-flow velocities at a high spatiotemporal resolution. Measurements are mostly performed non-invasively, although the imaging technique has been incorporated into invasive procedures, by mounting transducers at the tip of a catheter. The inspection results largely depend on the operator skills, and are limited by a relatively small field-of-view. Moreover, US is considerably prone to noise, which detracts from the image quality.

Color Doppler US relies on the Doppler shift to measure blood-flow velocities. The results can be quantified for a given angle of insonation. This angle between the ultrasonic beam and the blood-flow velocity is mostly unknown. In practice, the transducer is therefore positioned at a small angle, which the operator approximates to be zero. Doppler US inherently measures blood-flow velocities, and does not allow for acquisition of volumetric flow rates. In complex situations, specific contrast agents in the form of micro bubbles, may enhance the imaging results [41]. However, three-directional velocity information cannot be obtained.

Although computed tomography (CT) is one of the primary modalities for imaging of anatomical structures, blood-flow acquisition is limited. Typically, angiographic projections are inspected, based on an acquisition with an intravenously injected contrast agent. Instead, recent advances in rotational X-ray imaging enable the quantitative imaging of blood-flow information, such as velocities and the mean flow rate. Several approaches enable quantification, using iodine-based contrast agents, including bolus tracking and optic flow techniques [253]. Both CT and rotational X-ray imaging have the drawback of exposing the patient to harmful radiation, in combination with an administered contrast medium. This is preferably avoided for young patients.

Magnetic resonance imaging (MRI) is another established imaging modality, enabling non-invasive acquisition of both morphology and blood-flow information [185]. In comparison to US imaging, a larger field-of-view can be acquired, with a substantially better signal-to-noise ratio, at the expense of a lower spatiotemporal resolution. In addition, MRI enables imaging of structures that are hard to measure using US (e.g., pulmonary arteries [17]). Magnetic resonance flow imaging can be enhanced by means of contrast agents, which are typically Gadolinium based. However, a wide range of magnetic resonance flow imaging techniques do not require contrast administration. In particular, velocity-encoded phase-contrast (PC) MRI sequences enable the acquisition of blood-flow velocity information, for quantitative analysis. The capability to measure time-resolved volumetric flow information, largely independent of the operator, often outweighs the additional cost, especially in case of complex conditions. In this thesis, we focus on MRI-acquired unsteady volumetric blood-flow velocity data. The associated flow imaging techniques will be elaborated in chapter 3.

Clarification of ambiguous terms

The work presented in this thesis is concerned with blood-flow '**velocity**' data. The term velocity is used ambiguously in practice. To prevent confusion, we understand the physical quantity of velocity to be defined as *speed in a certain direction*, and hence consider velocity as a vector entity. Elsewhere, the term is frequently used to indicate the scalar-valued property speed. We will consistently use the distinct notions of velocity and speed throughout this thesis.

Terminology to define a vector is likewise used ambiguously. A vector is fully determined by its magnitude and direction, where the direction is specified by the order of two points on a line parallel to the vector. The term 'direction' is often interchanged with the term '**orientation**'. In the definition that we adhere, orientation is unsigned, and hence antiparallel vectors with different directions have the same orientation. In other sources, the term 'direction' is almost invariably used as a signed entity, while 'orientation' is typically the cause of ambivalence. Often the terms are used interchangeably, without a distinct separation of definitions. In some cases, however, opposite definitions are practiced.

2.3 Blood-flow analysis

2.3.1 Quantitative analysis

Physicians investigate both the spatiotemporal blood-flow velocity fields as well as derived measures [154]. They strive for a better understanding of the cardiovascular hemodynamics, and the interaction with surrounding biology. On a global scale, their interest is to obtain a patient-specific understanding of the efficiency of the cardiovascular circulation. However, blood-flow analysis also needs to assess blood-flow information on a smaller scale, as anomalous blood-flow characteristics may manifest themselves locally within the cardiovascular system.

Both globally and locally, quantitative information about the hemodynamics play a vital role in the analysis process. Convinced by the correlation between a quantitative measure and pathogenesis, these data can be used for clinical diagnosis. This is done under the assumption that the measurements are performed properly. Typical imaging modalities, such as PC-MRI, provide velocity information of the blood-flow field. Using this information, many parameters can be derived, or at least approximated. On the one hand, there are measures purely related to the bloodstream, such as *acceleration*, *flow rate*, *kinetic energy* [255], and *pressure* [232, 270]. On the other hand, there are measures that describe the relation with the surrounding biology, such as the *wall-shear-stress* (WSS) and *oscillating-shear-index* (OSI).

Derivation of various measures, e.g., mean flow rate or pressure [232], requires an accurate segmentation of the anatomical structures. In the two-dimensional case, this may be a relatively straightforward delineation of a vessel circumference. Either the inner vessel wall, or the blood-flow lumen may be segmented. In the three-dimensional case, manual segmentation becomes tedious and time-consuming. Hence, this process needs to be automated [237] (see section 5.2). For clinical use of quantitative measures, automated segmentation is an important task that needs to be addressed.

It is challenging to devise meaningful application-specific quantitative measures. It requires profound research, starting from preliminary case studies to large-scale clinical trials. In the literature, several examples show the use of quantitative measures for cardiovascular applications. For instance, a *speed* measurement may reveal stenotic regions. The *regurgitant volume* is a useful measure to inspect valvular regurgitation in the aorta [16]. Alternatively, left and right ventricular *stroke volumes* may be compared to measure aortic regurgitation. Stroke volumes may be employed to measure cardiac shunts, often associated with congenital conditions such as an atrial septal defect. Alternatively, the *pulmonary-to-systemic flow ratio* provides an indicator for the severity of the shunt [234]. In case of an aortic coarctation, which is easily diagnosed from morphology, the degree of narrowing can be determined quantitatively by the ratio between the *flow rates* in the ascending and descending aorta [16], as the residual flow will enter the branches of the aortic arch. Other measures are the *cardiac output*, *retained inflow*, *delayed ejection flow*, and the *residual volume* [57].

This listing is by no means exhaustive. Many more measures exist, while other application-specific indicators are yet to be found. The success of these measures mostly depends on the ease of measurement, and their diagnostic value and reliability. Besides these quantitative measures, also visual analysis is important in the process of understanding patient-specific hemodynamics.

2.3.2 Qualitative analysis

Visual analysis of time-resolved volumetric flow data on a slice-by-slice basis becomes a tedious and time-consuming task. While skilled physicians are able to mentally reconstruct a spatial image from 3D scalar data, this becomes significantly more difficult for 3D vector-valued data. A time-resolved volumetric blood-flow data set currently consists of twenty to thirty vector-valued volumes over time, which is virtually impossible to grasp for the human mind. Instead, a simplified and insightful representation is needed to gain understanding of the blood-flow behavior. A wide variety of visualization techniques has been proposed to convey the vital aspects of flow fields (see section 4.3). These techniques strive to depict the important flow patterns.

A good understanding of the blood-flow behavior in healthy volunteers is required, prior to detecting anomalous blood-flow patterns in patient data. An example from the work by Kilner et al. [113] is presented in figure 2.4, schematically depicting the aortic blood-flow patterns over time. Typically, blood flow in the ascending aorta shows a right-handed helix, which alters to a straight flow or a left-handed helix in the descending aorta. This helicity is largely due to the aortic geometry, while motion of the aorta also plays a role [101]. Also amongst healthy volunteers, there are considerable variations, related to factors such as age. While for young people blood moves from the aortic valve to the mid-descending aorta in one heart beat, this takes up to two or three heart beats for elderly people [15].




Figure 2.4 Schematic depiction of blood-flow patterns in the aorta. (a) Early systole shows acceleration in the ascending aorta. (b) During mid systole a secondary helical flow develops. A separation zone occurs in the distal arch. (c) End systole is characterized by combinations of rotational and recirculating secondary flows. Figure from the work by Kilner et al. [113].

Building on the knowledge of healthy blood-flow patterns, anomalous behavior can be detected, and related to the pathological case. In clinical research, it is therefore valuable to show and compare healthy and pathological flow patterns. For instance, Markl et al. presented a study of different aortic conditions, comparing pathological blood-flow to a healthy situation [151]. As an example, anomalous circular flow patterns are observed with an aortic aneurysm and coarctation.

Small alterations within the cardiovascular system considerably influence the hemodynamics. For instance, a sclerosed aortic valve leads to significant changes in the blood-flow patterns [153]. In early systole, the blood-flow is constrained, and follows an isolated flow channel along the outer curvature of the ascending aorta. Helical out-flow patterns may be reversed and amplified, and inner retrograde patterns may occur [153]. Similarly, nested helical flow patterns are observed in patients suffering from a bicuspid aortic valve, in combination with high speed blood-flow jets [91]. Characterizing these jets in an early stage may help to identify the risk for the development of an ascending aortic aneurysm.

Besides pathological blood-flow patterns, also flow behavior around a mechanical prosthesis is of importance. For instance, Kozerke et al. inspected the blood-flow patterns around a bi-leaflet aortic valve prosthesis, assessing its performance [122]. They observe a distinct jet pattern closest to the valve during early systole, together with adjacent retrograde flow. Furthermore, the qualitative analysis is also an important tool to gain understanding of the complex hemodynamics before and after different treatments or interventions. A recent review of both qualitative and quantitative analyses of 4D PC-MRI blood-flow fields was presented by Markl et al. [154].

2.4 Cardiovascular disease - adult conditions

2.4.1 Epidemiology

At present, CVD is the leading cause of death worldwide [267]. In Europe, the overall death rate was 332 per 100000 inhabitants, accounting for about 42% of the deaths in the year 2004. The total cost amounted to €192 billion [3]. In the United States, the overall death rate was 251 per 100000 inhabitants, accounting for 33.6% of the deaths in the year 2006. This amounts to an average of one deceased person every 39 seconds. The total cost involved exceeds any other diagnostic group, with a total of \$286 billion for the United States in the year 2007 [5].

The previous sections discussed the importance of the hemodynamics for the cardiovascular system, and hence for the progression of cardiovascular disease. To gain a thorough understanding, both a quantitative and a qualitative analysis of the blood-flow characteristics is required. The hemodynamics within various cardiovascular regions are clinically relevant. Such regions are for instance, the large thoracic arteries, the caval and pulmonary veins, and the valves [185]. In particular, cases with morphological deformations are of interest.

As depicted in figure 2.3, ischemic heart disease forms the largest diagnostic group of cardiovascular diseases. Ischemia entails a reduced blood supply, and hence blood flow is evidently important. The second largest group comprises cerebrovascular diseases, where blood flow is mostly investigated for cerebral aneurysms. The group of hypertensive diseases comes in many variations, all related to blood pressure. Many variations remain hard to detect, and blood-flow information can potentially provide new indicators. One example related to hypertensive disease will be elaborated in the next subsection. Hemodynamics likewise play a significant role in the development of diseases of the arteries and valves, as described in section 2.3.

For other CVD groups, such as cardiomyopathy and heart attack, the value of blood flow information is not directly apparent. However, it is fair to state that hemodynamics play an important role in the cardiovascular system, and hence provides great potential for the diagnosis and prognosis of CVD.

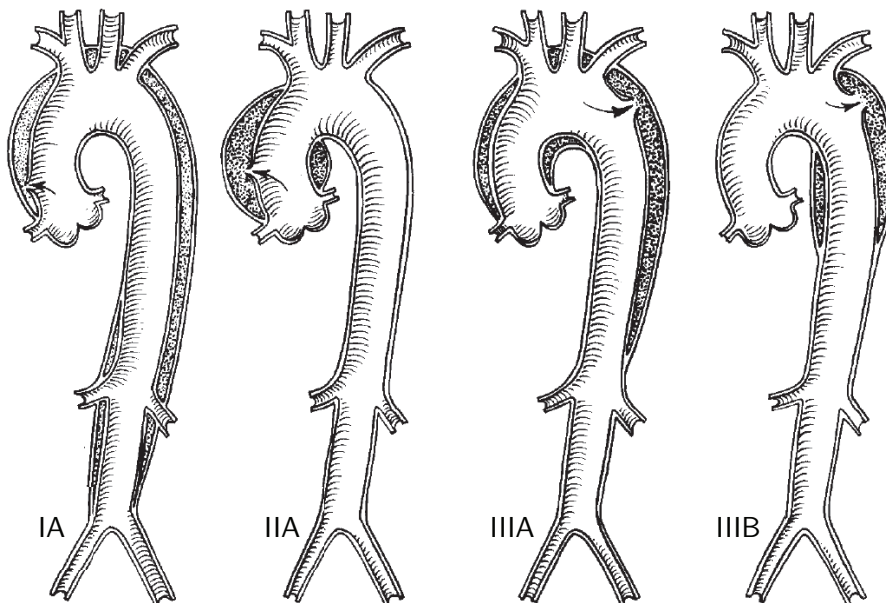


Figure 2.5 Anatomical categorization of an aortic dissection, according to the DeBakey system (I-III) and the Stanford system (A,B). (a) Type IA implies a tear in the ascending aorta, and extends to the descending aorta. (b) Type IIA implies a tear in the ascending aorta, and remain in the ascending aortic region. (c) Type IIIA implies a tear in the descending aorta, and extends to both the ascending and descending aortic regions. (d) Type IIIB implies a tear in the descending aorta, and is limited to the descending aortic region. Figure from the work by Larson et al. [137]

2.4.2 Pathology example

Aortic dissection is a common disorder of the aorta, with an incidence of five to thirty cases per million people per year [37]. A tear in the intimal layer of the aortic wall allows blood to enter the medial layer. This causes the intimal and medial layers to lacerate, forming a dissection flap that separates the original bloodstream from the secondary inter-layer stream, respectively known as the true and false lumina. The dissection commonly develops in the antegrade direction, while some cases extend retrograde from the tear location. Additional tear locations may develop after the inciting event. The initial tear location is often hard to detect, while this is an important indicator for treatment, together with the extent of the dissection.

Classification of an aortic dissection is commonly based on morphological characteristics of the pathology [109, 133]. In detail, the initial tear location and the propagation extent of the false lumen together define the type of dissection. Two systems are used in clinical practice, illustrated by figure 2.5. On the one hand, the *DeBakey* system classifies dissections according to three types. Type I involves an intimal tear in the ascending aorta that extends to the descending aorta, while type II is restricted to the ascending aorta. Type III comprises a tear in the descending aorta. On the other hand, the *Stanford* system classifies the pathology into group A or B, depending on the involvement of the ascending aortic region. Both systems strongly simplify the description by coupling the two main anatomical variables, limiting the value of the classification [133]. Furthermore, an aortic dissection is classified based on the identification time. A dissection detected within two weeks after the first symptoms is classified as an *acute* dissection. Otherwise, the dissection is said to be *chronic*.

The initial tear occurs at the area of greatest hydraulic stress [37], typically within ten centimeters of the aortic valve. While many risk factors are described in literature, still the inciting events that cause the initial laceration are largely unknown. A primary risk factor is systemic hypertension [137]. Disorders of the connective tissue, such as Marfan syndrome, are also shown to enlarge the risk of laceration. Furthermore, aortic valve disease may be associated with a dissection. Additionally, a chest trauma may result in an acute aortic dissection, which is the leading cause of incidence for young persons [60].

The risk of death from an untreated dissection in the first week is approximately 75%, increasing to 90% in the first month [265]. The largest risk is found in a rupture, leading to hemorrhaging into a body cavity. In this case the prognosis depends on the area of rupture. Other complications that often occur are vascular occlusion and retrograde dissection into the ascending aorta [55]. With no complications, treatment typically involves medication for blood pressure control, in particular for a chronic dissection of Stanford type B. Otherwise, a primary surgical repair is required.

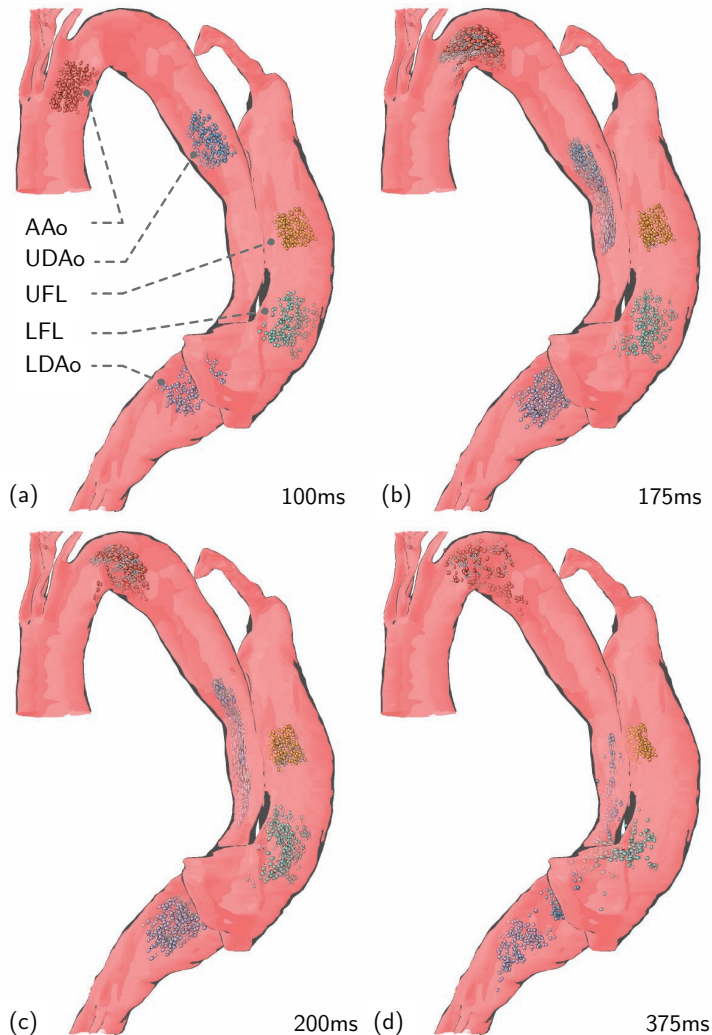


Figure 2.6 Particle trace depicting blood flow in a DeBakey type III, and Stanford type B aortic dissection. (a) Particles are seeded in the ascending aorta (AAo), the upper and lower descending aorta (UDAo, LDAo), and upper and lower false lumen (UFL, LFL). The seeding time is before the systolic peak. (b) Particles within the true lumen accelerate, while the blood-flow speed in the false lumen remains very low. (c) A circular flow pattern occurs in the aortic arch, while the particles seeded in the UDAo show a considerable acceleration. Particles seeded in the UFL region show a minimal retrograde displacement. (d) At the start of the diastolic phase, the circular flow pattern remains in the aortic arch, indicating inefficiency of the recirculation. Blood-flow speed in the false lumen region remains low, and relatively little displacement of the particles is observed.

For a chronic dissection, a surgical repair is typically performed when the enlargement is greater than 6.5 centimeters, following the guidelines for aortic aneurysms [55]. The surgical intervention may comprise a repair of the ascending or descending aorta, sometimes including the aortic arch or hemiarch. In other cases, an aortic root replacement is required [133]. In case of an acute proximal or distal aortic dissection with one or more complications, surgical treatment is also performed. Vascular repair is often performed by endovascular surgery, showing a good aortic remodeling for the acute case [200]. Alternatively, endovascular stent graft treatment may be performed [54]. In case of a vascular occlusion, aortic fenestration is mostly performed, reducing intraluminal pressure by creating a hole in the distal part of the false lumen [84].

In general, policies for acute dissections are aggressive for a Stanford type A, and selective for Stanford type B case. Over the past two decades, early and late survival after surgical repair has improved [218]. While there were less re-operations and follow-ups, intervention remains palliative in many cases [55]. Aortic repair still leads to high mortality rates (28% - 65%) and a high risk of paraplegia or stroke [36, 55].

Given the high mortality rates, and the uncertainties concerning the inciting event, further research is required to gain understanding of this complex pathological condition. Evidently, hemodynamics play a significant role in the pathogenesis, and hence blood-flow measurements can provide valuable insight [4], which may facilitate future treatment. Figure 2.6 shows an example of a qualitative inspection of blood flow in a patient suffering from an aortic dissection. Particles follow the PC-MRI acquired blood-flow field over time. More details concerning probing and visualization will follow later in this thesis.




Figure 2.7 Congenital heart disease prevalence, based on the work by Park [171]. The group of congenital heart diseases is subdivided into smaller groups, indicated by distinct colors. The groups comprise cyanotic, acyanotic, and obstructive diseases, as well as a group of false connections. For the majority of the congenital heart diseases, blood-flow analysis potentially provides valuable diagnostic and prognostic information.

2.5 Cardiovascular disease - congenital conditions

2.5.1 Epidemiology

Congenital heart disease (CHD) comprises a group of conditions that arise from an abnormal maturation of the heart and major blood vessels. CHD is the number one birth condition, with an approximate incidence of about 1 out of 150 live births in the United States in the year 2002 [5]. To put this into perspective, Down Syndrome occurs in about 1 out of 1000 live births. Of the births with congenital heart defects, approximately 9% suffered severe lesions, with a female predominance [149]. These severe anomalies are usually apparent, while minor defects are often asymptomatic and may resolve spontaneously. Hence, a true incidence figure is hard to obtain. Cardiovascular defects are the most common cause of infant death, accounting for more than 24% of deceased children with a birth defect [5]. In total, hospital costs amounted to \$2.6 billion in the United States in the year 2004 [5].

The fetal circulatory system, supported by the placenta, differs from the independent post-natal circulatory system. Within the first ten to fifteen hours, the fetal cardiovascular system changes under the influence of the increased alveolar pressure. Impediment of this process leads to an abnormal development of the heart, for which the cause in most cases remains unknown. Congenital heart defects have been related to chronic health issues of the mother, infection during pregnancy, and chromosomal abnormalities. Furthermore, the cause may be genetic [180] or environmental [67], for instance when the mother needs certain medication, or consumes drugs, such as alcohol or cocaine.

The group of congenital cardiovascular diseases can be divided into subgroups, as indicated by the distinct colors in figure 2.7. The largest group comprises acyanotic conditions. Cyanosis refers to the blue coloration of the skin, due to a lesser oxygen supply. For acyanotic conditions, the oxygenation level of the blood is normal, and hence the survival rate is higher [56]. The third largest group entails the cyanotic conditions, with the tetralogy of Fallot being the most common disease. The second largest group includes all obstructive conditions, such as stenoses. Furthermore, there is a group of false connections.

For most congenital heart conditions, hemodynamics play a significant role. Malformations of the cardiovascular system can be substantial, and hence the interaction of the blood with the surrounding biology affects pathogenesis. In particular for the youngest patients, development of the post-natal cardiovascular system heavily depends on the blood flow. Vascular remodeling is required in this stage. However, anomalous flow patterns may lead to malformations during the development of the anatomy. Colloquially speaking, the rule of thumb for congenital heart defects is: “No (blood) flow, no growth”. Consequently, treatment depends on the understanding of the patient-specific blood-flow dynamics.

2.5.2 Pathology example

To exemplify the hemodynamics for a CHD case, we follow a case study, presented by Valverde et al. [234]. The case report presents the congenital condition of an eight-year old girl, suffering from an atrial septal defect (ASD), and associated anomalous pulmonary venous return.

The asymptomatic ASD allows blood to flow between the left and right atria, and is occasionally associated with partial anomalous pulmonary venous return (PAPVR). In these cases, the pulmonary vein is incorrectly connected to the right atrium, or to veins carrying oxygen-poor blood, such as the vena cava. In the given case study, the right upper pulmonary vein shows a partial anomalous return to the superior vena cava. The PAPVR provides an additional source of left-to-right shunting, providing an increased risk of pulmonary artery hypertension [234].

To analyze the complex hemodynamics, 4D PC-MRI blood-flow measurements were performed. These measurements allow a direct quantification of the shunting, and an internal validation of the condition. The qualitative inspection provided new diagnostic insights, detecting the pulmonary venous return. Subsequently, particle tracing enabled inspection of the patient-specific blood-flow patterns. The anatomy of the case study was reproduced, and depicted in figure 2.8. In addition to the anatomy, also a global overview of the blood-flow is presented by means of arrows. Moreover, the particle tracing was reproduced using our visualization framework. These results are presented in figure 2.9. More details concerning probing and visualization will follow later in this thesis.

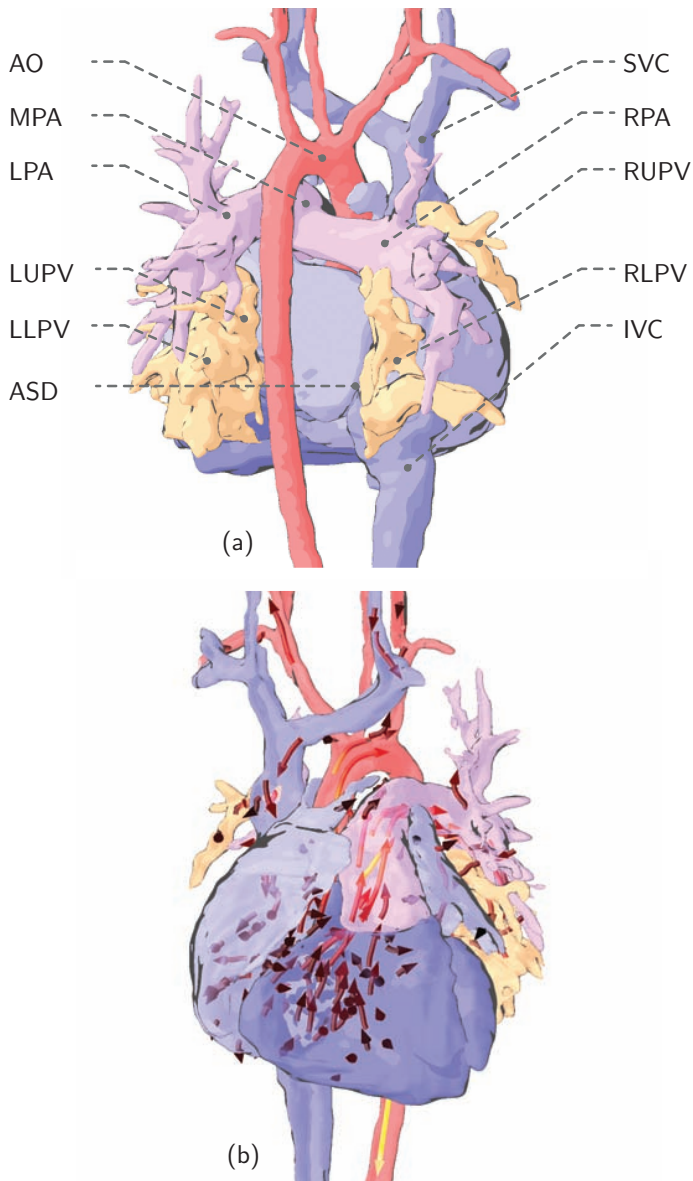


Figure 2.8 (a) Posterior view on anatomy segmentation of a CHD patient. There is a partial anomalous right upper pulmonary vein (*RUPV*) return to the superior vena cava (*SVC*). *AO* aorta; *MPA* main pulmonary artery; *LPA* left pulmonary artery; *LUPV* left upper pulmonary vein; *LLPV* left lower pulmonary vein; *ASD* atrial septal defect; *RLPV* right lower pulmonary vein; *IVC* inferior vena cava. (b) Anterior view of the segmented anatomy, with patharrows that indicate the flow at 375ms of the cardiac cycle. These figures are reproduced, inspired by work of Valverde et al. [234]. The visualization techniques are explained in detail in chapter 5.




Figure 2.9 Particle traces depict blood flow in a patient suffering from ASD with partial anomalous pulmonary venous return. (a) Particles are seeded in the superior vena cava (SVC), the inferior vena cava (IVC), the right upper pulmonary vein (RUPV) and the left lower pulmonary vein (LLPV). (b) Early diastole shows a filling of the ventricles. (c) Particles originating from the LLPV show a left-to-right shunt, and an initial formation of a circular flow pattern. (d) The circular pattern develops further. (e) A three-dimensional view with anatomical context. The orientation corresponds to the orientation of the planes. (f) A three-dimensional left-to-right view with anatomical context. These figures are reproduced, inspired by work of Valverde et al. [234]. The visualization techniques are explained in detail in chapter 7.

"I think that a particle must have a separate reality independent of the measurements. That is an electron has spin, location and so forth even when it is not being measured. I like to think that the moon is there even if I am not looking at it."

Albert Einstein

Data acquisition using
magnetic resonance

3

3.1 Magnetic resonance imaging

3.1.1 Spin physics

Magnetic resonance (MR) applications measure the density of selected atoms by exploiting the magnetic moment of their nuclei. Certain atomic nuclei have a physical property called spin, which have a magnetic moment. One of the most important atoms with this property for biomedical applications is hydrogen. The spin direction is a type of angular momentum, commonly conceived as a rotation of the particle around a specific axis (Fig. 3.1a). When placed in an external magnetic field, these tiny magnets align along the field, either parallel or anti-parallel (Fig. 3.1b left). Their individual behavior can only be described by quantum theory. However, the total magnetic moment of a group of nuclei, experiencing the same magnetic field strength, can be analyzed in a classical fashion. This macroscopic magnetic moment constitutes the net magnetization \vec{M}_0 (Fig. 3.1b right).

When the direction of the net magnetization differs from the main magnetic field \vec{B}_0 , a torque is exerted on the magnetization. This causes the magnetization vector to rotate around the main magnetic field. The resulting rotational motion, called precession, has an angular velocity proportional to the strength of the main magnetic field. The coefficient of proportionality is the gyromagnetic ratio γ , which depends on the nuclear spin, and hence varies for different nuclei. The net magnetization is mostly analyzed in a rotating frame of reference, in which case the transverse component of the net magnetization appears stationary.

Resonance occurs when nuclei are excited by the correct radio frequency (RF), which depends on their precession rate. This causes the magnetization vector to deviate from the main field. Consequently, the longitudinal component \vec{M}_{\parallel} of the magnetization vector will decrease, while the transverse component \vec{M}_{\perp} will increase (Fig. 3.1c). Thereafter, the tilted magnetization vector will spontaneously return to its equilibrium (Fig. 3.1d). This relaxation process exponentially returns the longitudinal component \vec{M}_{\parallel} to its equilibrium, with a time constant T_1 . The transverse component decays with the time constant T_2 , while rotating around the axis of the main magnetic field. This relaxation time is shorter in inhomogeneous magnetic fields, defined by the time constant T_2^* . The rotating transverse magnetization induces an electromagnetic current in a receiver coil, providing a measurable RF-signal.

Imaging techniques non-invasively measure the emitted RF-signal in living tissue. Since the human body primarily consists of water, the signal emitted by the abundance of hydrogen nuclei is typically measured. This signal is proportional to the density of the hydrogen nuclei, which varies per tissue type. Consequently, Magnetic Resonance Imaging (MRI) techniques enable the acquisition of anatomical images, by measuring the intensity of the radiation that is emitted from particular locations [163, 252].

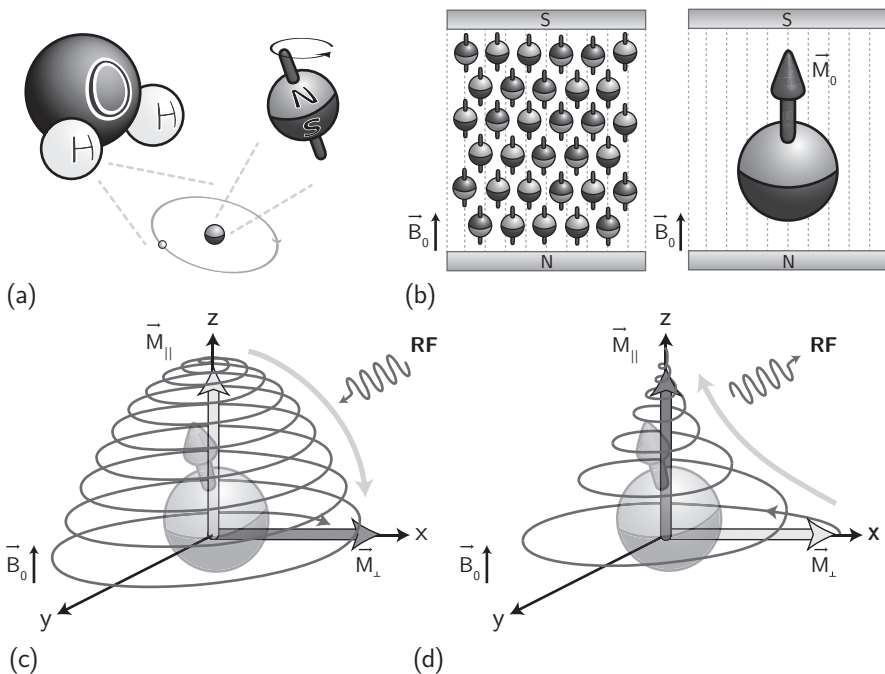


Figure 3.1 (a) Water molecules are abundant in the human body. The hydrogen nuclei, consisting of a single proton, have a physical property called spin. This causes the elements to behave as tiny magnets, with a north (N) and south (S) pole. (b) The hydrogen nuclei align parallel or anti-parallel with the main magnetic field \vec{B}_0 . For clarity, this schematic representation is based on the classical description, excluding quantum mechanical properties [155]. The net magnetization of a spin packet \vec{M}_0 can be analyzed in a classical fashion. (c) Excitation causes the net magnetization to deviate, resulting in a decrease of the longitudinal component $\vec{M}_{||}$, and an increase of the transverse component \vec{M}_{\perp} . (d) Subsequently, the magnetization spontaneously returns to its equilibrium state. During this relaxation, the transverse component can be detected for imaging [252].

3.1.2 Volumetric imaging sequences

Construction of an anatomical image requires localized measurements of the relaxation times and proton density, in order to obtain the desired contrast between soft tissues. The excitation of hydrogen nuclei should therefore occur at designated locations. This is achieved by superimposing the main magnetic field \vec{B}_0 with relatively small spatially varying magnetic fields, called gradient fields.

Spatial encoding is attained by imposing various gradient fields in sequence. First, an imaging slice is selected, comprising an RF excitation in the presence of a linear gradient field. The gradient field causes the nuclei to precess at distinct frequencies at different locations. The RF pulse subsequently controls which nuclei belong to the inspected slice.




Figure 3.2 The spin echo sequence. (a) A 90° pulse rotates the magnetization into the transverse plane. (b) Next, the magnetic moments dephase under the influence of both the phase encoding and frequency encoding gradient. (c) A 180° refocussing pulse inverts the magnetization. (d) This induces rephasing of the magnetic moments. The echo of transverse magnetization is measured when the maximum signal is attained [252].

Next, the positions within the slice require spatial encoding. Therefore, a phase shift is induced by a gradient field, perpendicular to the slice encoding gradient. This gradient field, called the phase-encoding gradient, influences the main field for a limited time span. While the gradient is applied, rows of spins will precess at slightly different frequencies, creating a phase offset. When the gradient is released, all spins again experience the same magnetic field, causing the spins to precess at identical frequencies, while the phase difference remains. This process is repeated for each column, at varying gradient field strengths.

Finally, a frequency-encoding gradient is employed, orthogonal to the slice-selection and phase-encoding directions. The signal is collected in the presence of this gradient, and contains the position-dependent contributions of all rows, encoded at distinct frequencies.

Series of RF and gradient pulses are commonly termed pulse sequences [13]. Repetitive execution of a pulse sequence implies a sampling of the domain of spatial frequencies, called k -space. From this space, the desired image data can be reconstructed by Fourier analysis, providing localized information, proportional to the proton density. A wide range of pulse sequences exists, emphasizing on either tissue contrast or acquisition speed, while limiting acquisition artifacts.

A fundamental group of pulse sequences relies on the spin echo or Hahn echo [80]. The sequence starts with a 90° pulse, rotating the longitudinal magnetization into the

transverse plane. Thereby, precession occurs under the influence of both the phase-encoding and frequency-encoding gradient fields, causing the magnetic moments of the different protons to dephase. A 180° refocusing pulse subsequently inverts the magnetization. The frequency-encoding gradient is again enabled, causing the precession to continue, and hence reversing the process of phase dispersal. After a time equal to the time elapsed between the two rotation pulses, the maximum signal is attained, which can now be detected by a receiver coil. The spin echo sequence is depicted in figure 3.2.

3.1.3 Blood flow imaging

The previous subsection explained the general concept of imaging by magnetic resonance. Throughout the description, the measured tissues were implicitly assumed to be stationary. However, motion of tissue occurs due to patient movement, and movement of anatomical structures. This leads to inconsistencies in the transverse magnetization, causing measurement artifacts, such as blurring and ghosting.

Besides patient movement and tissue displacement, also the flow of bodily fluids gives rise to motion artifacts. This is usually experienced as an adverse influence on the results. However, the inconsistencies in magnetization caused by fluid motion can be exploited to measure flow characteristics [185]. For instance, the speed of fluid displacement can be measured. The speed of bodily fluids varies by several orders of magnitude (Fig. 3.3). In particular, we focus on imaging of flowing blood, even though movement of cerebrospinal fluid (CSF) can be detected similarly.

Spins in motion experience the pulse sequence at various locations over time. The area of flow imaging sensitizes the pulse sequence for the spin movement. The area can be divided into two coherent groups of techniques. On the one hand, modulus contrast techniques utilize the magnitude difference between the transverse magnetization of moving spins and stationary spins. Inflow imaging is the most elementary modulus contrast technique, enhancing blood-flow regions based on the time-of-flight (TOF) effect. This effect is caused by spins without an RF history flowing into the imaging




Figure 3.3 Speed of fluid flow in the human body. Six orders of magnitude can be measured by various magnetic resonance flow imaging techniques. Figure is based on Potchen et al. [185].

slice. These spins carry full longitudinal magnetization into the slice, resulting in 'bright-blood' images. In contrast, black-blood techniques rely on spins flowing out of the imaging slice. After the slice is excited according to a spin-echo sequence, blood will move out of the slice, resulting in dark flow regions. An alternative modulus contrast technique relies on administered contrast media. Typically, a Gadolinium-based contrast agent is injected intravenously, enhancing flow regions by shortening the T_1 relaxation time.

On the other hand, phase-contrast techniques quantitatively measure flow, employing the phase difference between the transverse magnetization of moving spins and stationary spins. These techniques rely on subtraction methods, involving flow-insensitive data and flow-sensitized data. Intensities in the resulting images, representing the flow velocity, are proportional to the phase, enabling a quantitative analysis of the velocities. In this thesis, we focus on blood-flow data acquired by phase-contrast methods, which will be elaborated in section 3.2.

The qualitative inspection of the acquired flow data is mainly performed through projective views. This type of visual inspection is called MR angiography (MRA), dealing with high-contrast renditions that resemble x-ray angiograms. Modulus-contrast techniques are mostly employed for MRA, while phase-contrast techniques also suffice. For quantitative inspections, however, phase-contrast techniques prevail, while modulus-contrast techniques provide limited possibilities.

3.2 Phase-contrast acquisition

3.2.1 Velocity encoding

As described in the previous subsection, MRI is intrinsically sensitive to flow. Phase-contrast techniques exploit this sensitivity, measuring flow velocities by means of specific pulse sequences. Similar to spatial encoding, the velocity is encoded by means of sequences of well-chosen gradient fields, superimposed on the main magnetic field.

In order to derive a velocity-encoding pulse sequence, it is essential to define the position of hydrogen nuclei in time, since flow effectuates movement of nuclei. The position \mathbf{x} of these nuclei can be described by a Taylor series, expanded around the initial time $t = t_0$:

$$\begin{aligned}\mathbf{x}(t) &= \mathbf{x}(t_0) + \vec{v}(t_0) \cdot (t - t_0) + \frac{1}{2}\vec{a}(t_0) \cdot (t - t_0)^2 + \dots \\ &= \mathbf{x}_0 + \vec{v}_0 \cdot (t - t_0) + \frac{1}{2}\vec{a}_0 \cdot (t - t_0)^2 + \dots,\end{aligned}$$

where \mathbf{x}_0 is the initial position, \vec{v}_0 the initial velocity, and \vec{a}_0 the initial acceleration.

Higher order terms can be omitted, assuming constant acceleration in the measured time span [121]. Although acceleration measurements are possible, it is common to include only the first order term, assuming constant velocity throughout the measured time span. This assumption is reasonable when motion changes are small in relation to the temporal resolution.

Phase-contrast velocity encoding relies on the relation between the displacement of hydrogen nuclei and the phase shift of the transverse magnetization. Under influence of the main magnetic field, neither stationary, nor moving spins obtain a phase shift (Fig. 3.4a).

First, a linear gradient field \vec{G} is imposed on the main magnetic field \vec{B}_0 , which is inevitably subject to field inhomogeneities $\delta\vec{B}$. Along the gradient direction, nuclei will precess at distinct frequencies ω_0 , inflicting an incremental phase change. Next, the gradient field is released, causing the nuclei to precess at the same frequency again, while memorizing the phase shift. Hence, the phase shift ϕ at location \mathbf{x} is defined in the co-rotating frame at the time of measurement t_m as:

$$\phi(\mathbf{x}, t_m) = \int_{t_0}^{t_m} \omega_0 dt = \gamma \int_{t_0}^{t_m} \delta\vec{B}(t) + \vec{G}(t)\mathbf{x}(t) dt,$$

Nuclei in motion experience the gradient field at varying locations (Fig. 3.4b). Previously, we discussed that the position of moving nuclei can be described by a first order approximation of the Taylor series. We assume that the nuclei move with a nearly constant velocity \vec{v} in the direction of the gradient throughout the measurement. Without loss of generality, this results in the following phase shift for $t_0=0$:

$$\phi(\mathbf{x}, t_m) = \gamma \int_0^{t_m} \delta\vec{B}(t) dt + \gamma \int_0^{t_m} \vec{G}(t) \cdot (\mathbf{x}_0 + \vec{v}_0 t) dt.$$

Rewriting this equation shows the three terms that compose the total phase shift. In addition to the background phase ϕ_0 , the phase shift for both stationary spins and moving spins is proportional to the amplitude and duration of the gradient field.

$$\phi(\mathbf{x}, t_m) = \phi_0(\mathbf{x}, t_m) + \gamma \mathbf{x}_0 \int_0^{t_m} \vec{G}(t) dt + \gamma \vec{v}_0 \int_0^{t_m} \vec{G}(t) t dt.$$

This relation between motion and phase of the magnetization enables imaging of the flow. Therefore, susceptibility to spin displacements should be incorporated into conventional pulse sequences. Figure 3.4b shows that a single gradient lobe induces a phase shift on the spins in motion. However, also nuclei in stationary tissues acquire




Figure 3.4 Velocity encoding. (a) Neither stationary, nor moving spins, obtain a phase change under influence of the main magnetic field. (b) When a linear gradient field is enabled, stationary spins obtain a location dependent phase offset, while moving spins experience different parts of the gradient. (c) Inverting the gradient polarity ensures a zero phase shift for stationary spins, while the phase shift of moving spins remains linearly related to their velocity. Figure based on Lotz et al. [146].

a position dependent phase change. This offset can be reversed by a second gradient lobe with opposite amplitude $-\vec{G}$. This lobe is initiated at time T , with equal duration Δt , and the measurement time $t_m > T + \Delta t$. The total phase shift induced by the bipolar gradient waveform becomes:

$$\begin{aligned} \phi(\mathbf{x}, t_m) = & 2\phi_0(\mathbf{x}, t_m) + 2\gamma\mathbf{x}_0 \left(\int_0^{\Delta t} \vec{G}(t) dt + \int_T^{T+\Delta t} -\vec{G}(t) dt \right) \\ & + 2\gamma\vec{v}_0 \left(\int_0^{\Delta t} \vec{G}(t)t dt + \int_T^{T+\Delta t} -\vec{G}(t)t dt \right). \end{aligned}$$

Assuming a constant strength G of the linear gradient field over time, the equation can be rewritten through integral calculus as:

$$\begin{aligned} \phi(\mathbf{x}, t_m) = & 2\phi_0(\mathbf{x}, t_m) + 2\gamma\mathbf{x}_0 (G\Delta t - (G \cdot (T + \Delta t) - GT)) \\ & + 2\gamma\vec{v}_0 \left(\frac{1}{2}G\Delta t^2 - \left(\frac{1}{2}G \cdot (T + \Delta t)^2 - \frac{1}{2}GT^2 \right) \right), \end{aligned}$$

Consequently, the bipolar gradient suppresses the influence of the static spins, reducing the equation to:

$$\phi(\mathbf{x}, t_m) = 2\phi_0(\mathbf{x}, t_m) + \gamma\vec{v}_0\vec{M}_1,$$

where \vec{M}_1 denotes the first moment of the gradient time curve, defined as $\vec{M}_1 = G\Delta tT$. While the phase shift of stationary spins is nullified, the moving spins remain linearly related with the velocity and the first gradient moment (Fig. 3.4c).

For flow imaging, the bipolar gradient is included into pulse sequences as a flow sensitizing gradient waveform, which is depicted at the top of figure 3.4. For simplicity, the lobes of the gradient waveform are depicted adjacently, while they may well be disjoint.

Although the bipolar gradient reverses the phase shift of stationary spins, it cannot compensate for all background phase effects ϕ_0 , caused by inhomogeneities in the main magnetic field. These effects are reduced by subtraction of two consecutive flow measurements with inverse bipolar gradients, resulting in the final flow image.

3.2.2 Image formation

To reduce the influence of field inhomogeneities, two measurements with inverted polarities of the bipolar gradient waveform are performed along the same axis. Two distinct subtraction methods are commonly applied, based on the complex signal obtained from the Fourier image formation [12].

On the one hand, the phase-difference method subtracts the phase angles of the two individual measurements, defined by the arctangent of the complex signal (Fig. 3.5a). This is analogous to a pixel-wise subtraction of the two phase images. The resulting phase-difference images are often referred to as the phase images of phase-contrast angiography (PCA-P). They provide quantitative velocity information, describing both speed and direction. However, the phase-difference images contain uncorrelated subtraction noise, and are susceptible to phase wrapping artifacts, which will be described in subsection 3.2.5. The resulting phase-difference $\Delta\phi$ is devoid of background phase influences, and is defined as:

$$\begin{aligned}\Delta\phi &= \left(2\phi_0(\mathbf{x}, t_m) + \gamma\vec{v}_0\vec{M}_1^{(1)}\right) - \left(2\phi_0(\mathbf{x}, t_m) - \gamma\vec{v}_0\vec{M}_1^{(2)}\right) \\ &= \gamma\vec{v}_0\Delta\vec{M}_1.\end{aligned}$$

On the other hand, the complex-difference method subtracts the vectors of the acquired complex signals (Fig. 3.5a). These images are often referred to as the magnitude image of the phase-contrast angiography (PCA-M). The direction information is lost, and the intensities only indirectly relate to the flow speed. As a consequence, a quantitative analysis of the images is no longer possible. However, the complex difference method provides a high signal-to-noise ratio within the vessels [12].




Figure 3.5 (a) Field inhomogeneities are suppressed by subtraction of two individual measurements. Either the phase-difference, or the complex-difference is employed to form the images. (b) The velocities are measured by the phase of the transverse magnetization, with a scope of 180° . The minimum and maximum velocity is set by encoding speed v_{enc} .

Measurements of the transverse magnetization are very sensitive, and hence small changes can be measured. The velocity is measured according to phase shifts with a maximum scope of 180° , which compromises the measurement precision. The minimum and maximum flow speed that can be acquired without artifacts, is determined by the velocity encoding speed v_{enc} as:

$$v_{enc} = \frac{\pi}{\gamma \Delta M_1}$$

The velocity encoding speed needs to be configured prior to the acquisition of the flow data (Fig. 3.5b), and is typically defined in centimeters per second. The range of the imposed speed limit, for example $[-100 \text{ cm/s}, 100 \text{ cm/s}]$, corresponds to the phase extremities at respectively $-\pi$ and π radians.

Thus far, we have addressed flow imaging by velocity encoding of the fluid motion. However, since the velocity of moving nuclei is measured, this approach is not restricted to the flow of fluids. By selecting an appropriate velocity encoding speed, also the motion of tissue, such as the heart muscle, can be acquired quantitatively.




Figure 3.6 (a) Quantitative through-plane blood-flow velocities need to be acquired perpendicular to the flow direction. In this example, the flow is measured orthogonal to the ascending aorta. (b) Anatomical data is reconstructed, (c) in combination with phase-difference, and (d) complex-difference blood-flow information. A typical slice contains 256^2 pixels of about 1 mm^2 , over 25 phases of the cardiac cycle, with a temporal resolution of approximately 40 milliseconds.

3.2.3 Data acquisition

The previously described imaging techniques acquired flow velocities along a single axis. These unidirectional acquisition techniques are employed in clinical practice to obtain flow information on oblique slices. These 2D quantitative flow (qflow) measurements typically provide through-plane velocities. The acquisition of these velocities requires that the slices are orthogonal to the vessel cross-section. Prior to acquisition, the MRI technician positions the geometric basis of the slice to be perpendicular to a vascular structure (Fig. 3.6a). Anatomical information is reconstructed from the acquired complex signal, as depicted in figure 3.6b.

The through-plane blood-flow information is conventionally reconstructed by a phase-difference or complex-difference method, respectively depicted in figures 3.6c and 3.6d. Generally, the data is acquired throughout the cardiac cycle. In addition to through-plane velocities, also three-directional velocity vectors can be acquired on a slice. This requires repeated measurements in multiple directions.

In addition, volumes of three-directional blood-flow velocities can be acquired. Volumetric flow imaging requires measurements along three axes. Compensation for field inhomogeneities imposes a twofold measurement in each direction, leading to six separate measurements [175]. This is called a six-point acquisition method. The three-directional velocity vectors can subsequently be reconstructed from the individual Cartesian components. The six-point technique is a straight-forward and intuitive extension of the unidirectional flow imaging approach. However, redundant information is acquired, at the expense of valuable scan time.

To reduce the number of measurements, and hence to gain imaging time, a single reference measurement can be used for all three directions [175]. This reference measurement plus three orthogonal directions leads to a faster orthogonal four-point acquisition technique. In contrast to the six-point methods, the errors in the velocity component are correlated, and the noise gets an undesirable direction dependence.

Therefore, more advanced encoding schemes were devised. For instance, Pelc et al. [175] proposed a balanced four-point scan, sampling the Cartesian space on a tetrahedron. As a result, each measurement contributes to all velocity components. This encoding is often referred to as a Hadamard flow encoding, as the encoding table is reminiscent of the Hadamard matrix. Although the Hadamard transform has been employed for multi-slice imaging [214], there is no direct relation to flow imaging.

Recently, Johnson and Markl [104] proposed a five-point acquisition, adding a flow-compensated measurement to the balanced four-point technique. Their approach increases the signal-to-noise ratio with respect to the four-point balanced acquisition, at the expense of the temporal resolution.

3.2.4 Data characteristics

A typical 2D qflow slice comprises 256^2 pixels, with an isotropic pixel size of nearly 1 mm^2 . Each pixel location comprises either the through-plane blood-flow speed, or a three-directional velocity. The data is acquired throughout the cardiac cycle, and is approximately divided into 25 phases, with a temporal resolution of about 40 ms. The typical acquisition time is on the order of tens of seconds. This excludes the time required to position the slice perpendicular to the vessel structure. Hence, the total scan time can exceed the time needed for a volumetric flow acquisition [234].

A volumetric flow data set usually comprises $128 \times 128 \times 50$ voxels. The voxel size is in most cases slightly anisotropic at approximately $2.0 \times 2.0 \times 2.5 \text{ mm}$. Similar to planar flow imaging, anatomical information (Fig. 3.7a), as well as blood-flow information, is reconstructed. For the quantitative phase-difference data (Fig. 3.7b-3.7d), the acquired directions are combined per voxel location, creating a vector-valued velocity field. The complex-difference (Fig. 3.7e-3.7g) directions are typically averaged to form an approximate scalar-valued speed volume, suitable for angiographic projection. The temporal resolution is on the order of 50 ms, representing about 20 to 25 phases throughout the cardiac cycle. The acquisition time amounts to 15 to 20 minutes.

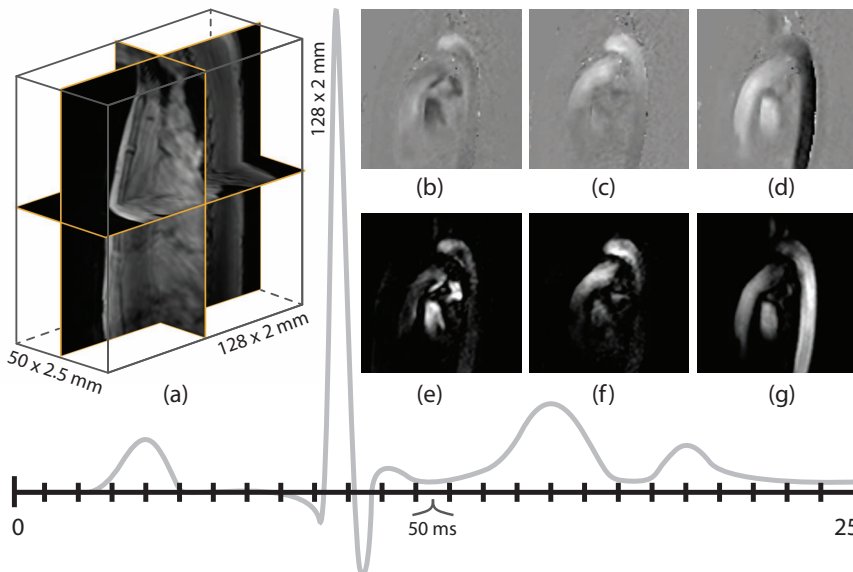


Figure 3.7 A typical volumetric flow data set consists of $128 \times 128 \times 50$ voxels, with an approximate size of $2.0 \times 2.0 \times 2.5 \text{ mm}$. (a) In addition to anatomical information, blood-flow velocities are acquired along three axes. The quantitative data, called PCA-P data, are reconstructed using a phase difference. (b) The x axis represents the blood flow in the right-left direction of the patient, (c) the y axis captures the anterior-posterior direction, (d) and the z axis the feet-head direction. Similarly, the complex difference, or PCA-M data, can be reconstructed in (e) right-left, (f) anterior-posterior, and (g) feet-head direction. The temporal resolution is approximately 50 ms, leading to 20 to 25 phases throughout the cardiac cycle, which is illustrated by an electrocardiogram (ECG).

Clarification of ambiguous terms

In the field of flow imaging, several terms are used ambiguously, and their intent should be interpreted from the contextual information. In particular, different meanings are assigned to the term '**phase**'. In phase-contrast flow imaging, the term is mostly used to describe the phase of the transverse component of the magnetization. Alternatively, the term phase is used to describe a time moment during the cardiac cycle. Again, this can be interpreted in two ways. On the one hand, a phase of the cardiac cycle can refer to a period that demarcates a functional part of the cardiovascular cycle, i.e., systole or diastole. This is also often referred to as a time interval. On the other hand, a phase of the cardiovascular cycle can refer to the equidistantly acquired time moments of the flow imaging sequence (Fig. 3.7).

The term '**magnitude**' is used frequently. Although its use is not ambiguous, the different applications can be confusing. For magnetic resonance imaging in general, the term magnitude typically refers to the length of the acquired complex signal, which is used to construct the anatomical images. However, for flow imaging the term also refers to magnitude images, reconstructed by a complex-difference approach. The magnitude in that case indicates the length of the vector difference between the complex signals, obtained from two measurements with opposite flow sensitizing gradient waveforms (Fig. 3.5a). Besides, the term magnitude can also refer to the length of a measured velocity vector, which equals the speed of the blood flow.

In clinical practice, imaging protocols combine flow acquisition sequences with additional information. Often, separate anatomical scans are acquired, for instance by means of a steady-state free precession (SSFP) sequence. These sequences provide higher tissue contrast, leading to better insight into the morphology. Additionally, other blood flow related information, such as bright-blood, black-blood, and contrast-enhanced flow images are often part of the imaging protocols (subsection 3.1.3). Together, these sequences provide a vast amount of valuable information, enabling insight into both morphology and function of the cardiovascular system.

3.2.5 Acquisition accuracy and artifacts

As described in the previous subsections, phase-contrast flow imaging enables non-invasive and quantitative measurement of blood-flow velocities. However, measurements come with imperfections, complicating a quantitative analysis. In particular, derived flow measures are sensitive to relatively small errors in the velocity measurement. Inaccuracies are caused by a combination of many factors, associated with the MRI hardware, imaging sequences and their parametrization, post-processing software, and patient movement. For most cardiac applications, the generally accepted objective is to acquire flow data with less than 10% error [72].




Figure 3.8 Artifacts in systolic aortic flow. (a) A bias field manifests itself as a slowly varying linear image gradient. The top row shows anatomy and conventional through-plane blood flow. In the bottom row, contrast is increased to convey the linear field gradient, which is even more apparent on a gelatin phantom [72]. (b) Flow data is affected by uncorrelated subtraction noise, (c) and phase-wraps when the actual speed transcends the encoding velocity speed v_{enc} .

The accuracy of flow measurements is largely influenced by the parametrization of the imaging sequence. With the abundance of parameters, it is hard to decide on an appropriate configuration, and hence the quality of the data depends on the expertise of the operator. In particular, the parametrization directly influences the spatial and temporal resolution, strongly affecting the measurement accuracy [77]. Especially the quantitative analysis of small vessels becomes cumbersome with a low resolution, and the associated partial volume effects [7]. For planar flow imaging, the operator also influences the accuracy by positioning the imaging plane, which should be orthogonal to the flow direction. In clinical practice, the tolerable error to estimate the blood flow is maximally 15° [146].

Besides user parametrization, various types of motion cause imaging artifacts. In subsection 3.1.3, we briefly touched upon motion induced artifacts. Besides movement of the patient, causes for tissue displacement are threefold. First, motion artifacts are induced by peristaltic motion. This movement can currently not be foreseen or monitored, and hence the acquisition cannot countervail this type of motion.

Secondly, contraction of the heart muscle leads to motion artifacts. These artifacts can be reduced by compensating the acquisition for cardiac motion. Therefore, a real-time electrocardiogram (ECG) is employed to trigger the measuring sequence. Alternatively, data can be acquired continuously during multiple cardiac cycles, before reordering into time slots. This is called retrospective gating.

In the third place, motion artifacts are caused by respiration. While short scans are often performed under breath hold, longer imaging procedures are affected by the respiratory motion. Breathing gives rise to movement of most organs, leading to severe blurring and ghosting. The impact of the respiratory motion can be largely suppressed by exploiting the relatively motionless period after exhalation. Therefore,

the movement is measured using either a pressure sensor, or a navigator echo. A navigator echo entails a pre-pulse sequence, imaging a narrow area perpendicular to a moving structure; typically the diaphragm. Detection of the high-contrast tissue interface reveals the respiratory motion.

In addition, flow measurements are subject to general MRI artifacts, largely due to hardware imperfections. A so-called bias field leads to background errors in the acquired images, caused by factors such as poor RF coil design, local flip angle variations and subject scanner interactions. Additionally, the superimposed gradient fields yield imaging inaccuracies. For instance, linear gradient fields contain higher-order terms, called concomitant gradients [14], resulting in phase errors. This effect is calculable, and can be corrected during reconstruction. Moreover, fast gradient switching induces eddy currents in the electromagnetic field, causing deviations in the gradient field. This causes background phase errors in the image, which manifest themselves as spatially and temporally slowly varying image gradients (Fig. 3.8a). These effects are difficult to predict, and hence challenging to correct [72, 196].

Furthermore, MRI is susceptible to noise. The amount of noise depends on various factors, including the strength of the main magnetic field as well as the thermal energy emitted by the human body. The conventional imaging noise follows a Rician distribution [78]. For imaging purposes, the tissue contrast is of importance, and hence the main interest lies in the relation between the signal and the noise. Typically, this is estimated through a signal-to-noise (SNR) metric. For flow imaging, it can be shown that the noise in flow regions depends on the velocity encoding speed, and is inversely proportional to the SNR of the corresponding magnitude image [145]. Hence, v_{enc} should be chosen as small as possible, while capturing the full dynamic range of the actual flow.

Moreover, flow data is affected by subtraction noise (Fig. 3.8b). The uncorrelated noise in the two measurements are subtracted, as described in subsection 3.2.2, which yields a random phase. This noise can be largely reduced by filtering with the absolute value of the difference signal.

Some artifacts occur only in flow data. For instance aliasing, or phase wrapping (Fig. 3.8c), erroneously introduces regions with opposite flow directions. As depicted in figure 3.5b, the phase of the transverse magnetization encodes the velocity with a scope of 180° . The velocity encoding speed v_{enc} determines the maximum and minimum speed that can be measured within that scope. Whenever the actual blood-flow speed transcends the velocity encoding speed, a phase wrap occurs. Several methods have been devised to correct these artifacts caused by a single phase wrap through postprocessing [51, 130, 132, 198, 269]. Resolving multiple phase wraps remains a challenging topic. For flow imaging, phase unwrapping is typically performed per acquired direction, which is valid for the 6-point method and the basic 4-point method. However, component-wise unwrapping of balanced volumetric flow encoding schemes is conceptually incorrect.

Another artifact found in flow imaging is called misregistration, causing blood-flow regions to be shifted with respect to the stationary tissue. This is due to the time between the phase encoding and frequency encoding gradients, which together determine a point in k -space [138]. Moving spins change location between the phase encoding and frequency encoding gradients, causing a misregistration of the blood-flow regions in the reconstructed image. These artifacts can be corrected by adding a bipolar gradient to each phase encoding gradient [252].

The flow imaging sequences are based on the assumption that the blood-flow velocities are constant at the time of measurement. Hence, the measurement provides proper results for laminar blood flow. Measurement of accelerated flows is less accurate and can cause undesirable artifacts. Accelerated flows can be found in pulsatile flows, for which flow imaging provides accurate results [77]. In addition, accelerated flows can be found in stenotic flows, or jets, which often induce phase wraps due to the high blood-flow speed. Moreover, accelerated flows are found in turbulent flow regions. Both jets and turbulent flows can cause so-called flow voids at relatively low spatial resolutions. A large velocity spread within a voxel leads to dephasing of the transverse magnetization, averaging out the emitted signal. To verify whether this behavior has occurred, black-blood scans are often employed to inspect the vessel delineation.

The time to acquire the time-resolved volumetric blood-flow data is relatively long, and hence physiological variations, such as respiration, blood pressure, and heart rate, have a significant impact. To a certain degree, the acquisition can be compensated for these variations through gating.

"The most enduring innovations marry art and science."

Steve Jobs

Technical background

4

This chapter is partly based on:

"Illustrative volume visualization using GPU-based particle systems" Roy van Pelt, Anna Vilanova, and Huub van de Wetering. In *IEEE Transactions on Visualization and Computer Graphics* (2010), 16(4):571-582.

4.1 Introduction

Contemporary humankind is overwhelmed with a vast amount of data, originating from various sources, and encompassing diverse characteristics. Visualization of these data facilitates interpretation, and hence enables understanding of the contained information. The term visualization, however, is ambiguous, and may refer to a mental process, a specific technique, the research field, or to the visual result [246]. Different disciplines will have a diverse philosophy and definition of the true significance of visualization. To set forth the notion of visualization within this thesis, we start at the most elementary level, providing the dictionary definition of the term ‘visualization’.

Vi-su-al-i-za-tion *noun*

\,vi-zhə-wə-lə-ˈzā-shən, \,vi-zhə-lə-, \,vɪzh-wə-lə-\

1. formation of mental visual images
2. the act or process of interpreting in visual terms or of putting into visible form

Source: Merriam-Webster

Visualization as a *process* indeed alludes to the formation of mental visual images. However, the visualization process in the context of this thesis is more closely related to the second interpretation. The contemplated notion of visualization comprises interpretation of large and often complex data, and abstraction of the results into a visual representation [33]. The visualization process is mostly considered from a technological perspective, dealing with a pipeline of algorithms, techniques and tools to achieve the goals. Metaphorically, visualization can be regarded as the optician of science, enabling disclosure of information that was previously unseen, or challenging to detect. Meanwhile, the aim is to convert non-visual data into effective images that are readable and recognizable [120]. Hence, visualization needs to surpass the opticians task, and take into account the human interpretation of visual information.

A major challenge is imposed by the visual channel that transfers the information from a data source to our brain. As the visualization process is generally effectuated by computer graphics, the monitor largely defines this visual channel, typically limiting the representation to a two-dimensional plane. Our brain is expected to mentally reconstruct the data dimensions, and interpret the provided information. Fortunately, the human visual system provides unique capabilities to detect features and patterns in a short time, and the human brain can rapidly interpret the provided quantity of information [246]. This process can be strengthened by exploiting knowledge of the visual system, and the way humans perceive their natural environment. Therefore, visualization builds upon a set of elementary visual cues, such as color, opacity and texture. Furthermore, higher level concepts such as shape, light, and context play a significant role in visual communication. Moreover, the notion of human convention may be employed by including customary symbols [254].

Visualization as a *research field* is maturing. Small innovations become notable, even if they might be considered incremental by non-experts [246]. The level of specialization in the field triggers debate about a possible theory of visualization [247]. Given the wide range of experimental results, theory could provide generalization, in accordance with perception theory. However, visualization is primarily a technical field that successfully yields solutions for a wide range of problems. Perhaps the field should strive for a consensus on design principles [1], instead of a comprehensive theory of visualization. The community could keep advancing effective and efficient approaches that are accepted as good practice [246]. An example is the visual information seeking mantra, defining a standardized process that starts with an initial overview, followed by zooming and filtering steps, prior to inspection of details [208].

Visualization *techniques* have become more effective with respect to visual perception, and amplifying application-specific cognition. To evaluate the effectiveness, current techniques are often validated with domain experts. Integration of the visualization process with the user's tasks is essential to answer the key questions at hand. This area currently draws increased attention within the visualization community, bridging the gap between the visualization expert and the target user group.

Additionally, visualization *techniques* have become more efficient. On the one hand, increased computational performance has led to a significant improvement concerning interactivity. User participation enables direct communication with the visualization. Interaction techniques, such as exploration and probing, enable the users to perform their tasks intuitively. In contrast to tedious parametrization, the interaction should be immediate and engaging. On the other hand, computational performance is required to meet the growing demand to visualize extremely large data sets. This demand arises from the increasing size of high dimensional and multivariate data [27], in which case inspection of the entire data becomes cumbersome, if not impossible.

Within the visualization *research field*, the prevailing taxonomy is based on data characteristics. The discipline of "scientific visualization" is concerned with physics-based data with an inherent spatial component, typically related to scientific applications. Instead, the area of "information visualization" generally deals with abstract, non-physical and non-spatial data [227]. While this distinction on a global level is convenient, the separation is not apparent for all situations. To classify the visualization field as a whole, other aspects play a role; e.g., the object of study, the user model, and the design model [227]. For instance, the rapidly emerging field of "visual analytics" is not categorized based on the involved data. Instead, the taxonomy is based on the types of distinct analytical reasonings about large abstract data. The boundaries between the disciplines are subject to debate, as the characteristics of the data partially coincide with the data used for "information visualization".

The visualization discipline is concerned with a multitude of application areas. For instance, visualization techniques may be applied for architectural purposes, terrain research, traffic control, weather prediction models, music exploration, photographic applications, and many more. In this thesis, the focus is on applications in the medical

domain. In chapter 2, we have elaborated on the variety of clinical imaging modalities that measure anatomical and functional information, providing a large amount of patient-specific data (see subsection 2.2.2).

Since the advent of radiographic applications, visualization has been a key component of the medical workflow. At the outset, x-ray radiography was the prevailing modality, and medical visualization entailed a mere light-box projection. With the arrival of tomographic imaging, and associated reconstruction methodologies, volumetric data became omnipresent in clinical practice. Consequently, medical visualization attained a more important role in the clinical routine. Elementary slice-based visualization techniques, such as multi-planar reformatting, are still commonplace. In addition, a wide range of projective visualization techniques have been developed, such as the ubiquitous maximum intensity projection (MIP). Moreover, direct volume rendering (DVR) has become a notable visualization technique for analyzing three-dimensional data, obtaining photo-realistic depictions of the anatomical data.

Furthermore, both planar and volumetric data have been used for image-guided interventions. The acquired data may be used to preoperatively plan the surgery, or may be registered in real-time to support the intervention. In addition, data sets from different imaging devices have been combined into multi-modal visualizations, exploiting the most valuable characteristics from the various modalities. In this way, anatomical and functional information can be combined into one comprehensive visualization.

From there on, the developments in medical visualization have not come to a standstill. Earlier in this subsection, we have stated the criterion that visualization techniques should effectively convey information within the data. Photo-realism may not

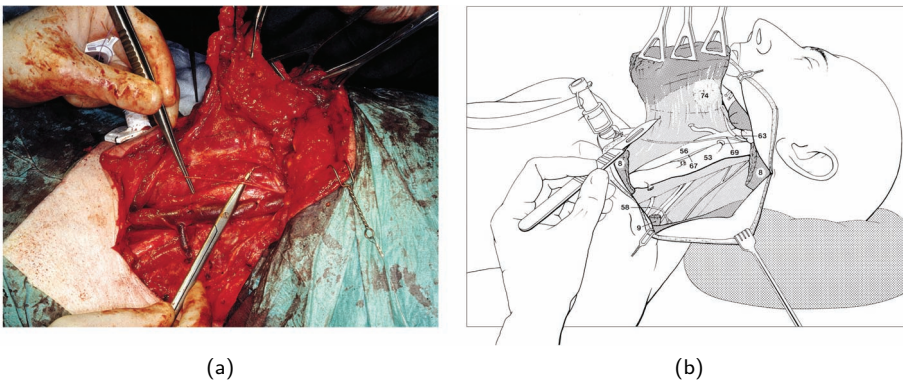


Figure 4.1 (a) A photograph of a medical intervention in the human neck area. (b) The same procedure is drawn up by a skilled medical illustrator. Superfluous information is omitted, and contextual information is introduced. The artistic freedom allows to convey the essence of the intervention, providing a much better insight into the procedure, which is important for educational purposes. Photo and illustration courtesy of R. Trompert - medical illustrator (<http://www.medical-art.nl/>).

in all circumstances provide the most effective solution. Instead, non-photorealistic models are adopted, mimicking artistic styles and techniques, inspired by traditional illustrations. Such illustrations emphasize important features, and omit superfluous information. Especially in the medical domain, valuable information is often communicated through illustrations, availing oneself of the effectual clarity of this medium. This is underlined for example by textbooks about human physiology, which mostly portray the information using illustrations, instead of photographs. Also, visual communication of medical procedures often relies on illustrations. An example is presented in figure 4.1. Instead of a photograph of the actual intervention, the illustration of the procedure elucidates the scene, omitting undesirable details, while providing the necessary context. In comparison to medical illustration, medical visualization is constrained by the underlying data. However, the expertise of illustrators, which is geared towards effective communication, is valuable in the process of designing visualization algorithms. The existing techniques in the field of illustrative visualization will be discussed in subsection 4.2.

Besides the developments in medical visualization techniques, advancements in medical imaging produce rapidly increasing amounts of data. From imaging a patient, various types of data with distinct characteristics may be acquired. Especially for MRI, each measured sequence provides different tissue contrasts, providing singular diagnostic value. Managing the large amount of data requires visualization techniques that fuse information sources, emphasizing the important features of the data sets.

Furthermore, advances in medical imaging have led to higher dimensional data, and three-dimensional data have become commonplace. Recent medical imaging techniques also enable acquisition of multi-valued data. Especially in the field of neuroimaging, diffusion in the brain is measured by diffusion weighted imaging (DWI). Depending on the number of directions, the acquired information can be modeled by tensors, referred to as diffusion tensor imaging (DTI), or higher order constructs, based on high-angular resolution diffusion imaging (HARDI). Visualization of these multi-valued data typically involves tensor glyphs [172, 187] and various types of tractography line renderings [174, 250], and remains a challenging topic for image analysis and visualization research.

Besides three-dimensional data, four-dimensional data are increasingly acquired, providing time-dependent volumetric information. Additionally, four-dimensional blood-flow velocity fields can be acquired in-vivo, as described in chapter 2 and 3. These unsteady vector fields require specific visualizations. Flow visualizations in general have been studied in depth, and the existing techniques in this research field are elaborated in subsection 4.3

4.2 Illustrative visualization

4.2.1 Motivation

The previous section briefly motivated the area of illustrative visualization, oftentimes referred to as non-photorealistic rendering. Photorealism often precludes effective depiction of information contained within the data. The aim of illustrative visualization is to generate images that contain the necessary information, without loss of comprehensibility. To this end, the visualization community investigates established techniques developed by illustrators.

The more extensive motivation for illustrative visualization is threefold. First, illustrative visualization may be employed to create *works of art* using a computer system. Computers have the ability to visualize different media with great precision, and they enable the automation of repetitive tasks. It remains questionable, however, if a computer system will ever be capable of generating art, and of making decisions like an artist. It seems unlikely that creativity can be automated. However, computer graphics certainly provides valuable support for modern-day artists.

Clarification of ambiguous terms

The term **non-photorealistic rendering** (NPR) is an area of computer graphics expressive styles, which are inspired by art. NPR was introduced as a consequence to the prevailing mind-set of pursuing photorealism. The main problem is that the term strives to demarcate an area of research, by telling what it does *not* contain. It can be compared to categorizing a 'non-impressionist' art style, comprising all art that is not impressionistic. Such definitions cannot grasp the essence of the field, and make it hard to categorize the respective research. Moreover, the term photorealism has been criticized, because it is perceived differently by artists and computer graphics researchers. The artist will refer to photorealism as a school of painting, whereas the computer scientist will refer to photorealism as a technique for generating images that are indistinguishable from reality. Even the rendering component of NPR has been criticized, because not all NPR techniques are based on a rendering process. Instead, lots of techniques are based on modeling or post-processing of images. Hence, the term rendering is loosely used in this context.

Alternatively, the term **illustrative visualization** has been used to describe computer supported visual abstractions, inspired by illustration. This term has also been subject to criticism, since both illustrations and visualizations aim for visual abstraction to achieve an effective visual representations of the data. Hence, some authors consider the term a tautology [191]. Although a certain overlap with related areas remains, the term intuitively demarcates an area in the scientific research field. In the remainder of this thesis we adopt this terminology, motivating our visualizations by traditional illustrations.

Besides the artistic motivation, also *aesthetics* provides a notable argument for illustrative visualization. A pleasing visual appearance is often valuable, for instance for marketing purposes. Illustrations abstract away unwanted details of reality, resulting in more appealing depictions. For example, brochures selling houses or kitchens often use artistic techniques to make the product look more attractive. Automation of these techniques serves a large commercial target group. Additional aesthetic imagery, for instance in newspapers, appeal to the observer, and may stimulate recognition and interest in the subject matter [42].

The key motivation for our purposes is visual *comprehension and clarity*. Often comprehensibility and legibility are far better in illustrations, compared to photographs or photorealistic renditions (e.g., Fig. 4.1). Especially for scientific visualization of physics-based data, it is important to obtain a thorough understanding of the underlying phenomena. Illustrative visualizations employ computer graphics to provide expressive visual abstractions, motivated by traditional illustrations [191].

Illustrative visualization is concerned with the synergy between technology and art, and likewise serves pragmatic and artistic goals [120]. Research in illustrative rendering is divided into two distinct movements. On the one hand, there is a notable group of researchers trying to *simulate* artistic media. Such systems are mainly based on models that simulate physical properties of the substrate, and the painting tools and dyes. Examples are physical simulations of water paintings [236], and ink drawings [34]. On the other hand, researchers have shown great interest in *imitating* artistic work. An art style is being modeled, without explicit definition of real-world characteristics. The illustrative visualizations presented in this thesis all follow this pragmatic approach.

Over the past decade, a variety of distinctive illustrative visualization techniques have been developed. A considerable number of these techniques focus on point- and line-based graphics, which will be elaborated in subsection 4.2.2. Traditional illustrations prove that line drawings are effective at visually conveying information. Effective line drawings are difficult to master from a computer graphics perspective. Additionally, various techniques are concerned with artistic illumination of the objects of interest. These techniques will be elaborated in subsection 4.2.3.

Besides these elementary illustrative visualization techniques, the field also proposes a number of higher level visual abstractions that effectuate comprehensive visualizations. For instance, the separation in focus and context regions is generally accepted as a good practice [23, 251]. Other smart visibility methodologies, inspired by traditional illustration, include semantics-driven illustrations [190] and exploded views [24].

A considerable advantage of computer graphics over traditional illustration is the ability to adjust the visual result interactively. Camera viewpoint, and lighting conditions can be changed in real-time. Furthermore, the user can vary a wide range of parameters to influence the visualization. Real-time interaction can easily become computationally expensive. A considerable amount of research in the visualization

community is therefore dedicated to computational performance. A good performance is valuable for the majority of visualization techniques, and this certainly holds for illustrative visualization. The desired level of abstraction for illustrative rendering generally requires additional processing. This rapidly becomes computationally demanding, especially when user parametrization is required to be interactive.

The motivation for illustrative visualization largely coincides with the motivation for visualization in general, both aiming to amplify cognition through comprehensive depictions. Hence, one might question whether all visualizations should become illustrative [191]. In the field of computer graphics, achieving photorealistic renditions will definitely remain a notable goal. For visualization, photorealism is of less importance. In some areas, including architectural and medical applications, photorealistic renditions will at least partly subsist. Many other areas will exploit the clarity of illustrative approaches. For instance, illustrative renditions are worthwhile to investigate in the medical domain, since the field already employs visual abstractions. An example is the bull's eye plot, which provides a visual model to communicate cardiac pathologies. Such abstractions constitute a valuable basis for illustrative visualizations [223].

4.2.2 Techniques - point and line

A substantial amount of illustrative visualization research is concerned with point- and line-based styles. Lines are considered the scaffold of illustrative visualization, and are mostly used to convey object silhouettes or contours. Such visual abstractions are known from traditional illustration, and have been proven to be effective [40]. This is largely due to the intrinsic capability of a human observer to perceive and interpret these features. For instance, a human observer will immediately recognize a human face from a few well chosen line segments, such as the Da Vinci caricature in figure 4.2. Furthermore, point- and line-based styles are used to mimic traditional pen-and-ink drawing techniques, conveying object shape by tone variation. In general, these techniques are particularly challenging to translate to computer graphics.



Figure 4.2 A caricature of Leonardo da Vinci emerges from a few well chosen lines.
(© www.cartoonstock.com)

Line drawings emphasize object boundaries. The outline that captures the very edge of an object is called a silhouette [99]. Somewhat confusingly, a silhouette may also refer to the filled image-space projection of an object [169]. In contrast to silhouettes, contours capture the object's outline including structural details. A contour is generally defined at regions where the surface turns away from the viewer and becomes invisible. By definition silhouette and contour extractions are view dependent.

Besides *image-based* filtering approaches [75], *object-based* methods exist that extract contours from polygonal models [50, 268] or volume data [28, 44, 272]. Various other feature line definitions have been investigated [74], including ridges and valleys [96], suggestive contours [49], and apparent ridges [106]. These techniques aim for correspondence with drawings by traditional illustrators [39]. Line renditions are increasingly used in combination with other styles to emphasize object shape [19, 238].

Alternatively, line primitives are used to produce tone variations, mimicking traditional pen-and-ink hatching. The hatches convey surface shape by means of combined stroke patterns, which follow the surface curvature. Single hatches and crossed hatches respectively provide a medium and dark tone for shading, while no hatches result in the brightest tone. Surface hatching has been implemented in *object-space*, generating the hatch stroke geometry [167]. Furthermore, hatching has been applied through procedural textures [186] in a *hybrid approach* that includes object-space information.

A related point-based shading technique is termed stippling. It employs density variations of dots. Stippling may be performed in *image-space* [204], which typically results in a lack of frame coherence. Alternatively, *object-based* approaches control the stipple density on a voxel basis [147]. Furthermore, frame coherence can be obtained with a *hybrid approach* [9].

To ensure interactive renditions, most of these illustrative techniques have been adapted for consumer graphics hardware, commonly referred to as the graphics processing unit (GPU). Contours [28, 160], hatching [167], and stippling [9, 147] techniques have been GPU accelerated. Few approaches combine these illustrative styles into a generic framework [272]. In particular, Busking et al. have presented an integrated framework, based on a generic particle system [29]. Their approach is flexible and configurable, and encompasses all previously mentioned pen-and-ink styles, independent of the screen resolution. In our work, we have exploited parallelism in modern consumer graphics hardware, to accelerate this particle-driven visualization framework [241]. This enables real-time interaction and parametrization, including iso-surface selection and user adjustments of the visual result. Moreover, we have extended the framework with ridge and valley lines [240].

Visualization results generated by our illustrative framework, are presented in figure 4.3. In accordance with the topic of this thesis, we present computer generated illustrations of the cardiovascular system. The figure includes contours, and pen-and-ink style techniques, such as stippling and hatching. Furthermore, an additional type of abstraction is presented by the exploded view in figure 4.3c. It clearly conveys the separate anatomical regions of the whole-heart segmentation. All images are rendered at interactive framerates, and can be parameterized in real-time [240].

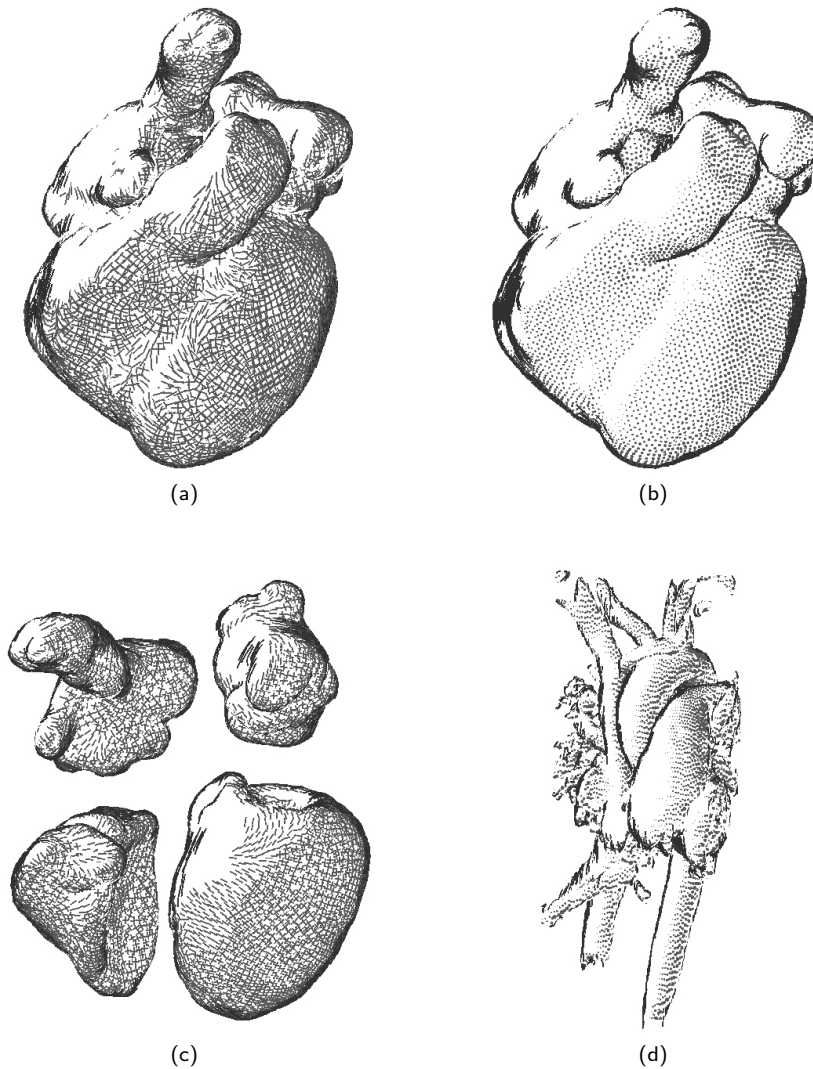


Figure 4.3 Interactive particle-based illustrative visualization, mimicking traditional pen-and-ink techniques [240]. (a) A whole-heart model, depicted with object-space contours and illumination based on hatches, (b) or stipple size. (c) Furthermore, the separate anatomical regions of the whole-heart segmentation are depicted using an exploded view. (d) Instead of a whole-heart segmentation, the blood-flow lumen from a 4D PC-MRI scan can be employed to approximate the cardiovascular morphology. The whole heart model was provided courtesy of dr. G. Hautvast - Philips Healthcare.

4.2.3 Techniques - shading and style

In addition to point- and line-based graphics, different shading techniques and styles have been proposed to illustratively convey object shape. In contrast to customary computer graphics approaches, these techniques do not aim for veridical conveyance of the object. Instead, the sole aim is to effectively communicate object shape, with minimal visual clutter.

In the early days of illustrative visualization, Gooch et al. proposed an illustrative shading model, tailored to imitate colored technical drawings [73]. Their shading approach produced different results as compared to the prevailing models that were aiming for realistic lighting. An example is the well-known Phong model [179], depicted in figure 4.4a. Instead, they simplified shading and illumination, conveying the essential shape characteristics. Similar techniques were later integrated with a volume rendering approach [148].

In accordance, toon shading, also referred to as cel shading, mimics the style of comic books or cartoons, as depicted in figure 4.4b. The term cel shading originates from traditional animation techniques, where shading was applied on transparent sheets of celluloid, or 'cels'. Cartoon images have a natural power of attraction. This is due to their simplicity, while preserving the essential information. The key characteristic of toon shading is the limited number of shades used to convey shape. Initial computer graphics techniques applied this type of shading to three-dimensional objects, employing a texture lookup that realized the desired number of shades [131]. Since then, several new variations and extensions have been presented [10, 159, 197, 215].




Figure 4.4 A prevailing aim within computer graphics is to closely emulate photorealistic illumination. (a) The Phong shading model [179] approximates realistic lighting; here presented on the whole-heart model. (b) Alternatively, cel shading provides a simplified illumination, employing a limited number of shades inspired by cartoons. (c) As a complementary technique, feature lines may be used to convey shape.

Additionally, several stylized rendering techniques have been proposed. Individual illustrative styles are quite diverse, and therefore hard to integrate into a single framework. Instead, stylized techniques emulate a wide range of illustrative styles, based on a sample sphere. Such a lit sphere captures color variations of an object as a function of the normal direction. A large variation of styles can be generated by introducing new lit spheres, often inspired by traditional illustration. Stylized techniques have been applied both to polygonal models [210], as well as volume data [25].

Although shading and style-based techniques adequately convey object shape, still these approaches are often combined with feature lines. While lines roughly emphasize the contours of the object, as depicted in figure 4.4c, shading elucidates shape details. Hence, these techniques are used complementarily, reinforcing shape conveyance.

4.3 Flow visualization

4.3.1 Motivation

Chapter 2 motivated the importance of understanding cardiovascular hemodynamics. Besides quantitative analysis, also visual exploration and inspection have proven valuable to gain insight into the intricate blood-flow behavior. In this thesis, we focus on unsteady volumetric blood-flow velocity fields, acquired by PC-MRI (see subsection 3.2.4). As described in subsection 2.3.2, the visual analysis of high-dimensional flow




Figure 4.5 Flow visualization of an artificial steady vector field. (a) A vector plot provides a direct visualization of the field. The presented arrows are based on 4000 sample positions, comprising a mere 2% of the number of voxels. (b) Alternatively, streamlines sparsely depict the same vector field, tracing lines from four seed points, positioned at critical points within the field. This visualization conveys the field structure more effectively. Images generated using ParaView (<http://www.paraview.org/>). The artificial flow field was provided courtesy of prof. C. Hansen - University of Utah.

data is very challenging, if not impossible, by means of conventional slice-by-slice inspection. The difficulties that arise with the qualitative analysis of these data are not restricted to the blood-flow application.

Many other applications are concerned with flow data, and likewise require visual analysis to gain understanding. Applications can be found in automotive and aerospace design, chemistry, meteorology, climate modeling and medicine. This list is by no means exhaustive, evidencing the need for specialized visualization techniques. Terminology to ascertain the type of flow field differs between application areas, and is often used interchangeably. Dimensionality of a flow field is simply indicated by the number of dimensions, while data with a common dimensionality have received specific terms, such as planar or volumetric. Flow fields over time are mostly termed to be unsteady, as opposed to non time-bound steady fields. Unsteady flow fields are also often named time-dependent, time-resolved or cine.

Traditionally, flow visualization researchers did not have the methods of computer graphics at their disposal. Instead, flow visualization was performed experimentally [183, 224]. On the one hand, materials were added to visualize the flow patterns. For instance, injection of dye or magnesium powder into the fluid flow reveals the dynamics. For gaseous fluid, smoke and oil droplets were employed. On the other hand, optical techniques visualize flow patterns based on changes in the optical refractive index. Light refraction changes in regions with large local variations in flow density. Using a light source, images can be generated with shadows and caustics. Different optical methods have been used, including shadowgraphs, interferometry, and Schlieren photography. Many of these techniques have inspired current graphics-based flow visualization techniques, as exemplified by a recent physics-based computer simulation of Schlieren photography [22].

Over the past two decades, graphics-based flow visualization has been an active field of research, yielding a variety of techniques to facilitate a qualitative analysis of flow data. Evidently, direct slice-based visualizations, such as the one in figure 3.7, are not effectual, especially for unsteady and volumetric flow fields. Other direct visualization approaches straightforwardly depict the directions of the vector field. In case of a hedgehog visualization, a flow vector is represented by a small line segment [117]. Alternatively, the flow directions are often depicted by small arrows, generally denoted as a vector plot [115].

While these techniques are effective for two-dimensional flow fields, both steady and unsteady, the situation becomes more challenging in the three-dimensional domain. Visual clutter and occlusion may dominate the visualization, as witnessed by the example in figure 4.5a. This can hamper interpretation of the flow fields. To depict volumetric flow fields, direct visualization techniques have been proposed, employing volume rendering based on raycasting [43, 62]. While these techniques reduce visual clutter to a certain extent, an expressive display of the flow field remains challenging. This is in part due to the involved transfer function specification, necessary to classify characteristics of the flow field.

Alongside direct flow visualization approaches, many other techniques have been proposed that aim to communicate the essential characteristics of flow fields. In many cases, specific aspects of the flow field are identified, enabling reduction of visual clutter. For instance, figure 4.5b shows a limited number of adequately seeded streamlines (see also subsection 4.3.5), resulting in a more comprehensible depiction in comparison to figure 4.5a. The remainder of this section provides an overview of recognized flow visualization techniques, adhering to the prevailing taxonomy in the field.

4.3.2 Texture-based graphics

Texture-based methods employ texture synthesis for flow visualization, providing dense representations of the flow data. The dynamics of the flow data are shown by means of a texture that is filtered along the vector field. A profound review of this group of techniques was performed by Laramée et al. [134].

The instigator of a wide variety of texture-based flow visualization methods was the spot-noise technique [243]. This method synthesizes a texture by distributing spots over the data set domain. Subsequently, the spots are deformed along the direction of the flow. Additionally, the renditions may be enhanced by animating the spots.

Thereafter, the line integral convolution (LIC) approach was introduced [30]. This method generates a similar dense representation of the flow field. It uses convolution of a random texture with pieces of streamlines as the filter kernels. The initial source texture typically consists of randomly generated noise, as depicted by figure 4.6a. The streamline-based filtering reveals the flow dynamics, as depicted by figure 4.6b for an artificial flow field.




Figure 4.6 Flow visualization using dense texture synthesis, based on the line integral convolution (LIC) method. (a) A random noise texture as input for the LIC method. (b) Locally filtering along the streamlines reveals the flow patterns of the artificial data set [100]. (c) Similar techniques were extended to three-dimensional flow fields, using sparse volumetric textures and volume clipping to reduce visual clutter. This example shows a simulation of a flow within the wheel casing of a car, as presented by Rezk-Salama et al. [193].

Both methods have been improved and extended. Performance of both spot noise and LIC has been improved by means of graphics hardware [46, 48, 102, 275]. For the LIC technique, adequate noise generation has been investigated [116], and the resulting textures have been combined with color mappings [216]. Furthermore, LIC has been extended to non-Cartesian grids [61]. Several other improvements concerned fast and coherent animation [61, 256, 257]. In addition, techniques with similar results have been introduced, such as physics-based smearing of virtual ink droplets [142].

Line integral convolution was furthermore extended to the three-dimensional domain [97, 98]. In three-dimensions, texture synthesis techniques suffer from visual clutter and occlusion, and depth perception is hard to achieve. A well-chosen selection of critical regions of interest is essential. Sparse noise textures and clipping techniques have been used to improve the visualization [193], as depicted by the simulated velocity field within the wheel casing of a car in figure 4.6c. Furthermore, three-dimensional LIC has been combined with colored dye injection to highlight local features [207], and texture advection was used to visualize the motion of colored particles [261].

Dense texture based techniques were occasionally applied to simulated and measured blood-flow fields. Typically, the blood-flow information is conveyed using a two-dimensional texture representation, either on a view-aligned plane [70], or mapped onto the vessel surface [136]. To the best of our knowledge, volumetric LIC has been applied only once to unsteady volumetric blood-flow velocity data [108].

4.3.3 Feature-based graphics

In contrast to direct and texture-based flow visualization, feature-based approaches aim for a higher level of abstraction. To this end, meaningful flow patterns are extracted from the flow data. The visualization subsequently shows the interesting characteristics, known as features, to the observer. A quintessential review paper on feature-based flow visualization was presented by Post et al. [184].

The goal of feature-based visualization techniques is to reduce the visual representation of intricate flow fields to simple primitives, such as points, lines or surfaces. These primitives represent the extracted features, yielding a simplified depiction of the flow phenomena. This leads to abstract renditions, presenting the flow in a concise and compact fashion [69]. Additionally, the flow field may be explored through features, for instance using a local probe [47]. The resulting renditions are generally effective for each specific application, and should be interpreted by domain experts.

Some ubiquitous features are relevant to the majority of flow applications. Notable examples are vortices, shock waves, and separation or attachment lines. Specifically for blood-flow velocity fields, important features mostly involve physical characteristics, such as speed, acceleration, flow rate, kinetic energy, relative and absolute pressure, vorticity or helicity. Other features may include characteristics of the surrounding biology, such as the wall-shear stress (see also subsection 2.3.1).

Extraction of features mostly relies on a set of mathematical methods, adopted from the image analysis field. For instance, derivatives of the flow field are frequently used in feature extraction techniques. In particular, the gradient matrix of a three-dimensional vector field, commonly denoted as the Jacobian matrix, is often employed. In addition, eigen analysis provides valuable information that is often applied to feature extraction. For example, eigen analysis of the Jacobian matrix indicates the direction of tangent curves through the flow [184].

Although most feature extraction techniques rely on specific mathematical foundations, various definitions of a single feature may coexist [69]. For instance, there is a wide range of definitions to identify a vortex. Each definition is appropriate within the application-specific context. However, non of these features provides a generic characterization for all flow fields, causing a feature ambiguity that may result in erroneous interpretations of the flow dynamics.

Closely related is the area studying topology-based methods. The elementary definition of topology describes “*a branch of mathematics, which is concerned with those properties of geometric configurations that are unaltered by elastic deformations, or homeomorphic*” (source: Merriam-Webster). In the context of flow visualization, the mathematical formalisms are concerned with the qualitative properties of flow systems, which should be preserved under continuous deformations. In contrast to mere feature extraction, the provided framework guarantees unambiguous results [69]. A profound review of topology-based techniques was presented by Laramée et al. [135].

Topology for flow visualization was first introduced by Helman and Hesselink [89]. They suggested to extract global flow patterns, defined with respect to the limit sets of streamlines. When the limit set consists of a single point, this is called a critical point. The flow in the neighborhood of such a critical point may be inspected by an eigen analysis of the Jacobian of the vector field. Depending on the real and imaginary parts of the eigen values, critical points are classified as sources, sinks,




Figure 4.7 A simulated flow field in a cylindrical container, involving a vortex breakdown bubble. In combination with restricted volume rendering, the flow topology is visualized based on critical points using singularity paths (yellow) and streamlines (blue), as presented by Tricoche et al. [229].

or saddle points. The type of critical point is often deduced from its index, which is based on the number of field rotations along a closed curve around the point [83, 141]. This approach was introduced by Poincaré, and is also known as the winding number.

Based on these critical points, the aim is to attain a topological skeleton that conveys the flow field, as exemplified by figure 4.7. Regions of different behavior can be discerned by means of particular streamlines or surfaces, called separatrices. These separatrices originate from saddle points in the flow field. Furthermore, other higher order constructs have been proposed as part of a topological skeleton [260].

Currently, a crucial drawback of topology-based methods, is the limitation to steady flows. Considerable effort has been invested to develop methods for unsteady flow fields [181]. Most methods focus on tracking critical points, which has been proven to yield incorrect results [177]. Alternatively, unsteady flow is considered from a Lagrangian point-of-view, analyzing fluid motion based on particles [65]. In particular, there is an increasing interest in Lagrangian coherent structures, enabling separation of dynamically distinct regions in unsteady flows. Especially, the finite-time Lyapunov exponent (FTLE) gained considerable attention [112, 123]. Alternatively, there are spacetime approaches, which consider time as an additional axis [225].

Furthermore, topology-based methods are inherently sensitive to noise, complicating their applicability to measured flow data, such as the unsteady PC-MRI blood-flow. Topology in this application is further restricted by the limited number of critical points, because sources and sinks by definition cannot occur in incompressible flows.




Figure 4.8 Different visualizations of clustered flow data. (a) Volume rendering of a time-resolved supernova simulation, hierarchically clustered in the spacetime domain [271]. (b) Curved stream arrows with a spot-noise texture show a hierarchically clustered artificial flow field [222].

4.3.4 Partition-based graphics

In contrast to feature-based techniques, partition-based flow visualizations convey the field characteristics based on a subdivision of the whole domain. Instead of pre-determining adequate features, which is considerably challenging in many applications, relevant partitions can be identified a posteriori. The partitioning relies on elementary characteristics of the flow field, resulting in a subdivision that captures the flow field structure. The structure reflects the whole flow field, while providing a sufficiently abstract representation to reduce visual clutter. A review of partition-based flow visualization techniques was presented by Salzbrunn et al. [199].

Partitioning of the flow domain is often performed by means of clustering. Adequate clustering reduces the complexity of the flow data, without neglecting possibly interesting small regions. Some methods require the definition of the number of clusters beforehand [53, 161]. Other approaches are fully data driven, mostly clustering the domain in a hierarchical fashion. Starting from individual flow vectors, hierarchical clustering may be performed by merging clusters, which is referred to as a bottom-up approach [31, 222]. Conversely, it is possible to iteratively perform splitting operations in a top-down order [88]. Furthermore, clustering may be performed without boolean merging or splitting, using a diffusion process to enhance correlations in the cluster set [66]. Given a set of clusters, different visualization techniques can be used to depict the flow field, as exemplified by figure 4.8.

All partition-based methods rely on carefully selected cost functions, defining which clusters correlate best. Most methods rely on elementary properties of the flow field [222], while some methods employ higher level concepts, such as the gradient, curl, divergence, or vorticity [161]. Alternatively, cluster measures may rely on streamline characteristics, such as curvature and associated bending energy [126].

A common drawback of partition-based techniques is the computational complexity, which scales poorly with the data dimensionality. Consequently, these techniques have been applied to relatively small-scale steady flow fields. For unsteady flow, some research has been carried out to perform the partitioning in the spacetime domain [271]. Our findings and visualization approaches will be described in section 5.3.

4.3.5 Geometry-based graphics

The area concerned with geometry-based flow visualization aims to reflect the characteristics of a flow field through discrete geometrical objects. This group of methods uses geometric primitives to directly communicate trajectories obtained from the flow field. In contrast, other areas have employed trajectories for texture generation or feature extraction. The trajectories are initiated from a set of carefully selected seed positions. Afterwards, a numerical integration through the flow field approximates the analytical solution. Consequently, geometry-based techniques are also

known as integration-based techniques, for which an elaborate review was presented by McLoughlin et al. [162].

Different types of trajectories exist, which individually represent different aspects of the flow field. A *streamline* represents the instantaneous tangent curve in a steady flow field $\vec{v}(\mathbf{x})$, for positions \mathbf{x} in the data domain. Initiated at the seed position \mathbf{x}_0 , the streamline trajectory for a given integration time τ is defined as:

$$d\mathbf{x}(\tau) = \vec{v}(\mathbf{x}(\tau)) d\tau, \text{ with } \mathbf{x}(0) = \mathbf{x}_0.$$

Alternatively, tangent curves can be obtained from unsteady flow fields. These curves, called *pathlines*, describe the path that a massless particle takes through a time-dependent flow field $\vec{v}(\mathbf{x}, t)$. For pathlines, the integration time t corresponds to the actual time; initiated at time t_0 . The pathline trajectory originates from a spatiotemporal seed position \mathbf{x}_0 , and is subsequently defined as:

$$d\mathbf{x}(t) = \vec{v}(\mathbf{x}(t), t) dt, \text{ with } \mathbf{x}(t_0) = \mathbf{x}_0.$$

In addition, flow field dynamics can be characterized by *streaklines*, which comprise a set of points seeded at identical spatial locations, and distributed uniformly in time. Streaklines can be observed experimentally, by continuous dye injection at a fixed position. Alternatively, points that form a line may start at the same time, with varying spatial locations. Such a *timeline* is subsequently advected through the flow field, analogous to a stretchable rope transported by the flow of a river. Details and in-depth formulations of these integral lines are provided by various authors [45, 260].

These trajectories have yielded multitudinous geometry-based visualization approaches. A broad range of techniques utilize *particles* to depict the trajectories, especially in pathline visualizations. Initial visualization research focussed on an accurate and efficient generation of the trajectories, increasingly using consumer graphics hardware [114, 125]. Meanwhile, the attention is targeted at adequate seeding strategies, which are generally challenging for all geometry-based visualization techniques [157, 244]. Thereafter, the visual representation of the particles was taken into account, experimenting with shape and enhanced depth perception [125, 211].

Particle animations effectively convey flow dynamics, and have been applied to a multitude of applications. Also for the blood-flow application, particles have been employed for both simulated and measured velocity data. The value of particle systems has been demonstrated in a research setting [11, 123, 234, 248, 263], and has been applied in commercial tooling [189]. An example is provided in figure 4.9a.

Alternatively, the trajectories can be represented directly by *line* primitives. Also for integral-line visualizations, placement of the seed positions largely determines

the effectiveness of the rendition. Several approaches have been proposed, aiming for an even distribution of lines to ease the interpretation of the line renderings [103, 139, 156]. Furthermore, various visualization approaches, including tubes and ribbons, enhance the appearance of the lines [203, 219], as exemplified by figure 4.9b. Additionally, illustrative approaches have been inspired by pen-and-ink drawings [59, 220]. Beside the appearance, improvements to enhance spatial relations and depth perception have been developed, such as line illumination [274] and halos [59]. Line visualizations have also been used to depict blood-flow [90, 129, 154].

In addition to particles and lines, trajectories can furthermore be conveyed through integral *surfaces*. By increasing the dimension of the seeding objects, the dimensionality of the geometric primitive will likewise increase. Hence, a surface primitive may be constructed from a group of adjacently seeded trajectories. Depending on the type of trajectories, streamsurfaces [143, 250], pathsurfaces [123, 201], timesurfaces and




Figure 4.9 Various geometry-based flow visualization approaches. (a) One time frame of a particle animation, depicted by illuminated spheres, in pathological blood-flow [239]. (b) Streamlines convey the numerically simulated flow field of a heat-driven cavity. Illustrative patterns on each line scale with the flow velocity [58]. (c) A streamsurface depicts the flow in a vortex breakdown bubble, using feature lines and illustrative shading [19]. (d) A pathsurface depicts the flow of a turbulent jet. The illustrative visualization uses color to distinguish between the front and back side of the surface. Stripe patterns are employed to visualize the individual pathlines on the surface [95].

streaksurfaces [26, 124] can be generated. Also for integral surfaces, seed placement highly influences the visual result. To explore the flow field, the interactive generation of the surfaces is beneficial. However, real-time computation of both the integral trajectories and the surface geometry is challenging [26, 123]. Surface renditions often suffer from self-occlusion, and hence different visualization approaches have been investigated, including surface perforation and transparency [68, 143]. Furthermore, illustrative approaches have been studied [19, 95], as exemplified by figure 4.9c-d.

All geometry-based methods rely on numerical approximations, and hence errors and uncertainty accumulate for longer trajectories. Visualizing this uncertainty is an open challenge in the visualization community [140]. From the previous discussion it follows that positioning of the seed objects is another major challenge for geometry-based methods. Although most issues are solved for two-dimensional flow fields, positioning in the three-dimensional domain remains challenging [162]. This especially holds true for unsteady flow data, in which case the seed objects need to be positioned in both space and time. In our work, we have investigated seeding for various geometry-based flow-visualization techniques. This will be elaborated in chapters 6 and 7.

4.4 Interactive visualization

4.4.1 Motivation

Interactivity is a key aspect for visualization techniques in general, and especially for exploratory approaches. Direct interaction with the visualization engages the user, who can directly steer the visual results. This involves adjustment of different user parameters as well as interaction with the virtual camera to inspect the scene.

This camera interaction is important for several reasons. On the one hand, the interaction enables visual inspection from different viewpoints. Therefore, different aspects of the data can be analyzed, and occlusion of objects can be circumvented. On the other hand, the camera interaction provides an important monocular depth cue. The motion parallax provides an important contribution to the perception of depth, together with, inter alia, perspective, light and shade.

Interactive visual inspection is particularly important for volumetric and time-resolved data. Besides inspection of the spatial relations, also the temporal behavior needs to be explored. Conventionally, animations are employed for this purpose. Although animations impose a considerable cognitive load, they also reveal the spatial relations between the presented objects. Repeated animation, together with interaction with the virtual camera, facilitates the qualitative analysis of the data.

The interactive parametrization and camera movements, as well as animation, require real-time actuation and rendering. In many cases, the performance of conventional graphics that use the central processing unit (CPU) is insufficient. Therefore, dedicated hardware has been employed to obtain interactive framerates. Nowadays, consumer graphics hardware provide a generic and powerful means to gain performance.

4.4.2 Consumer graphics hardware

Modern consumer graphics hardware provide a graphics processing unit (GPU), tailored for high-end graphics in the game industry. The GPU enables high-performance parallel processing of defined tasks. These tasks are carried out by a fixed graphics pipeline, which processes geometry and pixels in parallel.

Usage of the GPU relies on a single-program multiple-data (SPMD) architecture. A small program, called a shader program, defines fixed tasks to be carried out in parallel. The term shader originates from the initial intent of these programs to apply shading to different graphics primitives.

The graphics hardware and pipeline can furthermore be exploited to perform various computations in parallel, drastically improving the computational performance. This research field performs general computation on the GPU (GPGPU), often using clever detours. For example, costly sorting and searching algorithms can be performed embarrassingly parallel by using the GPU [240]. Throughout this thesis, the GPU was employed to achieve the desired performance and interactivity for all visualization techniques, and associated algorithms.

In the last few years, computational languages are emerging, facilitating general computations on the GPU. These languages, such as OpenCL or NVidia's TM CUDA, provide an abstraction to the programming of the graphics hardware. As a result, these languages make parallel processing of various algorithms on the GPU more accessible.

"What is real is not the external form, but the essence of things."

Constantin Brancusi

Blood-flow abstraction



This chapter is based on:

"Automated Segmentation of Blood-Flow Regions in Large Thoracic Arteries using 3D-cine PC-MRI Measurements." Roy van Pelt, Huy Nguyen, Bart ter Haar Romeny, and Anna Vilanova. In *International Journal of Computer Assisted Radiology and Surgery* (2012), 7(2):217-224.

"Visualization of 4D Blood-Flow Fields by Spatiotemporal Hierarchical Clustering." Roy van Pelt, Sander Jacobs, Bart ter Haar Romeny, and Anna Vilanova. In *Computer Graphics Forum* (2012), in press.

5.1 Temporal projections

5.1.1 Motivation

The foregoing chapters unveiled the challenges with respect to a qualitative analysis of flow fields. Especially for unsteady and volumetric flow data, visual clutter and occlusion adversely affect the analysis. Therefore, adequate abstraction of the flow data is necessary to enable an effective visualization.

Besides the variety of simplification techniques presented in section 4.3, abstraction of the flow data can also be achieved by a reduction of dimensionality. This can for instance be realized by projections along the time axis, providing a time-independent summary of the unsteady flow field. Throughout the work presented in this thesis, two distinct projections were employed. These projections will be elaborated in the remainder of this section.

5.1.2 Temporal maximum intensity projection

The first projection is termed the temporal maximum intensity projection (tMIP). The tMIP results in a new volumetric scalar-valued data set, representing a coarse static approximation of the cardiovascular structures. For each voxel position \mathbf{x} of this new volume, the maximum speed is determined along the time axis of the unsteady volumetric flow data $\vec{v}_{t_i}(\mathbf{x})$. For P cardiac phases, this is defined as:

$$\text{tMIP}(\mathbf{x}) = \max_{t_i}(\|\vec{v}_{t_i}(\mathbf{x})\|) \text{ for } i = 0, \dots, P - 1. \quad (5.1)$$

In the volume obtained by this temporal projection, voxels with a bright intensity indicate a position where flow with a substantial speed has occurred at least once during the cardiac cycle. Hence, such a voxel was certainly contained within the cardiovascular structure at least once during the heartbeat.

The resulting scalar-valued volume closely resembles a magnetic resonance angiogram, which enhances the morphological structures by contrast injection. Segmentation of the high intensity regions yields a luminal geometry of the cardiovascular structures. The tMIP volume can be inspected by means of a visual maximum intensity projection, as presented in figure 5.1a. A similar strategy to obtain a virtual angiogram from PC-MRI blood-flow data was previously presented by Hennemuth et al. [90]. Instead of a maximum operator, they employ a temporal averaging, which is less sensitive to noise. However, boundaries of the cardiovascular lumen are less sharply defined, which is disadvantageous for segmentation purposes.

5.1.3 Temporal mean orientation projection

In addition, we present an average of blood-flow orientations over time, referred to as the temporal mean orientation projection (tMOP). The tMOP provides a static summary of the blood-flow orientations for a full cardiac cycle. Therefore, voxel-wise mean orientation distributions are considered, defined by structure tensors [118]. A subset of a tMOP volume is depicted in figure 5.1b, in which case the orientation distributions are depicted by normalized ellipsoids.

The structure tensors are computed by a tensor product \otimes of each velocity \vec{v} , at all voxel positions \mathbf{x} in the data set. Subsequently, the temporal mean of these tensors is computed, resulting in a description of average orientation distributions, projected over P phases of the cardiac cycle. The tensor-valued tMOP volume is given by:

$$\text{tMOP}(\mathbf{x}) = \frac{1}{P} \sum_{t=1}^P (\vec{v}(\mathbf{x}, t) \otimes \vec{v}(\mathbf{x}, t)) \quad (5.2)$$

The use of structure tensors for the tMOP computation avoids direction cancelation, in contrast to averaging of the velocity vectors. This comes at the cost of losing directional information. For various purposes, however, the lack of directional information does not pose a problem, when only a global and steady description of the flow field is required.

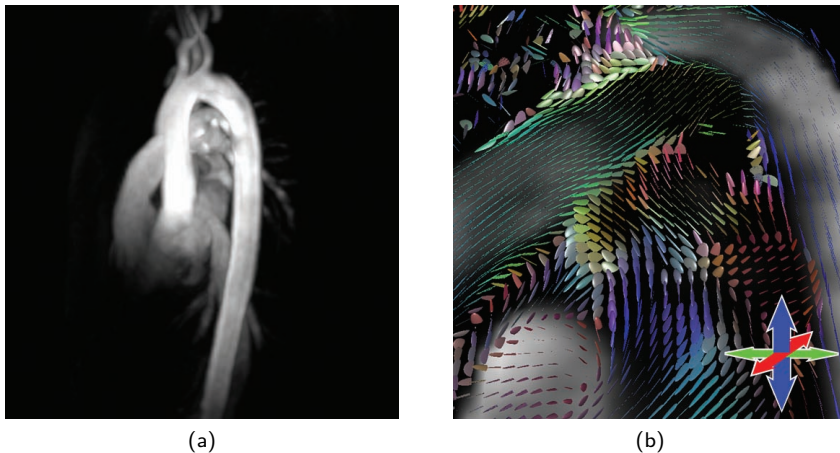


Figure 5.1 (a) A tMIP volume provides a rough approximation of the cardiovascular morphology, here depicted by a visual maximum intensity projection (see equation (5.1)). (b) The tMOP provides an average orientation field of the unsteady velocity data (see equation (5.2)). The orientation field is depicted by normalized ellipsoidal glyphs, with RGB color coding for xyz orientations. The magnitude of the principal directions are communicated by the gray-value intensities.

5.2 Segmentation

5.2.1 Motivation

Segmentation is concerned with a meaningful partitioning of data, leading to more abstract representations. In the context of medical data, this mostly entails automated extraction of anatomical structures. These segmented structures generally provide a scaffold for medical visualization, enabling the communication of anatomical context. Besides the benefits for visualization, accurate segmentation is likewise important for a quantitative analysis. This certainly holds true for the PC-MRI blood-flow data, where accurate segmentation of the vessel lumen is relevant.

In clinical practice, segmentations are largely obtained through manual delineation of the vessel circumference on two-dimensional slices. This is possibly repeated for several phases of the cardiac cycle. In case of velocity-encoded data, the blood-flow lumen is generally delineated, instead of the vessel morphology. Despite the intra- and inter-observer variability of these manual segmentations, this approach has shown to be sufficiently fast and accurate for practical purposes.

Extending the segmentation to a 3D-cine data sets is non-trivial. While segmentation in a 2D-cine case comprises a single circumference, segmentation of volumetric data involves a three-dimensional surface, as depicted in figure 5.2. Manual segmentation of 3D-cine blood-flow data is laborious and time-consuming, and hence not suitable for clinical practice. Therefore, an automated approach is needed to segment the lumen of time-resolved volumetric PC-MRI blood-flow data.

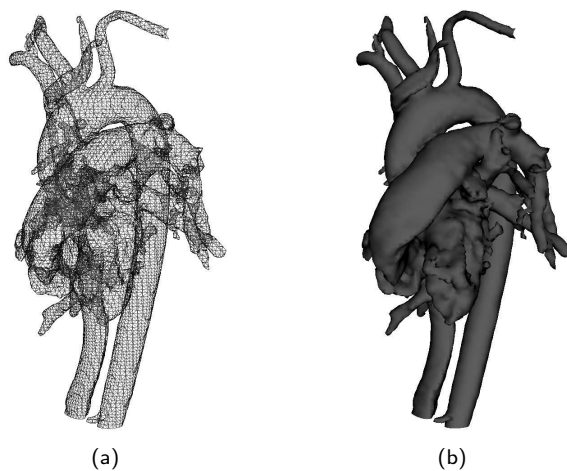


Figure 5.2 Surface geometry, segmented from a 4D PC-MRI acquired cardiovascular blood-flow field, depicted by (a) a wireframe representation, and (b) a Phong-shaded surface.

To this end, we introduce an automated global segmentation of the blood-flow lumen. We employ an active surface model, as opposed to thresholding, level-set, or active shape approaches. Active surfaces are a three-dimensional extension of active contours, which use energy-minimizing curves to extract features of interest in images [38, 111]. The energy to be minimized depends on the feature data, as well as the shape of the contour or surface.

An active surface benefits from topological stability, avoiding undesirable branches, splits, or holes in the surface. Moreover, active surfaces are generally more accurate than thresholding, and outperform level-sets in terms of computation time. On occasion, no high quality anatomical data is available along with the velocity-encoded data. Therefore, our segmentation approach relies on blood-flow data only. The segmentation does not approximate the vessel-wall morphology. It distinguishes blood-flow regions from surrounding stationary tissues, capturing the luminal geometry.

Our active surface approach uses the full velocity information, including the blood-flow direction. To the best of our knowledge, active surface segmentation based on full velocity information has not been employed for determining the blood-flow lumen before. Active surface models require an initial surface that roughly approximates the final segmentation result. We determine the initial approximate surface by iso-surface extraction from the tMIP volume, which was introduced in subsection 5.1.2. Although variation of the iso-threshold will provide different initial surfaces, this will not strongly affect the results of the active surface model.

Active surface models require a set of features that determine the forces that attract the surface. These features are introduced in subsection 5.2.2, prior to the description of the active surface model in subsection 5.2.3. The segmentation results were validated against manual segmentations, performed on 2D-cine PC-MRI blood-flow slices with higher spatial resolution. This validation and the results are presented in subsection 5.2.4.

5.2.2 Features

The attracting forces for our active surface model require features that emphasize the boundaries between blood-flow regions and stationary tissue. The features should be sensitive to the blood-flow speed and direction. In addition, we require that the features are spatially robust to noise.

Blood-flow speed (BFS) - The blood-flow speed is commonly inspected in angiograms, and is employed by the majority of segmentation approaches. BFS is defined as the length of velocity \vec{v} at each voxel position \mathbf{x} :

$$\text{BFS}(\mathbf{x}) = \|\vec{v}(\mathbf{x})\| = \sqrt{\vec{v}(\mathbf{x})_x^2 + \vec{v}(\mathbf{x})_y^2 + \vec{v}(\mathbf{x})_z^2}$$

This feature is not sensitive to the blood-flow direction, and is commonly computed without local averaging. Averaging the speed information in a local neighborhood improves noise-robustness, at the expense of contrast near the edges. Alternatively, the noise can be suppressed by masking with a complex-difference reconstruction, which was described in subsection 3.2.2.

Local Phase Coherence (LPC) - Besides speed, also direction provides valuable information for the segmentation process. Directions of the blood-flow field are said to be locally coherent, as opposed to blood-flow velocities near the boundaries.

Chung et al. [35] introduced the LPC metric, locally measuring the coherence as a scalar value for each voxel of the data set. They compute the *average angle* between a velocity vector and its direct neighbors. This is the sum of inner products of a unit-length velocity vector \hat{v} with the unit-length velocity vectors of the direct neighbors. The result is normalized by the number of velocity vectors M in the neighborhood.

$$\text{LPC}(\mathbf{x}) = \frac{1}{M} \sum_{u=-1}^1 \sum_{v=-1}^1 \sum_{w=-1}^1 \hat{v}(\mathbf{x}) \cdot \hat{v}(\mathbf{x}_x + u, \mathbf{x}_y + v, \mathbf{x}_z + w)$$

The LPC is based on inner products, which by definition results in sensitivity to blood-flow direction. Averaging the velocities in the neighborhood provides built-in robustness to noise. By normalizing the velocities, as proposed by Chung et al. [35], the LPC is not sensitive to the blood-flow speed.

Eigenvalue Coherence (EVC) - An alternative coherence metric was introduced by Solem et al. [212]. Instead of considering angles, they determine the coherence based on a local *distribution of orientations*. An orientation distribution is computed as the average structure tensor [76] of velocities in a local neighborhood. By default we use a 26-connected neighborhood. The structure tensor is the tensor product of a velocity vector with itself.




Figure 5.3 Image feature intensities. (a) blood-flow speed, (b) local phase coherence and (c) eigenvalue coherence. Within the blood-flow regions, intensities are high, with discontinuities on the boundaries towards stationary tissue.

To obtain the coherence metric, an eigenanalysis is performed on the average structure tensor \mathbf{T} at each voxel position. The eigenvectors span the basis of the orientation distribution. Intuitively, this can be thought of as an ellipsoid that represents the variation of orientations in a local neighborhood.

The coherence can be derived from the lengths of the eigenvectors of the local structure tensor. Hence, the local eigenvalue coherence metric is defined as a function of the first two ordered eigenvalues λ_1 and λ_2 :

$$\text{EVC}(\mathbf{T}) = \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right)^2$$

Due to the tensor product, this feature is only sensitive to orientation, discarding the sign of the velocity direction. It is plausible to assume that within a small blood-flow region there will be no directly opposing blood-flow directions. Therefore, orientation information is sufficient. Furthermore, EVC is sensitive to the blood-flow speed and is considerably robust to noise, because of the integration scale.

5.2.3 Active Surface Model

An active surface model deforms a surface, until the surface closely approximates the boundaries of interest in the provided data set. In comparison to other segmentation methods, active surfaces are computationally efficient, and can operate with sub-voxel accuracy. Moreover, active surfaces are topologically stable and generally robust to noise. However, active surface models suffer from numerical instability and may converge to local minima. The impact of these disadvantages is reduced if the model is provided with a well-estimated initial surface.

The initial surface was obtained by a marching cubes iso-surface extraction on the tMIP volume [144], providing a surface mesh. The surface is supplied with a potential energy, moving it towards edges in an attractor image. For our application, this attractor image, or energy image, is based on the features described in subsection 5.2.2. In addition, the surface is provided with an internal energy, which relies on the surface shape.

The active surface model aims to minimize the total energy functional, which is defined as the sum of the potential energy and the internal energy. Minimizing the total energy is equivalent to solving the corresponding Euler-Lagrange equation, which is performed by means of a gradient descent approach [2]. This approach results in a force balance between the internal forces \vec{F}_{int} and the potential forces \vec{F}_{pot} , defined for each node \mathbf{n} of the surface mesh as:

$$\vec{F}_{int}(\mathbf{n}) = w \cdot \vec{F}_{pot}(\mathbf{n}), \text{ with } \vec{F}_{int} = \alpha \cdot \vec{F}_e(\mathbf{n}) + \beta \cdot \vec{F}_r(\mathbf{n})$$

The *internal forces* restrain excessive surface deformations. The computation of these forces depends on surface derivatives. These derivatives are approximated using finite differences on the surface mesh nodes, in accordance with the work by Ahlberg et al. [2]. On the one hand, an elasticity force \vec{F}_e limits stretching of the surface. For each node, the average position of directly neighboring nodes determines the elasticity force, moving the node towards the neighborhood center. On the other hand, a rigidity force \vec{F}_r restrains inordinate bending. The averages of the first and second order neighboring nodes determine the rigidity force, moving nodes along the surface normal to restrain excessive bending. The stretching and bending constraints are imposed on all nodes of the surface mesh, yielding a smooth surface.

The *external forces* pull the surface towards the boundaries between blood-flow regions and stationary tissue. The direction and strength of these potential forces \vec{F}_{pot} are extracted from an attractor image, which is based on the various features introduced in subsection 5.2.2.

Figure 5.4a and 5.4b respectively depict BFS and EVC intensity profiles across a vessel. By incorporating the blood-flow direction, the EVC feature provides relatively sharp edges at the lumen boundaries. In contrast, the BFS shows a gradual intensity change. As a result, the EVC provides a better indication of the desired boundary locations. This likewise holds for the direction-based LPC feature.

As opposed to the feature volume, an attractor image requires maximal intensities at the desired boundaries, with a gradual intensity change towards these boundaries. Therefore, an elementary gradient-based edge detection is applied to the BFS and the EVC feature. By definition, the LPC feature is directly applicable as attractor image.

The potential forces are taken from the attractor image, for each node on the surface mesh. The surface is attracted towards the boundaries. To the best of our knowledge, coherence-based features have not been employed before, to define the external forces of an active surface model.




Figure 5.4 Schematic overview of the external forces. (a) The intensity profile across a vessel shows a slanted slope at the boundaries for the BFS. (b) By including direction information, the EVC feature defines sharper boundaries.

The active surface model is parameterized by three weighting parameters α , β and w . With these parameters, the user can steer the force balance, where α and β increase surface smoothness and w increases the influence of the blood-flow information. All parameters were scaled to lie within the $[0,1]$ interval. The active surface model is successively adjusted until the system reaches an equilibrium, or arrives at a maximum number of iterations.

5.2.4 Validation

For the validation, a number of 2D-cine PC-MRI slices are available for each 3D-cine PC-MRI data set. The slices are registered with the volumetric data, as illustrated in figure 3.6. These 2D-cine slices are acquired perpendicular to the vessel of interest, measuring the through-plane flow as a scalar value. These slices have a higher spatial and temporal resolution as compared to the 3D-cine blood-flow data. Therefore, the 2D-cine slices can be used as a local approximation of the ground truth. The validation was performed using four volunteer studies, inspecting two 2D-cine velocity-encoded slices per study. One slice was oriented perpendicular to the aorta, and one slice was oriented perpendicular to a region of the pulmonary artery.

Obtaining an approximate 3D ground truth segmentation is difficult. Acquisition of higher resolution volumetric flow data is not feasible with current MRI sequences. Moreover, manual segmentation of the 3D luminal geometry is tedious and time-consuming, and therefore, not feasible for larger validation studies.

The blood-flow lumen is first manually segmented on a 2D-cine slice, similar to segmentations performed in clinical practice. This results in a reference contour that captures the blood-flow region, described by a point set in patient coordinates. Next, the cross section of the active surface and the 2D-cine slice is determined, resulting in a point set that describes the cross section contour. Lastly, the correspondence between the two discrete contours needs to be computed. We have adopted the mean distance-to-closest-point (DCP) measure, which provides an intuitive distance in millimeters. As the relation between the contours is not one-to-one, the point correspondence is determined bi-directionally and subsequently averaged [86].

Table 5.1 provides an overview of the validation results. For each of the inspected regions, the distance between the automated and the manual segmentation is provided in millimeters, based on the DCP metric. To demonstrate the feasibility of our active surface approach, the optimal segmentation was obtained on a case-by-case basis. The unitless weighting parameters α , β and w were determined empirically.

The results show that the average distance between the active surface segmentation and the reference contour is generally less than 2.5 mm, which is smaller than the voxel size in the blood-flow data. All values larger than 2.5 mm are colored in red in table 5.1. This shows that both coherence measures, i.e., LPC and EVC, provide a better attractor image for the active surface model, and hence that the blood-flow direction information is worthwhile to incorporate in the segmentation. Furthermore, we observe no significant performance difference between the coherence measures.

Table 5.1 Results on volunteer aorta (Ao) and pulmonary artery (PA). The unitless parameters α , β and w , as defined in subsection 5.2.3, are in the $[0,1]$ range. The average contour distance is measured using the distance-to-closest-point in millimeters. Distances larger than the voxel size ($\approx 2.5 \text{ mm}^3$) are shown in red.

#	ROI	BFS (α, β, w)	LPC (α, β, w)	EVC (α, β, w)
1	Ao	1.80 (0.03, 0.05, 0.15)	1.55 (0.03, 0.10, 0.14)	1.29 (0.11, 0.06, 0.10)
2	Ao	3.04 (0.03, 0.07, 0.25)	2.07 (0.07, 0.03, 0.10)	1.78 (0.12, 0.10, 0.15)
3	Ao	2.92 (0.07, 0.01, 0.12)	2.44 (0.02, 0.01, 0.20)	1.95 (0.07, 0.04, 0.20)
4	Ao	2.31 (0.05, 0.03, 0.08)	2.01 (0.06, 0.03, 0.09)	1.85 (0.07, 0.04, 0.04)
1	PA (trunk)	3.17 (0.10, 0.10, 0.30)	1.38 (0.03, 0.10, 0.14)	1.47 (0.11, 0.06, 0.10)
2	PA (left)	2.78 (0.07, 0.05, 0.35)	1.81 (0.10, 0.03, 0.15)	1.88 (0.10, 0.03, 0.20)
3	PA (right)	2.99 (0.10, 0.02, 0.20)	1.76 (0.05, 0.03, 0.15)	1.97 (0.07, 0.04, 0.12)
4	PA (left)	1.97 (0.02, 0.02, 0.25)	1.64 (0.02, 0.01, 0.10)	1.52 (0.05, 0.03, 0.35)

5.2.5 Discussion and conclusions

To perform the active surface segmentation, an iso-threshold is required to extract the initial surface. In general, the iso-threshold is defined conservatively, imposing a reasonable amount of undersegmentation. This minimizes the chance that segmented structures overlap, due to partial volume effects. There is no need for a time-consuming search for the optimal iso-threshold. After determining the initial surface, the active surface segmentation performs automatically.

The results of the active surface approach are affected by the limited spatial resolution, and associated partial volume effects. Therefore, smaller arteries are challenging to segment. In addition, the active surface technique requires that the blood-flow speed is sufficiently fast to be captured by the initial threshold.

The presented validation, based on 2D-cine PC-MRI slices, is restricted to the regions that can be evaluated. For instance, bifurcations are difficult to assess quantitatively. However, qualitative inspection shows promising results for bifurcating regions. Segmentation of complex morphologies is supported by the well-defined initial surface and the topological stability of the active surface method, although further testing on different pathologies is required.

In conclusion, we have presented an active surface segmentation of the blood-flow lumen, using the full blood-flow velocity information. Active surface models are computationally efficient, in particular when the initial surface provides a good approximation of the end result. After roughly determining the initial surface by visual inspection, the algorithm executes on the order of seconds. Precomputation of the attractor images also takes merely seconds, and is required only once.

Four volunteer data sets were validated, locally assessing the quality of the segmentation under varying conditions. The results, summarized in table 5.1, show that the commonly used BFS feature performs poorly for the envisioned segmentation, in contrast to both coherence-based features.

In the future, the mesh quality can be improved. The internal forces are currently ill-defined at the volume boundaries, where the mesh is truncated. In addition, the influence of the implicit smoothing on the segmentation accuracy should be investigated further. Moreover, a larger evaluation study may be carried out, including volunteer and patient data. This should result in a global parameter set, suitable for the majority of data sets. Prior to such a study, it is worthwhile to investigate the automatic selection of the initial iso-threshold. The manual reference segmentations can be improved by including multiple expert segmentations per region.

5.3 Clustering

5.3.1 Motivation

Abstract representations of time-resolved volumetric blood-flow fields can be obtained through hierarchical clustering of the velocity data. Hierarchical clustering enables intuitive level-of-detail selection. Lower levels are closer to the original field, and higher levels provide a sparse overview of the velocity field, as depicted in figure 5.5. Typically, higher levels of abstraction yield comprehensive visualizations, reducing visual clutter and occlusion, while preserving important flow characteristics.

The method is fully data-driven, and does not rely on specific domain knowledge. Furthermore, no a priori selection of the number of clusters is needed. This is beneficial for our application, as the number of clusters depends on the dimensions of the data set, as well as the complexity of the blood-flow patterns. Hierarchical clustering is scalable to high-dimensional data, and hence suitable for time-resolved velocity fields.

Clustering methods require a dissimilarity measure. Two distinct dissimilarity measures were adapted for the spacetime domain, as described in subsection 5.3.2. Subsequently, we present our coarse hierarchical clustering approach in subsection 5.3.3. This method efficiently clusters unsteady velocity data in the spacetime domain.

Visualization of the cluster boundaries is not effective, as they have no explicit physical meaning. Instead, we propose a patharrow visualization to represent the spatiotemporal clusters. Several visualization approaches were investigated to convey the un-




Figure 5.5 Hierarchical clustering enables intuitive level-of-detail selection, without a priori knowledge of the data domain. Higher levels provide a sparse and abstract overview of the data.

steady nature of the blood-flow field, as described in subsection 5.3.4. The results were validated using artificial data. Both qualitative and quantitative validation of the clustering results will be described in subsection 5.3.5.

Although our methodology is mainly motivated by PC-MRI blood-flow measurements, no domain-specific knowledge is employed. The approach can therefore be applied generically to other unsteady vector fields.

5.3.2 Spatiotemporal dissimilarity measures

An elementary entity to represent a flow cluster is the average \bar{v} of all velocity vectors in a cluster. Subsequently, the dissimilarity between clusters can again be measured by an average. However, such an approach does not take into account position information, nor does it consider specific flow patterns. Hence, two more elaborate dissimilarity measures are used, which combine velocity and position information.

Elliptical dissimilarity - The elliptical dissimilarity measure, introduced by Telea and Van Wijk [222], includes spatial information using the positional average \bar{x} . An example clustering with this measure is depicted in figure 5.6b. The elliptical measure d_{ell} combines a velocity dissimilarity d_v , and a positional dissimilarity d_p in a linear fashion, controlled by a weighting parameter ω :

$$d_{ell}(\bar{x}_1, \bar{v}_1, \bar{x}_2, \bar{v}_2) = \omega d_p(\bar{x}_1, \bar{x}_2) + (1 - \omega) d_v(\bar{v}_1, \bar{v}_2).$$

The velocity dissimilarity d_v is similar to a Euclidean distance, with the exception that a higher cost is associated with flow deviations perpendicular to the average cluster flow. This anisotropic behavior is achieved by modeling the iso-error contours around the flow vector with an ellipse.




Figure 5.6 Hierarchical clustering yields different results, when varying dissimilarity measures are employed. (a) An example artificial vector field is clustered using: (b) the elliptical dissimilarity measure, with $\psi = 0.9$, and $\omega = 0.9$, and (c) the local linear expansion dissimilarity.

The positional dissimilarity d_p is defined accordingly. The iso-error contours describe the positional dissimilarity, modeled with an ellipse oriented along the velocity vector. The aspect ratio of the ellipse is defined by a parameter ψ .

For time-varying velocity data, positions are defined in spacetime, and hence include a time component: $\bar{\mathbf{x}}_{4D} = (x, y, z, t)$. The dimensional disparity is solved by extending the velocity component with a constant, transforming the component $\bar{\mathbf{v}}_{4D} = (u, v, w, 1)$ into the spacetime domain.

Local linear expansion dissimilarity - Instead of using the average positions $\bar{\mathbf{x}}$ and velocities $\bar{\mathbf{v}}$, local linear approximations of the flow field may be used to preserve flow patterns. A local linear expansion dissimilarity measure was presented by Carmo et al. [31], aiming for a linear model that estimates the velocities \vec{v}_i within each cluster. An example clustering with this measure is depicted in figure 5.6c. The linear model describes the relation between the voxel positions \mathbf{x}_i and the modeled velocity \vec{v}_i as:

$$\vec{v}_i = \mathbf{A}\mathbf{x}_i - \bar{\mathbf{v}},$$

where \mathbf{A} is a matrix describing the velocity gradients, estimated by a least squares fitting over N cluster voxels. For a cluster C , the total squared error ε is given by:

$$\varepsilon(C) = \sum_{i=1}^N \|\vec{v}_i - \vec{v}_i\|^2, \text{ where } \vec{v}_i \in C.$$

The dissimilarity measure d_{lin} is defined by the cost to merge two clusters, C_1 and C_2 , into a new cluster C_{merged} [31]:

$$d_{lin}(C_1, C_2) = \varepsilon(C_{merged}) + (\varepsilon(C_1) + \varepsilon(C_2)).$$

To the best of our knowledge, this measure has not been applied to time-varying velocity data before. To achieve this, we have extended matrix \mathbf{A} to include both spatial and temporal gradients, resulting in a 4×3 matrix that captures the relation between the spatiotemporal position \mathbf{x}_i , and the three-dimensional velocities \vec{v}_i .

5.3.3 Coarse hierarchical clustering

Prior to clustering, a few preprocessing steps are required. First, a tMIP volume is generated, as described in subsection 5.1.2. Next, an iso-threshold is applied to the tMIP volume, creating a voxel mask that includes all voxels considered for clustering. Subsequently, the blood-flow velocities, and the voxel mask, are used to perform the hierarchical clustering.

The hierarchy is described by a precomputed binary cluster tree, which provides the basis for our visualization. A binary cluster tree of the velocity data can be constructed by a top-down [88] or a bottom-up approach [222].

Top-down division can be performed in various ways, leading to different results. For our application, we employ a bottom-up approach, consistently generating the desired hierarchy. We initiate the clustering, starting from the data voxels, which are consequently defined as the leaves of the cluster tree. Iterative merging of individual clusters subsequently provides the tree nodes, and is based on dissimilarities between spatiotemporally neighboring clusters, as described in subsection 5.3.2.

Processing of high-dimensional data inherently leads to a significant computational cost. Thus, clustering methods are generally applied to relatively small two-dimensional and three-dimensional examples. To obtain an acceptable computational performance for actual measured data, the computational complexity should be reduced, without downscaling the data dimensions by plain subsampling.

A major performance bottleneck during the generation of the hierarchy, is the repetitive search for the smallest dissimilarity. This operation has a complexity of $O(n)$, and the number of iterations likewise scales with $O(n)$, where n is the number of clusters to be merged. The overall complexity therefore becomes $O(n^2)$ for a straightforward implementation. Especially the initial phase of the clustering is computationally expensive, as the number of clusters is still large.

Therefore, we introduce the coarse hierarchical clustering (CHC) approach. Instead of merging a single pair of clusters per iteration, as depicted in figure 5.7a, the CHC method merges multiple clusters per iteration, as depicted in figure 5.7b. An incremental dissimilarity threshold d_{th} is introduced, causing all cluster pairs with a dissimilarity smaller than this threshold to merge simultaneously. For each iteration, the threshold is increased with a factor Δd_{th} . This reduces the computational-complexity by lowering the number of iterations, as well as the number of dissimilarity computations.




Figure 5.7 (a) Schematic depiction of a bottom-up generation of a full cluster hierarchy. Tree nodes are merged, iteratively processing the two clusters with the smallest dissimilarity. (b) Schematic depiction of a coarse bottom-up generation of the cluster hierarchy, requiring half the number of iterations in comparison to a full cluster tree generation.

Table 5.2 Overview of the complexity of various container operations that are performed on the set of dissimilarities, associated with the current cluster set. These operations are part of the hierarchical cluster process, where n is the number of clusters to be merged.

Operation	Dynamic array	Linked list	Balanced tree
Lookup	$O(n)$	$O(n)$	$O(1)$
Insertion	$O(1)$	$O(1)$	$O(\log n)$
Deletion	$O(n)$	$O(1)$	$O(1)$

Experiments were performed with varying Δd_{th} values, and the results were visually compared with the full cluster tree. For our application, the algorithm approaches a linear complexity for a threshold $\Delta d_{th} = 0.01$. This threshold yields a good balance between computational cost and cluster quality, for both dissimilarity measures.

Merging of multiple clusters per iteration affects the resulting tree, and hence provides an approximation of the true data hierarchy. However, it is fair to assume that small variations in the order of clustering hardly affect the cluster tree structure. Whenever clusters within an iteration are neighbors, small deviations will occur. These deviations depend on the user-defined threshold increment Δd_{th} .

In addition to the algorithm complexity, performance of the clustering also depends on other aspects. For instance, we assume a substantial amount of background information in the data, which is discarded prior to processing. Furthermore, several implementation aspects greatly influence the overall performance; particularly the chosen data structure to store the dissimilarity values. The operations performed on the data container include dissimilarity insertion, deletion, and lookup of the minimum value. The operation complexity varies per container, as summarized in table 5.2.

On the one hand, we consider two unsorted containers: the dynamic array and the linked list. The dynamic array is discarded, based on the relatively expensive lookup and insertion complexity. On the other hand, we consider a sorted balanced tree. In comparison to the linked list, insertion is more expensive, due to an additional search operation. However, lookup and deletion are efficient.

We have implemented the clustering with these container types, employing the Standard Template Library (STL) for the C++ programming language. The implementation confirms the expected theoretical complexity $O(n^2)$ for the linked list, however the balanced tree did not achieve the complexity of $O(n \log n)$, and was actually close to a quadratic complexity. In practice, the best performance was obtained by a custom implementation of a linked list.

5.3.4 Visualization

Visualization of the generated clusters relies on a user-defined level-of-detail. Comprehensive visualizations are mostly obtained at the top-most levels. This level-of-detail selection is the only essential parametrization of the visualization, and should be ex-

plored by the user. Based on the selected detail level, the spatial structure of the clusters can be reconstructed for each time phase. The reconstructed clusters are commonly represented by labeled volumes.

Direct visualization of the cluster boundaries is not desirable, as they convey no evident physical meaning. Instead, we propose representative pathlines to convey the clusters. These lines are seeded at the center of each cluster. These cluster centers can be defined spatially, or in the spacetime domain.

The three-dimensional cluster centers can be generated from the reconstructed clusters per phase. An intuitive first approach would be to compute the center of mass. However, the cluster shape may be concave, and hence the center of mass may reside outside the cluster boundaries.

To guarantee that the cluster center is inside the cluster boundaries, we devise an approach that locates the point within the cluster C with the smallest distance to all N positions in that cluster. Using the Euclidean distance, the center position \mathbf{x}_m is defined by its index m as:

$$m(C) = \underset{i}{\operatorname{argmin}} \frac{1}{N} \sum_{j=1}^N \|\mathbf{x}_i - \mathbf{x}_j\|, \text{ where } \mathbf{x}_{i,j} \in C.$$

To avoid a disproportionate performance penalty, computing all mutual distances is avoided. For visualization purposes, the exact center is not required, and hence a stochastic sampling per cluster is performed before computing the center. Based on visual inspection, a minimum of hundred samples is sufficient for our application.




Figure 5.8 Simplified representation of the patharrow seeding. The horizontal axis shows the time phases, while the spatial domain is represented on the vertical axis. One phase is highlighted in light gray, and indicated with $t = 1$. (a) For 3D seeding, a spatial center is computed for each existing cluster, and short pathlines are traced interactively. (b) For 4D seeding, a spatiotemporal center is computed, and long pathlines are precomputed.

The clustering is performed in four dimensions, and hence the clusters reside in space-time. However, a physically meaningful definition of the cluster centers in 4D is non-trivial. This is due to the disparity between space and time, in terms of units and dimensionality. To prevent an unsubstantiated coupling between space and time, the temporal center is first selected using the median of the time spanned by the 4D cluster. Subsequently, the spatial center is computed, similar to a 3D cluster center.

Based on the cluster centers, integral lines are used to represent the clusters. From each cluster center, a single pathline is traced both forward and backward in time. The integral lines are depicted as arrows in the direction of the blood flow. These patharrows are rendered in real-time, and are composed from illuminated tuboid imposters, with arrowheads and end caps.

Starting from a center point, we devise two approaches to generate the patharrows:

3D-seeded arrows - Cluster centers defined in 3D have no explicit temporal coherence. The spatial centers are computed for all clusters that exist at the inspected phase, without considering the temporal extent of the cluster. This is demonstrated in figure 5.8a. This results in a relatively dense seeding, in comparison to the use of 4D centers. Naturally, the visual density also depends on the selected level-of-detail.

At the inspected phase, short patharrows are traced in real-time for each cluster center. The resulting renditions are not suitable for animation, due to the lack of temporal coherence between the 3D cluster centers. However, this approach provides a comprehensive overview of the structure of the blood-flow at a given phase.

4D-seeded arrows - Alternatively, the sparser set of 4D cluster centers can be employed to seed the pathlines in a temporally coherent fashion. In this case, long pathlines are precomputed and used to represent each cluster, as depicted in figure 5.8b.

Based on this set of precomputed pathlines, we present several visualization styles. First, the whole set of pathlines can be visualized directly using patharrows. This provides a static overview of the full spatiotemporally clustered flow data. To obtain a sufficiently sparse representation, a relatively high level-of-detail is required. The temporal characteristics can be regained by color mapping the time component, as depicted in figure 5.9a. Alternatively, the blood-flow speed can be encoded by color, as depicted in figure 5.9b, losing temporal information.

Because of the temporal coherence, the static lines are applicable for animation. During animation, we highlight parts of the lines that reside within a certain time window around the animated phase. One way is to maintain the direct visualization of the full set of patharrows, rendering it with a fully desaturated color. The parts within the sliding window around the animation phase are emphasized by color, as depicted in figure 5.9c. Throughout the animation, this sliding window moves along the patharrows, highlighting flow patterns over time.

In this case, the overall structure and temporal behavior are clearly conveyed, while visual occlusion impedes the effectiveness at lower detail levels. An alternative is to




Figure 5.9 Close-up of 4D-seeded patharrows. (a) A time color coding reveals the temporal character. (b) Color may also show the blood-flow speed. (c) The pathlines are used for animation. The color is desaturated, except for a sliding window around the animation phase. (d) Small patharrows may also represent the sliding window. Note the color correspondence between (b), (c), and (d).




Figure 5.10 Blood flow of a healthy volunteer depicted by 3D-seeded patharrows with anatomical context. (a) A cel shaded surface with contours based on curvature information. (b) Opacity modulation of the frontfaces ensures visibility of the patharrows.





Figure 5.11 (a) Aortic dissection with pathological flow, represented by 250 clusters using the elliptical dissimilarity measure. The 4D-seeded patharrows convey the spatiotemporal structure, revealing different blood-flow patterns. (b) A CFD simulation of a cerebral aneurysm, depicted by static patharrows, based on 4D seed positions. A high-speed in- and outflow is observed from the parent vessel. Recirculating patterns occur within the aneurysm, which are more clearly conveyed through animation. The wall-shear stress (WSS) is color coded on the surface.

depict the sliding window as a small patharrow, omitting the remainder of the pathline that resides outside of the sliding window, as depicted in figure 5.9d. Observe that these two approaches are essentially identical, with respect to the emphasized regions in figure 5.9c and figure 5.9d.

In addition to the flow visualization, we provide an anatomical context, as depicted in figure 5.10a. The context is based on an iso-surface, extracted from the tMIP volume. To avoid visual occlusion, frontfaces are rendered using a view-dependent opacity modulation [71]; exemplified by figure 5.10b. The backfaces are rendered with cel shading, using receding pastel colors, analogous to practice in medical illustrations. Additionally, the context contains contours, based on curvatures on the surface.

5.3.5 Validation

For the validation, three distinct data sets were considered. The first data set comprises a PC-MRI blood-flow measurement of a healthy volunteer (Fig. 5.10b). The second data set concerns a patient suffering from an aortic dissection (Fig. 5.11a). The slow flow in the false lumen is not captured by the tMIP, and is hence omitted by the clustering threshold. The third data set consists of a CFD simulation of the blood flow in a cerebral aneurysm (Fig. 5.11b), demonstrating the generality of the clustering approach. The temporal resolution has the same order of magnitude, while the spatial resolution is considerably better than for the PC-MRI examples.

Qualitative validation - To attain meaningful visualizations, several requirements are imposed on the spatiotemporal clusters. The clusters need to be spatially coherent, compact, and shaped along the blood-flow direction. The linear expansion dissimilarity d_{lin} yields compact clusters, which are spatially coherent, and capture patterns such as bifurcations. However, the cluster shape does not necessarily align with the flow direction. Instead, the elliptical measure d_{ell} shows compact clusters, which can be parameterized [222]. The positional part ensures spatial coherence, while the velocity dissimilarity leads to clusters that are shaped along the blood-flow. The preferred parameters were determined empirically: $\psi = 0.9$, and $\omega = 0.9$.

Blood-flow patterns were further inspected using the elliptical dissimilarity. As expected, the volunteer data primarily shows laminar flow behavior. In contrast, several anomalous patterns are detected for the pathological case, as depicted in figure 5.11a. The blood-flow interacts with the surrounding tissue, and hence relations between blood-flow patterns and the vessel wall are important. Figure 5.11b shows patharrows, in relation to the context surface, color coded with the wall-shear stress.

Quantitative validation - The noise robustness was evaluated quantitatively. The noise found in the velocity data follows a Rician distribution, with a signal-to-noise ratio ranging between 10 and 25. This can be approximated by Gaussian noise. The experiments are performed on artificial flow data, varying the standard deviation of the additive Gaussian noise distribution. The cluster results are compared with the ground truth using the adjusted Rand index [93]. For the linear expansion dissimilarity, we observe a varying adjusted Rand index, while the elliptical dissimilarity is stable.

Table 5.3 Performance results of clustering and visualization. For the elliptical measure d_{ell} ($\psi = 0.9$, $\omega = 0.9$), a full hierarchy generation and CHC approximation ($\Delta d_{th} = 0.01$) is presented. The framerates are based on a cluster set at 99.9% of the hierarchy.

	Healthy volunteer	Aortic dissection	Cerebral aneurysm
Dimensions	$144 \times 144 \times 50$	$144 \times 144 \times 60$	$144 \times 128 \times 64$
Resolution (mm)	$2.1 \times 2.1 \times 2.7$	$2.1 \times 2.1 \times 2.5$	$0.2 \times 0.1 \times 0.3$
Cardiac phases	20	25	50
Iso-threshold	5%	5%	5%
Max speed cm/s	180	240	250
Cluster phases	[1,12]	[1,12]	[1,10]
Cluster nodes	621.8K	744.7K	589.9K
Cluster d_{ell} full	08h:43m	12h:08m	07h:33m
Cluster d_{ell} CHC	58s	67s	53s
Cluster d_{lin} CHC	02h:20m	02h:02m	01h:04m
Cluster count	366	890	356
Framerate 3D seeded	23.7	19.8	16.6
Framerate 4D seeded	23.8	19.9	17.1
Figures	5.10	5.11a	5.11b

Furthermore, we have carried out a performance evaluation of both the clustering, as well as the rendering. The measurements were performed on a conventional computer system, using a dual core processor, with 6GB internal memory, and an NVidia GeForce 570 GTX graphics card. The selected voxels were clustered using different dissimilarity measures. The results of the evaluation are summarized in table 5.3.

For all data sets, the systolic and early diastolic phases are clustered in four dimensions, capturing the most important dynamics. Larger data sets currently lead to issues with the memory footprint. Alternatively, systole and diastole could be clustered separately.

For the elliptical measure d_{ell} , we observe significant improvements in the clustering computation time using the CHC approach, in comparison to a full tree generation (Table 5.3). The linear model dissimilarity d_{lin} is computationally more expensive, in particular for noise-prone PC-MRI data. All visualizations render at interactive framerates, including the anatomical context (Table 5.3). The level-of-detail was set to 99.9% of the hierarchy tree, rendering a relatively small number of nodes, close to the top of the tree. Lower levels lead to more clusters, resulting in lower framerates. All meaningful renditions achieve framerates well over fifteen frames per second.

5.3.6 Discussion and conclusions

Clustering - The presented data-driven approach generically clusters the unsteady vector data. The domain-specific knowledge is emerging, and hence no explicit knowledge exists to identify clusters in the unsteady blood-flow data. Therefore, we rely on properties of the blood-flow field.

Although the cluster boundaries represent regions of coherent velocities, they do not truly exist in the continuous flow field. They were not visualized to avoid false interpretation of the boundaries as physical entities in the field. Consequently, coverage of the patharrows may be hard to derive, especially in regions with opposing arrows.

The coherence depends on the dissimilarity measure. The definition of such a measure is complicated. A non-trivial weighting between space and time is required, as the units of these dimensions cannot be compared.

Visualization - Related work employs pathline visualizations, similar to our patharrows. Clinically-oriented papers show seeding throughout the cardiovascular anatomy [154]. However, the densely sampled seed-points are positioned randomly. Instead, the clustering provides a data-driven seeding approach, for a coarse to detailed view.

An important aspect of the visualization is the level-of-detail selection. The selection is currently global, and does not provide local detailed representations on demand. Construction of labeled cluster volumes is performed when the level-of-detail is changed. With the used data structures, real-time localized level-of-detail changes need precomputed cluster label volumes, leading to excessive memory use.

The temporal behavior of the intrinsically unsteady flow data is captured using static and animated patharrows. Although animations impose a considerable cognitive load [231], the human observer is sensitive to the apparent motion in flow. Using animated short patharrows (Fig. 5.9d), the patterns attract the attention. The structure of the flow field can be investigated in more detail using static patharrows (Fig. 5.9c).

In conclusion, we have presented a coarse hierarchical clustering technique for four-dimensional flow data. The CHC approach achieves a near linear computational complexity, in contrast to the quadratic complexity of a naive implementation. Two dissimilarity measures were adapted for the spacetime domain.

We introduced 3D-seeded and 4D-seeded patharrow visualizations, combined with an illustrative anatomical context. The 4D-seeded patharrows allow for a static representations and animations of the spatiotemporal information.

The clustering and visualization was applied to time-resolved volumetric blood-flow data, either measured through PC-MRI, or simulated by CFD. Different flow patterns could be distinguished. These patterns can be related to properties of the vessel wall.

In the future, the incremental threshold for the CHC approach could be adjusted during the clustering. Alternatively, parallelization of the clustering could further improve the performance. Interactive level-of-detail selection, together with associated visualization techniques, could provide details on demand.

The current visualizations depend on camera interaction and animation to reveal the spatial relations. Different depth cues, such as shadows and halos, may be investigated. Additional information could be visualized, such as the intra-cluster variance.

Moreover, it is worthwhile to investigate application-specific approaches. Specific dissimilarity measures could cluster regions with high vorticity, or low coherence of the blood-flow field. The clustering results and associated visualizations should be evaluated, measuring the value for clinical research, and potentially future diagnosis.

“When using comics [visualizations, ed.] as a teaching tool, there are two ways to fail: form apologizing for content, or content apologizing for form.”

Scott McCloud

Blood-flow inspection The thoracic arteries



This chapter is based on:

“Exploration of 4D MRI Blood-Flow Using Stylistic Visualization” Roy van Pelt, Javier Oliván Bescós, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, Bart ter Haar Romeny, and Anna Vilanova. In *IEEE Transactions on Visualization and Computer Graphics* (2010), 16(6):1339-1347.

6.1 Motivation

The introductory chapters of this thesis substantiated the need for interactive visualization techniques to analyze measured blood-flow data. In particular, section 4.3 addressed the difficulties involved with visualization of unsteady velocity fields, such as the PC-MRI blood-flow data.

Direct visualization in three-dimensions, such as hedgehogs, vector plots, or texture-based techniques, are often difficult to interpret, because of the considerable amount of visual clutter. Moreover, direct representations are typically susceptible to the uncorrelated noise that is present in the MRI measurement.

Instead, specialized visualization techniques are necessary, reducing the communicated information, while retaining essential characteristics of the flow. For instance, the quantity of conveyed visual information can be reduced by local inspection, communicating the blood-flow behavior in a delimited area only. Variation of the inspected area subsequently enables exploration of the intricate blood-flow velocity data.

6.2 Probing

6.2.1 Appearance

For local inspection of the blood-flow velocity data, we introduce a probe that enables physicians to inspect parts of the blood flow in the main vessels surrounding the heart. Generally, physicians know well which regions of the blood flow are relevant to analyze.

The probe comprises a local delineation of the vessel of interest. In particular, we focus on the main thoracic arteries. The appearance of the probe includes an elementary polyline, capturing the vessel circumference.

6.2.2 Positioning

Positioning of the vessel probe comprises a single mouse-click interaction. Using conventional viewpoint interaction, the physician is able to navigate through the anatomical structures and locate the region-of-interest. Subsequently, a mouse click on the depicted vessel structure initiates the computation of a cross-sectional plane.

After each mouse click, the depth buffer is queried at the window coordinates of the designated position, obtaining a three-dimensional position in the camera coordinate frame. After transformation to the patient coordinate frame, the orientation of the vessel is determined from the MRI acquired data, involving an eigen-decomposition of the structure tensor [118]. Subsequently, the cross-sectional perimeter is determined in a two-dimensional plane, based on a full-width at half-max edge detector [20].

This detector determines the best contour fit for a limited number of scales, robustly detecting vessels with varying diameters. The approach assumes tubular structures, and can not reliably probe cross sections of the heart chambers.

Generally, the vessel cross-section probing relies on anatomical data, acquired by MRI. In practice, anatomical data is acquired for one, or at most two phases of the cardiac cycle, impeding tracking of the cross section over time. Consequently, a static probe is employed for visual inspection. For current clinical research, often no anatomical data is acquired at all, saving valuable acquisition time. In these cases, a tMIP volume is used as an approximate static anatomy (see subsection 5.1.2).

6.3 Seeding

6.3.1 Seeding strategies

The vessel cross-section, defined by the probe, is employed as a seeding plane for geometry-based flow visualization. Distribution of the seeds imposes structure on the geometric primitives, such as lines or surfaces, and hence the seed positions lead to varying visual results. As the blood-flow patterns are of primary interest, often the aim is to reduce the effect of the seed distribution on the geometric primitives.

Associated with a seeding distribution, we propose so-called seeding templates, as depicted on the vertical axis in figure 6.1. These templates are used to change the seeding density along the vessel radius. Dependent on the pathology, physicians may be interested in the flow behavior within the vessel, or near the vessel wall. Therefore, seeding templates allow to focus on the flow behavior near the vessel center (central) or the vessel wall (peripheral), by changing the seeding density. Conceptually, one may understand these templates to be radial transfer functions, varying the seeding density based on the distance from the center of the vessel cross-section.

6.3.2 Seeding distributions

Figure 6.1 furthermore presents an overview of the seeding distributions that we have investigated. The radial, circular and rectilinear seeding strategies yield orderly structured seed distributions, transferring their character to the geometric primitives. Instead of the flow dynamics, merely the seeding structure is perceived when inspecting the resulting primitives. Therefore, seed positions are mostly distributed randomly.

There are various ways to generate a random distribution, yielding different flow-visualization results. An example is provided by the two distinct random distributions in figure 6.1. The first random distribution is based on a uniform sampling of the vessel cross-section, while for the second distribution seed points were randomly omitted from a rectilinear distribution.




Figure 6.1 Seeding approaches impose structure on the geometry-based flow visualizations. Different seeding distributions can be inspected. Furthermore, different seeding templates can be selected to differentiate the seed density with respect to the vessel center.

For analyzing the blood-flow velocity field, non-random strategies are valuable to investigate as well. For instance, the radial seeding strategy provides a denser seeding towards the center of the vessel cross-section. It is oftentimes valuable to emphasize the blood-flow behavior near the vessel center, as the blood-flow velocity profile typically has a peak velocity near the vessel center. Furthermore, we propose a circular seeding strategy, distributing the seeds according to equally-spaced concentric circles, creating a set of nested tubes after integration. While the direction of the integration is strongly affected by the seeding structure, the accumulated speed information can be observed from the circular profiles that arise at the other end of the primitives.

6.4 Visualization

6.4.1 Planar reformat

The probe captures a region-of-interest within a vessel, and therefore provides the basis for a local visualization. The first visualization technique is based on customary multi-planar reformatting (MPR). MPR planes are omnipresent in clinical practice, and are used to depict the individual components of the blood-flow velocity vectors, the blood-flow speed, or a color Doppler inspired through-plane flow component [150, 233].

Instead of using a full planar reformat, which slices the bounding box of the volume, we propose smaller reformats at the designated vessel cross-sections. This approach, depicted in figure 6.2a, intuitively relates the flow information to the anatomical context and limits the amount of visual information to the demarcated regions-of-interest. A similar approach was implied by Frydrychowicz et al. [64], presenting flow related parameters as an overlay to the vessel structure.




Figure 6.2 Multi-planar reformat visualization. (a) Integrated planes at the vessel cross-section positions. (b) Exploded view of planes at the vessel cross-section positions.

Unfortunately, the obliquely positioned planes adversely affect the analysis of the flow data. We propose a solution through an exploded view technique, regularly used in the field of medical and technical illustrations (e.g., figure 4.3c). This technique, presented in figure 6.2b, aligns the planes with the current view direction, and depicts them alongside the related vessel cross-sections. The original position of a plane is depicted by an indicative contour, accompanied by dashed connection lines, supporting visual correlation of the view-aligned plane and its original position and orientation.

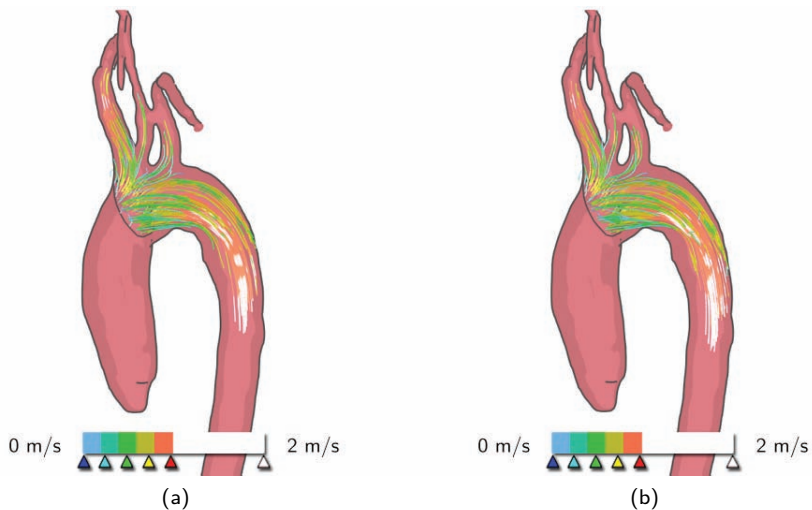


Figure 6.3 Integral line visualization based on random seed positions. (a) Streamlines. (b) Pathlines.

6.4.2 Integral lines

As described in subsection 4.3.5, tracing of line primitives is the prevailing geometry-based flow visualization method. The selected vessel-cross section is used as a seeding plane, utilizing the seeding approaches presented in section 6.3.

Tracing - Starting from the seed positions, different line primitives are traced using a fourth order Runge-Kutta integration scheme [188]. Line tracing is performed in real-time and can be parameterized interactively.

Streamlines represent the tangent curves of the velocity field at one point in time, as depicted in figure 6.3a. The streamlines are generated in real-time for each phase of the cardiac cycle, enabling physicians to inspect the temporal differences of the instantaneous flow-field structure. In addition, pathlines may be traced, as depicted in figure 6.3b. A pathline represents the trajectory of a massless particle through the flow-velocity field, enabling a physician to inspect the temporal flow behavior. A more extensive treatment of integral lines was provided in subsection 4.3.5.

As opposed to previous work, we devise an approach that generates restrained pathlines, which are generated for each phase of the cardiac cycle. Commonly, pathlines are traced throughout a large extent of the cardiac cycle, accumulating errors due to the numerical integration scheme. This certainly holds for the PC-MRI blood-flow data, with a relatively low temporal resolution of approximately fifty milliseconds, and with peak velocities of up to two meters per second. Moreover, long pathlines are more sensitive to the initial spatiotemporal seeding position.

Employing the probing approach, multiple cross sections can be interactively generated at fixed positions. For each cross section, streamlines or pathlines are traced with a restrained length, respectively integrating the lines in either the spatial or spatiotemporal domain. By placing sufficiently many cross sections and repeatedly tracing the line primitives for each phase, the total error decreases and the line primitives are generally less sensitive to seed positioning. For the pathlines, this results in a depiction of a set of confined particle trajectories for each phase, instead of a single set of very long trajectories throughout the cardiac cycle.

Line traces may also be used to inspect branching behavior of the blood-flow. For that purpose, Frydrychowicz et al. [63] proposed a reversed tracing approach, enabling inspection of the particle trajectory that arrives at the seed position under consideration. Examples of reversed-tracing are depicted in figure 6.4.

Visualization - A number of visualization techniques have been applied, to improve the perception of the line primitives, and the flow data they represent. Color is one of the most important visual cues to convey data characteristics. Therefore, we provide a set of predefined color maps, together with a transfer function editor, which allows to interactively change them. Typical color maps that have been included are the rainbow color map, as depicted in figure 6.4b, and a black-body color map. The most often inspected flow characteristic is the blood-flow speed, often depicted by

means of pseudocoloring. The blood-flow speed can be investigated by using linearly interpolated color maps, as well as regularly-banded color maps, as depicted in figure 6.3. The number of quantization steps can be selected interactively.

The rainbow color map is the de facto standard in current practice, using hue variations to distinguish blood-flow speed variations. This color map is generally applied without correcting for perceptual deficits [18]. Our transfer-function editor allows to interpolate color maps in both RGB color-space, as well as the perceptually more uniform CIELuv color-space [87]. Additionally, the editor allows the user to automatically approximate constant lightness in the CIELuv space.

Visual analysis of the blood-flow can be performed at different speed ranges, by interactively changing the color map. For example, physicians often inspect the existence of unexpected high-speed fluid streams, known as blood-flow jets. These jets can be easily detected by emphasizing the high-speed flows by choosing a salient color with respect to the chosen color map. An example of a color map that uses such a threshold is presented in figure 6.3. In addition, the transfer-function editor enables stepwise transparency modulation for different speed ranges.

Furthermore, the perception of the line primitives has been improved by super-sampling, anti-aliasing, and local illumination. The Phong reflection model is applied to the lines. The normal vector for each point of the line primitive is selected to be coplanar to the light direction and the local tangent direction [217].

In addition, we introduce animated highlights, emphasizing the particle trajectory that yields the line primitive under consideration. Using the probed cross sections, line primitives are traced for each phase of the cardiac cycle, spatially or spatiotemporally integrating flow profiles over time. Because a constant integration time is employed, a decrease of blood-flow speed over time will shorten the line primitives, which could be falsely interpreted as a reversal of the blood-flow direction. Hence, an animated visual cue is introduced to continuously indicate the blood-flow direction, as depicted in figure 6.5. The highlight is continuously animated, and is therefore not visually confused with the applied static illumination.

Additionally, the animated highlight provides a good indication of the relative blood-flow direction and speed between the line primitives. Assuming a forward tracing of the line primitives, the animated highlights initiate at the seed position in the cross-sectional plane. Subsequently, the highlights advect along the line primitive, fanning out due to mutual differences of the line primitives, which are based on variations in the blood-flow velocity field.

6.4.3 Arrows

The techniques presented in the previous subsections directly depict parts of the unsteady blood-flow velocity field. Other physically relevant and intuitive parameters can be derived from the blood-flow velocity field. Several informative parameters were listed in subsection 2.3.1.




Figure 6.4 Reverse integral line tracing. (a) Color per cross section, to distinguish the branching behavior. (b) Color determined by the line propagation, linearly mapping colors to the vertices from the seed at the cross section, until the end of the line.

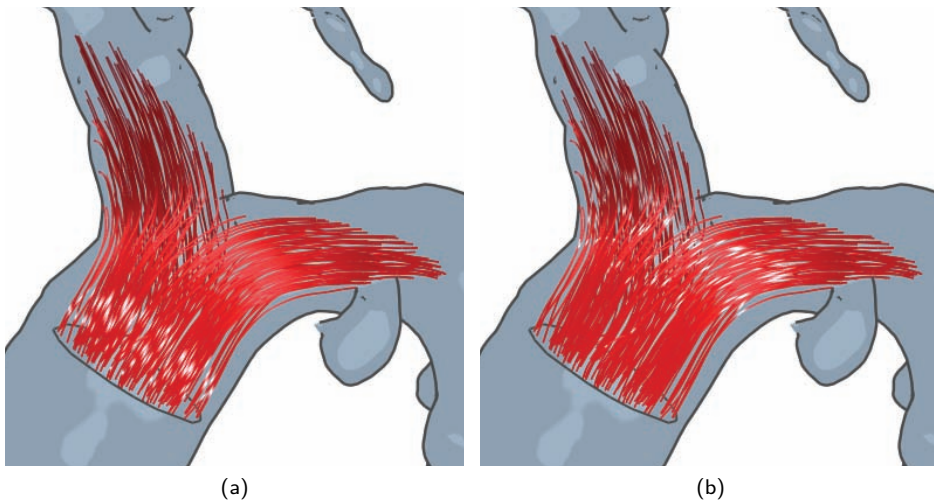


Figure 6.5 Integral lines with highlight animation, propagating along the flow direction. (a) Early propagation stage. (b) Late propagation stage.

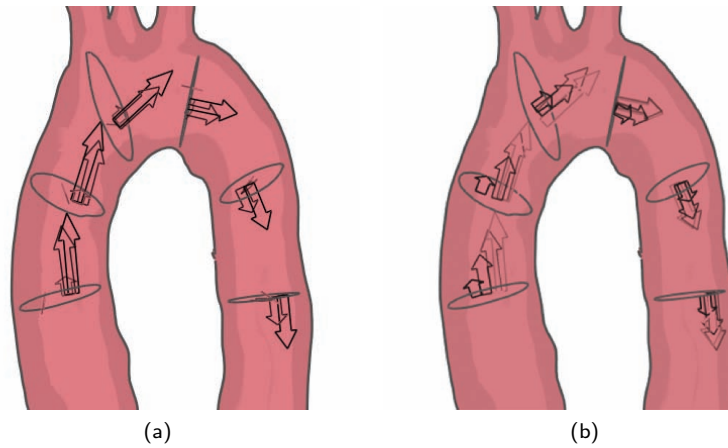


Figure 6.6 Flow-rate arrow trails. (a) Arrows at 140ms after the start of the cardiac cycle. (b) Arrows at 280ms after the start of the cardiac cycle.

An important derived parameter is the volumetric flow-rate, describing the volume of blood that passes through a vessel cross-section per second. This quantity, expressed in m^3/s or ml/s , is of particular interest whenever flow streams bifurcate. The flow-rates after the split should add up to the initial flow-rate value before the bifurcation. Hence, this parameter enables physicians to validate the branching of flow streams and to detect possible abnormalities, such as jets and leakages.

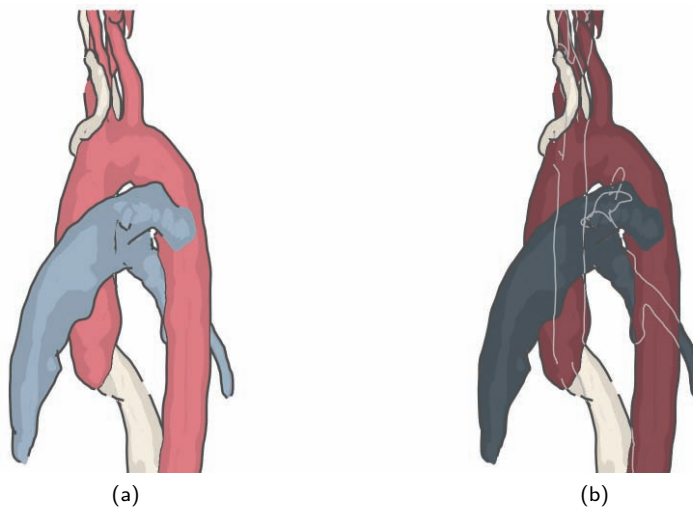


Figure 6.7 Illustrative anatomical context visualization. (a) Cel shaded inner vessel wall and occluding contours. (b) Hidden contours during viewpoint interaction.

Based on this parameter, we present an arrow visualization, as depicted in figure 6.6. The length of the arrow is determined by the flow-rate at the cross section. The flow-rate has a through-plane component and no direction. We use the origin and direction of the peak velocity within the cross-sectional area to determine the arrow positioning.

Commonly the volumetric flow-rate is inspected over time. To that end, the flow-rate arrows on each probed cross section can be animated. A motion trail of the arrows is included, capturing the temporal behavior. A motion-trail consists of a sequence of increasingly faded arrows, depicting flow-rates from the near past. Figure 6.6a depicts the increase in flow-rate after the contraction of the heart muscle, while figure 6.6b depicts the decrease of the flow-rate when the heart muscle relaxes.

6.4.4 Anatomical context

Exploration of the blood-flow data requires a patient-specific depiction of the anatomy, facilitating navigation to the regions-of-interest. With the primary interest in blood-flow dynamics, we allow an approximate representation of the anatomical structures. The anatomical context can be extracted from either a data set representing the anatomical structures, or a preprocessed tMIP volume.

The depiction of the anatomical structures includes a surface representation, typically extracted from a tMIP volume. A conventional segmentation technique, such as the marching cubes algorithm, suffices to extract the desired surface geometry from the tMIP data. For the figures in this chapter, third party software [273] was used to improve the automated segmentation. The vessel boundaries were manually retouched on a slice-by-slice basis. Furthermore, the vessel structures were separated, facilitating visibility changes of the structures and hence avoiding occlusion if necessary. With the focus on visualization, we have tolerated this laborious segmentation approach. Other methods endeavor an automated extraction of the full surface geometry from the MRI blood-flow data [107, 178], including our work presented in section 5.2 [237]. These methods typically segment the whole cardiovascular system, instead of the separate anatomical structures of interest.

Since the anatomy is presented as a contextual aspect of the visualization, details should be omitted without losing the necessary morphological information. Therefore, illustrative techniques are employed to depict the anatomical context, capturing the desired outline of the approximate anatomical representation.

Figure 6.7a depicts our anatomical context visualization, inspired by medical illustrations. The visualization is based on cel shaded surfaces [74], combined with superimposed occluding contours [50]. The surfaces continuously depict the inside of the vessel wall, by view-dependently rendering the rearmost triangles of the segmented structure. The cel shading provides the desired visual simplification, retaining a cue for visual depth and an outline of the morphological structure.

The outline of the anatomical structures is emphasized using contour lines. As described in subsection 4.2.2, extensive research has been conducted to investigate a variety of line structures that are effective at conveying shape. For the anatomical context, we use occluding contours, defined as the boundary between the visible and the hidden parts of the surface. These contours are generated by the set of points where the surface normal is perpendicular to the view direction. Complementary contour lines, such as suggestive contours, are deliberately omitted from the representation of the anatomical context. This context is a coarse approximation of the anatomical structure. Highlighting details of this structure would not provide relevant information to the physician.

The presented illustrative visualization of the anatomical context simplifies the representation, reducing depth perception and hence limiting insight into the spatial relations of the structures. In order to facilitate the vessel cross-section probing process, we introduce hidden contours that are visible during viewpoint interaction, as depicted in figure 6.7b. These hidden contours bypass occlusion and clarify spatial relations. The interaction is initiated when the left mouse button is pressed, and ends when it is released, leaving an uncluttered representation of the anatomy. After positioning the vessel cross-sections, the hidden contours may be occluded by the blood-flow renditions. However, they never impede the flow visualization.

6.5 Evaluation

As described in chapter 2, the PC-MRI blood-flow data provide a promising source of information for diagnosis and risk assessment of cardiovascular diseases. Nevertheless, many aspects of the blood-flow are still unknown. Visualization and interaction techniques support the analysis of the blood-flow data, which will lead to new insights and improve the understanding of the patient-specific hemodynamics.

To measure the value of the presented visual styles, we have performed an evaluation questionnaire with a group of four physicians that are actively involved with the acquisition of PC-MRI blood-flow data. The questionnaire is included in appendix B. Given the extent of the research field, this group of cardiologists and research fellows adequately represents the clinical cardiovascular blood-flow research community.

In the first part of the questionnaire, the depiction of the anatomical context is considered, as described in subsection 6.4.4. While inspecting the morphological structures, it is appreciated that the inside of the vessel wall is depicted view-dependently. While navigating around the vessel structures, three out of four respondents prefer an extensive illumination, currently presented by a diffuse lighting component. However, when the approximate anatomy is presented in the context of a flow visualization, simplified shading is preferred by all the physicians.

Furthermore, all respondents agree that the occluding contours greatly improve the

perception of the morphological structure. While the occluding contours are considered a necessary feature (Fig. 6.7a), most physicians also value the hidden contours (Fig. 6.7b). In particular those physicians who study anomalous cardiovascular structures, for instance in the context of congenital heart diseases, noticed an improved depth perception, and generally could grasp the complex morphology more easily.

In the second part of the questionnaire, the presented interaction and visualization techniques to analyze the blood-flow dynamics are considered. All the respondents agree that the automated cross-section probing, as described in section 6.2, is very intuitive, and saves valuable time during the data analysis. However, there is a need to interactively modify an automatically positioned cross section. In particular, it should be possible to move, angulate and scale the cross section.

Considering the blood-flow visualization aspects, as described in section 6.4, the questionnaire started with the planar reformat visualizations. All physicians in our study agree that reformat-based data analysis requires non-angulated and view-aligned cross-sectional planes. Therefore, the exploded planar reformat (Fig. 6.2b) is by all means preferred over the integrated planar reformats (Fig. 6.2a). In general, the respondents could very well relate the floating plane to the original location, guided by the dashed connection lines. Furthermore, occlusion can be sufficiently avoided, given the typically limited number of inspected planes, and the possibility to hide certain cross sections. The color Doppler inspired red-and-blue color coding is generally the preferred visualization method, although other representations could be tested.

In section 6.3, we have investigated a variety of seeding approaches, distinguishing between seeding distributions and seeding templates (Fig. 6.1). Considering these seeding approaches, the respondents were far from unanimous in many respects. As for the seeding distributions, the physicians indicate that it is worthwhile to investigate the effects, however they prefer not to vary the strategy. A common argument is that they prefer to use a single seeding strategy, for which they can get accustomed to the impact on the line primitives. In consequence, the physicians should be able to distinguish anomalous flow patterns more easily. Based on the same argument, three out of four respondents prefer the rectilinearly structured seeding over a random seeding, even if they are aware that this imposes a fixed structure on the line primitives.

None of the respondents valued the use of seeding templates. Even though some cardiovascular conditions require inspection of the blood-flow dynamics near the vessel wall or center, still there is a collective conception that changing the seeding template would falsely influence the perception of the resulting line primitives. The templates were intended to focus on the regions-of-interest, avoiding visual clutter and occlusion in the case of densely seeded line primitives. However, it is likely that this approach will not be taken up by the clinical research community.

Tracing of line primitives is currently considered to be the most effective visual style to analyze the blood-flow dynamics. For the data analysis, all respondents prefer to trace line primitives from multiple cross-section planes. For both streamlines and

pathlines, reversing the trace direction caused considerable doubt. Most physicians indicated that the reversed tracing seemed to be a promising feature, in particular for branching blood-flow streams. However, none of the respondents could indicate a direct application where they would benefit from this feature. Nevertheless, there was common interest to keep this feature available.

In the next part of the questionnaire, the physicians were asked to evaluate the appearance of the line primitives. Color coding of the line primitives can convey many aspects of the data. All physicians in the study agree that from the presented aspects, the blood-flow speed is the most valuable one. However, the possibility to inspect other parameters, such as a color coding of the line propagation, is highly appreciated.

With regards to the color coding, the rainbow color map was generally valued best to inspect the blood-flow speed. In particular, all respondents preferred the combination with a salient color for a range of high speed flow (Fig. 6.3). By qualitative inspection of the data, it was noticed that the salient color directly attracts the attention to high speed flow regions, while for other regions the blood-flow speed could be estimated fairly well. In addition, all participating physicians appreciated the possibility to interactively adjust the color map to inspect specific ranges of blood-flow speeds.

Illumination of the line primitives was generally considered to enhance the perception of the spatial relations between the line primitives. With respect to the animated highlight (Fig. 6.5), the physicians responded slightly more reserved. It was generally agreed that the highlight provides a profitable visual indicator. Indeed the highlight captures the direction and speed of the particle trajectories that yielded the line primitives. Moreover, relations between the line primitives became more apparent,




Figure 6.8 Exploded planar reformat and pathlines depicting pathological systolic blood-flow behavior in the ascending aorta.

potentially improving the understanding of complex flow patterns. The visual gain from the highlight was valued best in combination with the circular seeding strategy, which was however considered an improbable choice for current research purposes.

Subsequently, the questionnaire gauges the value of the presented flow-rate arrow trails. The physicians agree that the flow-rate arrow trail provides a clear indication of the temporal behavior of both the flow profile and the flow-rate at the considered cross-section. However, there is also the need for quantitative numbers. The arrow representation is considered a valuable visual indicator that should be accompanied by quantitative flow-rate information.

In the final part of the questionnaire, the physicians were asked to judge the overall performance of the framework. All respondents agreed that real-time parametrization is an important aspect for their data analysis tool, saving them valuable time. Interaction with the data and parametrization of the visual styles are performed in real-time. Consequently, the performance of our framework was positively valued. In particular, direct modification of the color coding, seeding density and the line-trace length were highly appreciated.

6.6 Discussion and conclusions

The figures in this chapter conveyed blood-flow behavior of healthy volunteers. Additionally, patient data was inspected, examining pathological blood-flow patterns such as depicted in figure 6.8. The color Doppler inspired red-and-blue color coding on the exploded planar reformat shows a through-plane regurgitant flow. In addition, the pathlines provide insight into the complex turbulent behavior over time. The behavior throughout the cardiac cycle can be inspected by animating the visual styles.

For all the techniques presented in this chapter, the computational power of modern graphics hardware was exploited. Consequently, the rendering performance of the visual styles, as well as the performance of computationally costly algorithms were improved considerably. In particular, tracing of the line primitives is performed in real-time by employing the geometry shader, similar to the approach by Köhn et al. [119]. As a result, the line primitives can be parameterized interactively. The system generally runs at interactive framerates of well over ten frames per second, even when multiple visual styles are combined. In order to achieve these framerates, the graphics hardware needs to implement the unified shading architecture.

In conclusion, we have discussed various visualization and interaction techniques that enable physicians to interactively explore unsteady volumetric PC-MRI blood-flow data. We have presented a number of visualization techniques to depict the blood-flow dynamics. All presented techniques rely on the probing method to interactively select vessel cross-sections. In addition, we have introduced an illustrative approach to depict an approximate anatomical context.

The blood-flow data is visualized through application-specific planar reformats, as well as streamlines and pathlines. These lines are generated for each phase of the cardiac cycle, based on various seeding distributions and templates. Properties such as the blood-flow speed can be conveyed by color coding the line primitives, interactively adjusting the color map. Additionally, we propose an animated highlight to continuously indicate the general blood-flow direction. Furthermore, we have presented flow-rate arrow-trails, depicting the flow-rate and peak-flow origin and direction throughout the cardiac cycle.

An evaluation questionnaire was performed with a group of domain experts, measuring the value of the presented visual styles. All respondents were hesitant about varying the seeding approach. Interactively changing the seed density, however, is a valuable feature in practice. The illuminated line primitives were considered to be the most effective visual style to inspect the blood-flow dynamics, for which the color coding of various parameters was a necessary aspect. The real-time parametrization and interactive color-map adjustments were highly appreciated. Additionally, the animated highlight and reversed tracing feature were considered valuable, even if their direct application in practice was not yet evident. The flow-rate arrow-trails provided a valuable visual indicator when combined with the appropriate quantitative information.

In the future, we see opportunities to include other visual styles that have proven successful for depicting line primitives. For example, the line perception could be enhanced by means of shadows [173] or halo effects [59, 156]. Moreover, the seeding approach will require further research. While the flow visualization literature generally proclaims random seeding strategies, physicians need recognizable patterns to which they get accustomed. To resolve this contrapositions, other approaches should be investigated. For instance, evenly spaced line primitives might be a viable solution [156, 250].

*"I am in blood
Stepp'd in so far, that, should I wade no more,
Returning were as tedious as go o'er."*

William Shakespeare - Macbeth

Blood-flow inspection The cardiovascular system



This chapter is based on:

"Interactive Virtual Probing of 4D MRI Blood-Flow" Roy van Pelt, Javier Oliván Bescós, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, Bart ter Haar Romeny, and Anna Vilanova. In *IEEE Transactions on Visualization and Computer Graphics* (2011), 17(12):2153-2162.

7.1 Motivation

The previous chapter presented a probing technique and associated blood-flow visualizations, tailored to the large thoracic vessels. This probe assumes a tubular structure, positioning a two-dimensional contour around the circumference of a vessel.

However, blood-flow patterns in the chambers of the heart are likewise important to inspect. Based on discussions with the involved physicians, we have extended the probing paradigm to enable inspection throughout the entire cardiovascular system. The motivation for interactive exploratory visualization of measured blood-flow velocity data remains the same, and was discussed in the introductory chapters.

In this chapter, we refer to the so-called virtual probe, whereby we distinguish between the vessel probe presented in chapter 6. This naming convention alludes to the fact that the probe in this chapter comprises a three-dimensional object, in a virtual manner positioned in the patient coordinate system.

7.2 Probing

7.2.1 Appearance

The shape of the virtual probe was envisioned to conform to the anatomical structures, and was simultaneously expected to be simple and easy to perceive. Abstraction of the cardiovascular morphology yields a cylinder as the elementary vascular shape, while a cone was determined to convey the elementary shape of a heart chamber. As a result, the probe is represented by half of a tapered cylinder, enabling visualization of both a cone and a cylinder.




Figure 7.1 (a) The virtual probe can be equipped with a parallel and orthogonal planar reformat, (b) and is automatically aligned to the orientation field, along view-direction \hat{e} .

The probe can be equipped with planes parallel and orthogonal to its long axis, conveying anatomical context. Figure 7.1a shows the probe with associated planes.

Occlusion of information within the probe is avoided by continuously orienting the open side of the cylinder towards the viewer. During interaction, the probe rotates around its long axis, directing the probe normal towards the viewpoint, as illustrated in figure 7.1b. The updated probe normal \hat{n} is defined as the projection of the view vector \hat{e} into the normal plane: $\hat{n} = (\hat{e} \cdot \hat{l})\hat{l} - \hat{e}$.

This approach introduces two issues. First, the view-alignment is ill-defined whenever the view vector is approximately parallel to the probe long axis. In practice, this does not pose a problem, since the flow information is not occluded in this case anyway. Second, the proposed view alignment may lead to counter-intuitive user interaction. This is due to the fact that the probe depends on the view direction, in addition to its position and orientation. Whenever the user interactively rotates the virtual probe, this rotation is restricted to the probe parallel plane, providing basic two-dimensional interaction. Moreover, constant updating of the probe to adhere to the view alignment causes a conflict with the user-specified rotation. We resolve this conflict by disabling the view alignment whenever the user rotates the probe, followed by an animation that gradually corrects the resulting misalignment.

7.2.2 Positioning

A new virtual probe can be added to the system using two distinct interaction approaches. On the one hand, we present a single-click ‘pen-stroke tool’. With this tool, the user creates a new line by dragging the mouse, creating a new virtual probe with the long axis identical to the user-defined line. On the other hand, we introduce a ‘two-click tool’, where the user selects the start and end point of the long axis in sequence.

Initial positioning of the probe is performed using a customary planar reformat, presenting either anatomical data or blood-flow velocity data. The planar reformat is continuously oriented towards the viewer. Different orthogonal views can be selected with respect to the patient, i.e., axial, sagittal, coronal, mimicking conventional slice-based data navigation. Besides, oblique reformats can be obtained by viewpoint interaction, using a trackball camera scheme.

Exact positioning along the view direction is challenging, while a tedious and time-consuming segmentation step is undesirable. Therefore, we introduce an automated fitting approach that approximately determines the depth position and orientation of the probe, using the underlying velocity data. Alternatively, a new probe could be positioned based on a MIP or DVR image. In contrast to positioning based on a planar reformat, no initial depth information is available in these cases. The depth information can likewise be resolved by the fitting approach, which is described in the next subsection.

7.2.3 Fitting

With common two-dimensional interaction techniques, it is straightforward to translate, rotate and scale the virtual probe in a three-dimensional environment. However, interactive adjustment of the position and orientation along the view direction is non-trivial. The depth and orientation of the probe are difficult to estimate visually, and segmentation of the anatomical structures is not preferred.

We aim to reduce the freedom of interaction with the virtual probe to two dimensions. This requires an automatic detection of position and orientation of the probe, along the view direction. We propose using the available velocity information of the unsteady blood-flow field, aligning the probe to this field. In particular, we introduce an automated fitting approach that aligns the orientation of the long axis of the virtual probe, to be tangential to the average local blood-flow orientation.

Preprocessing - Before precomputing the average blood-flow orientations, we address the uncorrelated noise that affects the velocity data. This undesirable spike-noise is caused by acquisition artifacts and the flow-field reconstruction process. We employ the vector median filter to largely remove this noise [8], inducing robustness for the fitting method. In addition, the filter moderately regularizes the velocity field without averaging, keeping tissue boundaries intact. The vector median filter VM is computed from vectors \vec{v} in a local neighborhood η surrounding the current position \mathbf{x}_i :

$$\text{VM}(\mathbf{x}_i) = \vec{v}(\mathbf{x}_{i_{\min}}), \text{ with } i_{\min} = \underset{i \in \eta}{\operatorname{argmin}} \left(\sum_{j \in \eta} \|\vec{v}(\mathbf{x}_i) - \vec{v}(\mathbf{x}_j)\| \right).$$

Objective function - Fitting the probe to the mean orientations of the unsteady velocity-field, described by the tMOP volume, can now be defined as an optimization problem, as depicted in figure 7.1b. The tMOP volume was described in subsection 5.1.3. The optimization process first requires a measure to compare the orientation of the probe long axis with the tMOP orientations, defining the objective function of the optimization problem.

To that end, coherence between orientations provides a suitable measure, defining the degree of alignment between orientation distributions in terms of direction and magnitude. In particular, we adopt the eigenvalue coherence (EVC) measure [178], as previously described in subsection 5.2.2.

Using the eigenvalue coherence measure, the degree of alignment between an arbitrary vector \vec{u} and the tMOP orientation at a position \mathbf{x} is then defined as the point-to-point coherence (PPC):

$$\text{PPC}(\mathbf{x}, \vec{u}) = \text{EVC}(\text{tMOP}(\mathbf{x}) + (\hat{u} \otimes \hat{u})), \text{ with } \hat{u} = \frac{\vec{u}}{\|\vec{u}\|}.$$

Subsequently, we integrate the point-to-point coherence PPC along the probe long axis to find the total line coherence (LC). The line coherence represents the alignment between the tMOP orientation field and the probe long axis, defined by the two outermost points \mathbf{p} and \mathbf{q} . The total line coherence provides the final objective function for the optimization method, and is defined as:

$$\text{LC}(\mathbf{p}, \mathbf{q}) = \left(\frac{1}{\Delta} \sum_{i=0, \Delta}^1 \text{PPC}(\mathbf{p} + i(\mathbf{q} - \mathbf{p}), \hat{l}) \right), \text{ with}$$

$$\Delta = \frac{s}{\|\mathbf{q} - \mathbf{p}\|} \quad \text{and} \quad \hat{l} = \frac{\mathbf{q} - \mathbf{p}}{\|\mathbf{q} - \mathbf{p}\|}.$$

The sampling distance Δ along the probe long axis is proportional to the resolution s of the tMOP orientation field.

Optimization - Lastly, the position and orientation of the optimal alignment between the probe and the tMOP field needs to be determined. The optimization space, spanned by the lines \mathbf{p}_{δ_1} and \mathbf{q}_{δ_2} along the view direction \hat{e} , must be searched to find the desired fit between the probe long axis and the orientation field. The probe fit optimization (PF) searches for this optimal alignment, and is defined as:

$$\text{PF}(\mathbf{p}_0, \mathbf{q}_0) = \underset{\delta_1, \delta_2}{\text{argmax}} (\text{LC}(\mathbf{p}_{\delta_1}, \mathbf{q}_{\delta_2})), \text{ with} \quad \begin{array}{l} \mathbf{p}_{\delta_1} = \mathbf{p}_0 + \delta_1 \hat{e} \\ \mathbf{q}_{\delta_2} = \mathbf{q}_0 + \delta_2 \hat{e} \end{array}.$$

As a proof of concept, we search the entire parameter space for maximum coherence between the probe long axis and the orientation field. The speed of searching the tMOP orientation field is improved by a branch-on-need octree [264] space partitioning. This results in a local fitting time in the order of one second, which is fast enough for our needs. An additional performance improvement may be obtained by employing a faster optimization method, such as gradient descent or simulated annealing.

7.2.4 Refinement

For the envisioned qualitative exploration, we explicitly perform neither global nor local segmentation. Instead, the presented initial positioning and fitting of the virtual probe provides a plausible approximation of the desired position and orientation, for which some further refinement may be necessary.

Similar to the initial positioning, full three degrees-of-freedom for interactive refinement of the probe would provide too much leeway to the target user. Instead, we restrict probe movement to a view-aligned plane, parallel to the probe long axis. This restriction narrows down the interaction to a manageable two degrees-of-freedom task, without sacrificing positioning possibilities. Especially in combination with the presented fitting approach, the navigation becomes much more practicable.

Within the parallel plane, the user can intuitively translate and rotate the probe. Translation is performed by dragging the object, while rotation is enabled through a dedicated handle on the virtual probe. Additionally, the base and top radius of the virtual probe can be adjusted.

7.3 Seeding

7.3.1 Seeding strategies

In section 6.3, seed points were positioned on the cross-sectional plane, based on the user selected probe position along the vessel. Similarly, the orthogonal plane attached to the virtual probe can be employed as a planar seeding basis. Interactive adjustment of the probe radii likewise changes the radius of the orthogonal plane, enabling physicians to manually fit the seeding plane to the inspected vessel or heart chamber. The orthogonal plane can furthermore be translated along the long axis of the virtual probe using an elementary two-dimensional interaction. This allows for rapid refinement of the seeding plane position.

Alternatively, the volume spanned by the virtual probe may be filled with seed points. Such a strategy leads to a volumetric distribution of seed points with a relatively large spatial dispersion. This is beneficial for a variety of geometry-based flow-visualization approaches. Volumetric seeding approaches require an adequate point distribution in a three-dimensional coordinate frame.

7.3.2 Seeding distributions

For planar seeding on the orthogonal plane of the virtual probe, the distributions described in subsection 6.3.2 could again be applied. However, based on the evaluation results, as described in subsection 6.5, we have restricted the seeding approach for the virtual probe.

For seeding based on the virtual probe, only the uniform distribution was adopted. Although uniformly distributed seed points are generally tightly packed, they are not necessarily evenly distributed. Visually, groups arise in the uniform random distribution, while certain areas are relatively sparsely covered. Therefore, we adapt a quasi-random Poisson-disk distribution [211], introducing a minimum distance between the seed points. This distance restriction results in a more evenly spaced distribution. Seeding templates are omitted altogether for the virtual probe.

The Poisson-disk distribution is computationally rather expensive. During exploration, the seed points are continuously updated, reflecting the new probe location. Therefore, interactivity of exploration may be affected with volumetric seeding, depending on the number of seed points. Consequently, a uniform seeding distribution is advisable during interaction.

7.4 Visualization

7.4.1 Particles

Based on the seeding distributions, described in section 7.3, we have investigated novel geometry-based flow-visualization approaches for the blood-flow application. Especially, particle visualizations are a recognized technique to depict flow fields, and were readily applied to convey the measured hemodynamics in previous work. Commonly, particles are depicted as spheres [248], or are otherwise represented by small integral lines [189]. Most approaches convey blood-flow speed through color, while direction information is captured by temporal cohesion. The small integral lines capture a history of the particle trajectory.

Methodology - Conventional approaches employ the available visual cues, such as color and shape, to capture merely the blood-flow velocity information. Instead, we propose an illustrative particle visualization, which captures the velocity information by means of shape, while saving the color cue for more elaborate blood-flow characteristics. To this end, we have extended the common sphere or integral line rendering. The approaches are inspired by illustrative techniques that are typically found in cartoons [164].

First, the particle spheres are deformed into ellipsoids, which convey speed and direction of the local velocity by means of their shape. This technique is often applied in cartoons, in which case the deforming behavior of a ball in a high-speed motion is replicated.

In addition, we enhance the perception of direction by means of speed lines, often used to depict motion in cartoons [105]. We apply two lines per particle, showing the particle trajectory for a limited amount of time. Using two lines has the advantage of capturing local reliability of the particle movement, indicating whether the particle resides within the blood-flow stream. In case the two speed lines are visually aligned, the particle fully resides within the flow field. Otherwise, the particle is probably at a position near the tissue boundary, as emphasized by the particles within the red dashed circle in figure 7.2a.

The thickness of the speed lines is adjusted according to the local speed, where higher speed results in thinner lines. This choice is specific to our application, where particles in high speed blood-flow commonly group together. Thinner lines for these densely-packed particles reduce visual clutter.

The novel combination of ellipsoids and speed lines captures the local blood-flow velocity, including speed and direction. Consequently, color coding of speed information is redundant, saving the color cue to convey more elaborate flow characteristics. Examples of such characteristics are vorticity or particle residence time. A more extensive list of characteristics was provided in subsection 2.3.1. Figure 7.2b provides a visual comparison between common sphere rendering and our illustrative approach.




Figure 7.2 Illustrative particles elongate according to the flow speed. (a) Aligned speed lines indicate that the particle resides within the flow field; otherwise the particle dwells near the tissue boundary (e.g., particles within the red dashed circle). (b) In contrast to traditional sphere renditions, our illustrative approach comprises ellipsoidal particles with speed lines.

Each seed point, and the associated particle movement, depends on the spatial location of the probe and the currently selected time frame. We present a ‘seed injection’ approach, analogous to the injection of contrast agent during MRI acquisition. A set of seed points - or ‘bolus’ - can be injected interactively during animation of the particle system. Figure 7.3 depicts three time frames of particle traces, originating from such a bolus. Alternatively, the animation can be paused to inject seed points at a specific time during the cardiac cycle.




Figure 7.3 A ‘bolus’ of illustrative particles moves through the cardiovascular system. In this example, the branching of the aortic arch towards the carotid arteries is depicted.

When particles leave the boundaries of the acquired volume, new seeds are commonly injected into the system at the original time of seeding. Alternatively, we have experimented with a continuous seeding approach, directly re-injecting seeds whenever particles vacate the volume. Each time a particle crosses the boundaries of the acquired volume, we visually fade out the particle, before creating new seeds.

Implementation - The presented particle system extensively uses modern consumer graphics hardware, providing real-time rendering. We employ a ping-pong buffer approach, alternating particles between two vertex buffer objects (VBOs). Additionally, each particle set is equipped with a second VBO, containing particle attributes such as particle age and local blood-flow velocity.

A step in the development of the particle system is divided into three stages that are executed for each render cycle. The first stage is named the *evolution* stage, controlling aging of particles in the system. A vertex shader with transform feedback is employed, managing particle birth and death in parallel.

Subsequently, particles are processed by the *advection* stage, displacing particles based on the unsteady velocity field. Displacement is based on a GPU-driven Runge-Kutta 4 integration scheme, approximating the spatiotemporal particle trajectory. Due to the moderate temporal resolution, extensive temporal supersampling is required.

Lastly, particles are depicted by the *rendering* stage. We employ geometry shaders for imposter rendering of the ellipsoids and speed lines. For each ellipsoid, a quad primitive is generated and rotated according to the projection of the local velocity vector [79, 187]. Instead of raycasting the ellipsoids, a more efficient normal-map texture is used, which delivers satisfactory results when combined with an anti-aliasing scheme.

The speed lines consist of short reverse-traced pathlines, in accordance with the pathlines introduced in subsection 4.3.5. Each speed line is generated in real-time by the geometry shader. The lines consist of view-aligned quad-strips extended with halos. The halos are generated during the fragment-shading stage [59].

7.4.2 Integral lines

In subsection 6.4.2, integral lines were introduced for visual exploration of the blood-flow behavior. For vessel probing, we proposed short integral lines, at multiple probing locations. These lines were seeded *dynamically*, i.e., the lines are initiated at a fixed spatial location, while varying the seed time with the current time frame. This provides an approximative depiction of the pulse-wave in the cardiovascular system.

For the cardiovascular system as a whole, longer line traces are often employed. The longer traces enable inspection of the particle trajectories throughout the cardiac cycle. This is considered valuable despite the rapidly accumulating uncertainty of the trace. Typically, these lines are seeded *statically*, in which case the lines are initiated from a fixed position in space and time.




Figure 7.4 Long pathlines generated using consumer graphics hardware. The lines are statically seeded in the heart chambers to capture aorta and pulmonary arteries. (a) Pathlines may be represented by shaded imposter tuboids, or (b) illuminated lines. The frames in (a) show larger versions of the pathlines, without and with halos and contours.

Methodology - To adhere to the exploratory nature of the virtual probe, long pathlines need to be generated in real-time, based on the set of seed positions. Therefore, the original hardware-driven approach to generate integral lines, as described in subsection 6.4.2, was extended to interactively generate long pathlines.

Moreover, we present an enhanced visual representation of the lines. Illustratively shaded imposter tuboids are employed (Fig. 7.4a), in contrast to illuminated lines (Fig. 7.4b). Perception of the spatial relations between the tuboid pathlines is improved by means of halos [59], and contours are applied to emphasize the structure of the pathlines. These additions are depicted in the magnification frames (Fig. 7.4a).

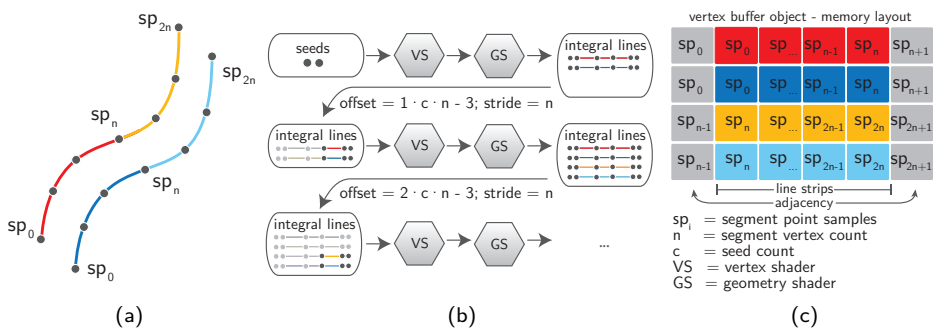


Figure 7.5 (a) Overview of the integral-line generation, using consumer graphics hardware. (b) Line-strip segments with adjacency information are iteratively processed by the GPU pipeline, (c) and stored to the line buffer, depending on their implicit topology.

Implementation - In previous work, various GPU-based integral-line trace approaches were proposed. For example, Köhn et al. [119] present a hardware-driven approach for fiber tracking. Their approach requires intermediate data transfers involving GPU memory. Alternatively, Cuntz et al. [45] present integral lines, employing the geometry shader and transform-feedback. The geometry shader allows the transformation of one input primitive into another one, which for line tracing amounts to a transformation from a point to a line primitive. Cuntz et al. observe that continuation of long line traces requires the last generated point of the previous segment, as well as the generated line segment itself. Since the geometry shader cannot return distinct primitive types, they forward the last position to the fragment-shading stage, storing the last positions to specific fragment locations.

In contrast, we present a GPU-driven approach for integral-line generation, omitting both intermediate buffer transfers and the unwieldy fragment-based storage process. We observe that vertex buffers, stored on the GPU side, may contain both geometry primitives and associated attributes, such as normal vectors and color information. Addressing these attributes requires a buffer offset, indicating the start of the specific attribute values. Instead of addressing specific attribute values, we employ the buffer-offset functionality to sequentially address segments of the line trace.

The first step of the iterative approach renders the provided seed points and generates the first line segments in parallel (Fig. 7.5a). Each generated segment consists of the maximum number for output vertices n . Subsequently, an offset pointer is set to address the last points of the previous segment in the current line buffer, providing the seed positions for the newly traced segment. This process is iterated, increasing the segment count c , until the desired trace length is achieved (Fig. 7.5b).

While generating the line primitives, adjacency information between segments should be included in the process. We efficiently store the line segments as line strips with adjacency information. Therefore, the geometry shader is set to output point primitives, as opposed to lines. The implicit topology of the lines can be manually generated by the geometry shader, resulting in a specific buffer memory layout (Fig. 7.5c).

Finally, the generated line-strip buffer can be rendered interactively, properly dealing with the adjacency information. For an illuminated-line representation, the tangential component is computed for each vertex, requiring the full adjacency information. Additionally, imposter rendering of tuboids requires both the tangential component, as well as the per-vertex binormal.

7.4.3 Integral surfaces

In addition to particles and integral lines, we have investigated the use of integral surfaces. Pathsurfaces are generated in real-time, conveying the temporal behavior of the blood flow. To the best of our knowledge, the value of integral surfaces for the inspection of blood-flow measurements has not been explored before.

Methodology - The pathsurfaces are illuminated according to the standard Blinn-Phong scheme, and are textured with alternating stripes. Especially for tube-shaped surfaces, the stripes emphasize the rotational motion around the centerline. The color of the stripes may convey various derived measures of the blood flow. Additionally, the propagation front is accentuated by an outline.

The virtual probe again provides a seeding basis for the real-time surface traces. The quality of the pathsurfaces is sensitive to the placement of a sparse number of seed points. Therefore, it is essential that all seed points reside within the blood-flow field. Seed points that are close to the vessel boundaries may cause erroneous pathsurfaces, due to partial volume effects. An interactive probe refinement allows the direct correction of the seed positions, whenever seeds are located outside the blood-flow regions.

Different seeding patterns lead to different integral-surface shapes. In detail, we have experimented with cross-shaped and tube-shaped pathsurfaces, respectively depicted by figures 7.6a and 7.6b. Additionally, multiple nested tube-shaped pathsurfaces are considered, as opposed to a single tube.

For all pathsurfaces, we propose a seeding approach similar to the one used for the dynamically seeded pathlines, presented in subsection 6.4.2. Each pathsurface trace is initiated at a fixed spatial location, while the seed time is adjusted to the currently inspected time frame.

Implementation - A range of approaches for integral-surface generation was presented in the literature. Garth et al. [68] presented an accurate generation of integral surfaces in unsteady flow fields, and focus on an adaptively-refined propagation of the surface front. They consider the separation between pathline and timeline generation, as depicted in figure 7.7. This characteristic is employed to approximate the integral surfaces in parallel.

Similarly, we exploit this intrinsic parallelism. First, pathlines are generated in parallel, as described in the previous subsection. Subsequently, the surface is triangulated per timeline in parallel, using its implicit topology. This approach enables the real-time generation of pathsurfaces, and hence facilitates an interactive exploration within the blood-flow field.

Corresponding to the presented pathline traces, adjacency information should be incorporated in the generation process. In particular, triangulation and per-vertex normals require spatiotemporal information of the direct neighborhood. As depicted in figure 7.7, this neighborhood information is exposed to the graphics pipeline. This is done by rendering the line primitive $(s_{p_i}, s_{p_{i+1}})$ with adjacency information, while providing neighboring lines (in red and green) as additional attributes. These line attributes are provided through appropriate buffer offsets, addressing the desired primitives within the buffer.

Various illustrative renditions of pathsurfaces were presented in recent work [19, 95].

Similar to Hummel et al. [95], we apply a stripe pattern to enhance the surface shape. The stripes are procedurally generated in the fragment shader, allowing interactive color mapping of various derived measures of the blood flow. The stripe pattern applied to tube-shaped surfaces is depicted in figure 7.6b. This is in contrast to conventional Phong-shaded surfaces, as depicted in figure 7.6c.

7.4.4 Anatomical context

Similar to the vessel probing approach, an anatomical context is required to facilitate navigation with the virtual probe. The previously introduced illustrative anatomical context, described in subsection 6.4.4, effectively conveys the morphology. However, this representation requires a segmentation of the anatomical structures, which is typically a tedious and time-consuming process, as discussed in section 5.2. This especially holds true for complex cases, with abnormalities in the morphology.

Instead of the illustrative surface renditions, we have investigated customary GPU-based raycasting with geometry intermixing and interactive volume clipping. No explicit segmentation is necessary, but a global classification of the anatomical structures is required that is based on a transfer function specification.

On the one hand, the volume renderer may use blood-flow speed information, or alternatively a tMIP volume, as input data. The speed information provides a cine representation of the blood flow, while the tMIP provides a static representation. Both these data sets contain strong boundaries between blood-flow regions and stationary tissues. Consequently, transfer function classification is manageable and the resulting iso-surfaces can be illuminated appropriately. An example of a volume rendered tMIP data set is depicted in figure 7.8a.

On the other hand, detailed anatomical scans are often acquired. These anatomical data can similarly serve as a contextual rendering, after spatiotemporal alignment with the 4D blood-flow data. Even though the anatomical data provide much richer information, they are likewise harder to classify. Moreover, these anatomical data are typically acquired at one or two time frames during the cardiac cycle, providing a static context. An example of an anatomical data set is depicted in figure 7.8b.

In order to correctly depict the spatial relations between the flow visualizations and the anatomical context, we intermix opaque geometry with semi-transparent volume raycasting. The intermixed representation, however, leads to unacceptable occlusion of the blood-flow visualizations.

Visibility of the blood-flow visualization is recovered by clipping the volume. Basic orthogonal clipping planes may be used for cutting, as depicted in figure 7.8c. Additionally, we employ the plane parallel to the long axis of the virtual probe as an interactive oblique clipping plane, as depicted in figure 7.8d. The positioned probe readily defines a region of interest, and clipping based on the probe is therefore likely to produce sensible views on the anatomy.




Figure 7.6 Pathsurfaces generated using consumer graphics hardware. (a) We present both cross-shaped surfaces, and (b) tube-shaped surfaces. The latter may be nested, approximately capturing the wave-front profile. A stripe pattern is employed to convey rotation, (c) which is in contrast to conventional Phong-shaded surfaces.

Alternatively, a gradient-based volume rendering can be used to reduce occlusion of the flow visualization. Regions with large gradient magnitudes roughly capture the contours of the morphology. Opacity modulation enables emphasis on these boundaries, while rendering other areas semi-transparently. Figure 7.9 shows an intermixing of illustrative particles with gradient-based volume rendering as anatomical context.




Figure 7.7 Overview of the generation of integral surfaces using consumer graphics hardware. Surfaces are generated with adjacency information, exploiting parallelism both in pathline tracing and triangle construction.




Figure 7.8 Anatomical context of an aortic dissection patient, conveyed by (a) raycasting of the tMIP data, (b) in comparison to raycasting of the anatomical data. Volume clipping is employed to avoid visual clutter while intermixing the presented blood-flow visualization with the context. (c) Either orthogonal planes, or (d) the probe parallel plane may be used to cut the volume. Since the intensities in the tMIP volume are relatively low at the morphological boundaries, the transfer function can be used to represent a wall, here shown in red. To exemplify the clipping, the interior of the anatomy is depicted in yellow, which would normally be set to fully transparent.

7.5 Evaluation

The initial conception for an exploratory virtual probe originates from discussions with physicians, actively involved with advancements in the emerging field of PC-MRI blood-flow acquisition. The value of the presented methods was measured by means of an evaluation questionnaire, which is included in appendix C, performed with four physicians. All respondents have in-depth general knowledge of CVD diagnosis, prognosis, and treatment. For the evaluation, we distinguish between cardiac-oriented and vascular-oriented respondents, based on the distinct ways of analyzing their data.

Positioning - First we consider the probe-positioning approach, based on the 'pen-stroke' and 'two-click' tools. While the differences are seemingly small, we have come to understand that the cardiac-oriented respondents prefer orthogonal views to select end-points of the probe centerline sequentially, which is only possible with the two-click tool. In addition, all respondents preferred a reduced degree-of-freedom for selecting the initial oblique reformat. Alternatively, three user-defined points may define the initial probe parallel plane, based on orthogonal slice views.

Fitting - Subsequently, we have evaluated the automated fitting. Regions of interest are commonly selected, based on anatomical scans. Our approach adds a view-dependent alignment of the probe according to the associated velocity field. This was positively valued by all the respondents, specifically for regions that are challenging to segment. Several respondents indicated that the fitting approach in practice may be used as an additional step, validating their anatomy-based positioning.

All respondents expressed the need to carefully control the visual probe location. Consequently, global positioning of the virtual probe based on MIP or DVR was not particularly desired. The approximate behavior of the presented fitting approach was not regarded disadvantageous. The close alignment with the average velocity field was considered adequate, especially with an interactive refinement of the probe location.

Refinement - The refinement of the virtual probe, was evaluated positively. The reduced degrees-of-freedom by the in-plane probe interaction was highly appreciated.

Visualization - All respondents preferred a semi-transparent glass style for the virtual probe, extended with contour lines. The continuous view alignment of the virtual probe was considered to be intuitive, and the associated parallel plane was greatly appreciated for showing anatomical context or additional blood-flow information. The orthogonal plane was valued positively for refining the probe location, even though the information is communicated from an oblique angle. Moreover, the orthogonal plane was appreciated as a seeding basis for the flow-visualization approaches.

Particles - The next part of the questionnaire addressed the various flow-visualization techniques, presented in section 7.4. All respondents agreed that the illustrative depiction of the particles enhances the perception of the *blood-flow recirculation*. The physicians could easily relate to this representation, and value the possibility to communicate other blood-flow measures by means of color. The fading cue for

disappearing particles was appreciated by two of the respondents. Remarkably, three out of four respondents indicated that other physicians may mistrust the cartoon-inspired style, questioning the reliability of the conveyed information. Nevertheless, the respondents themselves clearly appreciated the novel stylistic improvements.

Besides the visual representation, also seeding of the particles was evaluated. All respondents appreciated the injection analogy, mimicking the MRI contrast-agent behavior. Three out of four respondents questioned if they would use real-time injection of seeds, and prefer to pause the animation to inject the seeds. All respondents unanimously rejected continuous seeding of particles.

Integral lines - In previous work, the value of integral-lines for blood-flow analysis has been readily established. In the questionnaire, we evaluate several visual styles, as well as the virtual-probe based exploration.

All respondents preferred the tuboid representation over conventional illuminated lines. Two out of four respondents, however, note that the radius of the tuboids should be only slightly larger than a typical line rendering. Halos were again valued positively for depth perception. With respect to the contour rendering, the respondents were far from unanimous.

Altogether, the interactivity of the pathlines was considered the most valuable feature. All respondents indicated that interactive lines allow them to inspect the *blood-flow structure* in more regions within the same restricted time budget. Other tools typically require time-consuming local segmentation of vessel cross-sections and off-line precomputation of the integral lines.

Integral surfaces - The integral surfaces for blood-flow visualization caused some controversy between the respondents. The evaluation revealed a dichotomy between vascular-oriented and cardiac-oriented physicians. The vascular-oriented physicians responded positively to the surface visualizations, indicating the benefit to capture *blood-flow rotation* in large thoracic arteries. Nesting the tubular surfaces was said to reasonably approximate the wave-front profile. In contrast, cardiac-oriented physicians questioned the value of integral surfaces over the previously presented integral lines. They argue that within the heart chamber it is hard to define regions with rotation about a longitudinal axis.

Anatomical context - The last aspect of the visualization comprises the volume rendering of the anatomical context. All respondents appreciated the presented volume rendering, in particular based on the tMIP or blood-flow speed information. However, they prefer to use an anatomical scan for contextual information, involving a tedious and time-consuming transfer function classification. Therefore, planar-reformats currently suffice as anatomical context for general purposes, while volume rendering may occasionally serve as an aesthetically-pleasing broad overview.

Finally, all respondents positively valued clipping. In particular, orthogonal planes are appreciated for their simplicity. Using probe-based clipping was valued positively as

Table 7.1 Evaluation results of the blood-flow visualization and interaction techniques.

virtual probe	positioning	+	two-click tool, ortho-plane navigation
		±	pen-stroke tool
		-	view-aligned plane navigation
virtual probe	fitting	+	concept, accuracy, and speed
		+	planar reformat as basis
		-	MIP or DVR as basis
virtual probe	refinement	+	translation, rotation, scaling
	visualization	+	view-alignment, associated planes
	blood-flow visualization	particles	+
		+	contours, interactive seeding
		±	fading cue for leaving particles
		-	continuous seeding
blood-flow visualization	integral lines	+	lines, tuboids with halos, interactivity
		+	static and dynamic seeding
		±	contours
blood-flow visualization	integral surfaces	+	tubes with nesting, stripe pattern
		-	cross, only vascular oriented
blood-flow visualization	context	+	planar-reformat, raycasting, clipping
		±	probe-based clipping, tMIP data
		-	transfer function

a concept, and becomes worthwhile when the current interaction with the oblique clipping plane is getting less involved. In general, the presented GPU-based volume rendering with clipping was preferred over opacity modulated volume renderings, as included in competing tools.

The results of the evaluation questionnaire are presented in table 7.1. In this table, '+' indicates that the majority of physicians valued the aspect positively, while '-' indicates a predominantly negative judgment. This is also encoded in colors, respectively in green and red, while orange (or ±) indicates diverging judgments.

Performance - Besides the evaluation questionnaire, we have carried out a quantitative performance evaluation, using a conventional system with dual-core processor, with 6GB of internal memory and an NVidia GeForce 570GTX graphics card. For all renditions, a multi-sampling scheme was applied to avoid aliasing. The imposter particle visualization renders at a minimum of 40 frames per second (FPS), for the default of 200 particles. Adding particles did not cause a noticeable performance decrease. In contrast, the performance of the pathline rendering decreases linearly with the number of lines. Interactive frame rates are obtained for up to 700 lines. The pathsurfaces were generated from 40 seed points, providing a smooth surface. A single tube is rendered at about 45 FPS, while three nested tubes render at 30 FPS. For these results, rendering of the probe and a tMIP based context was included. The DVR can put a heavy load on the performance.

We have also evaluated computation times of the probe-fitting approach. The global fitting of the probe depends on the length of the ray-volume intersection. The fitting on average completes within two seconds, when octree optimization is applied. Instead, the local fitting generally finishes within one second. Using the octree, this can be reduced to merely dozens of milliseconds. In practice, increasing the octree depth by one or two levels yields satisfactory computation times and fitting results.

The fitting approach requires a tMOP volume, which takes about ten seconds to compute for an average data set of $150 \times 150 \times 50$ voxels and 25 phases. However, including median smoothing increases this computation time to twenty minutes. This time would reduce significantly with a parallel implementation.

7.6 Discussion and conclusions

The figures in this chapter convey blood-flow behavior of healthy volunteers, and patients suffering from an aortic dissection. In particular, we have employed the virtual probe on two distinct patient studies, involving an aortic dissection. An example of an illustrative particle visualization, conveying pathological flow in the aortic arch in combination with a volume-rendered anatomical context, is presented in figure 7.9.




Figure 7.9 Illustrative particles depict pathological flow in the aortic arch of a patient suffering from an aortic dissection. The anatomical context is conveyed by volume rendering through gradient-based raycasting, emphasizing the outlines of the morphology.

The evaluation, described in section 7.5, shows that the majority of the presented techniques were positively received by the physicians. The virtual probe, together with the presented flow-visualization techniques, enables a local inspection of the blood-flow information, limiting the scope to a specific spatiotemporal region. However, physicians have a priori knowledge of the anatomical regions to inspect. This follows from their experience with 2D-cine blood-flow acquisition. Interactive local exploration around the customary regions-of-interest easily provides a substantial improvement, while global inspection remains a major visualization challenge.

Due to the lack of segmentation, combined with integration inaccuracies, visual artifacts will occur. The integration primitives may cross morphological boundaries, because of the limited temporal resolution and partial volume effects. With this in mind, the physicians prefer to keep the outliers visible for clinical research purposes. A conventional velocity threshold to eliminate outliers is not appropriate, especially in the case of aortic dissection patients, where regions of slowly flowing blood are the primary regions-of-interest.

The exploratory character of the presented techniques heavily depends on the attained interactivity. Modern consumer graphics hardware is employed for the majority of the visualization approaches and related algorithms, enabling real-time visualization, parametrization and interaction. The graphics hardware should implement the unified shading architecture.

In conclusion, we have presented a variety of visualization techniques that facilitates interactive exploration of unsteady volumetric PC-MRI blood-flow data. In particular, we introduce a virtual probe as navigational object, providing a real-time seeding basis for various illustrative blood-flow visualizations.

All interactions with the virtual probe are reduced to two degrees-of-freedom, apart from conventional camera positioning in the three-dimensional scene. The hard-to-determine position and orientation of the virtual probe along the view direction is calculated using an automated fitting approach. The presented system incorporates various interactive flow-visualization approaches, inspired by illustrative techniques.

Both the virtual probe and the presented flow visualizations were evaluated with four physicians, who are actively involved in 4D MRI blood-flow acquisition. All respondents positively valued the virtual probe, although interaction with the initial view-aligned planar reformat is still considered to be somewhat cumbersome. Interaction with the virtual probe was perceived as intuitive. The presented fitting approach was greatly appreciated as a valuable addition to anatomy-based navigation.

In general, the visualization approaches were considered complementary, where particles are used to capture blood-flow recirculation, while pathlines capture temporal structures of the blood flow. The pathsurfaces were considered valuable to detect local rotations, which proved to be of interest to vascular-oriented physicians. To the best of our knowledge, pathsurfaces were not previously applied to measured blood-flow. Therefore, we have focussed on evaluating the use of pathsurfaces for blood-flow

analysis, provisionally omitting surface refinement. The quantitative evaluation shows that all components perform at interactive frame rates. This was greatly appreciated, and considered a necessary aspect for exploration.

The ability to explore the data without a predefined segmentation enables physicians to obtain a quick initial insight into the blood-flow dynamics. Based on the initial findings, the physician may judge on the acquisition quality, or may proceed with a thorough quantitative analysis. This analysis can now be performed locally, speeding up segmentation and computation of relevant blood-flow parameters.

In the future, other applications of the generic virtual-probe concept may be investigated. Thusfar, the virtual probe is employed as seeding basis for a variety of flow visualizations. Other types of probes may be defined. For example, we envision separate clipping probes, and probes that refine or cut segmentation surfaces. Furthermore, probes may be used to locally approximate quantitative information, such as vessel radii. An accurate quantitative analysis will require a local segmentation, in which case the probe could facilitate the navigation to specific regions-of-interest.

Furthermore, different visualization techniques for the anatomical context may be investigated. While the presented volume rendering provides satisfactory results, proper transfer function specification is still a time-consuming task.

*"You've got to push yourself harder.
You've got to start looking for pictures nobody else could take.
You've got to take the tools you have and probe deeper."*

William Albert Allard

Blood-flow inspection An ultrasound perspective



This chapter is based on:

"Time-resolved Volumetric Blood-Flow through Magnetic Resonance: A Doppler Ultrasound Perspective" Roy van Pelt, Javier Oliván Bescós, Eike Nagel, Valentina Puntmann, Bart ter Haar Romeny, and Anna Vilanova. To be submitted.

8.1 Motivation

Subsection 2.2.2 elaborated on the range of imaging modalities that enable acquisition of blood-flow information. Throughout this thesis, PC-MRI blood-flow data is used, providing time-resolved volumetric velocity fields, as introduced in section 3.2.

In clinical practice, color Doppler US remains the standard for blood-flow imaging, providing an appealing spatiotemporal resolution at relatively low cost. It effectively employs changes in frequencies, due to the Doppler effect, to detect and visualize the blood-flow dynamics, as illustrated in figure 8.1a. An example is shown in figure 8.1b.

Ultrasound imaging provides a limited field-of-view, and the quality of the acquired data depends on the operator skills. Moreover, US imaging provides limited flow information, and is intrinsically prone to noise. In contrast, PC-MRI provides full volumetric blood-flow information, i.e., an unsteady volumetric vector field, at the expense of a lower spatiotemporal resolution. Acquisition by PC-MRI is considerably more costly, and hence measurements are currently performed for complex pathologies.

Color Doppler US has proven to be an effective tool to assess pathological blood-flow patterns in clinical practice. For instance, Doppler US enables detection of regurgitant jets, which is valuable for the diagnosis of valvular heart disease. It is worthwhile to analyze the flow behavior for this application using PC-MRI data, availing ourselves of the advantages of PC-MRI over US, especially considering the volumetric acquisition.

Domain experts involved with Doppler US or MRI blood-flow measurements have extensive experience in interpretation of blood-flow data. They primarily rely on just a few distinct techniques for visual analysis. Therefore, we devise a visualization prototype for PC-MRI velocity fields, which is based on customary Doppler US visualization techniques. This prototype mimics Doppler US visualizations for PC-MRI data, and extends those with additional techniques that exploit the benefits of PC-MRI velocity data. Hence, we enable the domain experts to explore the advantages and disadvantages of both imaging modalities. Furthermore, we aim for synergy between the disciplines, which potentially leads to new insights into complex hemodynamics.

The simulation of US characteristics is mostly performed to teach young physicians to analyze US information. These calculations require sharp boundaries that can be related to acoustic properties through physical simulation [128, 192, 206]. In addition, simulated US data can be employed to register US to other modalities [92, 258, 259].

Instead, we mimic the US images, as integral part of a compound view for PC-MRI blood-flow data. In detail, we employ a virtual probe for interactive exploration, analogous to the ultrasound imaging probe. In a compound view, we use planar reformats to convey the simulated Doppler US velocity encoding, as well as velocities inspired by 2D-cine PC-MRI blood-flow data. The volumetric nature of the PC-MRI blood-flow data is reflected by a number of geometry-based flow visualization techniques. Moreover, we have extended the compound view with a novel velocity-field volume rendering, conveying the dynamics in three dimensions.

8.2 Probing

8.2.1 Appearance

To mimic Doppler US imaging, a virtual probe is employed for navigation through the volumetric blood-flow velocity field. Based on the position and orientation of the probe, we derive Doppler US images from the time-resolved volumetric MRI blood-flow data.

Instead of a photorealistic rendition of the virtual probe, the appearance is equivalent to the probe that was introduced in subsection 7.2.1. The probe consists of a semi-transparent half-open configurable tapered cylinder, and is equipped with two double oblique planes to represent contextual information, as depicted in figure 7.1a.

As in the previous chapter, the virtual probe is provided with a rotation handle, which appears as a green sphere on the right-hand side of the probe. For the US-inspired environment, the green sphere is furthermore used as an orientation marker that visually unites the different linked views. As demonstrated in figure 8.2, the green sphere indicates the positioning of the different visualizations with respect to the coordinate frame spanned by the virtual probe.

8.2.2 Interaction

As opposed to an externally positioned Doppler US probe, the virtual probe is embedded inside the volume. The initial positioning requires a three-point selection, in contrast to the positioning approaches presented in subsection 7.2.2. Placement of the three points is based on either anatomical or blood-flow data, presented on conventional orthogonal planes. The first two points determine the probe long axis, while the third point secures the binormal of the probe. By definition, the binormal is orthogonal to the probe long axis and the view-aligned normal vector.

The user can subsequently interact with the virtual probe to explore the underlying data. Similar to exploration with a hardware US probe, six degrees-of-freedom are provided by the virtual probe. In contrast to interaction with a tangible US probe, interaction with the virtual probe is performed through a standard computer mouse. Mostly simple two-dimensional interaction schemes are employed. The interaction with the virtual probe differs slightly from the approaches presented in the previous chapter.

Interaction with the probe comprises both translation and rotation. These interactions can be performed in the various viewports of the compound view, as depicted in figure 8.2. The viewports are fully linked, and hence interactions that affect the virtual probe are directly updated in all other viewports.




Figure 8.1 (a) Color Doppler ultrasound flow imaging is carried out using a probe, as illustrated on the left-hand side of the figure. This probe is aligned with the flow direction by a skilled operator. In the top row, a wave is transmitted to an object that is moving towards the probe, depicted by a moving blood cell. In contrast to stationary objects, the frequency of the returned wave is higher, due to the Doppler effect, as shown in the bottom row. Doppler ultrasound imaging employs these frequency changes to measure the blood-flow dynamics. (b) Conventionally, flows that move away from the probe are colored in blue, while flows towards the probe are colored in red [85]. The color Doppler ultrasound image was provided courtesy of P. Barbier (<http://www.echobyweb.com/>).




Figure 8.2 Based on the virtual probe, a compound view presents blood-flow information. This includes a Doppler US inspired planar reformat, a 2D-cine MRI inspired planar reformat, and a novel velocity-field volume rendering. In addition, a three-dimensional overview is provided, consisting of a comprehensive flow visualization, which includes anatomical context. Basic two-dimensional interactions are provided to navigate the virtual probe through the volumetric data. All degrees-of-freedom are obtained by in-plane translation, and rotation about the probe axes, comprising azimuth, roll, and pitch.

Translation is performed in the three-dimensional view, translocating the probe in its parallel plane. This restriction leads to an intuitive interaction, omitting complex three-dimensional translations. Convenient and intuitive probe rotation in three dimensions is more challenging. We distinguish three axes of rotation, providing azimuth, roll and pitch interactions. These interactions are performed by means of gestures in the views linked to the main three-dimensional overview. These gestures intuitively correspond to the rotation of the virtual probe.

For the probe-parallel plane, a horizontal gesture induces an azimuthal rotation around the plane normal, while a horizontal gesture pitches the plane. For the probe-orthogonal plane, only rotation around the probe long axis is required. Therefore, a roll is performed for both a horizontal and a vertical gesture. Similarly, a horizontal gesture induces a roll around the probe long axis for velocity-field volume rendering viewport.

To inspect the three-dimensional nature of the velocity-field volume rendering, a tilt interaction is introduced. In contrast to the aforementioned interactions, the location of the virtual probe is not affected. Instead, a vertical gesture tilts the camera position in the corresponding viewport, facilitating inspection of the three-dimensional projection.

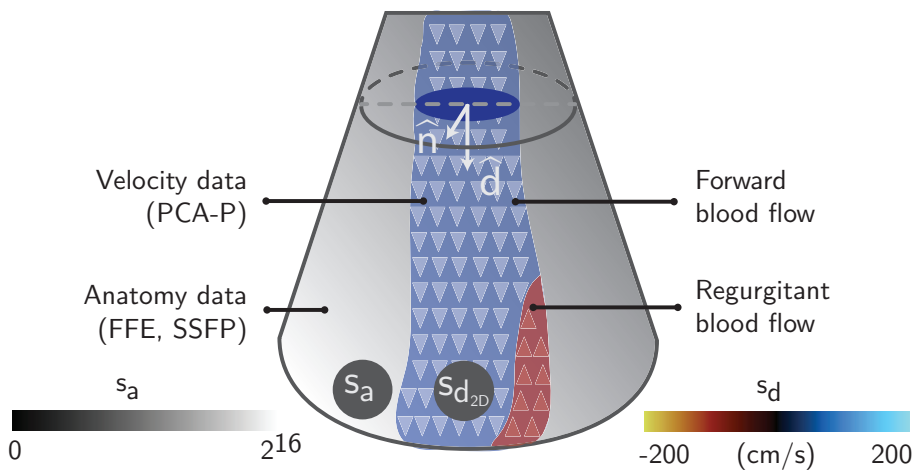


Figure 8.3 The probe-parallel plane provides a double oblique planar reformat, showing both anatomical (s_a) and blood-flow ($s_{d_{2D}}$) information. A speed threshold distinguishes the blood-flow regions, conveying in-plane velocities inspired by color Doppler US. The inspection direction \hat{d} is determined by the long axis of the virtual probe, and \hat{n} defines the plane normal.

8.3 Visualization

8.3.1 Mimicking ultrasound renditions

An integral part of the compound view is the Doppler US inspired planar reformat. The plane entails a double oblique reformat, based on the location and orientation of the virtual probe. It conveys both anatomy and blood-flow velocities, fusing MRI acquisition modalities. This plane is portrayed in the top-left view, as illustrated in figure 8.2. The presentation of the plane is in accordance with color Doppler US imaging, depicting the inspection direction from top to bottom.

A user-defined threshold on the blood-flow speed distinguishes between flow regions and stationary regions, as depicted in figure 8.3. Stationary tissue is depicted by anatomical data. This is typically the fast-field echo (FFE) anatomical data obtained with the flow acquisition. Alternatively, separate high-contrast SSFP data can be used, as described in subsection 3.2.4. On the US-inspired plane, the scalar values that represent anatomy s_a are equal to the values of the anatomical data. The values are displayed with the help of a gray-scale transfer function.

The blood-flow regions, demarcated by the speed threshold, rely on the measured PC-MRI blood-flow data. Based on the velocity information, we mimic conventional Doppler US renditions, requiring an inspection direction. In accordance to the direction of a physical US beam, the inspection direction for virtual exploration is defined by the long axis of the probe, pointing from base to top of the tapered cone.

First, the three-directional velocities are projected onto the oblique plane. Next, the sign of the angle between the normalized inspection direction \hat{d} and the projected velocity vectors is determined. The sign determines whether the probe is aligned parallel or anti-parallel with respect to the velocity field. According to Doppler US convention, flow that moves away from the virtual transducer is mapped to blue, while flow moving towards the probe is mapped to red. The scalar values that represent color Doppler $s_{d_{2D}}$ are mapped by a red-blue transfer function, as depicted in figure 8.4a. First, the velocity vector \vec{v} is projected to the probe-parallel plane. Next, the sign of the angle ζ between the projected velocity \vec{v}_p and the inspection direction \hat{d} , and the blood-flow speed determine the value $s_{d_{2D}}$ as:

$$\begin{aligned}\vec{v}_p(\mathbf{x}) &= \vec{v}(\mathbf{x}) - (\vec{v}(\mathbf{x}) \cdot \hat{n}) \times \hat{n} \\ \zeta(\mathbf{x}) &= \frac{\vec{v}_p(\mathbf{x})}{\|\vec{v}_p(\mathbf{x})\|} \cdot \hat{d} \\ s_{d_{2D}}(\mathbf{x}) &= \text{sgn}(\zeta(\mathbf{x})) \times \|\vec{v}_p(\mathbf{x})\|.\end{aligned}$$

Besides the conventional red-blue color mapping, clinical US imaging apparatus often provide a green emphasis in the visualization indicating either high-speed or turbulent

flows. Both approaches were incorporated into our visualizations. To emphasize jet flows, a one-dimensional red-blue transfer function is adapted to include green for high-speed forward flow. The notion of turbulence, however, is more challenging, as the term is often used incorrectly without intent.

Turbulent flow is characterized by chaotic and stochastic property changes. Transitional flow is observed with a Reynolds number greater than 2000, while turbulence is defined by a Reynolds number greater than 5000 [127]. However, Doppler US imaging does not carry out computations of the Reynolds number to determine turbulence. Instead, the temporal variance is considered, emphasizing rapidly changing flow regions [110, 209], which in practice pass for turbulent regions.

The temporal resolution of Doppler US is far superior to PC-MRI. Hence, the variance over time in PC-MRI data is not suitable to indicate anomalous regions. Instead, we emphasize abnormalities based on spatial variations, using a local coherence measure. We adopt the LPC measure, which was introduced in subsection 5.2.2. This measure indicates coherence between the flow directions in a local neighborhood. Low coherence indicates potential anomalous flow behavior, which is emphasized in green using a two-dimensional transfer function, as depicted in figure 8.4b.

Altogether, the US-inspired plane encodes blood-flow velocities using color. It communicates the direction in relation to the long axis of the virtual probe, in combination with the blood-flow speed. Although this results in a familiar view, the provided information concerning the blood-flow direction remains limited. Therefore, we have extended the visualization with an arrowhead overlay to add direction and speed information, enhancing traditional Doppler US imaging. The arrowheads are based on in-plane projections of the three-directional flow vectors. They are generated in real-time, based on regularly sampled seeds. The size of the arrowheads is scaled with the in-plane speed. The semi-transparent white arrowheads are superimposed on the plane, limiting interference with the underlying visualization, as shown in figure 8.4c.

Besides the arrowhead overlay, we propose a further addition to conventional Doppler US images. In many cases, the regurgitant flow is of interest. Although conventional visualizations highlight regurgitant areas using a red color scheme, the results are dependent on the orientation of the probe long axis. To reveal this dependency, we introduce a widget that enables inspection under a certain angle only.

Whenever the mouse cursor is hovering over the US-inspired plane, a circular widget appears, as depicted in figure 8.4d. Using the scroll wheel of the mouse, an angle can be set interactively, depicted by a cut-out that is similar to a pie-chart division. The regions where the angle between the projected velocities and the inspection direction \hat{d} is smaller than the user-defined angle are omitted from the visualization. This emphasizes areas that deviate from the forward direction by more than the set angle.




Figure 8.4 US inspired planes, based on the parallel plane of the virtual probe. (a) A conventional red-blue color map encodes antegrade and retrograde flow within a speed threshold. (b) In addition, incoherent areas can be emphasized by a green color, using a two-dimensional transfer function based on blood-flow speed and the local phase coherence (LPC). (c) To enhance a conventional US image, overlaid arrowheads can be used to show in-plane directions. (d) An angle widget enables inspection within a user-defined angle, focussing on regurgitant flows.

8.3.2 Enhancing ultrasound renditions

Besides the US-inspired planes, which mimic the color Doppler representation, several other visualizations are included in the compound view. We display an additional plane that shows through-plane velocities, similar to planar PC-MRI blood-flow acquisitions. Moreover, we exploit the volumetric nature of the PC-MRI velocity data using geometry-based flow visualizations, as well as a novel velocity-field volume rendering.

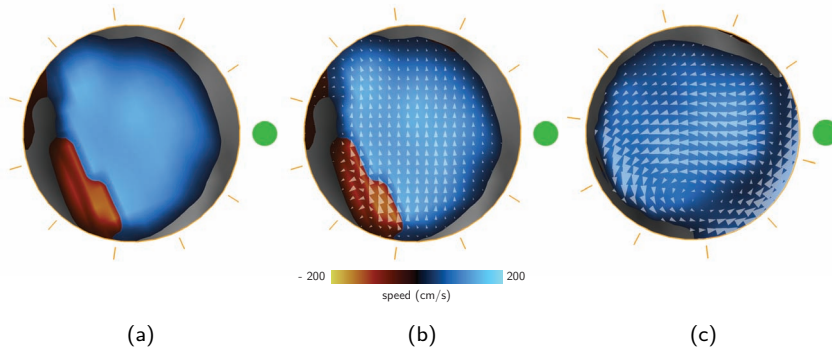


Figure 8.5 Through-plane velocities on the probe-orthogonal plane mimic planar PC-MRI acquisition. (a) The through-plane velocities are mapped by a conventional red-blue color map. (b) In addition, arrowhead overlays can be used to emphasize direction, showing deviations in regurgitant areas. (c) Arrowhead overlays reveal flow rotations. This example shows a projection of the right-handed helix in the ascending aorta, as the flow moves away from the probe, and hence towards the viewer.

2D-cine MRI flow inspired planar reformat - The second planar reformat is circular, and inspired by two-dimensional PC-MRI blood-flow imaging. Besides extensive unsteady volumetric blood-flow measurements, PC-MRI is often employed for planar acquisitions, providing through-plane velocities. To acquire meaningful through-plane velocity information, these planes need to be oriented perpendicularly to the blood-flow direction.

The virtual probe is equipped with an orthogonal plane, which is taken to display these planar blood-flow data. For inspection based on the probe-parallel plane, the long axis of the probe is required to be oriented parallel to the blood flow. Hence, the plane orthogonal to the probe long axis is perpendicular to the flow by construction.

Within the same speed threshold as before, the through-plane velocities are computed by projecting the three-dimensional velocity data \vec{v} to the plane normal, which corresponds to the Doppler inspection direction \hat{d} , as depicted in figure 8.3. The through-plane values s_t are defined as:

$$s_t(\mathbf{x}) = \vec{v}(\mathbf{x}) \cdot \hat{d}. \quad (8.1)$$

They are mapped through a red-blue transfer function to a visual representation, as depicted in figure 8.5a. Similar to the US-inspired planes, arrowhead overlays are introduced to enhance the perception of the blood-flow directions. In figure 8.5b the arrowheads show the in-plane blood-flow directions, including a regurgitant region. Furthermore, the in-plane directions, communicated by the arrowheads, enable the detection of rotations in a vessel. Figure 8.5c shows a projection of the right-handed helix in the ascending aorta.

Velocity-field volume rendering - While planes depict valuable characteristics of the blood flow and surrounding anatomical structures, the information remains restricted to a two-dimensional region. We introduce a novel volume rendering of the velocity field, adding a three-dimensional component to the compound view.

Inspired by three-dimensional US visualization, we propose a volume rendering of the velocities, communicating the volumetric blood-flow information. The shape of the virtual probe is employed as bounding geometry for the volume rendering, which comprises a raycasting approach [202]. Conventional volume rendering approaches are typically based on scalar-valued volumes. We aim for an holistic view of the velocity field [62]. A novel velocity-field volume rendering is introduced, which interactively projects blood-flow characteristics, based on the location and orientation of the virtual probe.

Similar to the Doppler US inspired planar reformat, the long axis of the virtual probe is used as inspection direction, which can be altered interactively. For each sample along a ray, the angle θ between the three-dimensional velocity field and the inspection direction is computed in real-time. Similar to the US inspired planar reformat, the sign of the angle θ is extracted to distinguish if the flow is forward or backward, relative to the inspection direction. In detail, the scalar values $s_{d_{3D}}$ used for the volume rendering are defined as:

$$\theta(\mathbf{x}) = \frac{\vec{v}(\mathbf{x})}{\|\vec{v}(\mathbf{x})\|} \cdot \hat{d}$$

$$s_{d_{3D}}(\mathbf{x}) = \text{sgn}(\theta(\mathbf{x})) \times \|\vec{v}(\mathbf{x})\|$$

The values within the speed threshold are composited front-to-back, and colors are mapped after data interpolation, using a red-blue transfer function. When interacting with the probe, the volume rendering is updated in real-time, depicting the flow characteristics from different angles by changing the location and orientation of the virtual probe.

The resulting renditions communicate blood-flow speed and the relative direction throughout the inspected region. This leads to a dense representation in three dimensions, as the lumen implicitly demarcates the morphology (Fig. 8.6a). However, we strive for insight into the dynamics in the interior of the flow field, and this information is occluded.

To reveal anomalies, different aspects of blood-flow velocity field can be inspected. In particular, excessive variations in the flow direction potentially indicate aberrant flow regions. Based on this observation, we investigate two distinct flow features, as part of the velocity-field volume rendering.

As first feature, the angle between the inspection direction \hat{d} and the velocity field \vec{v} is used. Instead of computing a single angle, the local variation of the angles is

computed by approximating the first-order derivative in a local neighborhood. As an alternative to this angle variation, we can employ the LPC measure [35], which is independent of the inspection direction.

For ease of use, all features are mapped to the color Doppler transfer function, as depicted in figure 8.6d. First, we determine the sign of the flow direction with respect to the inspection direction. Blood flows along the inspection direction are mapped to the right half of the transfer function, while regurgitant flow is mapped to left half. Then we determine the feature intensity, where a high angle variation or low coherence indicate high feature values. These are color mapped through the selected half of the transfer function. The angle derivative and coherence features are depicted respectively in figures 8.6b and 8.6c. An opacity transfer function is added, which by default suppresses lower feature values, and hence masks flow regions that appear laminar. As a result, anomalous flow patterns become the focus of the visualization, catching the attention of the user.

Anatomical context - Emphasis on anomalous blood-flow patterns is obtained by suppressing coherent flow regions. Consequently, anatomical context information provided by the luminal geometry is lost. An overlay is introduced to provide morphological context in addition to the velocity-field volume rendering. Instead of showing the anatomy in full detail, we outline the morphology as a spatial reference for the velocity-field volume rendering.

Separately acquired high-contrast anatomy scans would typically provide many details. For a context rendition only the rough global structure is required. We approximate the vessel lumen by calculating a tMIP volume, as described in subsection 5.1.2. Then, we project the maximum gradient magnitude of the tMIP volume. Lower values are suppressed by using transparency, while the high values are shown in white, providing a neutral color that is not contained in the Doppler color map. This maximum projection is overlay on the coherence-based velocity-field volume rendering, as presented in figure 8.6d.

Geometry-based flow visualization - As another part of the compound view, geometry-based flow visualizations are incorporated to fully exploit the volumetric PC-MRI data. The techniques presented in section 7.4 are adopted for this purpose, including anatomical context by means of volume rendering of the tMIP data.

An example of the compound view is depicted in figure 8.7. An illustrative particle visualization is shown in the 3D overview, while the probe-parallel and probe-orthogonal planes, as well as the velocity-field volume rendering, are depicted on the left-hand side. The regurgitant flow in the ascending aorta is clearly visible in both planes. It is depicted through contrasting colors, where red indicates the retrograde flow. Incoherent regions are emphasized in green, which in this example are found to be the boundary regions with regurgitant flows. The volumetric region of incoherence is indicated by the velocity-field volume rendering. Furthermore, the particle visualization reveals that some particles go astray, because of the regurgitant area.

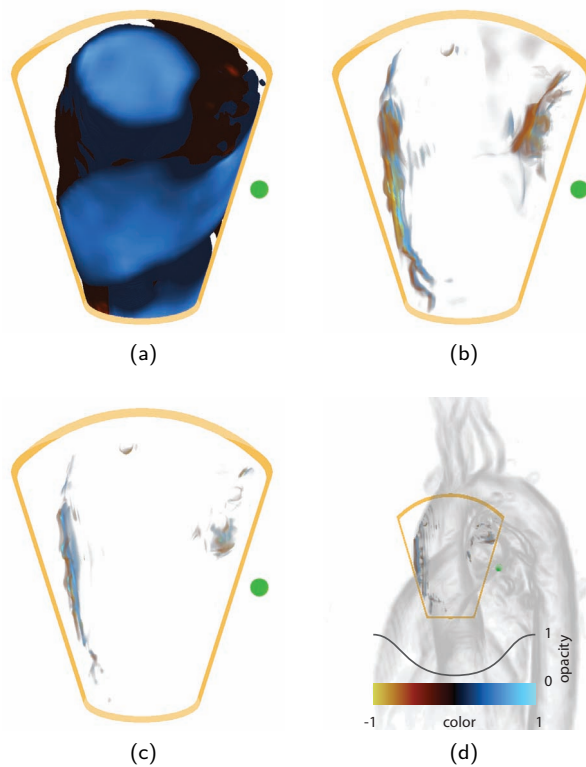


Figure 8.6 Velocity-field volume rendering (a) Without specific opacity modulation, velocity-field volume rendering, inspired by Doppler US, leads to a dense and cluttered visualization. (b) Alternatively, velocity-field volume rendering sparsely emphasizes abnormal flow. The variation of angles between the inspection direction and the blood-flow directions can be used, in combination with opacity modulation (see figure d). (c) Incoherent flow regions are highlighted, based on the LPC measure [35]. (d) The sparse representation in figures c and d lack anatomical context, which can be provided by a maximum gradient-magnitude projection of the tMIP data.

8.4 Evaluation

In the introduction of this chapter in section 8.1, the value of detecting regurgitant flow was motivated by the intention to investigate valvular heart disease. For this application, clinical blood-flow measurements are generally performed by US imaging. However, there is a growing interest to inspect volumetric blood-flow data for valvular anomalies, showing regurgitant flows, as well as blood-flow jets.

In clinical research, the current focus is to optimize parameters for the acquisition of blood-flow around the valves. In contrast to a full thorax blood-flow acquisition, as presented in figure 8.7, a higher temporal resolution is desirable, at the expense of the spatial resolution. A better temporal resolution facilitates the detection of blood-flow




Figure 8.7 Compound view of cardiovascular blood flow in the full thorax of a healthy volunteer. Conventional regurgitant flow occurs in the ascending aorta, superior to the aortic valve. This region is shown in the probe-parallel and the probe-orthogonal planes, as well as the velocity-field volume rendering. Some particles move astray, due to the regurgitant flow region.




Figure 8.8 Compound view of blood flow near the aortic valve of a healthy volunteer. To closely inspect the hemodynamics around the valves, the field-of-view is narrowed to gain temporal resolution. Instead of a volume rendered anatomical context, an cine anatomy slice is acquired. The velocity-field volume rendering shows incoherence in noisy regions surrounding the vessel. Some regions are strongly regurgitant. They are conveyed in yellow, based on the Doppler US inspired transfer function (Fig. 8.6d).

jets and regurgitation, even though PC-MRI acquisition is unable to outperform US imaging in this respect. A first volunteer PC-MRI scan provides a field-of-view that is limited to the valvular region, increasing the temporal resolution up to $14ms$. This is equivalent to approximately 50 phases per cardiac cycle. Cine anatomy slices are currently acquired to provide the necessary anatomical context.

This PC-MRI volunteer scan provides blood-flow velocities surrounding the aortic valve. An example visualization of the blood-flow at peak systole is depicted in figure 8.8, showing the different views. With the reduced field-of-view, particle traces in the 3D overview are no longer beneficial, while line traces could still communicate the structure of the velocity field. The hemodynamics can be further inspected through animation and interaction with the virtual probe.

Qualitative evaluation - Feedback was acquired from two physicians who are actively involved with the advancements in the emerging field of time-resolved volumetric PC-MRI blood-flow acquisition. They have particular interest in the blood-flow patterns involved with valvular diseases. The physicians indicate that they currently lack the knowledge on what to visualize and measure in the intricate blood-flow velocity data. To gain this knowledge, they indicate that interactive visual inspection is absolutely key. Based on demonstration movies, the physicians indicate that the presented interactions with the compound view are still slightly more cumbersome than operating an ultrasound probe. Nevertheless, the physicians emphasize that the presented interaction techniques are a step in the right direction. The combination of views with the ability to visualize blood flow similar to color Doppler ultrasound imaging is considered very intuitive. It is deemed likely that this approach will attract the attention of more involved physicians soon.

Performance - The combination of various visualizations affects the overall performance. Nevertheless, the compound view generally achieves interactive frame rates. The performance bottleneck is found in different volume projections, such as the velocity-field volume rendering, the maximum gradient-magnitude projection, and the volume rendering of the anatomical context. The performance of these techniques depends on the viewport size. By default, the layout of the compound view is divided into four viewports, as depicted in figures 8.7 and 8.8. Each of these viewports can be enlarged to a single full-size viewport.

First, we examine the performance of the individual components of the compound view, based on the velocity data from the full thorax acquisition of a healthy volunteer, as depicted in figure 8.7. The performance tests were carried out using a conventional computer system with a dual-core processor, 6GB internal memory and an NVidia GeForce 570GTX graphics card. The Doppler US inspired planar reformat and the 2D-cine MRI flow inspired planar reformat both render at approximately 58 frames per second, with a viewport size of 512×512 pixels. The velocity-field volume rendering performs at 29 frames per second without a an overlaid maximum gradient-magnitude projection for anatomical context. With this projection, the velocity-field volume rendering performs at a mere 8 frames per second. However,

Table 8.1 Rendering performance of the Doppler US inspired visualizations. All tests are performed using a viewport size of 512x512 pixels. In the compound view, smaller viewports are allocated to the individual visualizations, resulting in interactive framerates for the combination of views.

Visualization component	Framerate (FPS)
Doppler US inspired planar reformat	58
2D PC-MRI inspired planar reformat	58
Velocity-field volume rendering (without anatomical context)	29
Velocity-field volume rendering (with anatomical context)	8
3D flow visualization (without anatomical context)	≈ 30
Compound view (with anatomical context)	≈ 15

the velocity-field volume rendering is commonly active in a much smaller viewport, as part of the compound view. In this case, the overall interactivity of the system is not affected.

The performance of the geometry-based flow visualizations in the 3D overview varies, depending on the used technique. The corresponding performance evaluation was discussed in section 7.5. On average, these flow visualization techniques render at about 30 frames per second, and interactivity is maintained when combined with a volume rendering of the anatomical context.

Furthermore, we have measured the performance of the combined four components of the compound view, as depicted in figures 8.7 and 8.8. The complete compound view renders at approximately 15 frames-per second, enabling smooth animation and real-time interaction. An overview of the performance test is provided in table 8.1.

8.5 Discussion and conclusions

The described compound view communicates PC-MRI blood-flow information through various linked views, thereby addressing physicians involved with ultrasound and MRI blood-flow imaging. The diverse views enable detection of anomalous blood-flow patterns, communicating the information to physicians in a familiar manner. Visualization techniques from the established field of Doppler US imaging were adopted, and combined with expedient three-dimensional PC-MRI flow visualization techniques.

The compound view provides a valuable interactive synopsis of the blood-flow information for different cardiovascular applications. In particular, we have introduced the hemodynamics of patients suffering from valvular heart disease, typically associated with blood-flow jets and regurgitant flows. PC-MRI imaging of valvular anomalies is in its infancy, no pathological blood-flow data was available. However, volunteer studies show potential, and the visualizations facilitate the qualitative analysis.

Currently, the acquired field-of-view is narrowed, and planned around the valve of interest, in order to gain temporal resolution. As a consequence, the scope of exploration is still limited for this application, as well as the ability to present an anatomical context. In the future, acquisition techniques may improve, enabling a larger field-of-view, in combination with an better temporal resolution. Alternatively, the measured blood-flow data may be complemented with CFD simulated flow velocities.

In conclusion, we have presented a compound view that comprises diverse views, based on volumetric PC-MRI blood-flow velocity data. Besides geometry-based three-dimensional blood-flow visualizations, we introduce two planar reformats based on the position and orientation of a virtual probe. In particular, the probe-parallel plane constitutes the basis of an US-inspired visualization, simulating velocities along a Doppler US beam. These velocities are depicted through a red-blue transfer function. Additionally, the probe-orthogonal plane is used to mimic two-dimensional through-plane PC-MRI velocities, also color coded using a red-blue transfer function.

For both planes associated with the virtual probes, enhancements were proposed to the familiar visualization techniques. For the probe-parallel plane, we have presented an angle widget, enabling an interactive definition of the flow-angle range to be inspected. In a healthy volunteer, blood-flow along the inspection direction is typically laminar. These flows can be omitted from the visualization. Furthermore, both planes were equipped with arrowhead overlays, enhancing the perception of the in-plane blood-flow directions.

We have introduced a novel velocity-field volume rendering, using the shape of the virtual probe as a bounding geometry. Besides simulated Doppler US values, other features are used to emphasize potentially anomalous regions. In detail, the velocity-field volume rendering may utilize the local angle variation between the blood-flow velocities and the probe inspection direction. This dependency on the inspection direction mimics the Doppler US insonation. In this case, the results strongly depend on the placement of the probe. Instead, the velocity-field volume rendering may also use the local phase coherence, which is independent of the inspection direction. Large angle variations, or low coherence, indicate flow regions that are worthwhile to inspect. To preserve anatomical context, a maximum gradient-magnitude projection was employed.

All views are linked interactively, and corresponding positions can be verified by a green sphere that indicates the virtual probe location. The combined visualization renders at interactive framerates, and offers smooth animation of the hemodynamics.

Involved domain experts indicate that the presented visualization techniques can facilitate their research, aiming to understand what to visualize and what to measure in the intricate time-resolved blood-flow velocity data. The interactivity of the approach is deemed essential. Some further improvements on the interaction may be required. The compound view with the Doppler US inspired visualizations are considered intuitive, and valuable for the qualitative analysis.

In the future, the visualizations will be more extensively employed and validated for qualitative inspection of blood-flow behavior in patients suffering from valvular heart disease. This process will potentially lead to further improvements of the visualization techniques, in the context of the valvular application. Thereafter, domain experts envision similar visualization and interaction techniques to be incorporated into the acquisition process. In the case of valvular disease, this would enable the rapid detection of blood-flow jets in the three-dimensional domain. Subsequently, an accurate quantitative analysis should be performed along the jets, providing further insight into the condition of the valve.

"Nothing is so simple that it cannot be misunderstood."

Teague's Paradox

Discussion and conclusion

9

9.1 Contributions

In this thesis, we have introduced a variety of visualization and interaction techniques for the qualitative analysis of time-resolved and volumetric blood-flow velocity data. After the general introduction in **chapter 1**, the value of measuring and understanding cardiovascular blood-flow data was substantiated in **chapter 2**. This chapter addressed the benefits of blood-flow information in various medical conditions. We have mainly used quantitative blood-flow information, acquired by magnetic resonance imaging. The acquisition process was discussed in **chapter 3**.

Subsequently, **chapter 4** described the technical background, motivating the need for abstract flow visualizations of large and intricate blood-flow velocity fields. The provided overview of related work consisted of two parts, addressing relevant illustrative visualization methods, as well as established flow visualization techniques. The rationale provided in these chapters has given rise to the following key contributions:

Abstraction - In **chapter 5**, several approaches were introduced, through which we obtain simpler, more abstract representations of the intricate blood-flow data. First we have presented two temporal projection techniques, eliminating the time dimension to attain a static representation. On the one hand, we have described the tMIP volume, which provides a maximum projection of the blood-flow speed over time. Although the tMIP was introduced by others, we have demonstrated new uses of this valuable concept throughout our work. In addition, we have introduced the novel tMOP abstraction, comprising a static representation of the velocity-field orientations.

Secondly, another abstraction was introduced in terms of an application-specific segmentation technique that extracts the luminal geometry from the blood-flow velocity data. We have discussed an active surface approach that exploits the full velocity information, based on two distinct coherence measures. The extracted luminal geometry provides a reasonable approximation of the cardiovascular morphology, and can serve as an anatomical context for blood-flow visualizations.

Moreover, chapter 5 introduced abstraction of the blood-flow data by means of spatiotemporal hierarchical clustering. Two dissimilarity measures were employed to cluster the velocity data, which can then be visualized at a user-defined level-of-detail.

Visualization and interaction - Based on our hierarchical clustering approach, which utilizes the four data dimensions, i.e., space and time, we have devised a number of illustrative visualization approaches. Initiated from the cluster centers, we have introduced 3D-seeded and 4D-seeded patharrows, in combination with an illustrative anatomical context.

In **chapter 6**, we proposed various novel and enhanced visualization techniques, tailored towards inspection of thoracic arteries. An existing probing technique was adopted, enabling local inspection of the blood-flow velocities. We have introduced exploded planar reformats, dynamically seeded pathlines with animated highlights, and flow-rate arrow-trails, combined with an illustrative anatomical context.

Continuing on the previous chapter, **chapter 7** introduced a novel virtual probe, enabling exploration of the blood-flow field throughout the cardiovascular system. The cylindrical probe defines a volumetric region-of-interest, providing the basis for various innovative geometry-based visualization techniques. In detail, we have introduced illustrative particles, interactively generated long pathlines, and nested pathsurfaces. All visualizations were combined with an anatomical context, based on a real-time and gradient-based direct volume rendering of the blood-flow speed lumen.

The probing and visualization techniques proposed in chapters 6 and 7 enable a qualitative analysis of the blood-flow velocity data, facilitating detection of aberrant blood-flow patterns. To evaluate the effectiveness of the visualizations and interaction techniques, we have conducted interviews with several domain experts, based on structured questionnaires. The evaluation results were covered in both chapters.

Finally, in **chapter 8** we have proposed a visualization framework, inspired by color Doppler US imaging, and exploiting benefits of the PC-MRI velocity data. The established color Doppler US visualizations were mimicked, based on the volumetric PC-MRI blood-flow velocity data. In addition, we have simulated through-plane velocities, similar to planar PC-MRI blood-flow acquisitions. To communicate the true volumetric nature of the PC-MRI data, we have included geometry-based flow visualizations, as well as a novel velocity-field raycasting, which emphasizes anomalous flow regions. The visualization framework enables synergies between the different domain experts, and allows to compare the qualities of the respective imaging modalities.

9.2 Discussion

A thorough discussion was provided with each of the key chapters throughout this thesis. This section addresses the considerations from a birds-eye view, providing a critical assessment of the presented visualization research.

Methodological considerations - Abstraction: The variety of techniques presented in this thesis focussed on qualitative analyses of the measured blood-flow data. The PC-MRI acquisition provides blood-flow velocities that are in accordance with the actual physical speed and direction of the blood flow. A quantitative analysis is therefore also of great importance to gain insight into the hemodynamics. Throughout our work, we have carefully addressed the quantitative nature of the data, for instance by paying attention to the choice of color codings. In general, it is essential to communicate data characteristics that are not easily quantifiable, such as intricate blood-flow patterns.

For the reduction of visual clutter and occlusion, our abstractions of the cardiovascular anatomy and the blood-flow velocities have proven appropriate. Nevertheless, any type of abstraction should be addressed with some care. Simplifications inherently discard details, implying that potentially valuable information can be lost. It is essential that abstractions are underpinned by domain knowledge, and that experts are basically aware of the type of applied simplifications when interpreting the renditions.

Choosing the appropriate level of abstraction remains challenging, and depends on personal preference. In comparison to conventional visualizations in clinical practice, we have applied reasonably extensive abstractions with an illustrative character, opting for visual legibility. However, the level of abstraction may also be carried further. One can think of higher-level abstractions, such as straightening or unfolding of the vasculature, or exploded views of the cardiovascular system. It is worthwhile to incrementally reach an appropriate abstraction level, in consultation with the physicians.

Abstraction was enabled through segmentation of the luminal geometry, approximating the cardiovascular morphology. In this respect, we have evidenced that the segmentation accuracy improves when velocity information is incorporated into the segmentation process. Nevertheless, the current segmentation process is not viable for daily practice. Some limitations near the volume boundaries need to be resolved, and further research should show if a global parametrization of the model is possible. Additionally, a rigorous validation is needed. Segmentation as a research topic was not our primary focus. Instead, we have employed segmentation results to convey an anatomical context with the various flow visualizations.

Besides the segmented geometry, we have used gradient-based DVR to communicate the anatomical context. The associated transfer functions are currently set manually, comprising a time-consuming process that is not convenient for inexperienced users. An automated transfer function design seems feasible for the tMIP data, as it contains strong boundaries between the blood-flow lumen and the background. Transfer function automation remains challenging for anatomical MRI scans.

Abstraction was furthermore obtained by spatiotemporal hierarchical clustering of the blood-flow data. We have demonstrated that such an approach is promising, with an intuitive level-of-detail control. However, the notion of spacetime remains challenging to the human mind. While the physical interpretation of spatial cluster boundaries is already arguable, the virtue of spatiotemporal clusters is even more complicated to fathom. While some ad hoc choices were unavoidable, we have used the results of our four-dimensional clustering to visualize the unsteady blood-flow behavior.

Methodological considerations - Visualization: Various presented illustrative visualization approaches have proven effective, when judging from the acquired feedback. The applied abstractions enable communication of the essential information. However, for some physicians the illustrative character of the renditions gives rise to suspicions concerning the validity of the representation. Remarkably, most physicians personally appreciated the art-inspired approaches, while doubting that the opinion of their coworkers would be the same. Taking this into consideration, illustrative renditions may well be a suitable choice, although getting used to these methods might be necessary. Aesthetics play a role in the user experience, and foster sharing of the obtained findings amongst the physicians. A middle ground should be found by attenuating some of the emphatic illustrative characteristics, to gain trust with the skeptic observer. An example was given by the exploded planar reformats, locally communicating the acquired voxel information.

Considering the geometry-based flow visualizations, including the ensuing feedback, it becomes apparent that these techniques are largely complementary. Based on our findings, it seems improbable that a single visualization alone will be sufficient to allow clinical research to gain new insights into the blood-flow data.

Methodological considerations - Exploration: Exploration of the data should be performed with several different visualizations to grasp the intricate hemodynamics. Each visualization needs a certain amount of parametrization, taking up valuable time. For example, the suitable placement of seed positions in the spatiotemporal domain is a time-consuming process. In the spatial domain, physicians have a good understanding which regions to inspect. With respect to the temporal domain, seeding flow visualizations only at peak systole severely reduces the amount of communicated information, and hence increases the chance to miss important aspects.

The exploratory nature of the presented visualization techniques generally comes to full advantage through animations, in which case the temporal behavior of the blood-flow patterns is conveyed. After exploration and visual analysis, the findings need to be reported to fellow physicians. This often requires a different visualization approach, capturing the temporal behavior in static images, instead of an animation. Although reporting of the blood-flow information was not our focus, we have presented some techniques that are suitable. Illustrative particles, for instance, statically include a notion of displacement over time using the speedlines. Furthermore, we have presented a static representation of the spatiotemporal blood-flow, based on the cluster results.

Exploration of the velocity data is furthermore time-consuming because of the user interaction. For the different probing techniques, we have improved the ease of interaction, reducing the degrees-of-freedom. While some aspects can be automated, physicians generally want to make adjustments to get the desired results. This is especially true in the uncertain exploration phase. The search for the unknown, which is essential in clinical research, will take a considerable amount of time, and can only be facilitated by convenient exploratory tools. This laborious process will lead to indicators that can be used in more specific, and often task-driven, visualizations.

Technical considerations - Graphics hardware: To make strides in visualization research, a considerable amount of engineering work is necessary. For instance, we have exploited the traditional graphics pipeline, using general purpose computing on the GPU, to gain the necessary interactivity. In the last few years, there has been a rise of computational languages, facilitating hardware-driven parallel processing. These languages are equally well suited to speed up the algorithms, while rendering will continue to rely on the graphics pipeline.

Technical considerations - Software development: The main focus, however, remains on devising innovative visualization techniques, and hence the implementation work typically produces a relatively underdeveloped prototype. This unfortunately hampers early adoption of the visualization techniques, as well as more intensive and experience-driven feedback from the domain experts.

A solid prototype requires a user-friendly interface, with a limited number of parameters. Stable releases should be deployed regularly, in order to gain and incorporate the desired feedback. Moreover, flexibility in terms of data processing is needed, enabling input of various data types, and straightforward output to images or movie sequences. Furthermore, the prototypes should be able to deal with variations in data formats, since the velocity encoding process is not fully standardized. Especially usability and flexibility requirements stand out.

It takes a great deal of effort to fulfill such requirements. With limited time available, tradeoffs have to be made between these engineering tasks and the actual research. These considerations are characteristic for visualization research, which always operates on the frontier between computer graphics and a certain field of expertise.

Medical considerations - Validation: Throughout the collaboration, physicians involved with clinical research generally had a clear opinion on what they appreciate, and what they dislike about a particular visualization. It is typically not possible to query the expert about his explicit requirements for a visualization beforehand. The domain experts naturally cannot assess the possibilities of modern-day visualization techniques. Therefore, it is up to the visualization community to propose appropriate methods, familiarizing with the application-domain knowledge. This course of events especially holds for medical visualization, which is not always recognized in other visualization areas.

It is difficult to prove the effectiveness of a visualization. The value of the visualization should therefore be measured through validation. To this end, we have employed questionnaires, which enabled us to gain an understanding of the different opinions of the involved physicians in a compact time frame.

Medical considerations - Segmentation: The clinical research involved with PC-MRI blood-flow velocity data generally focusses on complex pathological conditions, with severe malformations of the cardiovascular anatomy. Besides the flow information, morphology of the anatomical structures remains important. Clinical research would greatly benefit from a fast and reliable segmentation of these complex cases, preferably utilizing different MRI sequences that include information about the morphology. Despite the technical challenges that remain, physicians currently require a serviceable tool. Automation of volumetric segmentations is necessary, and some specific user interventions are desirable to improve the segmentation results.

Medical considerations - Visualization: In many situations, generic viewing techniques are employed. Existing software to generate more extensive blood-flow visualizations often emerge from a medical setting. Hence, some well-established rules of thumb in the visualization field are often disregarded. For instance, we encountered inappropriate color mappings and seeding distributions. While the visualization techniques improve, a certain awareness of some perceptual concepts with the domain experts is beneficial in their research.

Ethical considerations: The various visualizations throughout this thesis conveyed blood-flow dynamics, based on different data sets, acquired from healthy volunteers and patients. The used MRI acquisition techniques involved very limited safety risks for the subjects, and hence we consider the measurements justifiable.

For all volunteer studies, written informed consent was obtained, prior to participation. The patient studies were approved by the St. Thomas and Guy's ethics committee. For patients suffering from an aortic dissection, the studies were carried out under the number 08/H0809/49. Studies of patients with congenital heart defect were carried out under the number 08/H0804/134.

9.3 Future research

Visualization research - In the methodological considerations, we have already touched upon several aspects that require further research, for instance with respect to the velocity field abstraction through hierarchical clustering. Currently, the clustering is not based on physical characteristics, and follows a purely mathematical approach. For the visualization, dealing with the notion of spacetime remains challenging. The current visualization depends on the animation of the clustered blood-flow, which provides an effective overview, prior to further exploration. In the future it is worthwhile to investigate other ways to communicate the four-dimensional clusters, providing detailed depictions of the blood-flow for further analysis.

From a visualization research perspective, it is worthwhile to investigate higher levels of abstraction to communicate the blood-flow dynamics. Recently, straightening of the aorta was presented to facilitate the comparison of the blood-flow velocities [6]. Other approaches are conceivable, such as unfolded and exploded views. Also concerning the illustrative character of the representation, further abstractions may be worthwhile to investigate. For reporting purposes, a comic-strip format is imaginable, if the key time frames and camera positions can be detected automatically or semi-automatically.

Another aspect for future work is to convey the reliability of the visualization. The blood-flow measurements involve a considerable amount of uncertainty. Furthermore, various types of uncertainty propagate through the post-processing pipeline, including filtering and visualization. Most visualizations communicate the information as being the truth, and omit these uncertainties. Especially for qualitative analyses in clinical research, inclusion of uncertainty information in the visualization is important. The uncertainties are often difficult to quantify and visualize, which provides an important challenge for future work.

For each visualization technique, it is essential to minimize the number of user parameters. Many parameters should be specified automatically, or at least initialized at a practicable value. A good initialization of these parameters requires domain knowledge, and hence advice from the involved physicians is necessary. In addition, tradeoffs between exploratory freedom and software automation need to be considered.

In the field of flow visualization, we observe an increased interest in Lagrangian coherent structures to describe time-dependent flow. In particular, the finite-time Lyapunov exponent gains considerable attention, and has been recently employed for PC-MRI blood-flow velocity fields [123]. Dealing with time in visualization of unsteady blood-flow velocity data is challenging in itself. In the future, this will become even more challenging, when multiple unsteady velocity fields are acquired over time. This would be valuable for patient follow-ups, as well as for post-operative assessments of the interventions. Time will be taken to the next level, and specific visualizations should be devised that enable the comparison of the consecutive blood-flow fields.

Flexibility is required in the provided visualization and interaction tools, as differently specialized experts typically need specific tools. There are interesting future opportunities in terms of interaction. In our work, we have focussed to reduce the degrees-of-freedom of interaction, which is valuable for use on a conventional computer system. Various other devices allow intuitive interaction with more degrees-of-freedom. For instance, the virtual probe could be operated in this way, simulating interaction with an ultrasound probe. Dedicated prototypes could be used by physicians with a distinct expertise, facilitating multi-disciplinary communication and comparison. Besides special three-dimensional navigation devices for computer systems, the gaming industry is also introducing various new interaction technologies that are worthwhile to investigate, such as the Nintendo™ Wii controller, and Microsoft's™ Kinect.

Clinical research: There are steady advances with respect to the PC-MRI blood-flow measurements, reducing acquisition time and enhancing quality. Besides improvement of the acquisition techniques, post-processing of the acquired velocity data remains an important topic of research. Various image processing methods are required to remove acquisition artifacts, such as phase wraps. Furthermore, extensive image analysis is necessary to identify features in the intricate blood-flow velocity data. Various techniques can be employed, such as flow field decomposition, or topological analysis. Furthermore, image analysis is needed to further validate the accuracy of the blood-flow acquisition techniques.

With a growing understanding of the hemodynamics, and the relevant characteristics for diagnosis and prognosis, an increasing amount of knowledge has to be incorporated into the algorithms. For instance, different MRI sequences provide diverse tissue contrasts, which may serve various purposes. In the future, prototypes and tools should cope with a fusion of these data sources, exploiting the strengths of each particular sequence. For instance, segmentation can benefit from such an approach, incorporating anatomical as well as blood-flow information to capture the morphology. Inclusion of domain specific knowledge requires a close and ongoing collaboration between clinical research, and the various involved technical disciplines.

Fluid simulations research: Besides blood-flow measurements, it is worthwhile to investigate patient-specific simulations of the hemodynamics. Simulated flow fields are not affected by noise, and allow for high spatiotemporal resolutions. Cooperation between the respective fields potentially leads to interesting new insights. Initially,

it is worthwhile to perform a cross-validation, investigating the similarities and differences between measurements and models. Neither measurements nor simulations capture the entire ground truth. However, the individual strengths can be exploited by combining the approaches. For instance, the measured flow can be introduced as a restrictive condition for the simulation, or the measured flow can be enhanced through local computational modeling. Moreover, simulations could be employed to predict hemodynamic changes in the course of a disease, based on an initial blood-flow velocity measurement. Alternatively, simulations can provide decision support for different types of interventions.

Personal outlook: In clinical research, the current focus is largely based on a quantitative analysis of hemodynamic characteristics. A quantitative analysis provides strong diagnostic or prognostic indicators. However, a large amount of information is contained in the intricate blood-flow patterns, which provide insight into the efficiency of the blood-flow circulation, and potentially provide early predictors of pathogenesis.

Especially for the time-resolved three-dimensional case, the variety of blood-flow patterns is still largely unknown, even in the case of healthy persons. Although many common patterns are known for the healthy case, variations with respect to age, gender and ethnicity are not completely understood. In addition, there is the vast number of pathological cases for which understanding of the blood-flow patterns is limited.

On a personal note, I would envision a knowledge database that provides easy access to the variety of blood-flow patterns, under diverse circumstances, preferably enriched with quantitative parameters. Such a database would constitute an atlas of well-known blood-flow patterns, much like widely used textbooks about the human physiology. Current visualization techniques, such as the work presented in this thesis, can provide the necessary exploratory tools to grasp the blood-flow patterns. The final representation of the blood-flow patterns may need further generalization and simplification, in order to communicate the hemodynamics to a larger audience. Therefore, manual illustrations or animations may be necessary.

The availability of such a knowledge about various blood-flow patterns would facilitate the ongoing investigation concerning the clinical relevance of the measured hemodynamics, and may accelerate the use of this information for diagnosis and prognosis.

9.4 Conclusion

In conclusion, we believe that the real-time illustrative visualization of time-resolved volumetric blood-flow data enables an effective communication of the hemodynamic behavior. By means of validation studies with the involved physicians we have exposed the advantages and weaknesses of the presented techniques. The physicians recognize the value of the exploratory visualizations for clinical research, and indicate that advanced interactive visualizations allow faster detection of pathologies. We were able to present exemplary pathological cases. Time will reveal what new insights can be obtained by means of interactive exploratory qualitative analyses.

"I used many times to touch my own chest and feel, under its asthmatic quiver, the engine of the heart and lungs and blood and feel amazed at what I sensed was the enormity of the power I possessed. Not magical power, but real power."

Stephen Fry

The cardiovascular system



Overview of the cardiovascular system

Throughout this thesis, we have focussed on the cardiovascular system. In particular, we were concerned with the heart and the main thoracic arteries. In visualization research, it is common to work closely with the domain experts. In our case we have closely collaborated with physicians in clinical research of blood-flow velocity data.

Inevitably, terminology from the medical domain has been used throughout this thesis, and many terms were assumed to be known. Since our work largely addresses a technical audience, several common terms in context of the cardiovascular system will be further introduced in this appendix. Besides the main anatomical structures, we also briefly describe the blood-flow circulation in the human heart.

Anatomical structures

The important anatomical structures of the heart and the thoracic arteries are depicted in figure A.1. The heart is composed of two ventricles, and two atria. The main thoracic arteries include the aorta, the pulmonary artery, and the vena cava.




Figure A.1 Human heart anatomy and blood-flow circulation. (© Alila07 | Dreamstime.com)

Blood-flow circulation

The arrows in figure A.1 indicate the direction of the blood flow in the heart. The red arrows indicate oxygen-rich blood, while the blue arrows indicate oxygen-poor blood. The course of the blood flow throughout the heart beat will be described, based on the numbers one to ten in the figure.

① First we follow the course of the oxygen-poor blood, which returns from the body through the superior and inferior vena cava, and enters the right atrium. Both atrioventricular (AV) valves are closed at this time. ② Next, the right AV valve opens, ③ and the right ventricle is filled with the deoxygenated blood. The period of relaxation of the heart is referred to as diastole. When the heart is fully relaxed, it is referred to as peak diastole. ④ Upon contraction of the heart muscle, the right AV valve closes, and the pulmonary valve opens. The blood is transported to the lungs for oxygenation through the pulmonary artery. The period of contraction of the heart is referred to as systole. When the heart is fully contracted, it is referred to as peak systole.

In parallel there is the course of the oxygen-rich blood. While the deoxygenated blood enters the right atrium, ⑥ oxygenated blood returns from the lungs, ⑦ and enters the left atrium through the pulmonary veins. ⑧ Subsequently in the diastolic phase, the left AV valve opens, and the left ventricle is filled with the oxygen-rich blood. ⑨ Upon contraction of the heart muscle, the left AV valve closes, and the aortic valve opens. In the systolic phase, the oxygen-rich blood is pumped into the aorta, which distributes the blood throughout the body. ⑩ For instance the branches of the aorta, respectively the brachiocephalic trunk, the left common carotid artery, and the left subclavian artery, supply blood to the upper part of the human body.

"The question is not what you look at, but what you see."

Henry David Thoreau

Appendix to Chapter 6

B

Illustrative Flow Visualization - Evaluation Questionnaire

Title:	Evaluation questionnaire
Name:	
Date:	

Introduction

With this questionnaire we aim to acquire formal feedback to our visualization prototype. This feedback will be part of a manuscript submission for the IEEE Vis 2010 conference (<http://vis.computer.org/VisWeek2010/>). The conference focusses on advances in visualization of scientific data, for which the medical field offers a great number of applications. Our manuscript will be classified as a so-called 'system' paper, which allows us to present a framework for specific tasks or applications. Such a manuscript requires an extensive description of design choices that were made during the development of the prototype and a comparison with existing tools that support a similar task. In our case we compare our results with the commercial software product GTFlow by a company called GyroTools (<http://www.gyrotools.com/>).

Our prototype, called *QFlow Explorer (QFE)*, aims to investigate new ways to understand the complex behavior of MRI acquired blood flow. Throughout the questionnaire we will assess various aspects of the visualization prototype. In particular we will focus on:

- Perception of the anatomical structures
- Perception of the blood flow (direction, speed, branching)
- Performance and interactivity

The questionnaire is divided into three sections. The first section comprises a series of questions to obtain your opinion about the visualization results of the current prototype, whereas the second section is concerned with the visualization results in comparison to GTFlow. In the last section we ask if you could think of use-cases for clinical diagnosis using QFE.

Throughout the questionnaire, each aspect of the prototype can be demonstrated.

Visualization prototype evaluation

Anatomy visualization

Visualizing blood flows requires an anatomical context. Currently QFE deals with the aorta, pulmonary arteries and the vena cava as anatomical context. These are manually segmented. In the following questions we ask for your opinion considering the visualization of the anatomy.

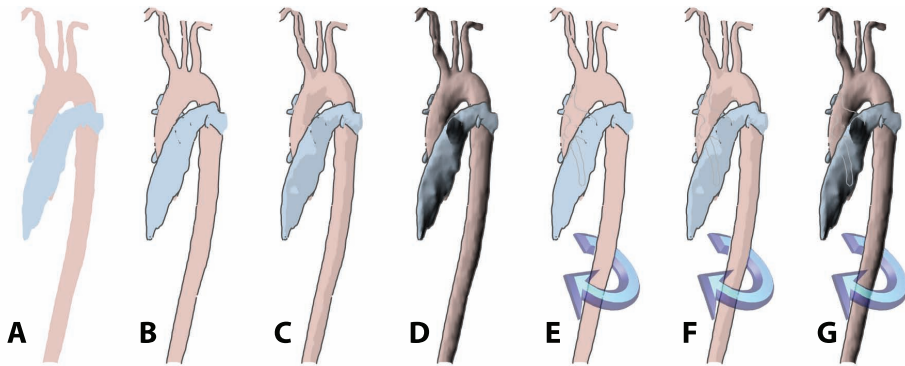


Figure B.1 Anatomical context visualization styles.




Figure B.2 Various planar visualizations of the phase reconstruction data

Q 01 Which image best depicts the <i>shape</i> of the structures, B.1A or B.1B?	A <input type="checkbox"/>	B <input type="checkbox"/>
Q 02 Which image best depicts the <i>shape</i> of the structures, B.1B or B.1C?	B <input type="checkbox"/>	C <input type="checkbox"/>
Q 03 Which image best depicts the <i>shape</i> of the structures, B.1C or B.1D?	C <input type="checkbox"/>	D <input type="checkbox"/>
Q 04 Which image best depicts the <i>spatial relations</i> of the structures, B.1A or B.1B?	A <input type="checkbox"/>	B <input type="checkbox"/>
Q 05 Which image best depicts the <i>spatial relations</i> of the structures, B.1B or B.1C?	B <input type="checkbox"/>	C <input type="checkbox"/>
Q 06 Which image best depicts the <i>spatial relations</i> of the structures, B.1C or B.1D?	C <input type="checkbox"/>	D <input type="checkbox"/>
Q 07 Which image best depicts the <i>spatial relations</i> of the structures, B.1B or B.1E?	B <input type="checkbox"/>	E <input type="checkbox"/>
Q 08 Which image best depicts the <i>spatial relations</i> of the structures, B.1C or B.1F?	C <input type="checkbox"/>	F <input type="checkbox"/>
Q 09 Which image best depicts the <i>spatial relations</i> of the structures, B.1D or B.1G?	D <input type="checkbox"/>	G <input type="checkbox"/>
Q 10 Which image best depicts anatomy as <i>context</i> (with flow), B.1A or B.1B?	A <input type="checkbox"/>	B <input type="checkbox"/>
Q 11 Which image best depicts anatomy as <i>context</i> (with flow), B.1B or B.1C?	B <input type="checkbox"/>	C <input type="checkbox"/>
Q 12 Which image best depicts anatomy as <i>context</i> (with flow), B.1C or B.1D?	C <input type="checkbox"/>	D <input type="checkbox"/>
Q 13 Are you looking at the inside (I) or outside (O) of the vessel?	I <input type="checkbox"/>	O <input type="checkbox"/>

Which anatomical visualization do you prefer from figure B.1?:						
B.1 A <input type="checkbox"/>	B.1 B <input type="checkbox"/>	B.1 C <input type="checkbox"/>	B.1 D <input type="checkbox"/>	B.1 E <input type="checkbox"/>	B.1 F <input type="checkbox"/>	B.1 G <input type="checkbox"/>

For the following set of questions, Q 14 up to Q 17, please consider the anatomical visualization you prefer.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 14 Considering only anatomy, how well can you interpret the <i>morphology</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 15 Considering only anatomy, how well can you interpret the <i>spatial relations</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 16 Considering only anatomy, how well can you <i>navigate</i> to areas of interest?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 17 Considering flow analysis, how well does the visualization provide <i>context</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In your opinion, what would be the <i>benefits</i> of the presented anatomical visualization?:

In your opinion, what would be the <i>drawbacks</i> of the presented anatomical visualization?:

Flow visualization

Planar reformat

There are several ways to depict the variety of aspects of 4D flow data. Quite often there is the need to validate the visualization techniques with the acquired data. QFE contains techniques to integrate a 2D view on the underlying data, incorporated in the 3D environment.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 18 Considering B.2A, how well can you interpret the underlying data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 19 Considering B.2A, how well can you validate the ring with the data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 20 Considering B.2A, how well is occlusion of important structures avoided?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 21 Considering B.2B, how well can you interpret the underlying data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 22 Considering B.2B, how well can you validate the ring with the data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 23 Considering B.2B, how well is occlusion of important structures avoided?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 24 Considering B.2C, how well can you interpret the underlying data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 25 Considering B.2C, how well can you validate the ring with the data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 26 Considering B.2C, how well is occlusion of important structures avoided?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In your opinion, would you benefit from more elaborate visualizations on 2D planes? If yes, can you give an example?

In your opinion, what would be the <i>benefits</i> of the presented planar visualization?:

In your opinion, what would be the <i>drawbacks</i> of the presented planar visualization?:

Seeding approaches

A common approach to visualize flow fields is to trace lines through the field. This tracing requires a set of start points, called seeds, positioned in the cross-cut of the vessel. Different positioning of these seeds leads to varying results of the visualization. In the following we ask for your preferences in terms of seed positions. A *seed distribution* determines the distribution of the seeds over the cross-cut area. A *seed template* redistributes the initial distribution to the vessel center or the vessel wall.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 27 Considering B.3A, how well do you judge this seeding distribution?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 28 Considering B.3B, how well do you judge this seeding distribution?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 29 Considering B.3C, how well do you judge this seeding distribution?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 30 Considering B.3D, how well do you judge this seeding distribution?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 31 Considering B.3E, how well do you judge this seeding distribution?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

(1 - not at all, 2 - sometimes, 3 - no opinion, 4 - often, 5 - always)	1	2	3	4	5
Q 32 When analyzing, would you benefit from varying the <i>seeding density</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 33 When analyzing, would you benefit from varying the <i>seeding distribution</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 34 When analyzing, would you benefit from varying the <i>seeding templates</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In your opinion, what would be the <i>benefits</i> of the presented seeding approach?:

In your opinion, what would be the <i>drawbacks</i> of the presented seeding approach?:

Line tracing

In the previous subsection we mentioned that tracing lines is a common technique to depict flow fields, such as the 4D flow data. The two most common approaches are:

- *Streamlines*: A single phase of the cardiac cycle in the 4D flow acquired data represents an instantaneous 3D blood flow field. Starting at a particular position in this volume, a line can be traced through this field, depicting the structure of the blood flow field at the phase under consideration.
- *Pathlines*: In contrast to streamlines, pathlines represent the temporal behavior over the cardiac phases. Suppose you could place a small object on the seed position and follow the objects trajectory, this would form a line which is called a pathline.

In this subsection these line types will not be distinguished. We will discuss properties that are common for both the stream- and pathlines.




Figure B.3 Various seeding distributions and templates used for line tracing

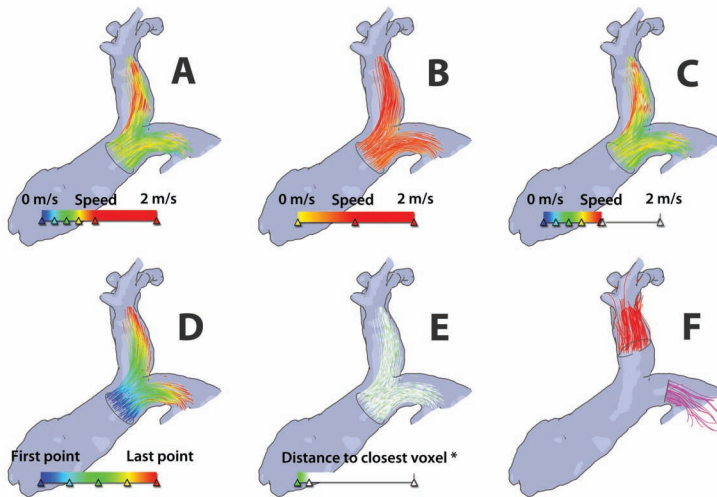


Figure B.4 Line color (* This color coding depicts the distance to the nearest voxel. If the distance is zero, the considered position on the line coincides with the position of the MR measurement and is not interpolated.)




Figure B.5 Reversed tracing of the flow

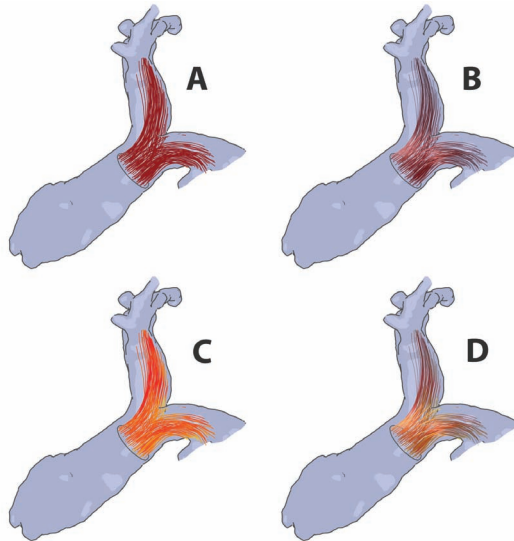


Figure B.6 Line illumination

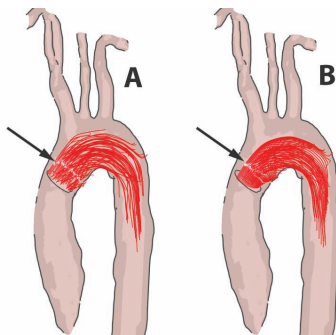


Figure B.7 Line highlights

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 35 Considering B.4A, how well / valuable do you judge this color coding?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 36 Considering B.4B, how well / valuable do you judge this color coding?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 37 Considering B.4C, how well / valuable do you judge this color coding?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 38 Considering B.4D, how well / valuable do you judge this color coding?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 38 Considering B.4E, how well / valuable do you judge this color coding?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 39 How well does color help to analyze the blood flow speed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

The lines can be illuminated by an artificial light source. The effect is especially apparent during interaction.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 40 For B.6A, how well do you perceive the <i>spatial relations</i> of the flow lines?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 41 For B.6B, how well do you perceive the <i>spatial relations</i> of the flow lines?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 42 For B.6C, how well do you perceive the <i>spatial relations</i> of the flow lines?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 43 For B.6C, how well do you perceive the blood flow <i>speed</i> (color coded)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 44 For B.6D, how well do you perceive the <i>spatial relations</i> of the flow lines?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 45 For B.6D, how well do you perceive the blood flow <i>speed</i> (color coded)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 46 How well does the <i>illumination</i> help to analyze the blood flow direction?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Inspired by work by Frydrychowicz et al. QFE includes the option to reverse the trace direction of the lines. This allows to investigate the origin of the flow for a given set of seed points at the cross-cut location.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 47 Considering B.5A, how well do you perceive the blood flow <i>branching</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 48 Considering B.5A, how well do you perceive the blood flow <i>stream relations</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 49 Considering B.5B, how well do you perceive the blood flow <i>branching</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 50 Considering B.5B, how well do you perceive the blood flow <i>stream relations</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

During 4D flow analysis, would you use the reversed tracing feature? If yes, can you give an example?

QFE traces new lines for each point in time from the fixed cross-cut area. As a result, the structure of the blood flow field can be inspected (instantaneous or in time) for each of the chosen locations. Shorter line traces are needed to grasp the flow field.

A possible drawback is that the *decrease* of blood flow *speed* can be falsely interpreted as regurgitant flow, since the lines become shorter over time. To overcome this problem, QFE incorporates a visual cue that indicates the flow direction for each traced line. This is an animated feature, please see the demonstration to observe the effect.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 51 For B.7A, how well does the highlight help to perceive the flow <i>direction</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 52 For B.7A, how well does the highlight help to perceive the flow <i>speed</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 53 For B.7A, how well does the highlight help to perceive the flow <i>branching</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 54 For B.7B, how well does the highlight help to perceive the flow <i>direction</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 55 For B.7B, how well does the highlight help to perceive the flow <i>speed</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 56 For B.7B, how well does the highlight help to perceive the flow <i>branching</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In your opinion, what would be the <i>benefits</i> of using the streamlines?:

In your opinion, what would be the <i>drawbacks</i> of using the streamlines?:

In your opinion, what would be the <i>benefits</i> of using the pathlines?:

In your opinion, what would be the <i>drawbacks</i> of using the pathlines?:

Flow-rate arrows

QFE present so-called *flow-rate arrows*. The position and direction of the arrows are determined by the maximum velocity, while the length represents the flow rate through the cross-cut area. The time behavior of the flow-rate arrows is depicted by fading away older arrows. This is called a *motion trail*.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 57 For B.8A, how well does the arrow represent the <i>flow direction</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 58 For B.8A, how well does the arrow represent the <i>flow rate</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 59 For B.8A, how well can you interpret the flow rate <i>behavior in time</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 60 For B.8B, how well does the arrow represent the <i>flow direction</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 61 For B.8B, how well does the arrow represent the <i>flow rate</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 62 For B.8B, how well can you interpret the flow rate <i>behavior in time</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In your opinion, what would be the <i>benefits</i> of using the flow rate arrows?:

In your opinion, what would be the <i>drawbacks</i> of using the flow rate arrows?:

Performance & interaction

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 63 During interaction, how well do the <i>anatomical structures</i> remain visible?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 64 During interaction, how well do the hidden lines show the spatial rel.?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 65 To what extent do you consider placement of the cross-cuts to be intuitive?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 66 How do you judge the performance (speed) of loading the data?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 67 How do you judge the performance (speed) of positioning the ring?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 68 How do you judge the performance (speed) of 3D navigation?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 69 How do you judge the performance (speed) of parametrization?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

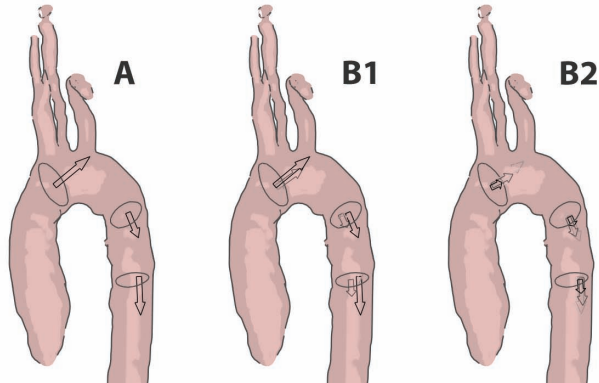


Figure B.8 (A) Flow-rate arrows at cardiac phase 4 (B1) Flow-rate arrows with motion trail at cardiac phase 4 (B2) As B1, at cardiac phase 7

Visualization prototype comparison

In this section we will compare QFE to the approach presented in the GTFlow product.

Anatomy visualization

Similar to the previous section, we first consider the visualization of the anatomy. Typically, GTFlow presents a full planar reformat of the flow data together with the stream- or pathlines. In GTFlow, the reformat is the common way to give a clue of anatomical context. QFE depicts a silhouette, based on the blood flow data, that gives a rough static approximation of the inner wall of the vessel. In combination, QFE includes planar reformats of the cross-cut areas, which present the blood flow data.




Figure B.9 Comparison: anatomical visualization

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 70 How do you judge the anatomical visualization of GTFlow (Fig. B.9 left)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 71 How do you judge the anatomical visualization of QFE (Fig. B.9 right)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>




Figure B.10 Anatomical visualization alternatives




Figure B.11 Comparison: seeding approach




Figure B.12 Comparison: streamlines




Figure B.13 Comparison: pathlines

GTFlow incorporates two alternatives to present an anatomical context. First it provides an iso-surface on the blood flow speed (Fig. B.10 left), which relates to the manual segmentation used in QFE. Alternatively, GTFlow presents a volume rendering of the blood flow speed (Fig. B.10 right).

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 72 How do you judge the iso-surface visualization of GTFlow (Fig. B.10 left)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 73 How do you judge the volume rendering of GTFlow (Fig. B.10 right)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Flow visualization

In this subsection we compare the flow visualization methods. GTFlow allows to vary the seeding density, based on a rectilinear grid to seed the volume.

As presented in the previous section, QFE contains various parameters to tune the seeding. In your opinion, are these options a benefit during the analysis of your 4D flow data set? Could you describe in a few words the main benefits and/or drawbacks, in comparison to the GTFlow approach?

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 74 How do you judge the streamlines of GTFlow (Fig. B.12 left)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 75 How do you judge the streamlines of QFE (Fig. B.12 right)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 76 How do you judge the pathlines of GTFlow (Fig. B.13 left)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 77 How do you judge the pathlines of QFE (Fig. B.13 right)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Performance & interaction

We also compare the performance and interaction possibilities of both systems. First we assess the interaction.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 78 How do you judge the placement of vessel contours in GTFlow?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 79 How do you judge the placement of vessel contours in QFE?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 80 How do you judge the 3D navigation in GTFlow?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 81 How do you judge the 3D navigation in QFE?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Lastly, we ask for your opinion on the performance of both systems. GTFlow first computes points and lines before visualizing them. This process repeats when changing a parameter of the visualization task, for instance the trace length. QFE uses graphics hardware to perform the computations of the line structures in real time.

(1 - very poor, 2 - poor, 3 - no opinion, 4 - good, 5 - very good)	1	2	3	4	5
Q 82 How do you judge the overall performance (speed) of GTFlow?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 83 How do you judge the overall performance (speed) of QFE?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 84 To what extend do you value the interactive parametrization?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Thank you for your valuable feedback!

*"It is better to debate a question without settling it,
than to settle a question without debating it."*

Joseph Joubert

Appendix to Chapter 7



Virtual Blood-Flow Probing - Evaluation Questionnaire

Title: Evaluation questionnaire

Name:

Date:

Introduction

With this questionnaire we intend to obtain formal feedback to our visualization prototype. The acquired feedback will be part of a manuscript submission for the IEEE Vis 2011 conference (<http://vis.computer.org/visweek2011/>). This conference focusses on advances in visualization of scientific data, including a wide variety of applications in the medical field. Our manuscript will be classified as an 'application/design' paper, describing the design choices for the development of the prototype and the potential benefits for the application domain. In addition, we compare with existing tools that support a similar task. In our case, we compare our results with the commercial software product GTFLOW by the company called GyroTools (<http://www.gyrotools.com/>). Besides, we compare to a recently introduced prototype by NVidia.

Our prototype, called *QFlow Explorer (QFE)*, investigates new ways to understand the complex behavior of MRI acquired blood-flow. The main goal of our prototype is to facilitate fast exploration the blood-flow, without the need for a full segmentation of the morphological structures. The qualitative exploration allows quick identification of regions with chaotic hemodynamics. In particular, we look for a generic approach that can be applied to both vessels and the heart. Moreover, the approach has to operate for abnormal morphological structures (e.g., congenital heart disease, dissection).

The questionnaire will assess various aspects of the visualization prototype. In particular we will focus on:

- Ease of virtual probe placement
- Ease of virtual probe interaction
- Perception of the blood flow (direction, speed, branching)
- Perception of the anatomical context
- Performance and interactivity

The questionnaire is divided into four sections. The first section comprises a series of questions to obtain your opinion about the interaction with the presented virtual probe. The second section is concerned with the various visualization approaches. Subsequently, we compare our prototype to existing tools. In the last section we ask for your view on the clinical relevance.

Throughout the questionnaire, each aspect of the prototype can be demonstrated.

Virtual probe interaction evaluation

Virtual probe placement

Before performing an extensive quantitative analysis of the blood-flow parameters, it is worthwhile to *explore* the blood-flow dynamics interactively. Therefore, we introduce a *virtual probe*, inspired by the operator-driven ultrasound acquisition. An example of the virtual probe, together with flow lines, is depicted in figure C.1a. The first step of the devised workflow, comprises the initial positioning of the virtual probe. In order to position the probe, we introduce a view-aligned slice.




Figure C.1 View-aligned slice




Figure C.2 Virtual probe positioning




Figure C.3 Probe positioning mockups

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 01 Figure C.1b: how do you judge the interaction in <i>orthogonal views</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 02 Figure C.1c: how do you judge the interaction in <i>oblique views</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 03 Figure C.1d: how do you judge the <i>fusion</i> of modalities?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 04 View-aligned slice: how well can you locate ROIs using the <i>t-mip</i> *?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 05 View-aligned slice: how well can you locate ROIs using <i>anatomy (ssfp)</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

**t-mip* - temporal maximum intensity projection: static summary of the maximum blood flow speed over time.

When the desired slice is determined, a virtual probe can be added by means of two different approaches.

1. *Two-click approach*:
Select the two click method (Fig. C.2a). While holding the CTRL key, click the two positions to create the probe (Fig. C.2b).
2. *Pen-stroke approach*:
Select the pen-stroke method (Fig. C.2a). While holding the CTRL key, click and drag to create the probe.

Subsequently, the depth and orientation along the view direction can be determined automatically, by fitting the probe to the blood-flow velocity field.

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 06 Probe positioning: how do you judge the <i>two-click positioning</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 07 Probe positioning: how do you judge the <i>pen-stroke positioning</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 08 Probe positioning: how do you judge the <i>concept of the fitting</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 09 Probe positioning: how do you judge the <i>accuracy of the fitting</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 10 Probe positioning: how do you judge the <i>speed of the fitting</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Alternatively, the initial placement of the probe may be performed on a volume rendering or MIP. The interactive slice positioning is omitted, at the expense of a decreased accuracy of the initial placement.

(1 - very inefficient, 2 - inefficient, 3 - indecisive, 4 - efficient, 5 - very efficient)	1	2	3	4	5
Q 11 Figure C.3a: how would you judge the <i>positioning using DVR</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 12 Figure C.3b: how would you judge the <i>positioning using a MIP</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Remarks concerning the probe placement (benefits / drawbacks):

Virtual probe interaction

A virtual probe has been positioned in the region of interest. While rotating around the scene, the probe continuously adjusts its orientation towards the viewer.

There are several ways to interact with the virtual probe. The probe can be translated and rotated in the presented slice. Besides, the length and the top and base radii of the probe can be adjusted. In addition, the probe can be augmented with an interactive slice orthogonal to the long axis.




Figure C.4 Virtual probe interaction

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 13 Figure C.4a: how do you judge the <i>in-slice translation</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 14 Figure C.4b: how do you judge the <i>in-slice rotation</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 15 Figure C.4c: how do you judge the <i>change of probe dims.</i> (length, radii)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 16 Figure C.4d: how do you judge the probe <i>orthogonal slice</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 17 How well can you <i>move</i> the probe to a particular region of interest?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Remarks concerning the probe interaction (benefits / drawbacks):

Blood-flow visualization

Virtual probe

The virtual probe is a navigational object, facilitating the exploration of the blood-flow field. Since the probe object is emphatically present in the visualization, we will assess the visual representation of the probe itself.

	a	b	c	d	e	f
Q 18 Figure C.5a..C.5e - Which style suits the application <i>best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Q 19 Figure C.5a..C.5e - Which style suits the application <i>second best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Q 20 Figure C.5c, C.5f - Which contour depiction has your preference?	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the probe visualization (benefits / drawbacks):

Particle tracing

Flow phenomena are commonly visualized using particle traces, where small virtual objects are directed by the velocity field. In our prototype, particles can be interactively injected and visualized. We present various ways to represent the particle traces, inspired by illustrative techniques.




Figure C.5 Virtual probe appearance

Seeding

Particles can only exist inside the acquired volume. Once a particle leaves the velocity field, the particle is newly injected into the system. As a result, a new particle trace starts at the seeding location at the actual time of the cardiac cycle. We call this approach *continuous particle seeding*. Alternatively, particle batches can be injected at fixed intervals, and particles leaving the volume will not be re-injected into the system. During the continuous particle trace, the user can decide to inject a new batch of particles at a desired point in time. All particles are removed, and a new set of particles is injected. This can be performed interactively, i.e., during a running animation.




Figure C.6 Interactive seeding

(1 - very ineffective, 2 - effective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 21 How do you judge <i>continuous particle seeding</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 22 Figure C.6 - How do you judge <i>interactive particle seeding</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the particle seedings (benefits / drawbacks):

Visualization

There are several ways to represent particles. Commonly, particles are represented by spheres (e.g., figure C.7c). We were inspired by illustrative techniques, where moving spheres are often presented in an elongated fashion. Similarly, we elongate the 'particles' based on the blood-flow speed (Fig. C.7d). Another illustrative technique is the motion line, depicting the speed and direction of movement (Fig. C.7e). In the subsequent, we focus on the flow visualizations, hence we temporarily omit anatomical context.

Q 23 Figure C.7a, C.7b - Which image best depicts the flow dynamics?	a <input type="checkbox"/>	b <input type="checkbox"/>	x <input type="checkbox"/>
Q 24 Figure C.7c, C.7d - Which image best depicts the flow dynamics?	c <input type="checkbox"/>	d <input type="checkbox"/>	x <input type="checkbox"/>
Q 25 Figure C.7c, C.7e - Which image best depicts the flow dynamics?	c <input type="checkbox"/>	e <input type="checkbox"/>	x <input type="checkbox"/>
Q 26 Figure C.7d, C.7e - Which image best depicts the flow dynamics?	d <input type="checkbox"/>	e <input type="checkbox"/>	x <input type="checkbox"/>
Q 27 Figure C.7f, C.7g - Which image best depicts the flow dynamics?	f <input type="checkbox"/>	g <input type="checkbox"/>	x <input type="checkbox"/>

x = *indecisive, please comment*

Particles leave the velocity field of the acquired volume. We present a fading cue for particles leaving the volume (see grey particles, figure C.7h).

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 28 Figure C.7h - How do you judge the <i>fading visual cue</i> for particles leaving the volume?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

During the particle trace, the probe can be repositioned. New particles will originate from the probe position.

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 29 How do you judge the <i>ease of exploration</i> using the probe?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 30 How do you judge the <i>speed of exploration</i> compared to slice-based views?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the particle traces (benefits / drawbacks):

Integral lines

Integral lines provide an alternative way of depicting flow fields. Instead of presenting separate particles that are directed by the flow, we depict the path that the particle has traveled over time.

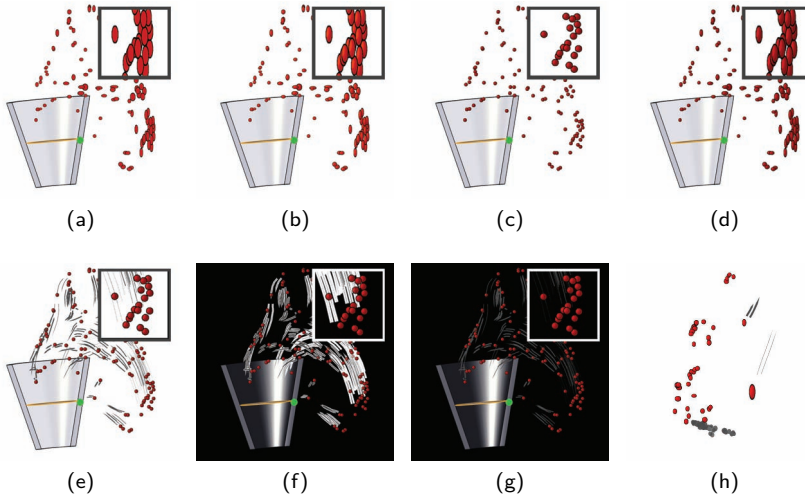


Figure C.7 Particle trace visualization

Q 31 Figure C.8a, C.8b - Which image best depicts the flow dynamics?	a <input type="checkbox"/>	b <input type="checkbox"/>	x <input type="checkbox"/>
Q 32 Figure C.8b, C.8c - Which image best depicts the flow dynamics?	b <input type="checkbox"/>	c <input type="checkbox"/>	x <input type="checkbox"/>
Q 33 Figure C.8a, C.8d - Which image best depicts the flow dynamics?	a <input type="checkbox"/>	d <input type="checkbox"/>	x <input type="checkbox"/>

x = indecisive, please comment

By moving the probe, integral lines can be inspected at any location.

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 34 How do you judge the <i>insight from the inspection using integral lines</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 35 How do you judge the <i>ease of exploration using integral lines</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 36 How do you judge the <i>speed of exploration</i> compared to slice-based repositioning?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the integral lines (benefits / drawbacks):

Integral surfaces

The flow visualization community has introduced *integral surfaces*. By connecting the paths that the particles have traveled, a surface is constructed. In order to keep the surfaces contiguous, the time that the 'particles' are allowed to travel is short.

	a	b	c	d
Q 37 Figure C.9a..C.9c - Which surface conveys the flow dynamics <i>best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	●
Q 38 Figure C.9a..C.9c - Which surface conveys the flow dynamics <i>second best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	●
Q 39 Figure C.9c..C.9d - Which visual style conveys the flow dynamics <i>best</i> ?	●	●	<input type="checkbox"/>	<input type="checkbox"/>




Figure C.8 Integral lines visualization




Figure C.9 Integral surface visualization

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 40 Figure C.9e - How do you judge the <i>perception of chaotic blood-flow regions</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 41 How do you judge the <i>insight from the inspection using integral surfaces</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 42 How do you judge the <i>ease of exploration using integral surfaces</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 43 How do you judge the <i>speed of exploration</i> compared to slice-based repositioning?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the integral surfaces (benefits / drawbacks):

Anatomical context

During probe positioning, the anatomical context is based on the view-aligned slice, attached to the virtual probe. For further inspection, we incorporate a volume visualization for the anatomical context. Both the anatomy, t-mip and blood-flow speed images can be used as a contextual visualization.

	a	b	c
Q 44 Figure C.10a..C.10c - Which modality conveys the anatomical context <i>best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 45 Figure C.10a..C.10c - Which modality conveys the anatomical context <i>second best</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

In order to mix the presented flow visualizations with the anatomical context (Fig. C.11a), we can 'cut away' or clip a part of the volume. This provides a view inside the anatomical structures. We present two ways to clip the volume. The first method clips a part of the volume based on an *orthogonal slice* (Fig. C.11b). The second method clips a part of the volume based on *the probe slice* (Fig. C.11c).

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 46 Figure C.11b - How well are the <i>morphological structures preserved</i> after clipping?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 47 Figure C.11c - How well are the <i>morphological structures preserved</i> after clipping?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 48 Figure C.11b, C.11c - How do you judge clipping as a means of looking into the anatomical structures?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 49 Figure C.12a..C.12c - How do you judge the <i>ease of using the probe slice for clipping</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 50 Figure C.12a..C.12c - How do you judge the <i>interactivity of the probe slice for clipping</i> ?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

(1 - very ineffective, 2 - ineffective, 3 - indecisive, 4 - effective, 5 - very effective)	1	2	3	4	5
Q 51 Figure C.12a - How well are the <i>flow dynamics</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 52 Figure C.12b - How well are the <i>flow dynamics</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 53 Figure C.12c - How well are the <i>flow dynamics</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 54 Figure C.12a - How well is the <i>morphology</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 55 Figure C.12b - How well is the <i>morphology</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Q 56 Figure C.12c - How well is the <i>morphology</i> conveyed?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Other remarks concerning the anatomical context (benefits / drawbacks):

Visualization prototype comparison

Mutual features

We compare our prototype to existing visualization tools. Currently, GTFlow by Gyrotools is the only commercial software available. Recently, also NVidia presented a prototype for blood-flow visualization.

We explicitly compare only the *exploratory and visualization functionality* of these packages, because this was the main focus of our prototype. Hence, we will not address approaches for quantitative analysis.

In the following, we provide a list of features that are present in both our prototype and any of the other tools on the market. Please provide your score for the features in general, using a score on a 1-5 scale (1 = very ineffective, 5 = very effective).

Feature	Insight	Importance
Probe
Particles
Integral lines
Volume rendering

In order to compare to other tools, please provide your score for each of these features per tool. Again, use a score on a 1-5 scale (1 = very ineffective, 5 = very effective). The numbers refer to the related figures.

Feature	GTFlow	Nvidia	QFE
Probe			
- positioning [C.13]	... [C.5]
- interaction	•	... [C.13]	... [C.5]
- style [C.13]	... [C.5]
Particles			
- seeding: planar	...	•	...
- seeding: volumetric	•
- interactivity: real-time seed location update	•	... [C.13]	... [C.7]
- style: spheres	•	... [C.13a]	... [C.7c]
- style: pathlets	... [C.14a]	•	... [C.7e]
- illumination	•	... [C.13a]	... [C.7c]
Integral lines			
- seeding: planar	...	•	...
- seeding: volumetric	•
- style: lines	... [C.14a]	... [C.13b]	... [C.8]
Volume rendering			
- anatomy	... [C.14a,b]	... [C.13b,c]	... [C.11]
- illumination	... [C.14a]	... [C.13c]	... [C.11]
- transparency	... [C.14b]	... [C.13c]	... [C.11]

Other remarks concerning differences with existing visualization tools:



Figure C.10 Anatomical context visualization

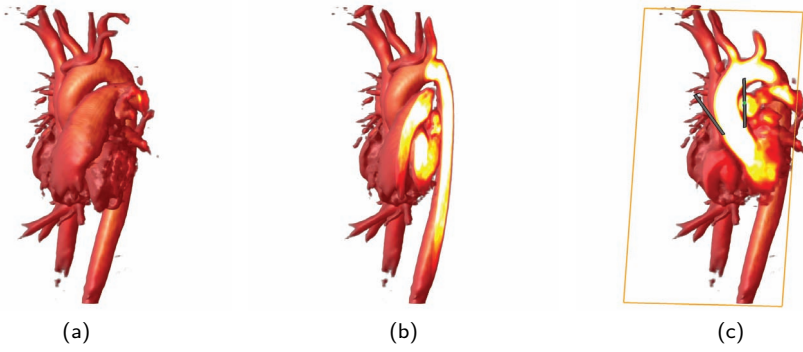


Figure C.11 Anatomical context clipping




Figure C.12 Anatomical context with flow visualization




Figure C.13 Nvidia blood-flow visualization




Figure C.14 Gyrottools GTFlow blood-flow visualization

Clinical relevance

<p>Would any of these techniques allow you to explore aspects of the data that you could not before? If yes, which techniques and in which context?</p>

<p>Which features are most beneficial to you, compared to other tools you have used?</p>
--

Thank you for your valuable feedback!

References

- [1] M. Agrawala, W. Li, and F. Berthouzoz. Design principles for visual communication. *Communications of the ACM*, 54(4):60–69, 2011. Cited on page 49.
- [2] J. Ahlberg. *Active Contours in Three Dimensions*. PhD thesis, Linköping University, 1996. Cited on pages 75 and 76.
- [3] S. Allender, P. Scarborough, V. Peto, and M. Rayner. European cardiovascular disease statistics, 2008. Cited on pages 12 and 19.
- [4] Y. Amano, T. Sekine, Y. Suzuki, K. Tanaka, R. Takagi, and S. Kumita. Time-resolved three-dimensional magnetic resonance velocity mapping of chronic thoracic aortic dissection: a preliminary investigation. *Magnetic Resonance in Medical Sciences*, 10(2):93–99, 2011. Cited on page 23.
- [5] American Heart Association. Heart disease and stroke statistics, 2010 update at-a-glance. <http://www.americanheart.org/statistics/> (last visited: 01-02-2012), 2010. Cited on pages 12, 19, and 24.
- [6] P. Angelelli and H. Hauser. Straightening tubular flow for side-by-side visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2063–2070, 2011. Cited on page 153.
- [7] H. Arheden, M. Saeed, E. Törnqvist, G. Lund, M. F. Wendland, C. B. Higgins, and F. Ståhlberg. Accuracy of segmented MR velocity mapping to measure small vessel pulsatile flow in a phantom simulating cardiac motion. *Journal of Magnetic Resonance Imaging*, 13(5):722–728, 2001. Cited on page 43.
- [8] J. Astola, P. Haavisto, and Y. Neuvo. Vector median filters. *IEEE*, 78(4):678–689, 1990. Cited on page 110.
- [9] A. Baer, C. Tietjen, R. Bade, and B. Preim. Hardware-accelerated stippling of surfaces derived from medical volume data. In *EG/IEEE EuroVis*, pages 235–242, 2007. Cited on page 55.
- [10] P. Barla, J. Thollot, and L. Markosian. X-Toon : an extended toon shader. In *Non-Photorealistic Animation and Rendering*, pages 127–132, 2006. Cited on page 57.
- [11] D. Bauer, R. Peikert, M. Sato, and M. Sick. A case study in selective visualization of unsteady 3D flow. In *IEEE Visualization*, pages 525–528, 2002. Cited on page 65.
- [12] M. A. Bernstein and Y. Ikezaki. Comparison of phase-difference and complex-difference processing in phase-contrast MR angiography. *Journal of Magnetic Resonance Imaging*, 1(2):725–729, 1991. Cited on page 38.
- [13] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI pulse sequences*. Elsevier, 2004. Cited on page 32.
- [14] M. A. Bernstein, X. J. Zhou, J. A. Polzin, K. F. King, A. Ganin, N. J. Pelc, and G. H. Glover. Concomitant gradient terms in phase contrast MR: analysis and correction. *Magnetic Resonance in Medicine*, 39(2):300–308, 1998. Cited on page 44.
- [15] H. G. Bogren and M. H. Buonocore. 4D Magnetic resonance velocity mapping of blood flow patterns in the aorta in young vs. elderly normal subjects. *Journal of Magnetic Resonance Imaging*, 10(5):861–869, 1999. Cited on pages 12 and 18.
- [16] H. G. Bogren and M. H. Buonocore. Complex flow patterns in the great vessels: a review. *Journal of Cardiac Imaging*, 15(2):105–113, 1999. Cited on page 17.
- [17] H. G. Bogren, R. H. Klipstein, R. H. Mohiaddin, D. N. Firmin, S. R. Underwood, R. S. O. Rees, and D. B. Longmore. Pulmonary artery distensibility and blood flow patterns: a magnetic resonance study of normal subjects and of patients with pulmonary arterial hypertension. *American Heart Journal*, 118(5):990–999, 1989. Cited on page 16.

- [18] D. Borland and M. R. Taylor. Rainbow color map (still) considered harmful. *IEEE Computer Graphics and Applications*, 27(2):14–17, 2007. Cited on page 97.
- [19] S. Born, A. Wiebel, J. Friedrich, G. Scheuermann, and D. Bartz. Illustrative stream surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1329–1338, 2010. Cited on pages 55, 66, 67, and 118.
- [20] H. Bouma, J. Oliván Bescós, A. Vilanova, and F. A. Gerritsen. Unbiased vessel-diameter quantification based on the FWHM criterion. In *Society of Photo-Optical Instrumentation Engineers*, volume 6512, pages 65122N.1–65122N.9, 2007. Cited on page 92.
- [21] M. Breeuwer, S. de Putter, U. Kose, L. Speelman, K. Visser, F. A. Gerritsen, R. Hoogeveen, R. Krams, H. van den Bosch, J. Buth, T. Gunther, B. Wolters, E. van Dam, and F. van de Vosse. Towards patient-specific risk assessment of abdominal aortic aneurysm. *Medical & Biological Engineering & Computing*, 46(11):1085–1095, 2008. Cited on page 15.
- [22] C. Brownlee, V. Pegoraro, S. Shankar, P. McCormick, and C. Hansen. Physically-based interactive Schlieren flow visualization. In *IEEE Pacific Visualization*, pages 145–152, 2010. Cited on page 59.
- [23] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. Illustrative context-preserving volume rendering. In *EG/IEEE EuroVis*, pages 69–76, 2005. Cited on page 53.
- [24] S. Bruckner and M. E. Gröller. Exploded views for volume data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1077–1084, 2006. Cited on page 53.
- [25] S. Bruckner and M. E. Gröller. Style transfer functions for illustrative volume rendering. *Computer Graphics Forum*, 26(3):715–724, 2007. Cited on page 58.
- [26] K. Bürger, F. Ferstl, H. Theisel, and R. Westermann. Interactive streak surface visualization on the GPU. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1259–1266, 2009. Cited on page 67.
- [27] R. Bürger and H. Hauser. Visualization of multi-variate scientific data. *Computer Graphics Forum*, 28(6):1670–1690, 2009. Cited on page 49.
- [28] M. Burns, J. Klawe, S. Rusinkiewicz, A. Finkelstein, and D. DeCarlo. Line drawings from volume data. *ACM Transactions on Graphics*, 24(3):512–518, 2005. Cited on page 55.
- [29] S. Busking, A. Vilanova, and J. J. van Wijk. Particle-based non-photorealistic volume visualization. *The Visual Computer*, 24(5):335–346, 2007. Cited on page 55.
- [30] B. Cabral and L. C. Leedom. Imaging vector fields using line integral convolution. In *ACM SIGGRAPH Computer Graphics and Interactive Techniques*, pages 263–270, 1993. Cited on page 60.
- [31] B. S. Carmo, Y. H. P. Ng, A. Prügel-Bennett, and G.-Z. Yang. A data clustering and streamline reduction method for 3D MR flow vector field simplification. In *Medical Image Computing and Computer-Assisted Intervention*, pages 451–458, 2004. Cited on pages 64 and 81.
- [32] M. A. Castro, C. M. Putman, and J. R. Cebal. Patient-specific computational modeling of cerebral aneurysms with multiple avenues of flow from 3D rotational angiography images. *Academic radiology*, 13(7):811–821, 2006. Cited on page 15.
- [33] E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *IEEE Information Visualization*, pages 69–75, 2000. Cited on page 48.
- [34] N. S.-H. Chu and C.-L. Tai. Moxi: real-time ink dispersion in absorbent paper. *ACM Transactions on Graphics*, 24(3):504–511, 2005. Cited on page 53.
- [35] A. C. S. Chung, J. A. Noble, and P. Summers. Vascular segmentation of phase-contrast magnetic resonance angiograms based on statistical mixture modeling and local phase coherence. *IEEE Transactions on Medical Imaging*, 23(12):1490–1507, 2004. Cited on pages 74, 139, and 140.
- [36] R. E. Clough, B. Modarai, J. A. Topple, R. E. Bell, T. W. G. Carrell, H. A. Zayed, M. Waltham,

- and P. R. Taylor. Predictors of stroke and paraplegia in thoracic aortic endovascular intervention. *European Journal of Vascular and Endovascular Surgery*, 41(3):303–310, 2011. Cited on page 23.
- [37] M. A. Coady, J. S. Ikonomidis, A. T. Cheung, A. H. Matsumoto, M. D. Dake, E. L. Chaikof, R. P. Cambria, C. T. Mora-Mangano, T. M. Sundt, and F. W. Sellke. Surgical management of descending thoracic aortic disease: open and endovascular approaches. *Circulation*, 121(25):2780–2804, 2010. Cited on page 21.
- [38] L. Cohen. On active contour models and balloons. *Computer Vision Graphics and Image Processing (CVGIP): Image Understanding*, 53(2):211–218, 1991. Cited on page 73.
- [39] F. Cole, A. Golovinskiy, A. Limpaecher, H. S. Barros, A. Finkelstein, T. Funkhouser, and S. Rusinkiewicz. Where do people draw lines? *ACM Transactions on Graphics*, 27(3):1–11, 2008. Cited on page 55.
- [40] F. Cole, K. Sanik, D. DeCarlo, A. Finkelstein, T. Funkhouser, S. Rusinkiewicz, and M. Singh. How well do line drawings depict shape? *ACM Transactions on Graphics*, 28(3):1–9, 2009. Cited on page 54.
- [41] D. Cosgrove. Ultrasound contrast agents: an overview. *European Journal of Radiology*, 60(3):324–330, 2006. Cited on page 15.
- [42] A. Cox. Capstone speaker: how editing and design changes news graphics. In *IEEE Visualization*, 2011. Cited on page 53.
- [43] R. Crawfis, N. Max, B. Becker, and B. Cabral. Rendering of 3D scalar and vector fields at LLNL. In *Supercomputing*, pages 570–576, 1993. Cited on page 59.
- [44] B. Csébfalvi, L. Mroz, H. Hauser, A. König, and M. E. Gröller. Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum*, 20(3):452–460, 2001. Cited on page 55.
- [45] N. Cuntz, A. Pritzkau, and A. Kolb. Time-adaptive lines for the interactive visualization of unsteady flow data sets. *Computer Graphics Forum*, 28(8):2165–2175, 2009. Cited on pages 65 and 117.
- [46] W. C. de Leeuw and R. van Liere. Divide and conquer spot noise. In *ACM/IEEE Supercomputing*, pages 19–32, 1997. Cited on page 61.
- [47] W. C. de Leeuw and J. J. van Wijk. A probe for local flow field visualization. In *IEEE Visualization*, pages 39–45, 1993. Cited on page 61.
- [48] W. C. de Leeuw and J. J. van Wijk. Enhanced spot noise for vector field visualization. In *IEEE Visualization*, pages 233–240, 1995. Cited on page 61.
- [49] D. DeCarlo, A. Finkelstein, S. Rusinkiewicz, and A. Santella. Suggestive contours for conveying shape. *ACM Transactions on Graphics*, 22(3):848–855, 2003. Cited on page 55.
- [50] D. DeCarlo and S. Rusinkiewicz. Highlight lines for conveying shape. In *Non-Photorealistic Animation and Rendering*, pages 63–70, 2007. Cited on pages 55 and 100.
- [51] C. Díaz and L. A. Robles. Fast noncontinuous path phase-unwrapping algorithm based on gradients and mask. *Progress in Pattern Recognition, Image Analysis and Applications*, 3287:116–123, 2004. Cited on page 44.
- [52] T. Doenst, K. Spiegel, M. Reik, M. Markl, J. Hennig, S. Nitzsche, F. Beyersdorf, and H. Oertel. Fluid-dynamic modeling of the human left ventricle: methodology and application to surgical ventricular reconstruction. *The Annals of Thoracic Surgery*, 87(4):1187–1195, 2009. Cited on page 15.
- [53] Q. Du and X. Wang. Centroidal Voronoi tessellation based algorithms for vector fields visualization and segmentation. In *IEEE Visualization*, pages 43–50, 2004. Cited on page 64.
- [54] H. Eggebrecht, U. Herold, O. Kuhnt, A. Schmermund, T. Bartel, S. Martini, A. Lind, C. K.

- Naber, P. Kienbaum, H. Kühn, J. Peters, H. Jakob, R. Erbel, and D. Baumgart. Endovascular stent-graft treatment of aortic dissection: determinants of post-interventional outcome. *European Heart Journal*, 26(5):489–497, 2005. Cited on page 23.
- [55] J. A. Elefteriades, C. J. Lovoulos, M. A. Coady, G. Tellides, G. S. Kopf, and J. A. Rizzo. Management of descending aortic dissection. *The Annals of Thoracic Surgery*, 67(6):2002–2005, 1999. Cited on pages 21 and 23.
- [56] P. Engelfriet, E. Boersma, E. Oechslin, J. Tijssen, M. A. Gatzoulis, U. Thilén, H. Kaemmerer, P. Moons, F. Meijboom, J. Popelová, V. Laforest, R. Hirsch, L. Daliento, E. Thaulow, and B. Mulder. The spectrum of adult congenital heart disease in Europe: morbidity and mortality in a 5 year follow-up period. The Euro Heart Survey on adult congenital heart disease. *European Heart Journal*, 26(21):2325–2333, 2005. Cited on page 24.
- [57] J. Eriksson, C. J. Carlhäll, P. Dyverfeldt, J. Engvall, A. F. Bolger, and T. Ebbers. Semi-automatic quantification of 4D left ventricular blood flow. *Journal of Cardiovascular Magnetic Resonance*, 12:1–9, 2010. Cited on page 17.
- [58] M. H. Everts. *Visualization of dense line data*. PhD thesis, Rijksuniversiteit Groningen, 2011. Cited on page 66.
- [59] M. H. Everts, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg. Depth-dependent halos: illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, 2009. Cited on pages 66, 105, 115, and 116.
- [60] C. R. Fikar and R. Fikar. Aortic dissection in childhood and adolescence: an analysis of occurrences over a 10-year interval in New York State. *Clinical Cardiology*, 32(6):e23–e26, 2009. Cited on page 21.
- [61] L. Forssell. Using line integral convolution for flow visualization: curvilinear grids, variable-speed animation, and unsteady flows. *Visualization and Computer Graphics*, 1(2):133–141, 1995. Cited on page 61.
- [62] T. Frühauf. Raycasting vector fields. In *IEEE Visualization*, pages 115–120, 1996. Cited on pages 59 and 138.
- [63] A. Frydrychowicz, R. Arnold, A. Harloff, C. Schlensak, J. Hennig, M. Langer, and M. Markl. In vivo 3-dimensional flow connectivity mapping after extracardiac total cavopulmonary connection. *Circulation*, 118(2):e16–e17, 2008. Cited on page 96.
- [64] A. Frydrychowicz, A. F. Stalder, M. F. Russe, J. Bock, S. Bauer, A. Harloff, A. Berger, M. Langer, J. Hennig, and M. Markl. Three-dimensional analysis of segmental wall shear stress in the aorta by flow-sensitive four-dimensional-MRI. *Journal of Magnetic Resonance Imaging*, 30(1):77–84, 2009. Cited on page 94.
- [65] R. Fuchs, J. Kemmler, B. Schindler, J. Waser, F. Sadlo, H. Hauser, and R. Peikert. Toward a Lagrangian vector field topology. *Computer Graphics Forum*, 29(3):1163–1172, 2010. Cited on page 63.
- [66] H. Garcke, T. Preußer, M. Rumpf, A. Telea, U. Weikard, and J. J. van Wijk. A phase field model for continuous clustering on vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):230–241, 2001. Cited on page 64.
- [67] V. Garg, I. S. Kathiriya, R. Barnes, M. K. Schluterman, I. N. King, C. A. Butler, C. R. Rothrock, R. S. Eapen, K. Hirayama-Yamada, K. Joo, R. Matsuoka, J. C. Cohen, and D. Srivastava. GATA4 mutations cause human congenital heart defects and reveal an interaction with TBX5. *Nature*, 424(6947):443–447, 2003. Cited on page 24.
- [68] C. Garth, H. Krishnan, X. Tricoche, T. Bobach, and K. I. Joy. Generation of accurate integral surfaces in time-dependent vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1404–1411, 2008. Cited on pages 67 and 118.
- [69] C. Garth and X. Tricoche. Topology-and feature-based flow visualization: methods and applications. In *Society for Industrial and Applied Mathematics (SIAM) Conference on Geometric*

- Design and Computing*, pages 1–21, 2005. Cited on pages 61 and 62.
- [70] R. Gasteiger, M. Neugebauer, O. Beuing, and B. Preim. The flowlens: a focus-and-context visualization approach for exploration of blood flow in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2183–2192, 2011. Cited on page 61.
- [71] R. Gasteiger, M. Neugebauer, C. Kubisch, and B. Preim. Adapted surface visualization of cerebral aneurysms with embedded blood flow information. In *EG Visual Computing for Biology and Medicine*, pages 25–32, 2010. Cited on page 88.
- [72] P. D. Gatehouse, M. P. Rolf, M. J. Graves, M. B. M. Hofman, J. J. Totman, B. Werner, R. A. Quest, Y. Liu, J. von Spiczak, M. Dieringer, D. N. Firmin, A. Rossum van, M. Lombardi, J. Schwitter, J. Schulz-Menger, and P. J. Kilner. Flow measurement by cardiovascular magnetic resonance: a multi-centre multi-vendor study of background phase offset errors that can compromise the accuracy of derived regurgitant or shunt flow measurements. *Journal of Cardiovascular Magnetic Resonance*, 12(5):1–5, 2010. Cited on pages 42, 43, and 44.
- [73] A. Gooch, B. Gooch, P. Shirley, and E. Cohen. A non-photorealistic lighting model for automatic technical illustration. In *ACM SIGGRAPH Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pages 447–452, 1998. Cited on page 57.
- [74] B. Gooch and A. Gooch. *Non-photorealistic rendering*. A. K. Peters, Ltd., 2001. Cited on pages 55 and 100.
- [75] B. Gooch, P.-P. J. Sloan, A. Gooch, P. Shirley, and R. Riesenfeld. Interactive Technical Illustration. In *Interactive 3D Graphics*, pages 31–38, 1999. Cited on page 55.
- [76] G. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Springer, 1995. Cited on page 74.
- [77] G. F. Greil, T. Geva, S. E. Maier, and A. J. Powell. Effect of acquisition parameters on the accuracy of velocity encoded cine magnetic resonance imaging blood flow measurements. *Journal of Magnetic Resonance Imaging*, 15(1):47–54, 2002. Cited on pages 43 and 45.
- [78] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995. Cited on page 44.
- [79] S. Gumhold. Splatting illuminated ellipsoids with depth correction. In *Vision, Modeling, and Visualization*, pages 245–252, 2003. Cited on page 115.
- [80] E. L. Hahn. Spin echoes. *Physical Review*, 80(4):580–594, 1950. Cited on page 32.
- [81] S. Hales. *Statistical essays: containing heamastaticks; or an account of some hydraulic and hydrostatical experiments made on the blood and blood-vessels of animals*. W. Innys and R. Maney, 1740. Cited on page 11.
- [82] M. Hall. *Observations on blood-letting: founded upon researches on the morbid and curative effects of loss of blood*. Sherwood Gilbert and Piper, 1836. Cited on page 10.
- [83] C. D. Hansen and C. R. Johnson. *The visualization handbook*. Academic Press, 2005. Cited on page 63.
- [84] G. G. Hartnell and J. Gates. Aortic fenestration: a why, when, and how-to guide. *Radiographics*, 25(1):175–189, 2005. Cited on page 23.
- [85] L. Hatle and B. Angelsen. *Doppler ultrasound in cardiology*. Lippincott Williams and Wilkins, 1985. Cited on pages 15 and 132.
- [86] G. L. F. Hautvast, S. Lobregt, M. Breeuwer, A. Vilanova, and F. A. Gerritsen. Automatic contour detection in short-axis cardiac cine MR data. *Journal of Cardiovascular Magnetic Resonance*, 7(1):323–324, 2005. Cited on page 77.
- [87] C. G. Healey. Choosing effective colours for data visualization. In *IEEE Visualization*, pages 263–270, 1996. Cited on page 97.
- [88] B. Heckel, G. Weber, B. Hamann, and K. I. Joy. Construction of vector field hierarchies. In *IEEE Visualization*, pages 19–25, 1999. Cited on pages 64 and 82.

- [89] J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *Computer*, 22(8):27–36, 1989. *Cited on page 62.*
- [90] A. Hennemuth, O. Friman, C. Schumann, J. Bock, J. Drexler, M. Huellebrand, M. Markl, and H.-O. Peitgen. Fast interactive exploration of 4D MRI flow data. In *Society of Photo-Optical Instrumentation Engineers*, volume 7964, pages 79640E:1–79640E:11, 2011. *Cited on pages 66 and 70.*
- [91] M. D. Hope, T. A. Hope, A. K. Meadows, K. G. Ordovas, T. H. Urbania, M. T. Alley, and C. B. Higgins. Bicuspid aortic valve : four-dimensional MR evaluation of ascending aortic systolic flow patterns. *Radiology*, 255(1):53–61, 2010. *Cited on pages 12 and 19.*
- [92] Y. Hu, H. Ahmed, C. Allen, D. Pendsé, M. Sahu, M. Emberton, D. Hawkes, and D. Barratt. MR to ultrasound image registration for guiding prostate biopsy and interventions. In *Medical Image Computing and Computer-Assisted Intervention*, volume 5761 of *Lecture Notes in Computer Science*, pages 787–794, 2009. *Cited on page 130.*
- [93] L. Hubert and P. Arabie. Comparing partitions. *Classification*, 2(1):193–218, 1985. *Cited on page 88.*
- [94] W. F. Hughes. *Fluid dynamics*. McGraw-Hill, 1999. *Cited on page 14.*
- [95] M. Hummel, C. Garth, B. Hamann, H. Hagen, and K. I. Joy. IRIS: illustrative rendering for integral surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1319–1328, 2010. *Cited on pages 66, 67, 118, and 119.*
- [96] V. Interrante, H. Fuchs, and S. Pizer. Enhancing transparent skin surfaces with ridge and valley lines. In *IEEE Visualization*, pages 52–59, 1995. *Cited on page 55.*
- [97] V. Interrante and C. Grosch. Strategies for effectively visualizing 3D flow with volume LIC. In *IEEE Visualization*, pages 421–424, 1997. *Cited on page 61.*
- [98] V. Interrante and C. Grosch. Visualizing 3D flow. *IEEE Computer Graphics and Applications*, 18(4):49–53, 1998. *Cited on page 61.*
- [99] T. Isenberg, N. Halper, and T. Strothotte. Stylizing silhouettes at interactive rates: from silhouette edges to silhouette strokes. *Computer Graphics Forum*, 21(3):249–258, 2002. *Cited on page 54.*
- [100] S. S. A. M. Jacobs. Hierarchical clustering and illustrative visualization of 4D MRI blood flow. Master’s thesis, Eindhoven University of Technology, 2011. *Cited on page 60.*
- [101] S. Jin, J. Oshinski, and D. P. Giddens. Effects of wall motion and compliance on flow patterns in the ascending aorta. *Journal of Biomechanical Engineering*, 125(3):347, 2003. *Cited on page 18.*
- [102] B. Jobard and G. Erlebacher. Hardware-accelerated texture advection for unsteady flow visualization. In *IEEE Visualization*, pages 155–162, 2000. *Cited on page 61.*
- [103] B. Jobard and W. Lefer. Creating evenly-spaced streamlines of arbitrary density. In *EG Visualization in Scientific Computing*, pages 43–56, 1997. *Cited on page 66.*
- [104] K. M. Johnson and M. Markl. Improved SNR in phase contrast velocimetry with five-point balanced flow encoding. *Magnetic Resonance in Medicine*, 63(2):349–55, 2010. *Cited on page 40.*
- [105] A. Joshi and P. Rheingans. Illustration-inspired techniques for visualizing time-varying data. In *IEEE Visualization*, pages 679–686, 2005. *Cited on page 113.*
- [106] T. Judd. *Apparent ridges for line drawing*. PhD thesis, Massachusetts Institute of Technology, 2007. *Cited on page 55.*
- [107] D. Kainmüller, R. Unterhinninghofen, S. Ley, and R. Dillmann. Level set segmentation of the heart from 4D phase contrast MRI. In *Society of Photo-Optical Instrumentation Engineers*, volume 6914, 2008. *Cited on page 100.*
- [108] B. Kainz, U. Reiter, G. Reiter, and D. Schmalstieg. In vivo interactive visualization of four-

- dimensional blood flow patterns. *The Visual Computer*, 25(9):853–862, 2009. Cited on page 61.
- [109] A. Karthikesalingam, P. J. E. Holt, R. J. Hinchliffe, M. M. Thompson, and I. M. Loftus. The diagnosis and management of aortic dissection. *Vascular and endovascular surgery*, 44(3):165–169, 2010. Cited on page 21.
- [110] C. Kasai, K. Namekawa, A. Koyano, and R. Omoto. Real-time two-dimensional blood flow imaging using an autocorrelation technique. *IEEE Transactions on Sonics and Ultrasonics*, 32(3):458–464, 1985. Cited on page 135.
- [111] M. Kass and A. Witkin. Snakes: active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988. Cited on page 73.
- [112] J. Kasten, C. Petz, I. Hotz, B. R. Noack, and H.-C. Hege. Localized finite-time Lyapunov exponent for unsteady flow analysis. In *Vision, Modeling, and Visualization*, pages 265–274, 2009. Cited on page 63.
- [113] P. J. Kilner, G.-Z. Yang, R. H. Mohiaddin, D. N. Firmin, and D. B. Longmore. Helical and retrograde secondary flow patterns in the aortic arch studied by three-directional magnetic resonance velocity mapping. *Circulation*, 88(5):2235–2247, 1993. Cited on page 18.
- [114] P. Kipfer, M. Segal, and R. Westermann. ÜberFlow: a GPU-based particle engine. In *Graphics Hardware*, pages 115–122, 2004. Cited on page 65.
- [115] R. M. Kirby, H. Marmanis, and D. H. Laidlaw. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *IEEE Visualization*, pages 333–341, 1999. Cited on page 59.
- [116] M. Kiu and D. Banks. Multi-frequency noise for LIC. In *IEEE Visualization*, pages 121–126, 1996. Cited on page 61.
- [117] R. Klassen and S. J. Harrington. Shadowed hedgehogs: a technique for visualizing 2D slices of 3D vector fields. In *IEEE Visualization*, pages 148–153, 1991. Cited on page 59.
- [118] H. Knutsson. Representing local structure using tensors. In *Scandinavian Conference on Image Analysis*, pages 244–251, 1989. Cited on pages 71 and 92.
- [119] A. Köhn, J. Klein, F. Weiler, and H.-O. Peitgen. A GPU-based fiber tracking framework using geometry shaders. In *Society of Photo-Optical Instrumentation Engineers*, volume 7261, pages 72611J:1–72611J:10, 2009. Cited on pages 104 and 117.
- [120] R. Kosara, F. Drury, L. E. Holmquist, and D. H. Laidlaw. Visualization criticism. *Computer Graphics And Applications*, 28(3):13–15, 2008. Cited on pages 48 and 53.
- [121] M. Kouwenhoven, M. B. M. Hofman, and M. Sprenger. Motion induced phase shifts in MR: acceleration effects in quantitative flow measurements— a reconsideration. *Magnetic Resonance in Medicine*, 33(6):766–777, 1995. Cited on page 35.
- [122] S. Kozerke, J. M. Hasenkam, E. M. Pedersen, and P. Boesiger. Visualization of flow patterns distal to aortic valve prostheses in humans using a fast approach for cine 3D velocity mapping. *Journal of Magnetic Resonance Imaging*, 13(5):690–698, 2001. Cited on page 19.
- [123] H. Krishnan, C. Garth, J. Guhring, M. A. Gulsun, A. Greiser, and K. I. Joy. Analysis of time-dependent flow-sensitive PC-MRI data. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–13, 2011. Cited on pages 63, 65, 66, 67, and 154.
- [124] H. Krishnan, C. Garth, and K. I. Joy. Time and streak surfaces for flow visualization in large time-varying data sets. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1267–1274, 2009. Cited on page 67.
- [125] J. Krüger, P. Kipfer, P. Kondratieva, and R. Westermann. A particle system for interactive visualization of 3D flows. *IEEE Transactions on Visualization and Computer Graphics*, 11(6):744–756, 2005. Cited on page 65.
- [126] A. Kuhn, D. J. Lehmann, R. Gasteiger, M. Neugebauer, B. Preim, and H. Theisel. A clustering-

- based visualization technique to emphasize meaningful regions of vector fields. In *Vision, Modeling, and Visualization*, pages 191–198, 2011. Cited on page 64.
- [127] P. K. Kundu and I. M. Cohen. *Fluid mechanics*. Elsevier Academic Press, 2008. Cited on pages 14, 15, and 135.
- [128] O. Kutter, R. Shams, and N. Navab. Visualization and GPU-accelerated simulation of medical ultrasound from CT images. *Computer Methods and Programs in Biomedicine*, 94(3):250–266, 2009. Cited on page 130.
- [129] J.-P. E. Kvitting, T. Ebbers, L. Wigström, J. Engvall, C. L. Olin, and A. F. Bolger. Flow patterns in the aortic root and the aorta studied with time-resolved, 3-dimensional, phase-contrast magnetic resonance imaging: implications for aortic valve-sparing surgery. *Journal of Thoracic and Cardiovascular Surgery*, 127(6):1602–1607, 2004. Cited on page 66.
- [130] M. Lachaise and R. Bamler. Minimum cost flow phase unwrapping supported by multibase-line unwrapped gradient. In *European Conference on Synthetic Aperture Radar*, pages 774–777, 2010. Cited on page 44.
- [131] A. Lake, M. Harris, M. Blackstein, and W. Disney. Stylized rendering techniques for scalable real-time 3D animation. In *Non-Photorealistic Animation and Rendering*, pages 13–22, 2000. Cited on page 57.
- [132] J. Langley and Q. Zhao. Unwrapping magnetic resonance phase maps with Chebyshev polynomials. *Journal of Magnetic Resonance Imaging*, 27(9):1293–1301, 2009. Cited on page 44.
- [133] S. L. Lansman, J. N. McCullough, K. H. Nguyen, D. Spielvogel, J. J. Klein, J. D. Galla, M. A. Ergin, and R. B. Griep. Subtypes of acute aortic dissection. *The Annals of Thoracic Surgery*, 67(6):1975–1978, 1999. Cited on pages 21 and 23.
- [134] R. S. Laramée, H. Hauser, H. Doleisch, B. Vrolijk, F. H. Post, and D. Weiskopf. The state of the art in flow visualization: dense and texture-based techniques. *Computer Graphics Forum*, 23(2):203–221, 2004. Cited on page 60.
- [135] R. S. Laramée, H. Hauser, L. Zhao, and F. H. Post. Topology-based flow visualization, the state of the art. In *Topology-based Methods in Visualization*, pages 1–19, 2007. Cited on page 62.
- [136] R. S. Laramée, B. Jobard, and H. Hauser. Image space based visualization of unsteady flow on surfaces. In *IEEE Visualization*, pages 131–138, 2003. Cited on page 61.
- [137] E. W. Larson and W. D. Edwards. Risk factors for aortic dissection: a necropsy study of 161 cases. *The American journal of cardiology*, 53(6):849–55, 1984. Cited on pages 20 and 21.
- [138] T. C. Larson, W. M. Kelly, R. L. Ehman, and F. W. Wehrli. Spatial misregistration of vascular flow during MR imaging of the CNS: cause and clinical significance. *American Journal of Roentgenology*, 155(5):1117–1124, 1990. Cited on page 45.
- [139] L. Li. *Advanced Flow Visualization*. PhD thesis, Ohio State University, 2007. Cited on page 66.
- [140] S. K. Lodha, A. Pang, R. E. Sheehan, and C. M. Wittenbrink. UFLOW: visualizing uncertainty in fluid flow. In *IEEE Visualization*, pages 249–254, 1996. Cited on page 67.
- [141] H. Löffelmann. *Visualizing local properties and characteristic structures of dynamical systems*. PhD thesis, Vienna University of Technology, 1998. Cited on page 63.
- [142] H. Löffelmann, A. König, and M. E. Gröller. Fast visualization of 2D dynamical systems by the use of virtual ink droplets. In *Spring Conference on Computer Graphics*, pages 111–118, 1997. Cited on page 61.
- [143] H. Löffelmann, L. Mroz, M. E. Gröller, and W. Purgathofer. Stream arrows: enhancing the use of stream surfaces for the visualization of dynamical systems. *The Visual Computer*, 13(8):359–369, 1997. Cited on pages 66 and 67.

- [144] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH Computer Graphics and Interactive Techniques*, pages 163–169. ACM, 1987. Cited on page 75.
- [145] J. Lotz, R. Döker, R. Noeske, M. Schütttert, R. Felix, M. Galanski, M. Gutberlet, and G. P. Meyer. In vitro validation of phase-contrast flow measurements at 3T in comparison to 1.5T: precision, accuracy, and signal-to-noise ratios. *Journal of Magnetic Resonance Imaging*, 21(5):604–610, 2005. Cited on page 44.
- [146] J. Lotz, C. Meier, A. Leppert, and M. Galanski. Cardiovascular flow measurement with phase-contrast MR imaging: basic facts and implementation. *Radiographics*, 22(3):651–671, 2002. Cited on pages 36 and 43.
- [147] A. Lu, C. J. Morris, J. Taylor, D. S. Ebert, C. Hansen, P. Rheingans, and M. Hartner. Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):127–138, 2003. Cited on page 55.
- [148] E. B. Lum and K.-L. Ma. Hardware-accelerated parallel non-photorealistic volume rendering. In *Non-Photorealistic Animation and Rendering*, pages 67–74, 2002. Cited on page 57.
- [149] A. J. Marelli, A. S. Mackie, R. Ionescu-Ittu, E. Rahme, and L. Pilote. Congenital heart disease in the general population: changing prevalence and age distribution. *Circulation*, 115(2):163–172, 2007. Cited on page 24.
- [150] M. Markl, F. P. Chan, M. T. Alley, K. L. Wedding, M. T. Draney, C. J. Elkins, D. W. Parker, C. A. Taylor, R. J. Herfkens, and N. J. Pelc. Time resolved three dimensional phase contrast MRI. *Journal of Magnetic Resonance Imaging*, 17(4):499–506, 2003. Cited on page 94.
- [151] M. Markl, M. T. Draney, M. D. Hope, J. M. Levin, F. P. Chan, M. T. Alley, N. J. Pelc, and R. J. Herfkens. Time-resolved 3-dimensional velocity mapping in the thoracic aorta: visualization of 3-directional blood flow patterns in healthy volunteers and patients. *Journal Of Computer Assisted Tomography*, 28(4):459–468, 2004. Cited on page 19.
- [152] M. Markl, J. Geiger, L. Herzer, B. Stiller, and R. Arnold. 3D aortic blood flow in patients with marfan syndrome: changes in hemodynamics and correlation with aortic geometry. In *Magnetic Resonance in Medicine*, volume 19, page 725, 2011. Cited on page 12.
- [153] M. Markl, A. Harloff, D. Föll, M. Langer, J. Hennig, and A. Frydrychowicz. Sclerotic aortic valve: flow-sensitive 4-dimensional magnetic resonance imaging reveals 3 distinct flow-pattern changes. *Circulation*, 116(10):e336–e337, 2007. Cited on page 19.
- [154] M. Markl, P. J. Kilner, T. Ebbers, M. Michael, K. Philip, and E. Tino. Comprehensive 4D velocity mapping of the heart and great vessels by cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 13(7):1–22, 2011. Cited on pages 17, 19, 66, and 90.
- [155] J. P. R. F. Marques. *Effects of dipolar fields in NMR and MRI*. PhD thesis, University of Nottingham, 2004. Cited on page 31.
- [156] O. Mattausch, T. Theußl, H. Hauser, and M. E. Gröller. Strategies for interactive exploration of 3D flow using evenly-spaced illuminated streamlines. In *Spring Conference on Computer Graphics*, pages 213–222, 2003. Cited on pages 66 and 105.
- [157] N. Max, R. Crawfis, and C. Grant. Visualizing 3D velocity fields near contour surfaces. In *IEEE Visualization*, pages 248–255, 1994. Cited on page 65.
- [158] S. McCloud. *Understanding Comics*. HarperCollins Publishers inc., 1993. Cited on page 3.
- [159] M. McGuire and A. Fein. Real-time rendering of cartoon smoke and clouds. In *Non-Photorealistic Animation and Rendering*, pages 21–26, 2002. Cited on page 57.
- [160] M. McGuire and J. F. Hughes. Hardware-determined feature edges. In *Non-Photorealistic Animation and Rendering*, pages 35–47. ACM Press, 2004. Cited on page 55.
- [161] A. McKenzie and S. Lombeyda. Vector field analysis and visualization through variational clustering. In *EG/IEEE EuroVis*, pages 29–35, 2005. Cited on page 64.

- [162] T. McLoughlin, R. S. Laramée, R. Peikert, F. H. Post, and M. Chen. Over two decades of integration-based geometric flow visualization. *Computer Graphics Forum*, 29(6):1807–1829, 2010. Cited on pages 65 and 67.
- [163] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince. *MRI - from picture to proton*. Cambridge University Press, 2007. Cited on page 30.
- [164] J. Meyer-Spradow, T. Ropinski, J. Vahrenhold, and K. Hinrichs. Illustrating dynamics of time-varying volume datasets in static images extracting dynamics from time-varying volume datasets. In *Vision, Modeling, and Visualization*, pages 333–340, 2006. Cited on page 113.
- [165] V. Mihalef, R. I. Ionasec, P. Sharma, B. Georgescu, I. Voigt, M. Suehling, and D. Comaniciu. Patient-specific modelling of whole heart anatomy, dynamics and hemodynamics from 4D cardiac CT images. *Interface Focus*, 1(3):286–296, 2011. Cited on page 15.
- [166] B. R. Munson, D. F. Young, and T. H. Okiishi. *Fundamentals of fluid mechanics*. Wiley, 2005. Cited on page 14.
- [167] Z. Nagy, J. Schneider, and R. Westermann. Interactive volume illustration. In *Vision, Modeling, and Visualization*, pages 497–504, 2002. Cited on page 55.
- [168] N. H. Naqvi and M. D. Blaufox. *Blood pressure measurement*. Parthenon publishing, 1998. Cited on page 11.
- [169] R. Otten, A. Vilanova, and H. van de Wetering. Illustrative white matter fiber bundles. *Computer Graphics Forum*, 29(3):1013–1022, 2010. Cited on page 54.
- [170] S. Pan, L. S. Kleppe, T. A. Witt, C. S. Mueske, and R. D. Simari. The effect of vascular smooth muscle cell-targeted expression of tissue factor pathway inhibitor in a murine model of arterial thrombosis. *Thrombosis and Haemostasis*, 92(3):495–502, 2004. Cited on page 12.
- [171] M. K. Park. *Pediatric cardiology*. Mosby, 2007. Cited on page 23.
- [172] T. H. J. M. Peeters, V. Prckovska, M. Almsick van, A. Vilanova, and B. M. ter Haar Romeny. Fast and sleek glyph rendering for interactive HARDI data exploration. In *IEEE Pacific Visualization Symposium*, pages 153–160, 2009. Cited on page 51.
- [173] T. H. J. M. Peeters, A. Vilanova, G. Strijkers, and B. M. ter Haar Romeny. Visualization of the fibrous structure of the heart. In *Vision, Modeling, and Visualization*, pages 309–317, 2006. Cited on page 105.
- [174] T. H. J. M. Peeters, A. Vilanova, and B. M. ter Haar Romeny. Interactive fibre structure visualization of the heart. *Computer Graphics Forum*, 28(8):2140–2150, 2009. Cited on page 51.
- [175] N. J. Pelc, M. A. Bernstein, A. Shimakawa, and G. H. Glover. Encoding strategies for three-direction phase-contrast MR imaging of flow. *Journal of Magnetic Resonance Imaging*, 1(4):405–413, 1991. Cited on page 40.
- [176] G. Pennati, F. Migliavacca, G. Dubini, and E. L. Bove. Modeling of systemic-to-pulmonary shunts in newborns with a univentricular circulation: state of the art and future directions. *Progress in Pediatric Cardiology*, 30(1-2):23–29, 2010. Cited on page 15.
- [177] A. E. Perry and M. S. Chong. Topology of flow patterns in vortex motions and turbulence. *Applied Scientific Research*, 53(3):357–374, 1994. Cited on page 63.
- [178] M. Persson, J. E. Solem, K. Markenroth, J. Svensson, and A. Heyden. Phase contrast MRI segmentation using velocity and intensity. *Scale Space and PDE Methods in Computer Vision*, 3459(4):119–130, 2005. Cited on pages 100 and 110.
- [179] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975. Cited on page 57.
- [180] M. E. Pierpont, C. T. Basson, D. W. Benson, B. D. Gelb, T. M. Giglia, E. Goldmuntz, G. McGee, C. A. Sable, D. Srivastava, and C. L. Webb. Genetic basis for congenital heart defects: current knowledge: a scientific statement from the American Heart Association Con-

- genital Cardiac Defects Committee, Council on Cardiovascular Disease in the Young: endorsed by the American Academy of Pediatrics. *Circulation*, 115(23):3015–3038, 2007. *Cited on page 24.*
- [181] A. Pobitzer, R. Peikert, R. Fuchs, and B. Schindler. On the way towards topology-based visualization of unsteady flow - the state of the art. In *Eurographics STARs*, pages 137–154, 2010. *Cited on page 63.*
- [182] R. Porter. *Blood & guts*. W. W. Norton & Company, 2004. *Cited on page 10.*
- [183] F. H. Post and T. van Walsum. Fluid flow visualization. In *Focus on Scientific Visualization*, pages 1–37, 1993. *Cited on page 59.*
- [184] F. H. Post, B. Vrolijk, H. Hauser, R. S. Laramée, and H. Doleisch. The state of the art in flow visualisation: feature extraction and tracking. *Computer Graphics Forum*, 22(4):775–792, 2003. *Cited on pages 61 and 62.*
- [185] J. E. Potchen, M. E. Haacke, J. E. Siebert, and A. Gottschalk. *Magnetic resonance angiography - concepts & applications*. Mosby, 1993. *Cited on pages 16, 19, and 33.*
- [186] E. Praun, H. Hoppe, M. Webb, and A. Finkelstein. Real-time hatching. In *Computer Graphics and Interactive Techniques*, page 581, 2001. *Cited on page 55.*
- [187] V. Prckovska and T. H. J. M. Peeters. Fused DTI/HARDI Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(10):1407–1419, 2011. *Cited on pages 51 and 115.*
- [188] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical recipes in C*. Cambridge University Press, 2nd edition, 1992. *Cited on page 96.*
- [189] K. Pruessmann, P. Boesiger, and S. Kozerke. Gyrotools GTFlow. <http://www.gyrotools.com/products/gt-flow.html> (last visited 01-02-2012). *Cited on pages 65 and 113.*
- [190] P. Rautek, S. Bruckner, and M. E. Gröller. Interaction dependent semantics for illustrative volume rendering. *Computer Graphics Forum*, 27(3):847–854, 2008. *Cited on page 53.*
- [191] P. Rautek, S. Bruckner, M. E. Gröller, and I. Viola. Illustrative visualization: new technology or useless tautology? *ACM SIGGRAPH Computer Graphics Newsletter*, 42(3):1–8, 2008. *Cited on pages 52, 53, and 54.*
- [192] T. Reichl, J. Passenger, O. Acosta, and O. Salvado. Ultrasound goes GPU: real-time simulation using CUDA. *Progress in Biomedical Optics and Imaging*, 10(37):1–10, 2009. *Cited on page 130.*
- [193] C. Rezk-Salama, P. Hastreiter, and C. Teitzel. Interactive exploration of volume line integral convolution based on 3D-texture mapping. In *IEEE Visualization*, pages 233–240, 1999. *Cited on pages 60 and 61.*
- [194] Y. Richter and E. R. Edelman. Cardiology is flow. *Circulation*, 113(23):2679–2682, 2006. *Cited on page 12.*
- [195] V. L. Roger, A. S. Go, D. M. Lloyd-Jones, R. J. Adams, J. D. Berry, T. M. Brown, M. R. Carnethon, S. Dai, G. de Simone, E. S. Ford, C. S. Fox, H. J. Fullerton, C. Gillespie, K. J. Greenlund, S. M. Hailpern, J. A. Heit, P. M. Ho, V. J. Howard, B. M. Kissela, S. J. Kittner, D. T. Lackland, J. H. Lichtman, L. D. Lisabeth, D. M. Makuc, G. M. Marcus, A. Marelli, D. B. Matchar, M. M. McDermott, J. B. Meigs, C. S. Moy, D. Mozaffarian, M. E. Mussolino, G. Nichol, N. P. Paynter, W. D. Rosamond, P. D. Sorlie, R. S. Stafford, T. N. Turan, M. B. Turner, N. D. Wong, and J. Wylie-Rosett. Heart disease and stroke statistics–2011 update: a report from the American Heart Association. *Circulation*, 123(4):210, 2011. *Cited on page 13.*
- [196] M. P. Rolf, M. B. M. Hofman, P. D. Gatehouse, K. Markenroth-Bloch, M. W. Heymans, T. Ebbers, M. J. Graves, J. J. Totman, B. Werner, A. C. van Rossum, P. J. Kilner, and R. M. Heethaar. Sequence optimization to reduce velocity offsets in cardiovascular magnetic resonance volume flow quantification - a multi-vendor study. *Journal of Cardiovascular Magnetic*

- Resonance*, 13(1):18, 2011. *Cited on page 44.*
- [197] S. Rusinkiewicz, M. Burns, and D. DeCarlo. Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics*, 25(3):1199–1205, 2006. *Cited on page 57.*
- [198] M. F. Salfity, P. D. Ruiz, J. M. Huntley, M. J. Graves, R. Cusack, and D. A. Beaugard. Branch cut surface placement for unwrapping of undersampled three-dimensional phase data: application to magnetic resonance imaging arterial flow mapping. *Applied optics*, 45(12):2711–2722, 2006. *Cited on page 44.*
- [199] T. Salzbrunn, H. Jänicke, T. Wischgoll, and G. Scheuermann. The state of the art in flow visualization: Partition-based techniques. In *Simulation and Visualization*, pages 75–92, 2008. *Cited on page 64.*
- [200] D. Sayer, M. Bratby, M. Brooks, I. Loftus, R. Morgan, and M. Thompson. Aortic morphology following endovascular repair of acute and chronic type B aortic dissection: implications for management. *European Journal of Vascular and Endovascular Surgery*, 36(5):522–529, 2008. *Cited on page 23.*
- [201] T. Schafhitzel, E. Tejada, D. Weiskopf, and T. Ertl. Point-based stream surfaces and path surfaces. In *Graphics Interface*, pages 289–296, 2007. *Cited on page 66.*
- [202] H. Scharsach. Advanced GPU raycasting. In *Central European Seminar on Computer Graphics*, volume 5, pages 67–76, 2005. *Cited on page 138.*
- [203] M. Schulz, F. Reck, W. Bartelheimer, and T. Ertl. Interactive visualization of fluid dynamics simulations in locally refined Cartesian grids. In *IEEE Visualization*, pages 413–553, 1999. *Cited on page 66.*
- [204] A. Secord. Weighted voronoi stippling. In *Non-Photorealistic Animation and Rendering*, pages 37–43, 2002. *Cited on page 55.*
- [205] N. Shahcheraghi, H. A. Dwyer, A. Y. Cheer, A. I. Barakat, and T. Rutaganira. Unsteady and three-dimensional simulation of blood flow in the human aortic arch. *Journal of Biomechanical Engineering*, 124(4):378–387, 2002. *Cited on page 15.*
- [206] R. Shams, R. Hartley, and N. Navab. Real-time simulation of medical ultrasound from CT images. In *Medical Image Computing and Computer-Assisted Intervention*, pages 734–741, 2008. *Cited on page 130.*
- [207] H.-W. Shen, C. R. Johnson, and K.-L. Ma. Visualizing vector fields using line integral convolution and dye advection. In *Volume visualization*, pages 63–71, 1996. *Cited on page 61.*
- [208] B. Shneiderman. The eyes have it : a task by data type taxonomy for information visualizations. In *Visual Languages*, pages 1–8, 1996. *Cited on page 49.*
- [209] I. A. Simpson, L. M. Valdes-Cruz, D. J. Sahn, A. Murillo, T. Tamura, and K. J. Chung. Doppler color flow mapping of simulated in vitro regurgitant jets: evaluation of the effects of orifice size and hemodynamic variables. *Journal of the American College of Cardiology*, 13(5):1195–1207, 1989. *Cited on page 135.*
- [210] P.-P. J. Sloan, W. Martin, A. Gooch, and B. Gooch. The lit sphere : a model for capturing NPR shading from art. In *Graphics Interface*, pages 143–150, 2001. *Cited on page 58.*
- [211] J. S. Sobel, A. S. Forsberg, D. H. Laidlaw, R. C. Zeleznik, D. F. Keefe, I. Pivkin, G. E. Karniadakis, P. Richardson, and S. Swartz. Particle flurries: synoptic 3D pulsatile flow visualization. *Computer Graphics and Applications*, 24(2):76–85, 2004. *Cited on pages 65 and 112.*
- [212] J. E. Solem, M. Persson, and A. Heyden. Velocity based segmentation in phase-contrast MRI images. In *Medical Image Computing and Computer-Assisted Intervention*, pages 459–466, 2004. *Cited on page 74.*
- [213] F. Sotiropoulos and I. Borazjani. A review of state-of-the-art numerical methods for simulating flow through mechanical heart valves. *Medical & Biological Engineering & Computing*, 47(3):245–256, 2009. *Cited on page 15.*

- [214] S. P. Souza, J. Szumowski, C. L. Dumoulin, D. P. Plewes, and G. H. Glover. Simultaneous multislice acquisition of MR images by Hadamard-encoded excitation. *Computer Assisted Tomography*, 12(6):1026–1030, 1988. Cited on page 40.
- [215] M. Spindler, N. Röber, R. Döhring, and M. Masuch. Enhanced cartoon and comic rendering. In *Eurographics*, pages 141–144, 2006. Cited on page 57.
- [216] D. Stalling and H.-C. Hege. Fast and resolution independent line integral convolution. In *ACM SIGGRAPH Computer Graphics and Interactive Techniques*, pages 249–256, 1995. Cited on page 61.
- [217] D. Stalling, M. Zöckler, and H.-C. Hege. Fast display of illuminated field lines. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):118–128, 1997. Cited on page 97.
- [218] L.-M. Stevens, J. C. Madsen, E. M. Isselbacher, P. Khairy, T. E. MacGillivray, A. D. Hilgenberg, and A. K. Agnihotri. Surgical management and long-term outcomes for acute ascending aortic dissection. *The Journal of Thoracic and Cardiovascular Surgery*, 138(6):1349–1357, 2009. Cited on page 23.
- [219] C. Stoll, S. Gumhold, and H.-P. Seidel. Visualization with stylized line primitives. In *IEEE Visualization*, pages 695–702. IEEE, 2005. Cited on page 66.
- [220] P. Svetachov, M. H. Everts, and T. Isenberg. DTI in context : illustrating brain fiber tracts in situ. *Computer Graphics Forum*, 29(3):1023–1032, 2010. Cited on page 66.
- [221] C. A. Taylor and C. A. Figueroa. Patient-specific modeling of cardiovascular mechanics. *Annual review of Biomedical Engineering*, 11:109–134, 2009. Cited on page 15.
- [222] A. Telea and J. J. van Wijk. Simplified representation of vector fields. In *IEEE Visualization*, pages 35–42, 1999. Cited on pages 63, 64, 80, 82, and 88.
- [223] M. Termeer, J. Oliván Bescós, M. Breeuwer, A. Vilanova, F. A. Gerritsen, and M. E. Gröller. CoViCAD: comprehensive visualization of coronary artery disease. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1632–1639, 2007. Cited on page 54.
- [224] The Japan Society of Mechanical Engineers. *Visualized flow*. Pergamon Press, 1990. Cited on page 59.
- [225] H. Theisel and H.-P. Seidel. Feature flow fields. In *EG VisSym*, pages 141–148, 2003. Cited on page 63.
- [226] Y. Tokuda, M.-H. Song, Y. Ueda, A. Usui, T. Akita, S. Yoneyama, and S. Maruyama. Three-dimensional numerical simulation of blood flow in the aortic arch during cardiopulmonary bypass. *European Journal of Cardiothoracic Surgery*, 33(2):164–167, 2008. Cited on page 15.
- [227] M. Tory and T. Möller. Rethinking visualization: a high-level taxonomy. In *IEEE Information Visualization*, pages 151–158, 2004. Cited on page 49.
- [228] A. Treuille, A. Lewis, and Z. Popović. Model reduction for real-time fluids. *ACM Transactions on Graphics*, 25(3):826–834, 2006. Cited on page 14.
- [229] X. Tricoche, C. Garth, G. Kindlmann, E. Deines, G. Scheuermann, M. Ruetten, and C. Hansen. Visualization of intricate flow structures for vortex breakdown analysis. In *IEEE Visualization*, pages 187–194, 2004. Cited on page 62.
- [230] H. Tucker. *Blood work, a tale of medicine and murder in the scientific revolution*. W. W. Norton & Company, 2011. Cited on page 10.
- [231] B. Tversky, J. B. Morrison, and M. Betrancourt. Animation: can it facilitate? *International Journal of Human-Computer Studies*, 57(4):247 – 262, 2002. Cited on page 90.
- [232] J. M. Tyszka, D. H. Laidlaw, J. W. Asa, and J. M. Silverman. Three-dimensional, time-resolved (4D) relative pressure mapping using magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 12(2):321–329, 2000. Cited on page 17.
- [233] S. Uribe, P. Beerbaum, T. S. Sørensen, A. Rasmusson, R. Razavi, and T. Schaeffter. Four-

- dimensional (4D) flow of the whole heart and great vessels using real-time respiratory self-gating. *Magnetic Resonance in Medicine*, 62(4):984–992, 2009. Cited on page 94.
- [234] I. Valverde, J. Simpson, T. Schaeffter, and P. Beerbaum. 4D phase-contrast flow cardiovascular magnetic resonance: comprehensive quantification and visualization of flow dynamics in atrial septal defect and partial anomalous pulmonary venous return. *Pediatric cardiology*, 31(8):1244–1248, 2010. Cited on pages 17, 25, 26, 27, 41, and 65.
- [235] A. E. van der Hulst, J. J. M. Westenberg, L. J. M. Kroft, J. J. Bax, N. A. Blom, A. de Roos, and A. A. W. Roest. Tetralogy of Fallot: 3D velocity-encoded MR imaging for evaluation of right ventricular valve flow and diastolic function in patients after correction. *Radiology*, 256(3):724–734, 2010. Cited on page 12.
- [236] T. Van Laerhoven, J. Liesenborgs, and F. Van Reeth. Real-time watercolor painting on a distributed paper model. In *Computer Graphics International*, pages 640–643, 2004. Cited on page 53.
- [237] R. F. P. van Pelt, T. Q. H. Nguyen, B. M. ter Haar Romeny, and A. Vilanova. Automated segmentation of blood-flow regions in large thoracic arteries using 3D-cine PC-MRI measurements. *Journal of Computer Assisted Radiology and Surgery*, 7(2):217–224, 2012. Cited on pages 17 and 100.
- [238] R. F. P. van Pelt, J. Oliván Bescós, M. Breeuwer, R. E. Clough, M. E. Gröller, B. M. ter Haar Romeny, and A. Vilanova. Exploration of 4D MRI blood flow using stylistic visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1339–1347, 2010. Cited on page 55.
- [239] R. F. P. van Pelt, J. Oliván Bescós, M. Breeuwer, R. E. Clough, M. E. Gröller, B. M. ter Haar Romeny, and A. Vilanova. Interactive virtual probing of 4D MRI blood flow. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2153–2162, 2011. Cited on page 66.
- [240] R. F. P. van Pelt, A. Vilanova, and H. van de Wetering. Illustrative volume visualization using GPU-based particle systems. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):571–582, 2010. Cited on pages 55, 56, and 68.
- [241] R. F. P. van Pelt, A. Vilanova, and H. Wetering van de. GPU-based particle systems for illustrative volume rendering. In *Volume and Point-based Graphics*, pages 2–9, 2008. Cited on page 55.
- [242] A. A. van Steenhoven, P. C. Veenstra, and R. S. Reneman. The effect of some hemodynamic factors on the behaviour of the aortic valve. *Journal of Biomechanics*, 15(12):941–950, 1982. Cited on page 12.
- [243] J. J. van Wijk. Spot noise texture synthesis for data visualization. *ACM SIGGRAPH Computer Graphics and Interactive Techniques*, 25(4):309–318, 1991. Cited on page 60.
- [244] J. J. van Wijk. Flow visualization with surface particles. *IEEE Computer Graphics and Applications*, 13(4):18–24, 1993. Cited on page 65.
- [245] J. J. van Wijk. Bridging the gaps. *IEEE Computer Graphics and Applications*, 26(6):6–9, 2006. Cited on page 4.
- [246] J. J. van Wijk. Views on visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):421–432, 2006. Cited on pages 48 and 49.
- [247] J. J. van Wijk, C. Ware, C. Demiralp, and D. H. Laidlaw. Panel on theories of visualization. In *IEEE Visualization*, pages 1–4, 2011. Cited on page 49.
- [248] S. Venkataraman. 4D visualization of cardiac flow (NVIDIA). http://www.nvidia.com/content/GTC-2010/pdfs/2009_GTC2010.pdf (last visited 01-02-2012), 2010. Cited on pages 65 and 113.
- [249] J. A. Vierendeels, K. Riemsdijk, E. Dick, and P. R. Verdonck. Computer simulation of intra-ventricular flow and pressure gradients during diastole. *Journal of Biomechanical Engineering*,

- 122(6):667–674, 2000. *Cited on page 15.*
- [250] A. Vilanova, G. Berenschot, and C. van Pul. DTI visualization with streamsurfaces and evenly-spaced volume seeding. In *EG VisSym*, pages 173–182, 2004. *Cited on pages 51, 66, and 105.*
- [251] I. Viola, A. Kanitsar, and M. E. Gröller. Importance-driven volume rendering. In *IEEE Visualization*, pages 139–146, 2004. *Cited on page 53.*
- [252] M. T. Vlaardingerbroek and J. A. den Boer. *Magnetic resonance imaging*. Springer, 1999. *Cited on pages 30, 31, 32, and 45.*
- [253] I. Waechter, J. Bredno, R. Hermans, J. Weese, D. C. Barratt, and D. J. Hawkes. Model-based blood flow quantification from rotational angiography. *Medical Image Analysis*, 12(5):586–602, 2008. *Cited on page 16.*
- [254] C. Ware. *Information visualization - perception for design*. Morgan Kaufmann, 2000. *Cited on page 48.*
- [255] T. A. Wåxnäs, E. Heiberg, J. Togert, M. Larsson, H. Arheden, and M. Carlsson. Quantification of left ventricular kinetic energy using 4D flow MRI. *Journal of Cardiovascular Magnetic Resonance*, 12(1):O94–O95, 2010. *Cited on page 17.*
- [256] R. Wegenkittl and M. E. Gröller. Fast oriented line integral convolution for vector field visualization via the internet. In *IEEE Visualization*, pages 309–316, 1997. *Cited on page 61.*
- [257] R. Wegenkittl, M. E. Gröller, and W. Purgathofer. Animating flow fields: rendering of oriented line integral convolution. In *Computer Animation*, pages 15–21, 1997. *Cited on page 61.*
- [258] W. Wein, A. Khamene, D.-A. Clevert, O. Kutter, and N. Navab. Simulation and fully automatic multimodal registration of medical ultrasound. In *Medical Image Computing and Computer-Assisted Intervention*, pages 136–143, 2007. *Cited on page 130.*
- [259] W. Wein, O. Kutter, A. Aichert, D. Zikic, A. Kamen, and N. Navab. Automatic non-linear mapping of pre-procedure CT volumes to 3D ultrasound. In *IEEE International Symposium on Biomedical Imaging*, pages 1225–1228, 2010. *Cited on page 130.*
- [260] T. Weinkauff. *Extraction of topological structures in 2D and 3D vector fields*. PhD thesis, University of Magdeburg, 2008. *Cited on pages 63 and 65.*
- [261] D. Weiskopf. Fields by texture advection via programmable per-pixel operations. In *Vision, Modeling, and Visualization*, pages 439–446, 2001. *Cited on page 61.*
- [262] J. J. M. Westenberg, S. D. Roes, S. Hammer, P. J. van den Boogaard, N. A. Marsan, J. J. Bax, J. H. C. Reiber, A. de Roos, and R. J. van der Geest. Flow assessment over all heart valves simultaneously using 3D velocity-encoded MRI with retrospective valve tracking. In *International Society for Magnetic Resonance in Medicine*, volume 17, pages 317–317, 2009. *Cited on page 12.*
- [263] L. Wigström, T. Ebbers, A. Fyrenius, M. Karlsson, J. Engvall, B. Wranne, and A. F. Bolger. Particle trace visualization of intracardiac flow using time-resolved 3D phase contrast MRI. *Magnetic Resonance in Medicine*, 41(4):793–799, 1999. *Cited on page 65.*
- [264] J. Wilhelms and A. van Gelder. Octrees for faster isosurface generation. *ACM Transactions on Graphics*, 11(3):201–227, 1992. *Cited on page 111.*
- [265] K. C. Woo and J. I. Schneider. High-risk chief complaints I: chest pain - the big three. *Emergency Medicine Clinics of North America*, 27(4):685–712, 2009. *Cited on page 21.*
- [266] World Heart Organization. The global burden of disease: 2004 update. http://www.who.int/healthinfo/global_burden_disease/ (last visited 01-02-2012), 2008. *Cited on page 12.*
- [267] World Heart Organization. Fact sheet cardiovascular diseases. <http://www.who.int/mediacentre/factsheets/fs317/> (last visited 01-02-2012), 2011. *Cited on page 19.*

- [268] X. Xuexiang and Y. He. An effective illustrative visualization framework based on photic extremum lines (PELs). *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1328–1335, 2007. *Cited on page 55.*
- [269] G.-Z. Yang, P. Burger, P. J. Kilner, S. P. Karwatowski, and D. N. Firmin. Dynamic range extension of cine velocity measurements using motion-registered spatiotemporal phase unwrapping. *Journal of Magnetic Resonance Imaging*, 6(3):495–502, 1996. *Cited on page 44.*
- [270] G.-Z. Yang, P. J. Kilner, N. B. Wood, S. R. Underwood, and D. N. Firmin. Computation of flow pressure fields from magnetic resonance velocity mapping. *Magnetic Resonance in Medicine*, 36(4):520–526, 1996. *Cited on page 17.*
- [271] H. Yu, C. Wang, and K.-L. Ma. Parallel hierarchical visualization of large time-varying 3D vector fields. In *ACM/IEEE Supercomputing*, pages 24:1–24:12, 2007. *Cited on pages 63 and 64.*
- [272] X. Yuan and B. Chen. Illustrating surfaces in volume. In *IEEE Visualization*, pages 9–16, 2004. *Cited on page 55.*
- [273] P. A. Yushkevich, J. Piven, C. Hazlett, H. Smith, G. Smith, R. Ho, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006. *Cited on page 100.*
- [274] M. Zöckler, D. Stalling, and H.-C. Hege. Interactive visualization of 3D-vector fields using illuminated stream lines. In *IEEE Visualization*, pages 107–113, 1996. *Cited on page 66.*
- [275] M. Zöckler, D. Stalling, and H.-C. Hege. Parallel line integral convolution. In *Parallel Computing*, pages 975–989, 1997. *Cited on page 61.*

List of figures

2.1	History of hemodynamic research	11
2.2	Disease mortality	12
2.3	Cardiovascular disease mortality	13
2.4	Aortic blood-flow patterns	18
2.5	Dissection example: types	20
2.6	Dissection example: blood flow	22
2.7	Congenital heart disease prevalence	23
2.8	Congenital heart disease example: anatomy	26
2.9	Congenital heart disease example: blood flow	27
3.1	Spin physics	31
3.2	Spin echo pulse sequence	32
3.3	Speed of fluid flow in the human body	33
3.4	Velocity encoding	36
3.5	Subtraction methods and velocity encoding speed	38
3.6	PC-MRI acquisition - planar: through-plane flow	39
3.7	PC-MRI acquisition - volumetric: velocity field	41
3.8	PC-MRI acquisition - artifacts	43
4.1	Medical photograph versus illustration	50
4.2	Caricature of Leonardo da Vinci	54
4.3	Illustrative visualization of the cardiovascular anatomy	56
4.4	Illustrative visualization: toon shading and contours	57
4.5	Flow visualization: vector plot versus streamlines	58
4.6	Flow visualization: texture-based	60
4.7	Flow visualization: feature-based	62
4.8	Flow visualization: partition-based	63
4.9	Flow visualization: geometry-based	66
5.1	Temporal projections	71
5.2	Active surface segmentation	72
5.3	Active surface segmentation features	74
5.4	Active surface segmentation external forces	76
5.5	Clustering level-of-detail selection	79
5.6	Clustering dissimilarity measures	80
5.7	Clustering strategies	82
5.8	Patharrow seeding	84
5.9	Patharrow visualization	86
5.10	Patharrow visualization with context: volunteer	86
5.11	Patharrow visualization with context: pathology	87

6.1	Seeding approaches	94
6.2	Flow visualization: planar reformat	95
6.3	Flow visualization: integral lines	95
6.4	Flow visualization: integral lines reverse	98
6.5	Flow visualization: integral lines highlight	98
6.6	Flow visualization: flow-rate arrows	99
6.7	Flow visualization: anatomical context	99
6.8	Flow visualization: pathology	103
7.1	Virtual probe: planes and fitting	108
7.2	Flow visualization: illustrative particles	114
7.3	Flow visualization: illustrative particle bolus	114
7.4	Flow visualization: illustrative pathlines	116
7.5	GPU implementation: integral lines	116
7.6	Flow visualization: integral surfaces	120
7.7	GPU implementation: integral surfaces	120
7.8	Context visualization: anatomy	121
7.9	Flow visualization: pathology example	125
8.1	Color Doppler ultrasound	132
8.2	Overview of the ultrasound probe interaction	132
8.3	Overview of the ultrasound-inspired MPR	133
8.4	Flow visualization: color Doppler US inspired plane	136
8.5	Flow visualization: 2D-cine MRI flow inspired plane	137
8.6	Flow visualization: velocity-field volume rendering	140
8.7	Volunteer thorax example	141
8.8	Volunteer valve example	141
A.1	The cardiovascular system	158
B.1	Anatomical context	163
B.2	Planar visualizations	163
B.3	Seeding distributions and templates	166
B.4	Integral line color	166
B.5	Integral line reversed tracing	167
B.6	Integral line illumination	167
B.7	Integral line highlights	167
B.8	Flow-rate arrow trails	170
B.9	Anatomical context comparison	170
B.10	Anatomical context alternatives	171
B.11	Seeding comparison	171
B.12	Streamlines comparison	171
B.13	Pathlines comparison	171

C.1	View-aligned slice	175
C.2	Virtual probe positioning	175
C.3	Virtual probe positioning mockups	175
C.4	Virtual probe interaction	177
C.5	Virtual probe appearance	178
C.6	Flow visualization interactive seeding	178
C.7	Flow visualization: particle trace	180
C.8	Flow visualization: integral lines	181
C.9	Flow visualization: integral surfaces	181
C.10	Context visualization: anatomy	184
C.11	Context visualization: anatomy clipping	184
C.12	Context visualization: anatomy clipping and blood flow	184
C.13	Flow visualization: Nvidia	185
C.14	Flow visualization: Gyrotools GTFlow	185

List of tables

5.1	Segmentation results	78
5.2	Clustering container complexities	83
5.3	Clustering results	89
7.1	Geometry-based flow visualization: evaluation results	124
8.1	Color Doppler US inspired flow visualization: performance results . . .	143

List of abbreviations

AAo	Ascending Aorta
Ao	Aorta
AHA	American Heart Association
ASD	Atrial Septal Defect
AV	Atrioventricular
BFS	Blood-Flow Speed
B-TFE	Balanced Turbo Field Echo
CFD	Computational Fluid Dynamics
CHC	Coarse Hierarchical Clustering
CHD	Congenital Heart Disease
CIE	Commission Internationale de l'Eclairage
CL	Computing Language
CPU	Central Processing Unit
CSF	Cerebrospinal Fluid
CT	Computed Tomography
CUDA	Compute Unified Device Architecture
CVD	Cardiovascular Disease
DCP	Distance to Closest Point
DTI	Diffusion Tensor Imaging
DVR	Direct Volume Rendering
DWI	Diffusion Weighted Imaging
ECG	Electrocardiogram
EVC	Eigenvalue Coherence
FFE	Fast Field Echo
FPS	Frames Per Second
FTLE	Finite-Time Lyapunov Exponent
GPGPU	General-Purpose computing on the GPU
GPU	Graphics Processing Unit
GS	Geometry Shader
HARDI	High-Angular Resolution Diffusion Imaging
ICD	International Classification of Diseases
Im	Imaginary part of the complex plane
IVC	Inferior Vena Cava
LC	Line Coherence
LDAo	Lower Descending Aorta
LFL	Lower False Lumen
LIC	Line Integral Convolution
LLPV	Left Lower Pulmonary Vein
LPA	Left Pulmonary Atery
LPC	Local Phase Coherence

LUPV	Left Upper Pulmonary Vein
MIP	Maximum Intensity Projection
MPA	Main Pulmonary Artery
MPR	Multi-Planar Reformat
MR	Magnetic Resonance
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
NPR	Non-Photorealistic Rendering
OSI	Oscillating Shear Index
PAPVR	Partial Anomalous Pulmonary Venous Return
PC	Phase-Contrast
PPC	Point-to-Point Coherence
PCA	Phase-Contrast Angiography
PCA-M	Phase-Contrast Angiography Magnitude
PCA-P	Phase-Contrast Angiography Phase
PDA	Patent Ductus Arteriosus
Re	Real part of the complex plane
RF	Radio Frequency
RGB	Red Green Blue
RLPV	Right Lower Pulmonary Vein
ROI	Region Of Interest
RPA	Right Pulmonary Atery
RUPV	Right Upper Pulmonary Vein
SNR	Signal-to-Noise Ratio
SPMD	Single-Program Multiple-Data
SSFP	Steady-State Free Precession
STL	Standard Template Library
SVC	Superior Vena Cava
tMIP	Temporal Maximum Intensity Projection
tMOP	Temporal Mean Orientation Projection
TOF	Time-of-Flight
UDAo	Upper Descending Aorta
UFL	Upper False Lumen
US	Ultrasound
VBO	Vertex Buffer Object
VM	Vector Median
VS	Vertex Shader
VSD	Ventricular Septal Defect
WHO	World Health Organization
WSS	Wall Shear Stress

List of symbols

\vec{a}	acceleration, with direction and magnitude
\mathbf{A}	matrix describing velocity gradients
\vec{B}_0	main magnetic field
C	set of cluster voxels
\hat{d}	ultrasound inspection direction of unit-length
d_{ell}	elliptical dissimilarity measure
d_{lin}	local linear expansion dissimilarity measure
d_p	positional dissimilarity
d_{th}	dissimilarity threshold
d_v	velocity dissimilarity
\hat{e}	view direction of unit-length
\vec{F}_e	elasticity force
\vec{F}_{int}	internal forces
\vec{F}_{pot}	potential forces
\vec{F}_r	rigidity force
G	linear gradient field constant strength
\vec{G}	linear gradient field
\hat{l}	probe long-axis direction of unit-length
m	voxel index of the cluster center position
M	number of velocity vectors in a local neighborhood
\vec{M}_0	net magnetization
$\vec{M}_{ }$	longitudinal magnetization
\vec{M}_{\perp}	transverse magnetization
\vec{M}_1	first order gradient moment
n	maximum number of output vertices
\mathbf{n}	node on surface mesh
\hat{n}	probe normal of unit-length
N	number of voxels in a cluster
\mathbf{p}	spatial position at the center of the bottom of the probe
\mathbf{p}_0	projection of \mathbf{p} on the view plane
\mathbf{p}_{δ_1}	the line \mathbf{p}_0 to \mathbf{p}
\mathbf{p}_0	projection of \mathbf{p} on the view plane
\mathbf{q}	spatial position at the center of the top of the probe
\mathbf{q}_0	projection of \mathbf{q} on the view plane
\mathbf{q}_{δ_2}	the line \mathbf{q}_0 to \mathbf{q}
P	number of cardiac phases
s	spatial resolution of tMOP volume
s_a	anatomy value
$s_{d_{2D}}$	simulated color Doppler value in 2D

$s_{d_{3D}}$	simulated color Doppler value in 3D
s_p	point sample on a line segment
s_t	simulated through-plane speed value
t	time
t_m	time of measurement
T	time second lobe of a bipolar gradient
\mathbf{T}	structure tensor
T_1	spin-lattice relaxation time
T_2	spin-spin relaxation time
T_2^*	spin-spin relaxation time, including magnetic field inhomogeneities
\vec{v}	velocity, with direction and magnitude
\vec{v}_p	velocity vector projection
\hat{v}	velocity direction with unit-length
\bar{v}	velocity vector average in a local neighborhood
\tilde{v}	velocity vector modeled using a linear expansion model
v_{enc}	velocity encoding speed
w	weighting parameter potential forces
\mathbf{x}	spatial position
$\bar{\mathbf{x}}$	spatial position average in a local neighborhood
\mathbf{x}_m	center position of a cluster
α	weighting parameter elasticity force
β	weighting parameter rigidity force
γ	gyromagnetic ratio
$\delta\vec{B}$	magnetic field inhomogeneities
Δd_{th}	dissimilarity threshold increment
ΔS	complex difference
$\Delta\phi$	phase difference
ε	error of cluster merging
ζ	angle between velocity direction and inspection direction in 2D
η	local spatial neighborhood
θ	angle between velocity direction and inspection direction in 3D
λ_i	sorted eigenvalues, with $i \in \{1, 2, 3\}$
π	mathematical constant: ratio of circle's circumference to its diameter
τ	integration time
ϕ	phase shift
ϕ_0	background phase
ψ	weighting parameter aspect ratio positional dissimilarity
ω	weighting parameter elliptical dissimilarity
ω_0	precession frequency in co-rotating frame
\times	multiplication
\otimes	outer product
\cdot	inner product

*“Statler: Boo!
Waldorf: Boooo!
S: That was the worst thing I've ever heard!
W: It was terrible!
S: Horrendous!
W: Well it wasn't that bad.
S: Oh, yeah?
W: Well, there were parts of it I liked!
S: Well, I liked a lot of it.
W: Yeah, it was GOOD actually.
S: It was great!
W: It was wonderful!
S: Yeah, bravo!
W: More!
S: More!
W: More!
S: More!”*

The Muppet Show - Waldorf & Statler

Dankwoord

Het lijkt zo kort geleden: het begin van vier uitdagende jaren onderzoek. Na een onstuimige zoektocht naar het juiste onderwerp viel de keuze op visualisatie van cardiovasculaire bloedstroming. Nu, ruim 170 miljoen hartslagen verder, na grofweg 12 miljoen liter bloed te hebben rondgepompt, is dit het resultaat: mijn proefschrift.

De weg naar dit proefschrift bestond uit een enerverende mix van bezigheden: programmeren, experimenteren, artikelen schrijven en beoordelen, conferenties en cursussen bezoeken en studenten begeleiden. Een promotie biedt de vrijheid om zelfstandig te werken aan je onderzoek: je bent manager van je eigen project. Het behaalde resultaat zou echter niet mogelijk zijn geweest zonder actieve samenwerkingen, intensieve begeleiding en de onvoorwaardelijke steun van familie en vrienden.

Allereerst wil ik mijn copromotor, dr. Anna Vilanova, bedanken voor haar toegewijde begeleiding. Beste Anna, met jouw vertrouwen, tomeloze inzet en kritische blik wist je mijn werk naar een hoger niveau te tillen. Onze talloze discussies, zowel inhoudelijk als persoonlijk, waren van onschatbare waarde voor het totstandkomen van dit proefschrift, en tevens een belangrijke inspiratiebron voor mij. De deur stond altijd voor me open, van je kantoor, maar ook bij je thuis als we in het holst van de nacht weer naar Londen vertrokken. Als er voor bedankt een superlatief zou bestaan, dan was dat hier van toepassing: *moltes gràcies* per tot!

Daarnaast wil ik uiteraard mijn beide promotoren, prof. Bart ter Haar Romeny en prof. Marcel Breeuwer, van harte bedanken voor de mogelijkheden die ze mij geboden hebben. Bart, bedankt voor je toewijding en grenzeloze enthousiasme, en uiteraard voor de kansen die je me hebt geboden om mezelf te ontwikkelen in een wetenschappelijke, maar ook warme en gemoedelijke omgeving. Marcel, bedankt voor het faciliteren van mijn project binnen Philips Healthcare en voor je kritische feedback, waarbij je altijd feilloos industrie, klinici en academici vertegenwoordigde.

Natuurlijk wil ik ook alle BMIA collega's bedanken voor de gezellige tijd. De sfeervolle Ardennen weekendjes, de ongedwongen filmavonden, barbecues en borrels, maar ook de prettige dagelijkse werksfeer maken deze jaren onvergetelijk. Een aantal mensen wil ik in het bijzonder vernoemen: Bram, -jonguh-, bedankt voor je humor, je fotografisch vernuft en de memorabele uitstapjes, zelfs al eindigden die soms in een moshpit. Ralph, promotie partner, bedankt voor de ongeremde discussies over visualisatie, uitlopend op vele andere thema's naarmate de avonden, en daarmee het promillage, vorderden. Marieke, bedankt voor je pogingen tot Tilburgs dialect, je bezielende aansturing in de keuken, en de gezelligheid, zowel bij Philips, als op de universiteit. Alessandro, amico mio, thanks for the many great cozy evenings. You're a wonderful host, narrator, and cook. Neda, thank you for your enthusiasm and energy, and for being the truly adorable structured scatterbrain you are. Vesna, one-day-younger lady, thank you for being a vivacious and bright personality. I relished

the homey dinners, and will never forget the fun we had at the many events. Paulo, I can't even begin to list all common interests we have. I really enjoyed sharing our views on music, movies, photography, gadgets, and more. Tim, bedankt voor je spontaniteit, soms heerlijk onbezonnen, maar altijd oprecht. Daarnaast natuurlijk ook Hans, Frans, Evgeniya, Laura, Luc, en alle andere (ex-)BMIA leden: bedankt voor de onvergetelijke tijden.

Verder wil ik graag Remco, alias dr. Don Neville, Andrei en Andrea van harte bedanken voor de fijne samenwerking in verschillende projecten, en natuurlijk ook voor de gezelligheid. Karin, bedankt voor je hulp in de voorbereidingen van mijn verdediging. Ook de studenten die ik heb mogen begeleiden ben ik uiteraard erg dankbaar voor hun bijdrage aan dit proefschrift. Huy, Sander, Andres, Yidan en Geert, thanks a lot! Ik heb jullie met veel plezier begeleid, en jullie harde werk heeft geleid tot schitterende resultaten.

Bovendien wil ik graag mijn collega's van Philips Healthcare bedanken voor hun waardevolle bijdrage aan mijn promotie. Prof. Frans Gerritsen, hartelijk dank voor de mogelijkheden die je hebt geboden om me te ontplooiën in de healthcare industrie. Hubrecht en Javier, bedankt voor de plezierige en stimulerende out-of-the-box brainstormsessies, en het delen van jullie rijke technische ervaring. Maurice, de redder in nood op het gebied van visualisatie en engineering: OrakelGL lijkt me een geschikte benaming. Daarnaast ook Fred, Kees, Stewart, Pierre, Rutger, Ursula, Jeroen, Raymond, en andere Philips collega's die hebben bijgedragen aan mijn promotie: bedankt.

In de afgelopen jaren heb ik met veel plezier deel uitgemaakt van de actieve en creatieve visualisatiegemeenschap. Allereerst prof. Jack van Wijk, hartelijk dank voor de inspirerende inzichten en de bijdrage aan mijn onderzoek. Huub, bedankt voor de wetenschappelijke discussies en de waardevolle inbreng in mijn master werk. A very special thanks goes to Eduard Gröller and his visualization crew. Meister, I shout thy name, and kneel before thee once more, in reverence and out of sincere gratitude. Thank you so much for your wisdom, your zealous reviews, your inspiring idiosyncratic ways, your true hospitality and the memorable social events. Also thanks to your 'disciples', Stefan, Peter, Martin, and many others. Furthermore, I'd like to thank the visualization group of prof. Bernhard Preim. Bernhard and Rocco, thank you for the valuable and exciting collaboration. Mathias, Steffen, Alexandra, thank you for the lively discussions and great times at conferences. Also I want to thank the Bergen crew, headed by prof. Helwig Hauser, for repeatedly including 'this Eindhoven hitchhiker' in the conference hotel sharing. Åsmund, Veronica, Johannes, Endre, Julius, and the rest, we've shared some great experiences all over the globe. Also, Markus, Johanna, and Philipp, many thanks for all the great conversations.

Additionally, I had the pleasure to work with skillful clinical partners, whose contribution has proven invaluable for my thesis. From King's College Londen, prof. Tobias Schaeffter, prof. Eike Nagel, dr. Rachel Clough, dr. Isra Valverde, and dr. Valentina Puntmann, thank you very much for all the help and the effort you have put into my project. Your guidance in the clinical domain, the extraordinary 4D blood-flow data

sets, and the constructive feedback have been of vital importance for the successful outcome of my promotion. This is really highly appreciated. Thanks again!

Tevens wil ik graag een aantal betrokken illustratoren hartelijk danken voor hun bijdrage. Van de Academie Beeldende Kunsten Maastricht wil ik graag Rogier Trompert en Jacques Spee bedanken. Jacques, bedankt voor de inspirerende discussies over het illustreren van medische beelden. Rogier, bedankt voor de schitterende afbeelding die ik heb mogen opnemen in mijn proefschrift. Tenslotte wil ik graag Koert Stavenuiter bedanken. Koert, ontzettend knap hoe je van mijn onbeholpen potloodkrabbels een perfecte strip hebt weten te creëren. Bedankt voor de geanimeerde gesprekken over striptekeningen, en uiteraard voor het schitterende stripje voor mijn proefschrift.

Naast het werk prijs ik mij gelukkig met mijn vrienden, waarmee ik de afgelopen jaren heb genoten van terrasjes, etentjes, bioscoopjes, feestjes, concerten en festivals. Sander, bedankt dat je, samen met Ellen, mijn paranimf wil zijn. We hebben in onze jarenlange hechte vriendschap al veel plezier gemaakt: trommelen, poolen, darten, karten, niks is te gek! Samen met Joyce bedankt voor de leuke tijden en jullie belangstelling. Bart, maatje, bedankt voor je briljante humor, en de geweldige tijden in allerlei werelden vol wonderen, van de Efteling tot de Everglades. Ad, bowl dude, bedankt voor onze vele gezellige en immer sympathieke gesprekken onder het genot van een klein biertje. Louis, Brigitte, Francis, Peter, Daan, Monica, bedankt voor de rockin' good times tijdens concerten en festivals, maar ook voor de gezellige middagen en avonden. Bonny, Edwin, Anneke, Jeroen, Kristel en Sander, bedankt voor alle gezelligheid. Mayke, bedankt voor de vele bosritjes, en voor alle goede zorgen voor Nellie. In de drukke tijden van mijn promotie kon ik altijd op je rekenen. Allemaal bedankt voor jullie steun en begrip in de 'druk druk druk' perioden. De laatste 'druk' hebben jullie nu in je handen!

Beste Lia en Jo, bedankt voor jullie gastvrijheid en de interesse in mij en mijn onderzoek. Sanne, Jelmer, Inge en Ralf bedankt voor jullie oprechte belangstelling en de gezelligheid. Lieve familieleden, Dimphie en Jack, bedankt voor jullie interesse en de hulp tijdens mijn 'intrede bij de Populieren'. Debby, Michael, Miranda en Jeffrey natuurlijk bedankt voor jullie belangstelling in mij en de ontwikkelingen tijdens mijn promotie. Oma, ut buukske is aaf! Bedankt voor uw betrokkenheid en interesse in al mijn bezigheden. Lieve pap en mam, bedankt voor jullie onvoorwaardelijke liefde. Zonder het (zelf)vertrouwen, de steun, en de mogelijkheden die jullie me hebben gegeven had ik dit nooit kunnen bereiken. Ik kan me geen betere ouders wensen: jullie staan altijd voor me klaar! Lieve Ellen, bedankt voor jouw liefde en luisterend oor, in het bijzonder in de tumultueuze laatste-loodjes-maanden. We stimuleren elkaar om het beste uit onszelf te halen, zowel in het werk als in ons privéleven. Ik hoop nog vele uitdagingen samen met jou te trotseren.

Curriculum vitae

Roy van Pelt was born on September 22, 1982, in Tilburg, the Netherlands. He received his secondary education at Theresia Lyceum Tilburg in 2000. Subsequently, he started a Bachelor in Information Technology at Fontys University of Applied Sciences in Eindhoven, which he completed in 2004. Thereupon, he started his studies in Computer Science and Engineering, specializing on Embedded Systems, at Eindhoven University of Technology, where he obtained his Master of Science degree in 2007.



His master's project was a collaboration between the Visualization group of the department of Mathematics and Computer Science, and the Biomedical Image Analysis group of the department of Biomedical Engineering at Eindhoven University of Technology. His project focussed on real-time illustrative visualization techniques, using the computational power of modern consumer graphics hardware.

Subsequently, he started his Ph.D. project at the Biomedical Image Analysis group of Eindhoven University of Technology, in strong collaboration with Philips Healthcare in Best. The project focussed on medical visualization and real-time interaction techniques for comprehensive exploratory analysis of unsteady volumetric blood-flow fields, measured by magnetic resonance.

During his Ph.D. project, he performed an externship at the 'Institut für Computergraphik und Algorithmen' of Vienna University of Technology, Vienna, Austria, where he worked on illustrative visualization techniques for blood-flow visualization. Furthermore, he performed an externship at King's College London and St. Thomas' Hospital, London, United Kingdom, where he investigated PC-MRI blood-flow acquisition and associated clinical applications.

Since April 2012, he is employed as a post-doctoral researcher at the Biomedical Image Analysis group of Eindhoven University of Technology.

List of publications

Journal papers

- 2012 **R.F.P. van Pelt**, S.S.A.M. Jacobs, B.M. ter Haar Romeny, A. Vilanova. "Visualization of 4D Blood-Flow Fields by Spatiotemporal Hierarchical Clustering", in *EG Computer Graphics Forum (CGF)*, in press.
- R.F.P. van Pelt**, H. Nguyen, B.M. ter Haar Romeny, A. Vilanova. "Automated Segmentation of Blood-Flow Regions in Large Thoracic Arteries using 3D-cine PC-MRI measurements", in *Journal of Computer Assisted Radiology and Surgery (JCARS)*, 7(2):217-224.
- 2011 **R.F.P. van Pelt**, J. Oliván Bescós, M. Breeuwer, R.E. Clough, M. E. Gröller, B.M. ter Haar Romeny, A. Vilanova. "Interactive Virtual Probing of 4D MRI Blood-Flow", in *Journal of IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 17(12):2153-2162.
- 2010 **R.F.P. van Pelt**, J. Oliván Bescós, M. Breeuwer, R.E. Clough, M. E. Gröller, B.M. ter Haar Romeny, A. Vilanova. "Exploration of 4D MRI Blood-Flow using Stylistic Visualization", in *Journal of IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 16(6):1339-1347.
- R.F.P. van Pelt**, A. Vilanova, H. van de Wetering. "Illustrative Volume Visualization using GPU-based Particle Systems", in *Journal of IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 16(4):571-582.

Conference papers

- 2012 **R.F.P. van Pelt**, J. Oliván Bescós, E. Nagel, V. Puntmann, B.M. ter Haar Romeny, and A. Vilanova “Time-resolved Volumetric Blood-Flow through Magnetic Resonance: A Doppler Ultrasound Perspective”, to be submitted.
- A. Fuster, **R.F.P. van Pelt**, R.H.J. Fick, G.G.H. Claassen, B.M. ter Haar Romeny, H.C. van Assen, L.M.J. Florack. “3D Saddle Point Detection and Applications in Cardiac Imaging”, in *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI)*.
- 2010 **R.F.P. van Pelt**, J. Oliván Bescós, M. Breeuwer, M. E. Gröller, B.M. ter Haar Romeny, A. Vilanova. “Exploring Hemodynamics by Raycasting 4D MRI Flow”, in *Proceedings of Advanced School for Computing and Imaging (ASCI)*.
- 2009 **R.F.P. van Pelt**, A. Vilanova, H. van de Wetering. “Illustrative Volume Visualization using GPU-based Particle Systems”, in *Proceedings of Advanced School for Computing and Imaging (ASCI)*.
- 2008 **R.F.P. van Pelt**, A. Vilanova, H. van de Wetering. “GPU-based Particle Systems for Illustrative Volume Rendering” (best paper), in *Proceedings of Volume and Point-Based Graphics (VPBG)*, pages 89-96.

PhD portfolio

International conferences

2012	Conference	IEEE/EG EuroVIS	Vienna	oral
2011	Conference	IEEE VISweek	Providence	oral
	Conference	IEEE/EG EuroVIS	Bergen	poster
	Conference	CARS	Berlin	poster
2010	Conference	IEEE VISweek	Salt-Lake City	oral
	Symposium	Visual Comp. in Medicine	Düsseldorf	attending
	Symposium	Visual Comp. in Medicine	Münster	attending
2009	Conference	ECR	Vienna	attending
2008	Conference	IEEE/EG EuroVIS	Eindhoven	volunteer
	Conference	ACM SIGGRAPH	Los Angeles	attending
	Conference	IEEE/EG VPBG	Los Angeles	oral

Other meetings

2011	Symposium	MISP	Rotterdam	oral
	Conference	Biomedical Engineering	Egmond aan Zee	poster
	Conference	BME research day	Eindhoven	poster
2010	Colloquium	BMIA	Eindhoven	organizing
	Conference	ASCI	Veldhoven	oral
	Conference	BME research day	Eindhoven	oral
2009	Conference	ASCI	Zeewolde	oral
	Conference	BME research day	Eindhoven	poster
2008	Conference	BME research day	Eindhoven	poster
	Meeting	COMRADE	Johnsbach	attending

Externships


2009	KCL and St. Thomas' Hospital	London
	Vienna University of Technology	Vienna

Courses

2010	Course	Supervising Master Students	PROOF
2009	Course	Knowledge-Driven Image Segmentation	ASCI
	Course	Advanced Pattern Recognition	ASCI
2008	Course	Measuring Features	ASCI
	Workshop	Science and Intellectual Property	PROOF
	Course	Writing Articles and Abstract	PROOF
	Course	Planning and Communication	PROOF
	Course	Intercultural Cooperation	PROOF
	Course	Tutor Training	TU/e
	Workshop	MeVisLab Introduction	MEVIS

Education

2011	Supervising	Master's project	S.S.A.M. Jacobs
	Supervising	Master's project	A.M. Arias
	Supervising	Internship	G.G.H. Claassen
	Supervising	Internship	Y. Guo
	Supervising	Basic image processing	
	Tutoring	Case study pathology	
2010	Supervising	Master's project	T.Q.H. Nguyen
	Supervising	Basic image processing	
2009	Supervising	Basic image processing	
2008	Supervising	Skill lab BioMIM	
	Tutoring	Case study pathology	

The image features a circular flow visualization composed of multiple parallel paths of white lines with colored oval markers. The markers transition in color from red at the bottom to yellow at the top. A central white rectangular box with a red border contains text. The background is a light gray with a subtle circular pattern.

Modern magnetic resonance imaging techniques enable acquisition of time-resolved volumetric blood-flow velocity data. With these data, physicians aim for newfound insight into the intricate blood-flow dynamics. This conceivably leads to improved diagnosis and prognosis of cardiovascular diseases, as well as a better assessment of treatments and risks. We facilitate the time-consuming and challenging process of visual analysis of these unsteady, multi-dimensional and multi-valued data by comprehensive exploratory visualization techniques, tailored to communicate blood flow in the heart and the thoracic arteries. We introduce abstraction techniques to reduce the abundance of information contained in the data. Interactive exploration is enabled by probing tools, selecting regions-of-interest that serve as a basis for our real-time illustrative visualizations. Based on evaluation studies with the involved physicians, we believe that real-time visual exploration of blood-flow data facilitates the qualitative analysis.