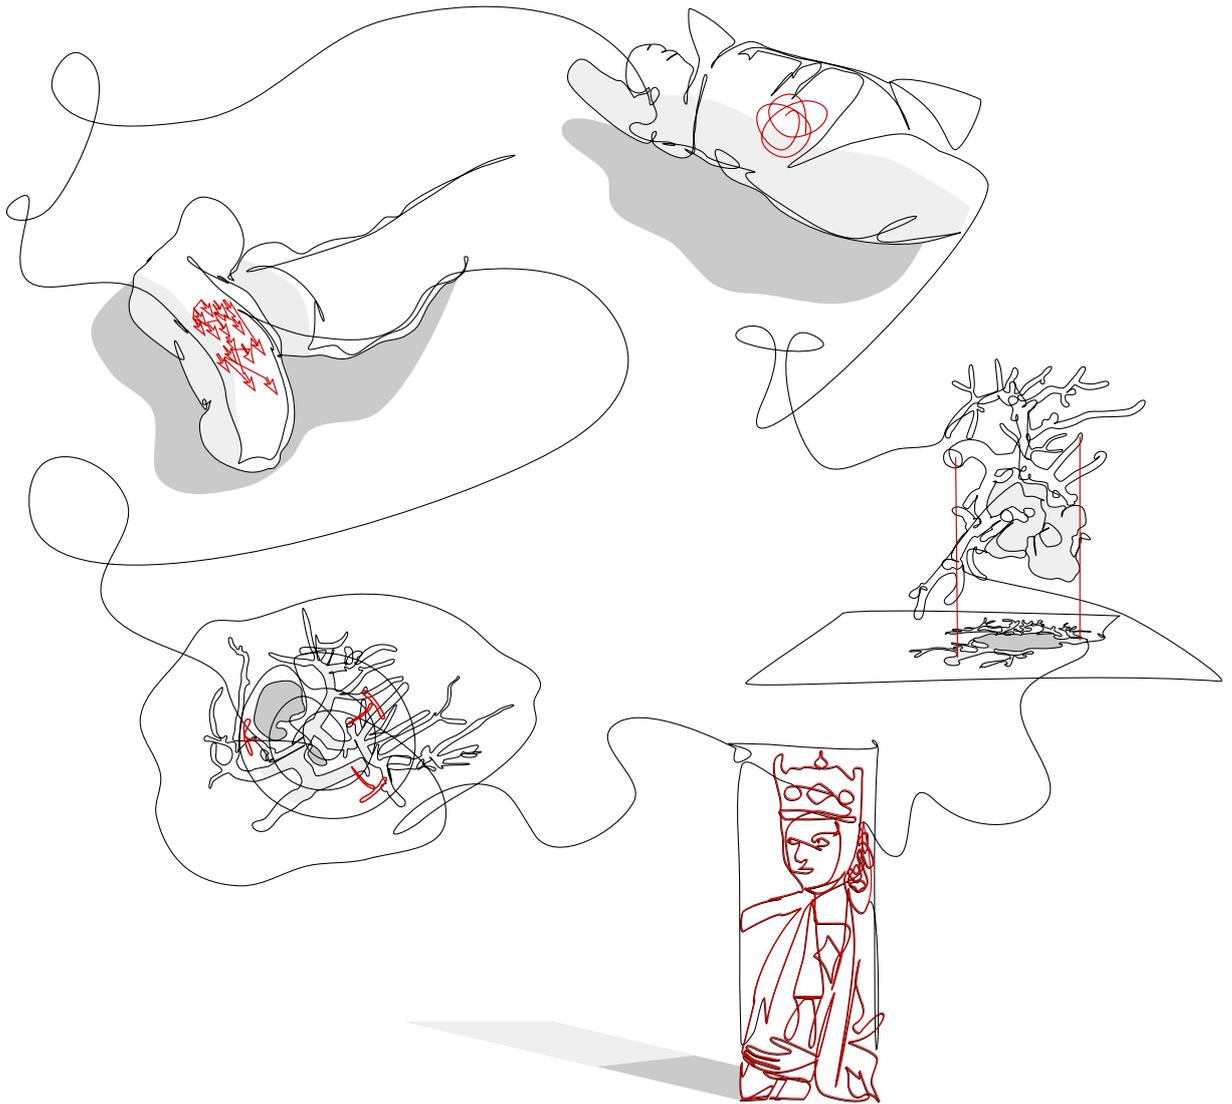


KAI RENÉ HARTMUT LAWONN

VISUAL ENHANCEMENT OF FOCUS STRUCTURES
IN SELECTED APPLICATIONS





VISUAL ENHANCEMENT OF FOCUS
STRUCTURES IN SELECTED
APPLICATIONS

HABILITATIONSSCHRIFT

zur Erlangung der Venia Legendi für
Computervisualistik

angenommen durch die Fakultät für Informatik
der Otto-von-Guericke-Universität Magdeburg

von DR.-ING. KAI RENÉ HARTMUT LAWONN

geb. am 04.10.1985, in Berlin

Gutachter

Prof. Dr. Bernhard Preim
Prof. Dr. Eduard Gröller
Prof. Dr. Anders Ynnerman

Magdeburg, den 02. November 2017

Für Editha

ABSTRACT

This thesis presents selected applications with the aim of visually enhancing focus structures. It covers four different applications with a clear problem statement. The solution is anything but simple. Thus, four applications will be presented and analyzed, complemented by their respective solutions. Novel visualization techniques to enhance focus structures will be presented. Furthermore, two overviews of the current state of the art are described, which cover the fields of illustrative visualization techniques and multimodal data visualization.

ZUSAMMENFASSUNG

Die Habilitation umfasst verschiedene Anwendungsgebiete mit dem Ziel der visuellen Verdeutlichung bestimmter Fokusstrukturen. Es werden vier verschiedene Anwendungen vorgestellt. Diese Anwendungen haben eine einfache Problemstellung aber die Lösung ist mitunter sehr kompliziert. Zur Detektion als auch zur visuellen Hervorhebung werden Techniken in den verschiedenen Anwendungen präsentiert. Dazu werden aktuelle Überblicke über die illustrative Visualisierung sowie der Visualisierung von multimodalen Daten vorgestellt.

*The proof of a high education is
the ability to speak about
complex matters as simply as possible.*

— Ralph Waldo Emerson

DANKSAGUNGEN

BERNHARD mein größter Dank gilt vor allem dir und deiner großartigen Unterstützung; für dein offenes Ohr und für die vielen Diskussionen über weitere Ideen zu neuen Veröffentlichungen. Ich freue mich auf weitere Jahre mit erfolgreichen Kooperationen!

An dieser Stelle bedanke ich mich bei allen Kollegen, mit denen ich das Vergnügen hatte, gemeinsam zu publizieren. **Monique:** wie lange saßen wir wohl abends immer im Büro und hatten nur die Deadline im Auge? Und jedes Mal ermutigten wir uns gegenseitig, weiter zu machen. Es macht so viel Spaß mit dir zusammenzuarbeiten. Wir sind ein sehr gutes Team und ich freue mich auf viele neue Projekte in der Zukunft! **Noeska:** thanks for the great time in Delft and thanks for such nice collaborations. I hope that we always find the time to skype and to talk about further projects. Team MedVis! **Paul:** das waren schon sehr lustige und produktive Zeiten am ISG. Ich hoffe, wir treffen uns bald mal wieder. **Alexandra, Maria:** vielen Dank für die Diskussionen über die Durchführung korrekter Studien und für die vielen Auswertungen. **Nils:** ich bin froh, dass wir zusammenarbeiten. Es macht unglaublich viel Freude und ich lerne jedes Mal aufs Neue dazu! Auf viele weitere Paper. Für die schönen VIS-Paper möchte ich mich auch bei **Sylvia** bedanken. Die erste Visualisierung der Wanddicke – das war ein tolles Paper! **Patrick:** das erste gezeichnete Gefäß in 2D und damals dachten wir schon über 3D nach. Und nun hast du es einfach in 3D übertragen, super umgesetzt! **Christian:** wie oft saßen wir zusammen und haben über den DFG-Antrag beratschlagt und diskutiert? Ich wünsche mir, dass wir auch in Zukunft über weitere spannende Projekte nachdenken können und das nicht nur im Büro, vielleicht auch mal in einer ganz anderen Atmosphäre, einer Bar vielleicht? **Anna:** I am very grateful that you gave me the opportunity to join your group in Delft. I had a lot of fun and I enjoyed my stay very much. Thanks also for your advice and help. **Erik:** vielen Dank, dass ich mit dir das Projekt der Putzritzungen bearbeiten durfte. Es hat sehr viel Spaß gemacht und ich freue mich schon auf die weitere Zusammenarbeit. **Klaus:** Ich danke dir für die umfassenden Diskussionen und das Brainstormen. In dieser Zeit habe ich viel gelernt. **Niels:** is it HW-o'clock? Thanks for having a great time in Delft. First week and we produced an awesome movie. Dann möchte ich

noch meinen Dank an **Oliver, Samuel** und **Uta** für die vielen produktiven Diskussionen richten. Für die oftmals lustigen Stunden am ISG geht mein Dank an **Tobias, Uli, Steffen, Nico, Ben** und (der zweite) **Ben**. **Tobias**: Vielen Dank, dass du mir die Möglichkeit gegeben hast, nach Paris zu kommen. Auch wenn wir uns nicht immer einig waren, wenn wir ein Paper zusammengeschrieben haben, muss ich doch zugeben, dass ich sehr vieles gelernt habe. Ich würde sogar so weit gehen und sagen, dass ich durch dich Paper mittlerweile anders schreibe oder einen anderen Blick habe. Vielen Dank für die tolle Zusammenarbeit.

Last but not least bedanke ich mich natürlich bei meiner Familie. **Susanne, Michael, Melissa, Hartmut, Christa** für euer Vertrauen und eure offenen Ohren. Mein letzter großer Dank gehört **Daniela** und **Editha**. Ich bin euch sehr verbunden für eure Unterstützung und die Fähigkeit, mich zum Lachen zu bringen und für eure Zeit, mir zuzuhören und Ratschläge zu geben :).

Kai Lawonn, 02. November 2017

CONTENTS

I	INTRODUCTION	1
1	INTRODUCTION	5
II	PRELIMINARIES	7
2	PRELIMINARIES	11
2.1	Differential Geometry Background	11
2.1.1	Discrete Differential Geometry	15
2.1.2	Isolines on Surface Meshes	18
III	VISUALIZATION CONCEPTS	21
3	SURFACE-BASED ILLUSTRATIVE RENDERING FOR VISUALIZATION	25
3.1	Concept of Illustrative Visualization	25
3.2	General Requirements	27
3.2.1	Smoothing	27
3.2.2	Robustness	28
3.2.3	Refinement	28
3.2.4	Frame Coherence	28
3.3	Low-level Illustrative Visualization Techniques	29
3.3.1	Silhouettes and Contours	29
3.3.2	Feature Lines	31
3.3.3	Hatching	40
3.3.4	Stippling	45
3.3.5	Illustrative Shading	47
3.4	Evaluation	52
3.5	Applications	56
3.6	Discussion & Future Work	58
3.7	Concluding Remarks	60
4	MULTIMODAL MEDICAL DATA VISUALIZATION	65
4.1	Medical Background	68
4.1.1	Computed Tomography (CT)	68
4.1.2	Magnetic Resonance Imaging (MRI)	68
4.1.3	Ultrasound	70
4.1.4	Modalities from Nuclear Medicine	70
4.1.5	Hybrid Scanners	71
4.2	Clinical Workflow and Requirements	73
4.2.1	Clinical Workflow	73
4.2.2	Requirement Analysis	74
4.3	Visualization Techniques	76
4.3.1	Basic Representation	77
4.3.2	Smart Visibility	78
4.3.3	Focus-and-context visualization	82

4.3.4	Summary	85
4.4	Multimodal Rendering and Interaction Techniques . .	86
4.4.1	Rendering Techniques	86
4.4.2	Interaction Techniques	90
4.5	Applications	92
4.5.1	Medical Research	92
4.5.2	Diagnosis	94
4.5.3	Treatment Planning	96
4.5.4	Discussion	104
4.6	Conclusion and future challenges	106
4.7	Summary	108
IV	SELECTED APPLICATIONS	111
5	ILLUSTRATIVE MULTI-VOLUME RENDERING FOR PET/CT SCANS	115
5.1	Medical Background	116
5.1.1	Requirement analysis	116
5.2	Method	117
5.2.1	Context Visualization	118
5.2.2	Focus Visualization	119
5.2.3	Contextual Visualization of the Surrounding Fo- cus	120
5.3	Implementation	120
5.4	Results	121
5.5	Evaluation	123
5.5.1	First evaluation	125
5.5.2	Final evaluation	126
5.6	Discussion	128
5.7	Conclusion	129
6	OCCLUSION-FREE BLOOD FLOW VISUALIZATION	133
6.1	Related Work	134
6.2	Medical Background	136
6.3	Requirement Analysis	138
6.4	Basic Visualization Algorithms	140
6.4.1	Cutaway Technique	140
6.4.2	Jump Flooding Algorithm	141
6.4.3	Pathline Animation	141
6.4.4	Centerline Extraction for Slice View Analysis .	142
6.5	Visualization Design for Integrated Depiction of Vessel Morphology and Blood Flow	142
6.5.1	Pathline Visualization	142
6.5.2	Jump Flooding Details	144
6.5.3	Cutaway Illustration	145
6.5.4	Slice View	145
6.5.5	Vessel Illustration	146
6.5.6	Depth Attenuation	147

6.5.7	Framework	147
6.6	Comparison with Related Approaches	149
6.6.1	Near-Wall Visualization	149
6.6.2	Wall Thickness Visualization	151
6.6.3	Summary	151
6.7	Qualitative Evaluation	152
6.7.1	Questionnaire	152
6.7.2	Results and Discussion	154
6.8	Expert Feedback	155
6.9	Conclusion and Future Work	157
7	VISUALIZATION AND EXTRACTION OF CARVINGS FOR HER- ITAGE CONSERVATION	161
7.1	Related Work	162
7.2	Data Acquisition	165
7.3	Requirement Analysis	166
7.4	Extraction and Generation of Carvings	167
7.4.1	Carve Extraction and Ranking	167
7.4.2	Carve Generation	171
7.5	Evaluation with Domain Experts	174
7.6	Results	176
7.7	Comparison to Alternative Methods	178
7.8	Discussion & Future Work	181
8	IMPROVED DEPTH-PERCEPTION WITH SUPPORTING LINES	185
8.1	Related Work	185
8.1.1	Improving Spatial Perception for Medical Visu- alization	186
8.2	Medical Background	188
8.3	Method	189
8.3.1	Hatching - ConFIS	190
8.3.2	Distance Illustrations	190
8.3.3	Depth Cues	191
8.4	Evaluation	194
8.5	Discussion and Results	195
8.6	Results	199
8.7	Discussion	200
9	IMPROVED DEPTH-PERCEPTION WITH SUPPORTING AN- CHORS	207
9.1	Method	207
9.1.1	From Voxels to 3D Anatomical Models	207
9.1.2	Visualization Approach	208
9.1.3	Candidate Selection	208
9.1.4	Optimal view-dependent Candidate Picking	209
9.1.5	Hatching	213
9.1.6	Real-time Hatching with Isolines	215
9.2	Implementation	215
9.3	Quantitative Evaluation	217

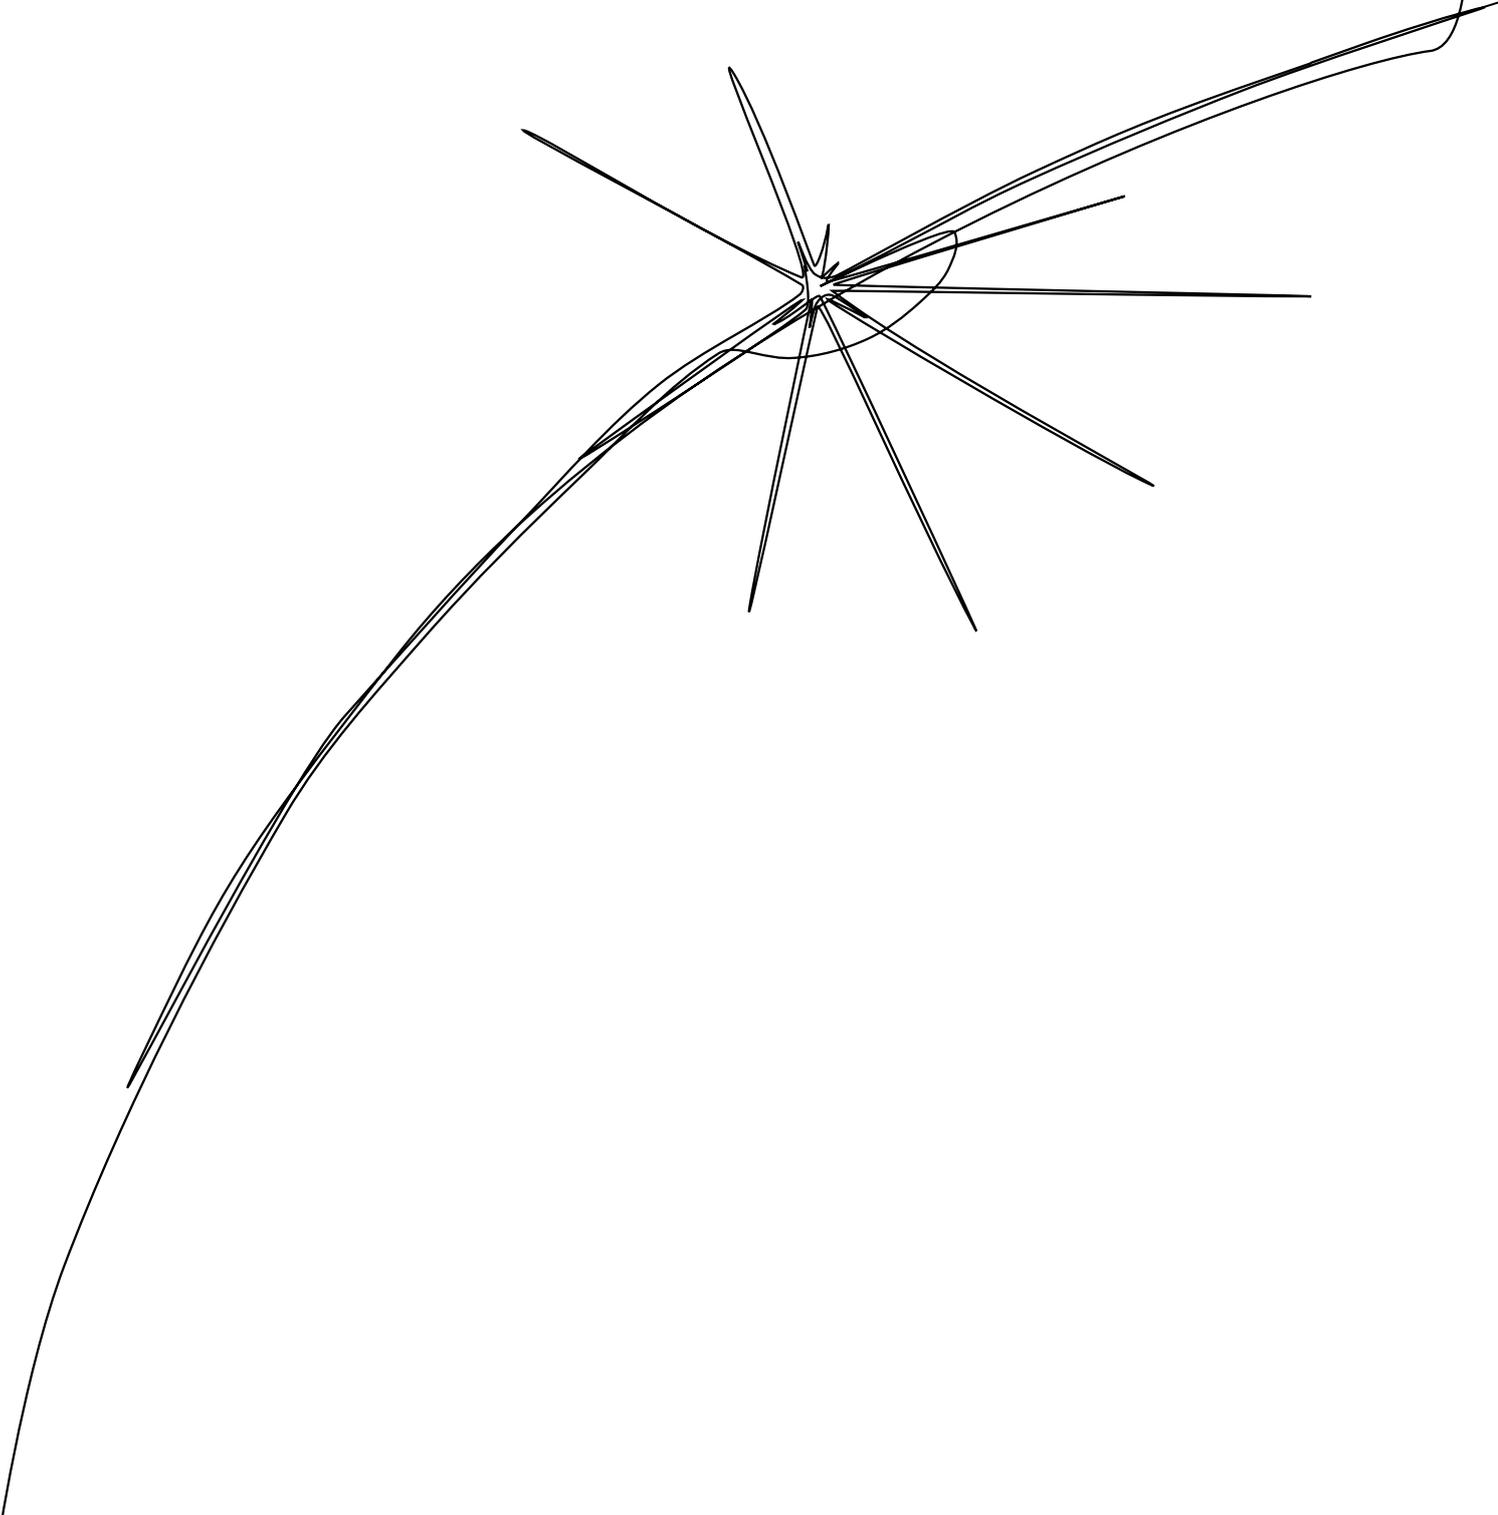
9.4	Results	218
9.5	Discussion and Future Work	220
V	CONCLUSION	225
10	CONCLUSION	229
	BIBLIOGRAPHY	231
	LIST OF PUBLICATIONS	250

Part I

INTRODUCTION

1

INTRODUCTION



INTRODUCTION

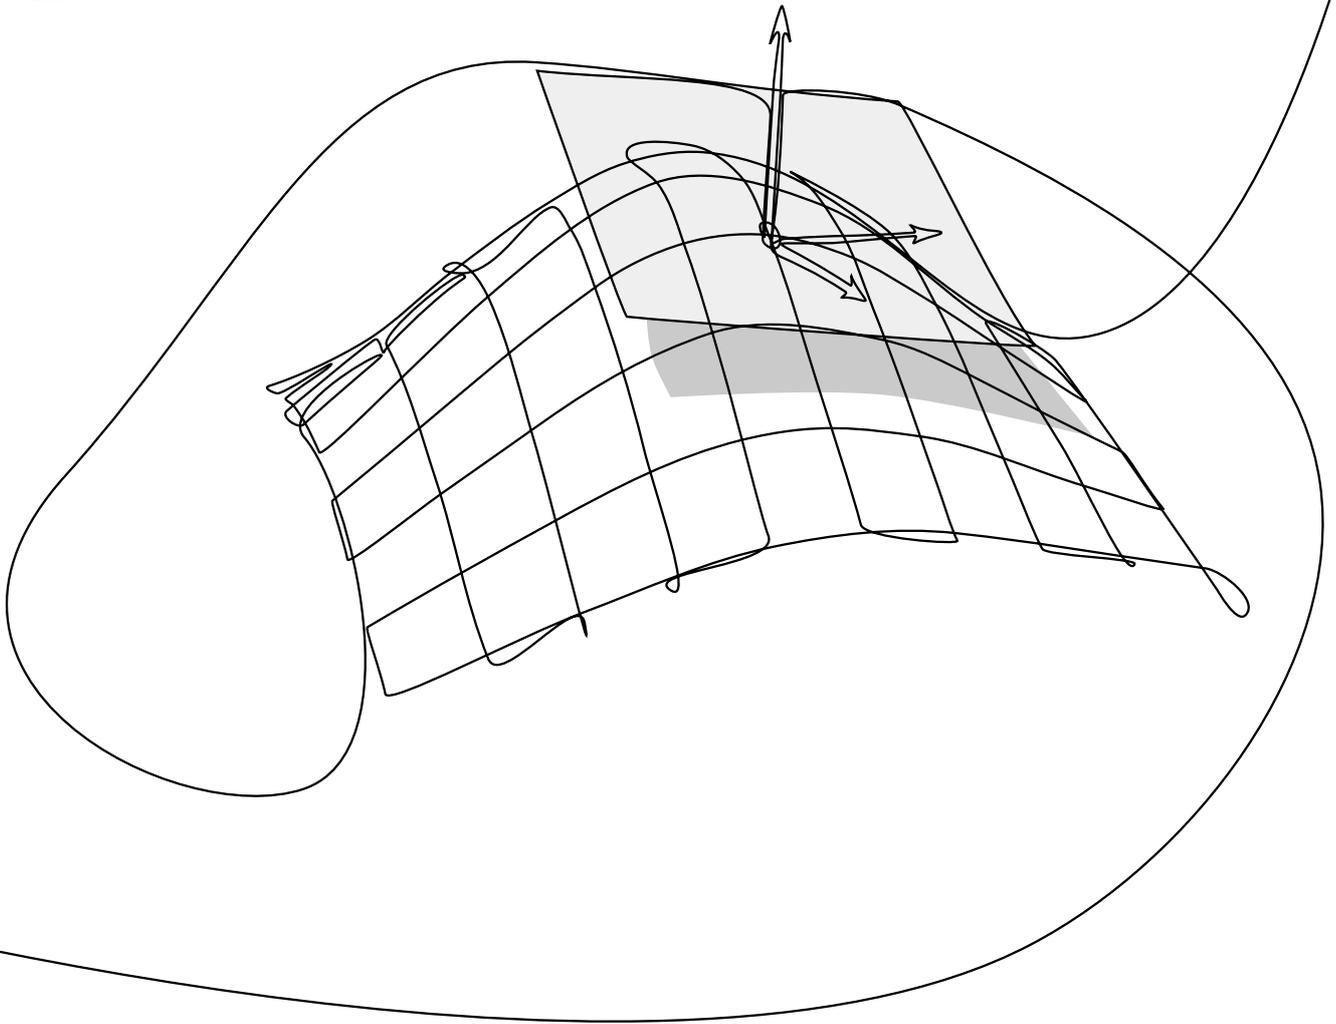
VISUALIZATION techniques provide us with a multitude of ways to answer various questions that were raised by the exploration and analysis of datasets. The objectives vary from the simple depiction of the data with the goal of showing phenomena and connections to the visualizations that were derived from the data. A very common goal, with a problem that can be simply formulated, is to display certain features that are already part of the data. Features that should be visualized are referred to as *focus structures*. In this regard, the formulation of the problem to detect or enhance focus structures may be simple, but the solution may be more difficult. In other words, if we have a dataset and the question is to display interesting features, the algorithm strongly varies from the underlying data and the problem itself. Even though it may be easy to visually perceive the interesting structures, it may be tough to algorithmically detect those focus structures. In this habilitation, I focus firstly on the question of what could be done to detect focus structures, and secondly how these features could be visually enhanced. For this, I examined focus structures in selected applications with the goal of emphasizing these structures. This led to various application fields with problems that may be easily formulated, but have complex solutions. First, this thesis provides the reader with a mathematical background, see Part ii. It serves as a mathematical prerequisite to the fundamentals of differential geometry and discrete differential geometry. Part iii provides a survey about various rendering techniques. This survey covers illustrative visualization methods and provides an overview of line drawings and shading concepts. In addition to the state of the art, an overview of multimodal data visualization will be provided. This covers the combination of different visualization techniques to display objects that were acquired by multiple acquisition modalities. Initial examples of the visualization of focus structures will be presented. The penultimate Part iv presents various applications with the goal of visually enhancing focus structures. This part introduces applications and presents solutions that detect and enhance these features. Finally, Part v holds the conclusion of this thesis and discusses future work.

Part II

PRELIMINARIES

2

PRELIMINARIES



This chapter is partly based on:

Kai Lawonn, Ivan Viola, Bernhard Preim, and Tobias Isenberg.

“A Survey of Surface-Based Illustrative Rendering for Visualization.”

In: Computer Graphics Forum (2017), submitted.

THIS chapter provides the preliminaries of differential geometry¹, which is necessary for the illustrative rendering of surface geometries. This chapter gives the necessary terms to implement the presented methods. First, the basic terms are introduced and specific mathematical concepts like curvature and directional derivatives are explained. Furthermore, the Laplace–Beltrami operator is introduced, which is also crucial for illustrative rendering. This chapter is inspired by differential geometry books [Car76; Car92; Küh06].

2.1 DIFFERENTIAL GEOMETRY BACKGROUND

2.1.0.1 Basic Prerequisites

A parametric surface patch is a map $f: I \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ such that the derivatives $\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}$ are linearly independent at each point $\mathbf{p} \in f(I)$. Formally, the tangent space at $\mathbf{u} \in I$ is defined as

$$\mathcal{T}_{\mathbf{u}}f := \left\{ \alpha \cdot \frac{\partial f}{\partial x_1}(\mathbf{u}) + \beta \cdot \frac{\partial f}{\partial x_2}(\mathbf{u}) \mid \alpha, \beta \in \mathbb{R} \right\}. \quad (1)$$

As f is defined to be an immersion it follows that the tangent space is two-dimensional. The tangent space is a subspace of \mathbb{R}^3 and can therefore be represented by a normal vector. Thus, a normalized normal vector \mathbf{n} is defined as a function $\mathbf{n}: I \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$:

$$\mathbf{n}(\mathbf{u}) = \frac{\frac{\partial f}{\partial x_1}(\mathbf{u}) \times \frac{\partial f}{\partial x_2}(\mathbf{u})}{\left\| \frac{\partial f}{\partial x_1}(\mathbf{u}) \times \frac{\partial f}{\partial x_2}(\mathbf{u}) \right\|}, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm. The normal vector has the following property:

$$\langle \mathbf{v}, \mathbf{n}(\mathbf{u}) \rangle = 0, \quad \text{for all } \mathbf{v} \in \mathcal{T}_{\mathbf{u}}f. \quad (3)$$

Here, $\langle \cdot, \cdot \rangle$ denotes the Euclidean dot product: $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i$ with $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. The map \mathbf{n} maps every point on the parametric surface to a point on the unit sphere S^2 .

2.1.0.2 Shape Operator

To analyze the change of the normal in certain directions, the derivative of the Gauss map, which is only defined if f is at least twice

¹ A similar chapter was presented in my dissertation “Illustrative Visualization of Medical Data Sets”. The current chapter was newly written.

differentiable needs to be determined. Thus, for now parametric surfaces should be at least twice differentiable. The change of the normal \mathbf{n} is provided by the Jacobian matrix $J\mathbf{n}$:

$$J_{\mathbf{u}}\mathbf{n} = \begin{pmatrix} \frac{\partial \mathbf{n}}{\partial x_1}(\mathbf{u}) & \frac{\partial \mathbf{n}}{\partial x_2}(\mathbf{u}) \end{pmatrix}. \quad (4)$$

The Jacobian can be considered as a linear map that maps every point $\mathbf{w} \in I$ to the tangent space $(J_{\mathbf{u}}\mathbf{n})(\mathbf{w}) \in \mathcal{T}_{\mathbf{u}}f$. The linear map $(J_{\mathbf{u}}\mathbf{n})(\mathbf{w})$ is just a matrix multiplication of the Jacobian with the point $\mathbf{w} = (w_1, w_2)$, i.e., $(J_{\mathbf{u}}\mathbf{n})(\mathbf{w}) = w_1 \cdot \frac{\partial \mathbf{n}}{\partial x_1}(\mathbf{u}) + w_2 \cdot \frac{\partial \mathbf{n}}{\partial x_2}(\mathbf{u})$. To show that $(J_{\mathbf{u}}\mathbf{n})(\mathbf{w}) \in \mathcal{T}_{\mathbf{u}}f$ is true, it is shown that $\langle (J_{\mathbf{u}}\mathbf{n})(\mathbf{w}), \mathbf{n} \rangle = 0$ is valid. The normal is defined such that $\langle \mathbf{n}, \mathbf{n} \rangle = 1$, deriving this with x_i gives:

$$\begin{aligned} \frac{\partial}{\partial x_i} \langle \mathbf{n}, \mathbf{n} \rangle &= \frac{\partial}{\partial x_i} 1 \\ 2 \left\langle \frac{\partial \mathbf{n}}{\partial x_i}, \mathbf{n} \right\rangle &= 0. \end{aligned}$$

Therefore, the derivative $\frac{\partial}{\partial x_i} \mathbf{n}$ is perpendicular to \mathbf{n} and therefore an element of the tangent space.

Furthermore, the Jacobian of f is defined, which can also be interpreted as a linear map from I to the tangent space. As f is an immersion, it yields that $J_{\mathbf{u}}f$ has a full rank and, thus, is a bijective map. Therefore, the inverse of $J_{\mathbf{u}}f$ is well-defined. Then, a linear map is defined

$$S = -(Jf)\mathbf{n}(Jf)^{-1} \quad (5)$$

which is called the *shape operator* (also *Weingarten map* or *second fundamental tensor*). The shape operator maps the derivative of f to the derivative of \mathbf{n} , precisely:

$$S\left(\frac{\partial}{\partial x_i} f\right) = \frac{\partial}{\partial x_i} \mathbf{n}. \quad (6)$$

Importantly, the shape operator is a self-adjoint linear map, i.e., $\langle S\mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{v}, S\mathbf{w} \rangle$. To show this, $\langle \mathbf{n}, \frac{\partial f}{\partial x_i} \rangle = 0$ is derived with respect to x_i :

$$\left\langle \frac{\partial \mathbf{n}}{\partial x_j}, \frac{\partial f}{\partial x_i} \right\rangle + \left\langle \mathbf{n}, \frac{\partial^2 f}{\partial x_i \partial x_j} \right\rangle = 0.$$

This yields the equation:

$$\left\langle \frac{\partial \mathbf{n}}{\partial x_j}, \frac{\partial f}{\partial x_i} \right\rangle = -\left\langle \mathbf{n}, \frac{\partial^2 f}{\partial x_i \partial x_j} \right\rangle. \quad (7)$$

To prove that S is a self-adjoint operator, the basis $\frac{\partial f}{\partial x_i}$ is used:

$$\begin{aligned} \left\langle S\left(\frac{\partial f}{\partial x_i}\right), \frac{\partial f}{\partial x_j} \right\rangle &= \left\langle -\frac{\partial \mathbf{n}}{\partial x_i}, \frac{\partial f}{\partial x_j} \right\rangle = \left\langle \mathbf{n}, \frac{\partial^2 f}{\partial x_i \partial x_j} \right\rangle \\ &= \left\langle -\frac{\partial \mathbf{n}}{\partial x_j}, \frac{\partial f}{\partial x_i} \right\rangle = \left\langle S\left(\frac{\partial f}{\partial x_j}\right), \frac{\partial f}{\partial x_i} \right\rangle. \end{aligned}$$

Interestingly, self-adjoint operators have real eigenvalues and if the eigenvalues $\lambda_1 \neq \lambda_2$ differ, then the eigenvectors $\mathbf{v}_1 \perp \mathbf{v}_2$ are perpendicular:

$$\lambda_1 \langle \mathbf{v}_1, \mathbf{v}_2 \rangle = \langle S\mathbf{v}_1, \mathbf{v}_2 \rangle = \langle \mathbf{v}_1, S\mathbf{v}_2 \rangle = \lambda_2 \langle \mathbf{v}_1, \mathbf{v}_2 \rangle.$$

The equation is only true if $\lambda_1 \neq \lambda_2$ holds. Note that the shape operator can directly operate on the 3D space with a three-dimensional vector in the tangent space, as well as the 2D space with the coefficients of the basis.

2.1.0.3 Curvature

The curvature measures how strongly the surface is bend at a certain point in a specific direction. Let $\mathbf{w} \in I$ be a direction and $S_{\mathbf{u}}(\mathbf{w}) \in \mathcal{T}_{\mathbf{u}}f$, then $\langle S_{\mathbf{u}}(\mathbf{w}), \mathbf{w} \rangle$ defines a quadratic form and the result is called the curvature. S has two real eigenvalues κ_1, κ_2 with their corresponding eigenvectors $\mathbf{k}_1, \mathbf{k}_2$. Let $\kappa_1 \geq \kappa_2$ then the eigenvector \mathbf{k}_1 maximizes the quadratic form $\langle S_{\mathbf{u}}(\mathbf{w}), \mathbf{w} \rangle$. The eigenvectors $\mathbf{k}_1, \mathbf{k}_2$ are called the principal curvature directions. They yield the direction where the surface is bend the strongest/least. Recall that if the curvatures differ $\kappa_1 \neq \kappa_2$ the principal curvature direction are mutual perpendicular.

2.1.0.4 Gradient and Directional Derivative

Previously, it was shown that the direction of the derivative has an influence on the outcome. Therefore, to analyze local variations in certain directions, the directional derivative is explained. The directional derivative is applied to a scalar field $\varphi: I \rightarrow \mathbb{R}$, which is defined on the set I , but can also be defined on the surface. The question arises how does the scalar field change if it is derived in direction $\mathbf{v} \in \mathcal{T}_{\mathbf{u}}f$ at the point $f(\mathbf{u}) = \mathbf{p}$. If φ is differentiable at \mathbf{p} , the directional derivative can be written as:

$$D_{\mathbf{v}}\varphi(\mathbf{p}) = \langle \nabla\varphi(\mathbf{p}), \mathbf{v} \rangle,$$

where ∇ denotes the gradient. To determine the directional derivative, the gradient on the surface needs to be explained. The gradient of a scalar field at the point $f(\mathbf{u})$ is a vector, which lies in the tangent space $\mathcal{T}_{\mathbf{u}}f$ and points in direction of the greatest slope. Again, the gradient can be written as a 3D vector in the tangent space or as a 2D vector, which represents the coefficients of the tangent space basis given by $\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}$. The gradient of φ in the domain I can be easily determined by the partial derivatives with respect to the coordinates:

$\nabla_I\varphi = \begin{pmatrix} \frac{\partial \varphi}{\partial x_1} \\ \frac{\partial \varphi}{\partial x_2} \end{pmatrix}$. To apply the gradient as a 3D vector in the tangent space, it needs to be restricted to the tangent space by claiming:

$$\nabla_{\mathcal{T}_{\mathbf{u}}f}\varphi = \alpha \cdot \frac{\partial f}{\partial x_1} + \beta \cdot \frac{\partial f}{\partial x_2}. \tag{8}$$

Then, the dot product of the gradient with the basis vectors in the space I must be the same as the dot product of the gradient with the basis vectors in the tangent space:

$$\langle \nabla_I \varphi, \mathbf{e}_i \rangle = \langle \nabla_{\mathcal{T}_u f} \varphi, \frac{\partial f}{\partial x_i} \rangle \quad (9)$$

$$\frac{\partial \varphi}{\partial x_i} = \langle \nabla_{\mathcal{T}_u f} \varphi, \frac{\partial f}{\partial x_i} \rangle. \quad (10)$$

This yields an equation system, which can be used to determine the coefficients α, β :

$$\begin{pmatrix} \frac{\partial \varphi}{\partial x_1} \\ \frac{\partial \varphi}{\partial x_2} \end{pmatrix} = \underbrace{\begin{pmatrix} \langle \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_1} \rangle & \langle \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2} \rangle \\ \langle \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2} \rangle & \langle \frac{\partial f}{\partial x_2}, \frac{\partial f}{\partial x_2} \rangle \end{pmatrix}}_{g:=} \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$$

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = g^{-1} \begin{pmatrix} \frac{\partial \varphi}{\partial x_1} \\ \frac{\partial \varphi}{\partial x_2} \end{pmatrix}.$$

The matrix g is called the *metric tensor* and describes the metric distortion, e.g., stretching, shearing, which is induced by the map f and g^{-1} is denoted as the inverse matrix of g .

2.1.0.5 Laplace–Beltrami Operator

The Laplace–Beltrami operator is defined as the *divergence* of the *gradient*. The gradient of a scalar field $\varphi: I \rightarrow \mathbb{R}$ can be written as:

$$\nabla \varphi = \sum_{i,j=1}^n \left(g^{ij} \frac{\partial \varphi}{\partial x_j} \right) \frac{\partial}{\partial x_i}, \quad (11)$$

where g^{ij} denotes the i, j -th matrix entry from g^{-1} and $\frac{\partial}{\partial x_i}$ means the basis in the n -dimensional tangent space.

The divergence is a measure of how much more flow leaves a specific region than flow enters. Formally, the divergence of a vector field $\Phi: I \rightarrow \mathbb{R}^n$ with a given metric tensor g is given by:

$$\operatorname{div} \Phi = \frac{1}{\sqrt{|g|}} \sum_{i=1}^n \frac{\partial}{\partial x_i} \left(\sqrt{|g|} \Phi_i \right), \quad (12)$$

with $|g| = |\det g|$.

With the definition of the gradient and the divergence, the Laplace–Beltrami operator $\Delta \varphi$ of a scalar field $\varphi: I \rightarrow \mathbb{R}$ on a surface has the following form:

$$\Delta \varphi = \operatorname{div} \nabla \varphi = \frac{1}{\sqrt{|g|}} \sum_{i,j=1}^n \frac{\partial}{\partial x_i} \left(\sqrt{|g|} g^{ij} \frac{\partial \varphi}{\partial x_j} \right). \quad (13)$$

2.1.1 Discrete Differential Geometry

In Section 2.1 some background information on differential geometry were given. In this section, an introduction on how the operators and measures from the continuous differential geometry can be adapted to polygonal meshes is presented. The following notation is used in the remainder of this thesis. Let $M \subset \mathbb{R}^3$ be a triangulated surface mesh with vertices $i \in V$ and its associated positions $\mathbf{p}_i \in \mathbb{R}^3$, edges $E = \{(i, j) \mid i, j \in V\}$, and triangles $T = \{(i, j, k) \mid (i, j), (j, k), (k, i) \in E\}$. The normalized normal vector at vertex i is written as \mathbf{n}_i . If nothing else is mentioned, normal vectors at vertices are used. Furthermore, $\mathcal{N}(i)$ denotes the neighbors of i such that for every $j \in \mathcal{N}(i)$, $(i, j) \in E$ holds. Furthermore, for a triangle, this notation is used: given a triangle $t = (i, j, k)$ with vertices $\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k$, and the edges are defined as $\mathbf{e}_{ij} = \mathbf{p}_i - \mathbf{p}_j$.

2.1.1.1 Voronoi Area

Mostly, the important discrete differential geometry measures can be determined for every triangle. Afterwards, the question arises how to compute the measures for every vertex. Intuitively, one idea is to use measures of the incident triangles and weight them based on the underlying triangle. One way to weight them is by using the Voronoi area. Given is a triangle $t = (i, j, k)$, which is divided in three regions $\mathcal{R}_i, \mathcal{R}_j, \mathcal{R}_k$ with $\mathbf{p}_i \in \mathcal{R}_i, \mathbf{p}_j \in \mathcal{R}_j, \mathbf{p}_k \in \mathcal{R}_k$. The regions are called *Voronoi regions*, if for every point \mathbf{p} inside the triangle and $\mathbf{p} \in \mathcal{R}_m$ with $m \in \{i, j, k\}$ follows that $\|\mathbf{p} - \mathbf{p}_m\| \leq \|\mathbf{p} - \mathbf{p}_{\{i, j, k\} \setminus m}\|$. The area of the Voronoi region can be computed according to Meyer et al. [MDSBo2]:

$$\mathcal{A}_t(i) = \begin{cases} \frac{\|\mathbf{e}_{ji} \times \mathbf{e}_{ki}\|}{4} & \text{if triangle } t \text{ is obtuse at } \mathbf{p}_i \\ \frac{\|\mathbf{e}_{ji} \times \mathbf{e}_{ki}\|}{8} & \text{if triangle } t \text{ is obtuse} \\ \frac{1}{8} \left(\|\mathbf{e}_{ij}\|^2 \cdot \cot(\mathbf{e}_{jk}, \mathbf{e}_{ki}) + \|\mathbf{e}_{ki}\|^2 \cdot \cot(\mathbf{e}_{ij}, \mathbf{e}_{jk}) \right) & \\ \text{otherwise.} & \end{cases} \quad (14)$$

2.1.1.2 Normals

The normals can be calculated for every triangle as well as for every vertex. For the triangle $t = (i, j, k)$, the normal is the cross product of the edges: $\mathbf{n}_t = \frac{\mathbf{e}_{ji} \times \mathbf{e}_{ki}}{\|\mathbf{e}_{ji} \times \mathbf{e}_{ki}\|}$. Note, that it must be assured that the orientation is consistent in the surface mesh. The normal for every

vertex can be obtained by a weighted accumulation on the incident triangle normals:

$$\mathbf{n}'_i = \sum_{t \in T, i \in t} \omega_t \mathbf{n}_t.$$

The natural choice to use $\omega_t = 1$ or $\omega_t = \text{area}(t)$ leads to discontinuities in the shading. Using the Voronoi area $\omega_t = \mathcal{A}_t(i)$ or the angle $\omega_t = \angle(\mathbf{e}_{ji}, \mathbf{e}_{ki})$ work quite well. The normal at i still needs to be normalized: $\mathbf{n}_i = \frac{\mathbf{n}'_i}{\|\mathbf{n}'_i\|}$.

2.1.1.3 Discrete Curvature

The calculation of the principal curvature directions and their curvatures can be determined by fitting higher order polynomials to the mesh [CP05; GI04] or by calculating the normal curvatures along the edges and then estimate the shape operator [CS92; HS03; MDSB02; PKSP+01; Tau95a]. This section focusses on another category of methods that estimates the shape operator directly [ACSDL+03; CSM03; Ruso4; HP11], see also Váša et al. [VVPS+16] for an analysis of curvature estimations. Here, the curvature estimation is provided according to [Ruso4]. First, the shape operator is determined for every triangle and afterwards for every vertex. To formalize this, an orthonormal basis for each triangle needs to be created. Given the triangle $t = \{i, j, k\}$ the basis $(\mathbf{x}_t, \mathbf{y}_t)$ is constructed by:

$$\mathbf{x}_t := \frac{\mathbf{e}_{ij}}{\|\mathbf{e}_{ij}\|}, \quad \mathbf{y}_t := \frac{\mathbf{x}_t \times (\mathbf{e}_{kj} \times \mathbf{x}_t)}{\|\mathbf{x}_t \times (\mathbf{e}_{kj} \times \mathbf{x}_t)\|}. \quad (15)$$

Applying the property that the shape operator yields the change of the normal in a certain direction leads to the following equation system:

$$S \begin{pmatrix} \langle \mathbf{x}_t, \mathbf{e}_m \rangle \\ \langle \mathbf{y}_t, \mathbf{e}_m \rangle \end{pmatrix} = \begin{pmatrix} \langle \mathbf{x}_t, \mathbf{n}_m \rangle \\ \langle \mathbf{y}_t, \mathbf{n}_m \rangle \end{pmatrix}, \quad (16)$$

with $m \in \{(i, j), (j, k), (k, i)\}$ and $\mathbf{n}_{ij} = \mathbf{n}_i - \mathbf{n}_j$. This results in an overdetermined system, which can be approximated by the least square method. Next, the shape operator S is calculated for each vertex of the mesh. Given an orthonormal basis $(\mathbf{x}_i, \mathbf{y}_i)$ for the vertex i , the basis $(\mathbf{x}_t, \mathbf{y}_t)$ of the incident triangle is rotated such that the basis are coplanar. Given the shape operator in the triangle basis $S_t = \begin{pmatrix} e_t & f_t \\ f_t & g_t \end{pmatrix}$, the matrix elements can be obtained by multiplying a combination of the 2D basis vectors $(1, 0), (0, 1)$ from left and right, e.g., $f_t = (1 \ 0)S_t \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. The matrix entries in the new basis $(\mathbf{x}_i, \mathbf{y}_i)$ are given by

$$e_i = \begin{pmatrix} \langle \mathbf{x}_t, \mathbf{x}_i \rangle \\ \langle \mathbf{y}_t, \mathbf{x}_i \rangle \end{pmatrix}^T S \begin{pmatrix} \langle \mathbf{x}_t, \mathbf{x}_i \rangle \\ \langle \mathbf{y}_t, \mathbf{x}_i \rangle \end{pmatrix},$$

The other entries are determined by analogous calculations. Thus, for every incident triangle of a vertex, the shape operator is re-expressed with the vertex' basis. Then, the shape operator is weighted by the Voronoi area of its corresponding triangle and sum the shape operators. The resulting shape operator is divided by the sum of the weights yielding a shape operator for every vertex. The eigenvalues and eigenvectors are the curvatures and principal curvature directions.

2.1.1.4 Discrete Gradient

In this section, the concept of the gradient is adapted to calculate it on surface meshes. First, the gradient is determined for every triangle and then it is computed for every vertex. Given is a scalar field φ on the surface that yields a value for every vertex i : $\varphi(i)$. For simplification, it is written as $\varphi_i := \varphi(i)$. The gradient of the triangle $t = \{i, j, k\}$ can now be determined by either constructing a basis for t and then building a linearized 2D scalarfield that coincides with $\varphi_i, \varphi_j, \varphi_k$ at the position of the vertices according to the basis. The gradient of the scalar field can then be calculated straightforwardly. Directly, the gradient is determined on the triangle yielding a 3D vector in \mathbb{R}^3 . The gradient results from:

$$\nabla\varphi_t = (\varphi_j - \varphi_i) \frac{(\mathbf{p}_j - \mathbf{p}_i)^\perp}{2A_\Delta} + (\varphi_k - \varphi_i) \frac{(\mathbf{p}_k - \mathbf{p}_i)^\perp}{2A_\Delta}, \quad (17)$$

where A_Δ denotes the area of the triangle and \perp stands for a counter-clockwise rotation by 90° in the triangle plane, see [BKPA+10]. The gradient per vertex is determined by transforming the basis of incident triangles to the basis of the vertex tangent space. Then, the gradients are weighted according to the Voronoi area of the associated vertices. Finally, the accumulated vector is divided by the sum of weights, recall Section 2.1.1.3.

2.1.1.5 Discrete Laplace–Beltrami Operator

For an overview of discrete Laplace–Beltrami operators commonly used in graphics, are state of the art reports by Sorkine [Sor05] and by Patané [Pat16] is recommended. In general the Laplace–Beltrami operator of a scalar field φ on surface meshes can be written by:

$$\Delta\varphi_i = \sum_j w_{ij} (\varphi_j - \varphi_i). \quad (18)$$

There exist different weights which lead to different Laplace operators [WMKGo7]. Mostly, the weights only operate on the neighbors. In the following, the commonly used weights are introduced.

Combinatorial: The combinatorial Laplace–Beltrami operator is determined by:

$$w_{ij} = \begin{cases} 1, & \text{if } (i, j) \in E \\ 0, & \text{otherwise.} \end{cases}$$

Uniform: Taubin [Tau95b] proposed the uniform Laplace–Beltrami operator. The weights are determined by the number of neighbors of \mathbf{p}_i :

$$w_{ij} = \begin{cases} \frac{1}{N(i)}, & \text{if } (i, j) \in E \\ 0, & \text{otherwise.} \end{cases}$$

Floater’s mean value: Floater [Flo03] suggested the mean value weights:

$$w_{ij} = \begin{cases} \frac{\tan(\delta_{ij}/2) + \tan(\gamma_{ij}/2)}{\|\mathbf{p}_i - \mathbf{p}_j\|}, & \text{if } (i, j) \in E \\ 0, & \text{otherwise.} \end{cases}$$

Cotangent weights: MacNeal [Mac49] proposed the cotangent weights:

$$w_{ij} = \begin{cases} \frac{\cot(\alpha_{ij}) + \cot(\beta_{ji})}{2}, & \text{if } (i, j) \in E \\ 0, & \text{otherwise.} \end{cases}$$

Belkin weights: Belkin [BSW08] determined the weights over the whole surface:

$$\Delta\varphi_i = \frac{1}{4\pi h^2(\mathbf{p}_i)} \sum_{\Delta_k} \frac{A(\Delta_k)}{3} \sum_{j \in \Delta_k} e^{-\frac{\|\mathbf{p}_i - \mathbf{p}_j\|^2}{4h(\mathbf{p}_i)}} (\varphi_j - \varphi_i),$$

where $A(\Delta_k)$ denotes the area of the triangle Δ_k and h corresponds intuitively to the size of the neighborhood.

2.1.2 Isolines on Surface Meshes

In Section 3.3.2, an important category of illustrative visualization techniques named the *feature lines* is described. These lines mostly assume an underlying scalar field on the surface mesh. The location of the lines are given by the zero-crossing of the scalar field φ , i.e., the loci of points \mathbf{p} with $\varphi(\mathbf{p}) = 0$. To obtain a more tessellation-independent result, it is essential to not restrict the lines to the edges than rather allow to draw entirely on the mesh even inside a triangle. For this in every triangle the sign of the scalar field with respect to the corresponding vertices is checked. If only one sign is positive (and the rest is negative) or only one sign is negative (and the rest is positive) then a zero-crossing inside the triangle occurs. First, the position of the zero-crossing on the two edges is determined and then they are

connected with a line. Let φ_i be the value at vertex i , which is the only one that is negative (or positive). The zero-crossing on the edge \mathbf{e}_{ij} is determined by first determining α for which $\varphi_i + \alpha(\varphi_j - \varphi_i) = 0$ is fulfilled: $\alpha = \frac{\varphi_i}{\varphi_i - \varphi_j}$ and then plugging this into $\mathbf{p}_i + \alpha(\mathbf{p}_j - \mathbf{p}_i)$, which yields the position on the edge:

$$\mathbf{p}' = \mathbf{p}_i + \frac{\varphi_i}{\varphi_i - \varphi_j}(\mathbf{p}_j - \mathbf{p}_i), \quad (19)$$

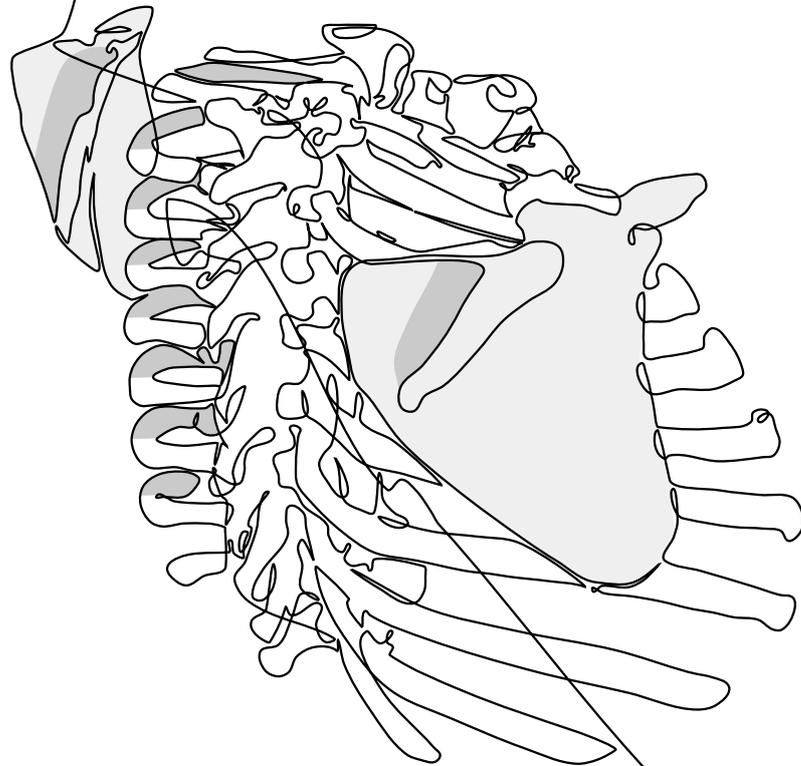
where the zero-crossing occurs. Analogously, this calculation can be carried out on the other edge. Afterwards, both points are then connected.

Part III

VISUALIZATION CONCEPTS

3

ILLUSTRATIVE RENDERING



This chapter is partly based on:

Kai Lawonn, Ivan Viola, Bernhard Preim, and Tobias Isenberg.

“A Survey of Surface-Based Illustrative Rendering for Visualization.”

In: Computer Graphics Forum (2017), submitted.

SURFACE-BASED ILLUSTRATIVE RENDERING FOR VISUALIZATION

MODERN rendering and visualization techniques provide us with a plethora of ways to turn 3D datasets into visual representations. Many of them use novel approaches to depict different data aspects. While most of these techniques are useful and have their application cases, e. g., enhancing details on the surface [RBD06; VPBG+11; SK14], the subgroup of illustrative visualization techniques focuses on learning from and/or being inspired by centuries of experience of scientific illustration. Such techniques aim to provide meaningful, expressive, and sometimes simplified depictions of a problem, a scene, or a situation. For example, illustrative techniques can introduce intentional abstraction for reducing visual clutter that allow the viewers to focus their attention on one or more points of interest in the visualization. Illustrative visualization techniques are mainly intended to be used in medical applications and to generate material for textbooks, but also applications in the natural sciences and in engineering have been discussed in the past.

However, while the previous definition of illustrative visualization [RBGV08] related it to traditional visual arts, it has been lacking a strong distinction from the general characteristics of visualization. Therefore, this chapter proposes a more comprehensive definition that amplifies the unique characteristics of the discussed subfield.

It makes an attempt to provide a structured overview of this subfield of visualization by surveying the approaches that are considered to generate illustrative visualizations. Yet, the field is still rather large and encompasses many visualization methods. In this overview the focus is on illustrative rendering techniques for 3D surface shapes. Specifically, it contributes with:

- a definition of illustrative visualization based on the literature,
- an overview of illustrative visualization concepts, and
- a discussion of open problems in illustrative visualization and a perspective on future research directions.

3.1 CONCEPT OF ILLUSTRATIVE VISUALIZATION

Visualization is a method to “convey salient information about [...] data” by means of visual encoding [HJ05], ultimately to “amplify cognition” [CMS99]. Over the past decades, visualization researchers

have been successful in finding numerous ways to effectively represent data aspects visually in wide variety of application domains. The field of visualization shares, however, the goal of the communication of information and insights about complex phenomena with the field of traditional illustration. Consequently, as the necessary computational and representational tools were developed within computer graphics,¹ visualization researchers took inspiration from traditional illustration to be able to generate *illustrative visualizations* [BCPV+12; Ise15; Ise16; LMo2b; RBGV08; VHE10].

In addition to this inspiration from the traditional craft, illustrative visualization techniques are characterized by *abstraction* and *simplification* [BHIM+05; Ise16; LMP13; LP16; RBGV08; VHE10], by emphasis of the important/relevant and suppression of the context [BHIM+05; LMP13], therefore by leading to clearer [BCPV+12], more expressive [BCPV+12; LP16; LMo2b; VHE10] depiction, and consequently more effective visualization [RBGV08; VHE10]. Thus, it allows visualization creators to use more visual variables [Ise15], use several layers of information [Ise15; Ise16], reduces clutter and occlusion [BCPV+12], and improves the perception of shape and depth [BCPV+12; Ise15; PBCI+16]. It thus also relies on insights from perception and cognition [BHIM+05; VHE10], ultimately to show the relevant information [BHIM+05], to allow people to better understand what they see [BCPV+12; HWRR+10; LP16], to better explore data [BCPV+12; Ise16], to better gain knowledge from it [BCPV+12; Ise16; VHE10], and to better communicate insights from data exploration [Ise16; VHE10]. Based on these notions, following definition is provided:

Definition 1. *An illustrative visualization is a visualization technique for measured, simulated, and modeled data that — inspired by traditional illustration — uses established illustration principles such as abstraction and emphasis, layered organization, dedicated use of multiple visual variables, or the support of perception of depth and shape — with the goal to create more effective, more understandable, and/or more communicative data representations than it is possible with other visualization techniques.*

Nevertheless, the question remains whether illustrative visualization is a “new technology” or simply a “tautology” [RBGV08] because is simply a characteristic of ‘good visualization’—after all, its goals and approaches are arguably shared with many if not most visualization techniques. Rautek et al. [RBGV08] thus answer that it is the latter, arguing that illustrative visualization will ultimately be an aspect of virtually all visualization work. Generally this notion is true, but seeing illustrative visualization as only a “useless tautology” [RBGV08] is selling it short. The focus on a dedicated reflection of principles that have long been established in traditional visualization

¹ Specifically as part of the computer graphics sub-field of non-photorealistic rendering [GG01; SS02; KCWI13; RC13].

and the exploration of how they are best applied to visualizations of “real” data embedded in a problem-solving situation. It allows illustrative visualization to ultimately advance the field of visualization as a whole.

It is thus essential to understand the state of the art of the field. Several surveys have already summarized the research for sub-domains of visualization including flow visualization [BCPV+12], the visualization of brain connectivity [Ise15], as well as a general tutorial [VGHB+05]. This overview focusses specifically on the rendering part for visualization of surface-based data as it results, for example, from iso-surface extraction (e.g., in physical simulations), from segmentation (e.g., in medical datasets), or from dedicated surface models (e.g., in the visualization technical objects and processes). The extraction and application of sparse line illustrations, surface-filling marks, and illustrative surface shading will be described.

3.2 GENERAL REQUIREMENTS

The illustrative visualization of surface meshes leads to several crucial requirements, to achieve the high quality results typically expected from illustration-like visuals. In particular, the following requirements need to be fulfilled: smoothing meshes, robustness, filtering the result, and frame-coherent behavior.

3.2.1 *Smoothing*

Most illustrative visualization techniques use higher order derivatives, such as gradient calculation or curvature estimation. To obtain reasonable values, the techniques require sufficiently smooth surface meshes. This property cannot be expected if the used data were acquired by a laser scanner or tomographic devices, such as industrial or medical CT. Discontinuities in the surface mesh may lead to distracting results. A smoothing algorithm to address these issues should thus keep prominent surface features, but suppress insignificant details [SCBW14; WYPW+15; ZDZB+15]. Objects that were acquired by tomographic devices may inherit staircasing artifacts in the generated surface. In such cases a special smoothing algorithm should be applied to reduce the staircases [MLKW+13; WZYW+15]. Depending on the illustrative visualization technique it may be sufficient to smooth the surface normals as the algorithm works on the normals only. If working on a scalar field a smoothing approach that suppresses noise in the scalar field may be sufficient.

3.2.2 *Robustness*

The result of the visualization technique should not significantly depend on the quality of the triangles or the tessellation of the surface mesh. In case the algorithm is applied on the same model with first high-quality triangles, i.e., mostly with equilateral triangles and second with arbitrary underlying, i.e., obtuse triangles can also occur, the result should not significantly differ. Furthermore, the result should also not deviate significantly if the number of triangles is strongly reduced, as it is often necessary to achieve faster computation times. Moreover, the algorithm should also handle holes or local noise. Even with these defects the rendering result should be reasonable and similar to the same mesh without these defects.

3.2.3 *Refinement*

Even smoothing the surface may lead to annoying or distracting illustrative visualization results. This may be caused by small local irregularities such as noise. One possibility to restrict unwanted results is to give the user the possibility to change the result manually. User-defined thresholds may be a way to control the outcome, if the illustrative visualization technique comes with a quantitative measure that influences the result. Some algorithms determine a scalar field, which measures the *importance* to apply an illustrative approach. Then thresholding can restrict the results to higher values. Another possibility may be to let the user directly change the result by drawing on the surface or to change parameter, which influences the result locally. In essence, it is crucial to give the user opportunities to directly change the visualization result.

3.2.4 *Frame Coherence*

Illustrative visualization techniques need to be frame-coherent because the resulting visualizations are typically explored in interactive contexts. This means that user interactions with the surface mesh (e.g., rotations, zooming, etc.) must not lead to sudden changes of the visual representation because these can be annoying and can disrupt the analysis. For example, if the surface is rendered with, e.g., lines or points and these primitives suddenly appear during the interaction, users are likely to shift their attention automatically to these regions. Instead, the appearance should be constant or changes should be introduced in a smooth manner, both during an interaction or in an animation.

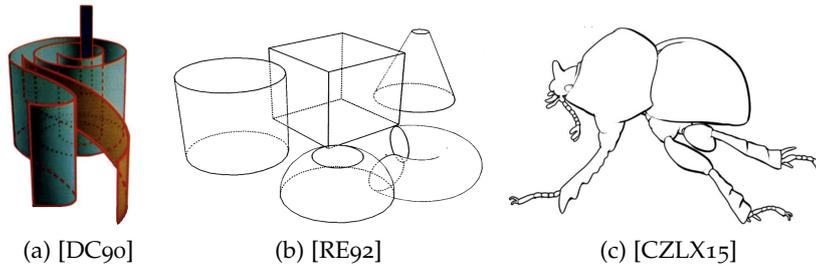


Figure 1: Contour examples with varying styles.

3.3 LOW-LEVEL ILLUSTRATIVE VISUALIZATION TECHNIQUES

Many of the fundamental principles discussed in Section 3.1 that are essential for illustrative visualization can be realized by applying one or more from a number of low-level techniques, see Rautek et al. [RBGV08]. These low-level illustrative visualization techniques include the creation of sparse line drawings based on *silhouettes and contours* (Section 3.3.1) and *feature lines* (Section 3.3.2), the use of marks on the surface including *hatching* (Section 3.3.3) and *stippling* (Section 3.3.4), and the application of *illustrative shading methods* (Section 3.3.5). They are based on different degrees of abstraction and this section discusses them starting from the most abstracting ones to the ones with less abstraction.

3.3.1 *Silhouettes and Contours*

Silhouettes, (occluding) contours, and feature lines (for the latter see Section 3.3.2) are part of the group of sparse line drawings.² They restrict the depiction to only a few lines that can potentially be stylized and have been used for centuries in illustrations.

3.3.1.1 *Silhouette and Contour Detection*

Among the lines that are used for sparse line drawings, the silhouette is defined as the illustration of an object’s outline, i. e., the border of the object to the background. This definition can be traced back to Étienne de Silhouette, the finance minister of the French king Louis XV, who is often connected to the paper-cut shadow profile portraits which were popular at his time. The contour (of a completely smooth object), on the other hand, is defined as the loci of all the points for

² A general overview of sparse line drawings is provided by Rusinkiewicz et al. [RCDF08], the state of the art reports by Hertzmann [Her99], Isenberg et al. [IFHS+03], and Li et al. [LZLX+12], and a discussion of the advantages of stylization by Al-Rousan et al. [ARSK15].

which the normal vector \mathbf{n} and the view vector \mathbf{v} are mutually perpendicular:

$$\langle \mathbf{n}, \mathbf{v} \rangle = 0. \quad (20)$$

This contour is a very important cue for the understanding of a shape’s surface. For example, perceptual studies [Mar76] confirmed that the first stage of visual perception includes the detection of the contours. Because the contour, at the same time, does not provide information about the shape of the surface itself it provides a strong abstraction of a shape—yet one that can typically still be recognized.

Techniques for the detection and illustration of contours can roughly be divided into three categories [Her99; IFHS+03]:

- *image-based techniques,*
- *object-based techniques,* and
- *hybrid techniques;*

see the examples in Figure 1 and a summary of existing techniques in Figure 7 which also includes approaches not mentioned in the text. The first among them, image-based algorithms, detect discontinuities of the pixel values, i. e., they operate on the rendered image in the view plane. Some techniques also make use of the z-buffer to find relevant discontinuities [ST90; MBC02]. Object-based methods, on the other hand, employ the geometry with their primitives, e. g., triangles and vertices, and use the 3D coordinates as well as additional information such as the normals to detect contours. It is important to note that these object-based approaches, unlike the other two groups, also need a dedicated processing of the visibility of the resulting contours [IFHS+03]. By using the surface mesh, the contours can be generated as a geometrical object. The generated object can then be further processed for additional stylization or information [HSC12; LLPH15]. Hybrid methods combine aspects of both approaches, e. g., by first performing operations on the object space and then using image-space operations [Rus89; RE92], see Figure 1b, or plain rendering [GSGS+99].

For surface-based illustrative visualization the state of the art is the extraction of a sequence of edges that represent the contours to facilitate further stylization. Initially, this sequence was extracted as those edges from the surface mesh for which the sign of the dot product of the view vector with the normals of the incident triangle normals changes [Her99] which unfortunately leads to artifacts [IHS02]. Hertzmann and Zorin [HZ00] thus interpreted the surface mesh as the approximation of a smooth surface and extracted the *sub-polygon* approximation of the smooth surface’s contour by detecting zero-crossings of the dot product between the interpolated vertex normals with the view vector, and then connecting these zero-crossings

across the triangles of the mesh. This process is just as fast as the mesh-based edges, but leads to strokes with much higher quality. A process for the correct computation of the contour based on these sub-polygon contours exists as well [BHK14], but does not facilitate a computation at interactive frame rates.

3.3.1.2 *Stylizing Silhouettes and Contours*

Pioneering work for the use of stylization in the field of illustrative rendering was done by Appel [App67] as well as Dooley and Cohen [DC90]. In particular, they derived the visibility of lines and then applied various line styles to encode spatial relations. For example, Dooley and Cohen illustrated hidden lines as dashed lines as shown in Figure 1a. This general concept was later replicated and realized numerous times, for example by Saito and Takahashi [ST90] (for image-based line extraction), Strothotte et al. [SPRS+94], Markosian et al. [MKGT+97], and Gooch et al. [GSGS+99]—all focusing on slightly different aspects and different line styles.

3.3.2 *Feature Lines*

In addition to silhouettes and contours, sparse line illustrations also comprise feature lines—lines that characterize particular features on the surface of the object that are not necessarily characterized by (potential) changes in visibility [LP16]. These feature lines are similarly important to be able to understand the shape of the depicted objects and are typically placed in regions where discontinuities occur. For example, strong changes of curvature or a strong rise of the illumination values may warrant the use of feature lines. In general, feature lines can be divided in two main classes:

- view-independent feature lines
 - illumination-independent
 - illumination-dependent
- view-dependent feature lines.

3.3.2.1 *View-Independent Feature Lines*

View-independent feature lines are determined based on features derived from surface information (independent from the orientation and location of the surface) or from light sources. Thus, this category can be further divided into illumination-independent and illumination-dependent techniques. The advantage of the illumination-independent lines over the illumination-dependent or the view-dependent lines are that they can typically be computed in a pre-processing stage and only their visibility has to be determined at run-time, which makes

the overall computation faster. View-dependent and illumination-dependent feature lines, in contrast, have the advantage that they can be more stably computed and are derived based on the features that are perceived by a viewer from a certain vantage point, thus making them generally the accepted choice for sparse line drawings [DFRS03; JDA07; CGLB+08; LBSP14; BLSP15; LP16]. Only for objects with mathematically ideal features (e. g., illustrative visualizations of technical models) the two classes produce equivalent output [PBCI+16], which makes the use of view-independent lines a better choice in this case.

ILLUMINATION-INDEPENDENT FEATURE LINES

CREASE LINES To detect such feature lines one can use a definition inspired by that for contours: Instead of looking at the set of points where $\langle \mathbf{n}, \mathbf{v} \rangle = 0$ holds, the algorithm for *crease lines* looks for edges where the dihedral angle, i. e., the angle between two incident triangles, exceeds a user-defined threshold:

$$\langle \mathbf{n}_t, \mathbf{n}_{t'} \rangle \geq \tau, \quad (21)$$

where $\mathbf{n}_t, \mathbf{n}_{t'}$ are the normals of neighbored triangles and τ is the user-defined threshold. The value $\tau = \cos \alpha$ denotes the cosine of the dihedral angle which should be exceeded to display the edge, independent of any viewing conditions. This method, however, can only detect features in the surface mesh if they are characterized by large angles between adjacent triangles—features on smooth surfaces that are represented by highly-tessellated meshes cannot be found. Moreover, it is difficult and often impossible to find a well-suited angle α that leads to an adequate set of feature lines, even for a single mesh. For example, for meshes with strong as well as weak features this method may not be able to depict both features, especially if the mesh suffers from noise.

RIDGES AND VALLEYS A more stable approach is thus to treat meshes as approximations of smooth surfaces and to then derive features based on properties of these smooth surfaces, in particular based on the local curvature. Interrante et al. [IFP95] described such *ridges and valleys* (RV) for volume data, an approach that was later adapted to surface meshes by Ohtake et al. [OBS04]. The definition of *ridges and valleys* is based on the local curvature values with $|\kappa_1| \geq |\kappa_2|$ and the principal curvature direction \mathbf{k}_1 which corresponds to κ_1 . The *ridges and valleys* are then defined as the loci of points where the directional derivative of the curvature κ_1 in direction of the principal curvature direction \mathbf{k}_1 reaches an extremum:

$$D_{\mathbf{k}_1} \kappa_1 = 0. \quad (22)$$

Depending on the extremum and the sign of the curvature, ridges and valleys are distinguished: According to two constraints, the sets of points are called

$$D_{\mathbf{k}_1} D_{\mathbf{k}_1 \kappa_1} \begin{cases} < 0, & \text{and } \kappa_1 > 0: \text{ ridges} \\ > 0, & \text{and } \kappa_1 < 0: \text{ valleys.} \end{cases} \quad (23)$$

Intuitively, one can think of a plane that cuts the surface along the principal curvature direction and looks for an extremum with respect to the curvature. The point with the local highest curvature value is then defined as a ridge or valley point. To filter the lines, which may occur at noisy regions, a user-defined threshold is employed. For every connected series of lines, an integral is determined which measures the curvature along the line. If the magnitude of the integral surpasses the user-defined threshold the line is drawn, otherwise it is discarded. Because the ridge and valley lines are of 3rd order derivatives, however, this method is susceptible to noise and requires a smoothed mesh. Small discontinuities may lead to erroneous results of the derivatives, thus leading to visually unpleasant results. The ridges and valleys are, nevertheless, of good quality if the mesh has strong features and is guaranteed to be smooth.

Several approaches to extract salient features similar to ridge and valley lines exist. Watanabe and Belyaev [WB01], for example, use an approximation of the mean curvature, a non-linear averaging of the curvature maps, a histogram-based curvature extrema filtering, and a skeletonization procedure. They used a simplification to determine salient curvature extrema triangles. This approach was later improved to a fully automatic approach by Belyaev and Anoshkina [BA05]. Their method consists of two steps: first a smoothing of normals with a non-linear diffusion filter and, second, an application of a Canny-like non-maximum suppression using a hysteresis threshold operation. Yoshizawa et al. [YBS05; YBYS07; YBYS08], finally, presented various techniques to detect robust ridge and valley lines on surface meshes based on a better estimation of the shape operator, yielding more reliable curvature measures.

DEMARCATING CURVES Kolomenkin et al. [KST08] presented *demarcating curves* (DEM) as the points of maximum curvature derivative:

$$\langle \mathbf{w}, S\mathbf{w} \rangle = 0 \quad \text{with } \mathbf{w} = \arg \max_{\|\mathbf{v}\|=1} D_{\mathbf{v}}\kappa. \quad (24)$$

S is the shape operator and κ the curvature along the direction \mathbf{v} . The direction \mathbf{w} that satisfies this property can analytically be determined as the root of a 3rd order polynomial, but is consequently susceptible to noise. This noise, however, can be reduced by discarding line parts where the curvature derivative in the gradient direction is lower than a user-defined threshold.

RELIEF EDGES Kolomenkin et al. [KST09] presented another approach to detect characteristic lines. They assume that the surface mesh consists of a (smooth) base mesh with a height field. The *relief edges* are then defined as the edge of the height field. The height field of a surface S can locally be expressed as

$$S(\mathbf{p}) = \frac{1}{2}\langle \mathbf{p}, S\mathbf{p} \rangle + \frac{1}{2}C(\mathbf{p}, \mathbf{p}, \mathbf{p}), \quad (25)$$

where C denotes a $2 \times 2 \times 2$ rank-3 tensor defined as $C := (D_u S \ D_v S)$. A smooth step edge function E :

$$E(\theta, \alpha, u, v) = \frac{1}{6}\alpha(u \cos(\theta) + v \sin(\theta))^3, \quad (26)$$

is fitted to the height field S with the local coordinates u, v and the edge intensity α . *Relief edges* are then defined as

$$\langle \mathbf{w}, S\mathbf{w} \rangle = 0 \quad \text{with} \quad (\theta, \alpha) = \arg \min \int (E - S)^2 \rho \, d\rho d\omega, \quad (27)$$

where $\mathbf{w} = (\cos(\hat{\theta}), \sin(\hat{\theta}))$. Note that θ is written as the minimization argument, but used $\hat{\theta}$ as the solution. The condition were simplified for a relief edge; Kolomenkin et al. [KST09] describe a detailed derivation of the relationship between θ and $\hat{\theta}$. Similar to the demarcating curves, this approach uses 3rd order derivatives.

ILLUMINATION-DEPENDENT FEATURE LINES

PHOTIC EXTREMUM LINES Xie et al. [XHTS+07] presented *photic extremum lines* (PELs) to illustrate the surface mesh. While all previously described feature lines are based on discontinuities of the surface mesh, PELs analyze discontinuities in the illumination. These feature lines are motivated by the importance of illumination as a shape cue or visibility and are thus based on the shading of the surface. PELs lines are defined as those locations where the variations of the shading reaches a maximum based on Lambertian reflectance $f := \langle \mathbf{n}, \mathbf{l} \rangle$, with \mathbf{n} being the normal and \mathbf{l} being the light vector (equivalent to the view vector \mathbf{v} for headlight setups). Using the normalized light gradient $\mathbf{w} = \frac{\nabla f}{\|\nabla f\|}$ PELs are then defined as:

$$D_{\mathbf{w}}\|\nabla f\| = 0 \quad \text{and} \quad D_{\mathbf{w}}D_{\mathbf{w}}\|\nabla f\| < 0. \quad (28)$$

In this case the PELs only depend on a single light source and the result can be improved by adding additional light sources. For example, additional local light sources can reduce erroneous lines in noisy regions. Similar to ridges and valleys, the filtering is done by measuring the strength of the integral along a line with respect to the magnitude of the light gradient. A line is drawn if the integral exceeds the user-defined threshold, otherwise it is discarded. Zhang

et al. [ZHS10] later improved the computation of PELs to achieve real-time rendering. While the computation of the PELs also relies on 3rd order derivatives, it is only necessary to smooth the normals to obtain reasonable results on noisy surfaces.

LINES VIA ABSTRACTED SHADING The line generator technique (LAS) proposed by Lee et al. [LMLH07] determines view-dependent regions around potential ridge and valley lines and is based on two passes. The first pass computes the shading, for example based on Lambertian reflectance $f := \langle \mathbf{n}, \mathbf{l} \rangle$. The second pass then identifies image pixels that form ridges and valleys of the image intensities due to the shading. Then a second-degree polynomial $f(x, y) = a_0x^2 + 2a_1xy + a_2y^2 + a_3x + a_4y + a_5$ is fitted to the shading values at each fragment. In practice, the authors use nine (image-independent) sample points (x_i, y_i) , arrange them in a 3×3 grid, and construct the matrix X which is made up of the rows $(x_i^2 \ 2x_iy_i \ y_i^2 \ x_i \ y_i \ 1)$. The matrix $H = (X^T X)^{-1} X^T$ can then be calculated as a pre-processing step. At run-time, the authors determine the matrix $A = HT$ where T is a matrix that consists of the shading values of the sample points. The matrix A yields the coefficients of the fitted second-degree polynomial. Then, the function f can be written as a quadratic form:

$$Q(\mathbf{x}) = (\mathbf{x} - \mathbf{c})^T \underbrace{\begin{pmatrix} a_0 & a_1 \\ a_1 & a_2 \end{pmatrix}}_M (\mathbf{x} - \mathbf{c}), \text{ with} \quad (29)$$

$$\mathbf{c} = -\frac{1}{2} M^{-1} (a_3 \ a_4)^T. \quad (30)$$

The eigenvalues and eigenvectors of the matrix M serve as principal curvatures and principal curvature directions on the image. The *abstracted shading lines* are then defined as the curve through \mathbf{c} in direction of the lower curvature value. As such, with a headlight lighting setup, the line generator extracts pixels that contribute to contours and will, in part, reproduce the suggestive contours generator. The visualization designer through the use of different lighting, however, has the freedom of defining where the lines will depict more or less details. As an image-based technique, the second pass automatically handles the depicted level of detail. Interestingly, in contrast to other feature line approaches this method is not a technique where lines are identified as zero-crossings. Therefore, this method uses first order derivatives, but if the principles curvatures and the directions would be determined by the Hessian matrix on the shaded image and the ridge and valley line would have been determined as isolines, the technique would be of third order.

LAPLACIAN LINES *Laplacian lines* (LL) were introduced by Zhang et al. [ZHXX+11]. They determine the Laplacian of the Lambertian reflectance $f := \langle \mathbf{n}, \llbracket \cdot \rrbracket \rangle$ and then find the zero-crossings:

$$\Delta f = 0 \quad \text{and} \quad \|\nabla f\| \geq \tau, \quad (31)$$

Noisy lines can be filtered with a user-defined threshold τ which measures the magnitude of the light gradient. The *Laplacian lines* are of 3rd order, but similar to the PELs one can restrict the smoothing to the normals to get a better result for noisy surfaces. An advantage of this method is that, instead of determining the Laplacian of the actual shading, it can be calculated on the normals: $\Delta f = \langle \Delta \mathbf{n}, \mathbf{v} \rangle$. This yields a simplified pre-processing step and increases the frame rates during the interaction.

DIFFERENCE OF GAUSSIANS Inspired by their use in image processing [MH80], Zhang et al. [ZXYH+12] adapted the difference of Gaussians (DoG) concept to surface meshes for the depiction of characteristic lines. Their main idea is to apply two different Gaussian kernels $G_{\sigma_e}, G_{\sigma_r}$ to the illumination $f := \langle \mathbf{n}, \mathbf{l} \rangle$ of the surface. The final image is then obtained by calculating the difference by subtracting the smoothed illumination results from each other, formally:

$$H(f, \sigma_e, \sigma_r, \tau) = G_{\sigma_e}(f) - \tau \cdot G_{\sigma_r}(f). \quad (32)$$

Zhang et al. define the Gaussian of the illumination as:

$$G_{\sigma}(f)|_{\mathbf{x}} = \frac{1}{2\pi\sigma} \int f(\mathbf{y}) \exp\left(-\frac{\|\mathbf{x}-\mathbf{y}\|}{2\sigma^2}\right) d\mathbf{y}. \quad (33)$$

The strength of a feature line at a point is then defined as:

$$D = \begin{cases} 1 & \text{if } H > 0 \\ 1 + \tanh(\phi H) & \text{otherwise} \end{cases} \quad (34)$$

where ϕ determines the sharpness of the rendered result and D determines the fragment's color. Zhang et al. simplified the Gaussian with $G_{\sigma}(f)|_{\mathbf{x}} \approx \mathbf{v}(\mathbf{x}) \cdot G_{\sigma}(\mathbf{n})|_{\mathbf{x}}$ and used σ_e as the average edge length and $\sigma_r = \sqrt{1.6}\sigma_e$. The computation of the Gaussians is, therefore, just a pre-processing step. Zhang et al. also extended the difference of Gaussians to an anisotropic version. Here, they determine a local parametrization $h(\mathbf{u}, \mathbf{v})$ such that \mathbf{u} points in direction of the principal curvature direction \mathbf{k}_2 with $|\kappa_1| \geq |\kappa_2|$ being the minimal absolute value and \mathbf{u} being orthogonal to \mathbf{v} . Then, they smooth the normals along \mathbf{v} , yielding a visually more pleasing result. In contrast to other feature line approaches, the DoG method detects (linear) regions rather than lines defined by zero-crossings. Due to the use of pre-processing and the simple subtraction step, however, this method is faster than the other approaches.

3.3.2.2 View-Dependent Feature Lines

In contrast to the view-independent lines, other feature lines are computed based on the given viewing (and sometimes illumination) conditions. A large number of lines have been suggested which will be reviewed below.

SUGGESTIVE CONTOURS DeCarlo et al. [DFRS03] introduced *suggestive contours* (SC) as the first view-dependent feature line method. They describe features that are occluding contours in nearby view-points, and that naturally extend occluding contours in the 2D view plane. Two equivalent definitions exist: The first is based on the surface normal \mathbf{n} , the view vector \mathbf{v} which points towards the camera, and the projected view vector on the tangent space $\mathbf{w} = (\mathbf{Id} - \mathbf{nn}^\top)\mathbf{v}$. Suggestive contours are then defined as the loci of points where $\langle \mathbf{n}, \mathbf{v} \rangle$ reaches a minimum in direction of \mathbf{w} :

$$D_{\mathbf{w}} \langle \mathbf{n}, \mathbf{v} \rangle = 0 \quad \text{and} \quad D_{\mathbf{w}} D_{\mathbf{w}} \langle \mathbf{n}, \mathbf{v} \rangle > 0. \quad (35)$$

The other definition is based on the radial curvature κ_r . Given the principal curvature directions $\mathbf{k}_1, \mathbf{k}_2$ with curvatures κ_1, κ_2 , the projected view vector \mathbf{w} can be written as a linear combination of the principal curvature directions $\mathbf{w} = \alpha \mathbf{k}_1 + \beta \mathbf{k}_2$. These coefficients then yield the radial curvature $\kappa_r = \alpha \kappa_1 + \beta \kappa_2$ and the *suggestive contours* can be defined as the loci of points where

$$\kappa_r = 0 \quad \text{and} \quad D_{\mathbf{w}} \kappa_r > 0. \quad (36)$$

is fulfilled. Filtering can be applied based on a user-defined threshold which tests if the radial curvature exceeds the given threshold. *Suggestive contours* thus use 2nd order derivatives, making them less susceptible to noise than lines based on 3rd order derivatives.

Several authors have adjusted and extended the concept to date. For example, DeCarlo et al. [DFR04] themselves explored the real-time computation of the lines, while McGuire and Hughes [MH04] discussed a hardware-based implementation that also facilitates line stylization. Jeong et al. [JNLM05] and Ni et al. [NJLM06] have presented an approach controlling the depicted amount of line detail by adjusting the underlying mesh in a progressive fashion based on the current view. Goodwin et al. [GVH07], finally, addressed line stylization issues as mentioned above. Other authors have explored the application of suggestive contours specifically to illustrative visualization applications. For example, Burns et al. [BKRF+05] demonstrated the computation for volumetric models as used in medical visualization, while Lawonn et al. [LGP14a] employed them for vessel visualization.

HIGHLIGHT LINES DeCarlo and Rusinkiewicz [DR07] extended the concept of suggestive contours and added two new classes of

feature lines, *suggestive highlights* and *principal highlights* (HL), that provide shape cues analogous to shading. The *suggestive highlights* are defined as the loci of points where $\langle \mathbf{n}, \mathbf{v} \rangle$ reaches a positive maximum in the direction of \mathbf{w} :

$$\kappa_r = 0 \quad \text{and} \quad D_{\mathbf{w}}\kappa_r < 0. \quad (37)$$

Note that, in contrast to Equation 36, this definition evaluates where the second condition is negative. The *principal highlights*, on the other hand, are defined as strong positive maxima of $\langle \mathbf{n}, \mathbf{v} \rangle$ in the direction of $\mathbf{w}_{\perp} := \mathbf{n} \times \mathbf{w} / \|\mathbf{n} \times \mathbf{w}\|$. For this purpose the authors define the radial torsion τ_r as $\langle S(\mathbf{w}_{\perp}), \mathbf{w} \rangle = \tau_r \|\mathbf{w}\|$ and use the principal curvature direction \mathbf{k}_1 with $|\kappa_1| \geq |\kappa_2|$ as follows:

$$\langle \mathbf{k}_1, \mathbf{w} \rangle = 0 \quad \text{and} \quad D_{\mathbf{w}_{\perp}}\tau_r < 0. \quad (38)$$

A user-defined threshold is then used to discard lines whose derivative is lower than this value. Like suggestive contours, *suggestive* and *principal highlights* use 2nd order derivatives and are thus less susceptible to noise than lines based on 3rd order processing.

APPARENT RIDGES Judd et al. [JDA07] presented *apparent ridges* (AR), an extension to the ridges and valley concept that uses a view-dependent curvature term. Formally, apparent ridges are the loci of points at which the view-dependent curvature assumes an extremum in the view direction:

$$D_{\mathbf{t}'}\kappa' = 0 \quad \text{and} \quad D_{\mathbf{t}'}D_{\mathbf{t}'}\kappa' < 0. \quad (39)$$

In this definition, the sign of κ' is always positive, and the authors use the sign of curvature (defined on the mesh) to distinguish between ridges and valleys. A ridge occurs whenever Equation 39 is fulfilled and the curvature on the mesh is negative, while a valley is detected if the curvature is positive. The view-dependent curvature is defined by a projection operator P that maps points on the surface mesh to the view plane. Then, κ' is defined as

$$\kappa' = \max_{\|P(\mathbf{x})\|=1} \|S(\mathbf{x})\| \quad (40)$$

and \mathbf{t}' is the corresponding vector. Like in the approaches discussed before, the authors filter out undesired lines with a user-defined threshold for the view-dependent curvature. The major advantage of *apparent ridges* is that they combine the advantages of static ridge and valley lines with the benefits of view-dependent features, depicting strong and thus salient changes in the surface mesh as observed by the viewer/camera. Their disadvantage is the use of 3rd order derivatives, which can lead to cluttered or noisy results for insufficiently smooth meshes.

Table 1: Major feature line methods. The derivative order are stated as well as if the method is view-dependent, illumination-dependent, and if the method can be applied to animated surfaces in real-time.

Name	Order	V.-dep.	I.-dep.	Anim.
Occluding Contours	1	yes	no	yes
Crease Lines	1	no	no	yes
Ridges & Valleys	3	no	no	no
Demarcating Curves	3	no	no	no
Relief Edges	3	no	no	no
Suggestive Contours	2	yes	no	yes
Highlight Lines	2	yes	no	yes
Apparent Ridges	3	yes	no	no
Photic Extremum Lines	3	no	yes	yes
L. Abstracted Shading	1 (3)	no	yes	yes
Laplacian Lines	3	no	yes	no
Difference of Gaussian	3	no	yes	yes

3.3.2.3 Summary

The line concepts discussed in Section 3.3.1 and Section 3.3.2 differ in their specific characteristics, Table 1 provides an overview. Occluding contours are essential for illustrative sparse line visualizations, while a good choice of feature lines is recommended—in typical cases view-dependent ones. For most of these concepts the respective authors have demonstrated their suitability for illustrative visualization, for example suggestive contours for terrain models [NJLM06], medical objects [JNLM05; NJLM06; LGP14a], biological illustrations [GVH07], and even medical volume scans [BKRF+05], highlight lines for mathematical shapes [DR07], apparent ridges for mechanical shapes and medical objects [JDA07], photic extremum lines for mathematical shapes and volumetric medical data [XHTS+07], abstracted shading for terrain visualizations [LMLH07], and Laplacian lines for medical and biological surface models [ZHXX+11]. Other line concepts exist as well but, like Sousa and Prusinkiewicz’s [SP03] approach, are often geared more toward an artistic representations of shapes instead of efficiently communicating object shape.

The specific choice of lines depends on the application and the line properties. For example, some line concepts have noise issues with non-smooth surfaces which makes them less suited for surfaces that are directly derived from segmentations or 3D scans. Other feature line methods cannot easily be applied to animated and deforming objects because surface curvature and the curvature derivatives cannot easily be computed in real-time for deforming objects. To address

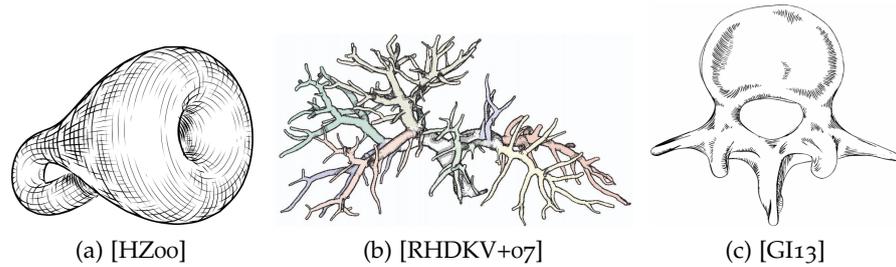


Figure 2: Three hatching examples, (a) as an image-based hatching, (b) as a texture-based hatching, and (c) as an object-based hatching

this last issue, Kalogerakis et al. [KNSM+09] presented an approach that learns a mapping from a set of animation parameters to surface curvatures for a deforming surface mesh. With this model they are able to predict the changing curvatures as well as their derivatives, based on which they derived a fast algorithm to extract feature lines of deforming objects at runtime.

3.3.3 Hatching

Hatching is another category of line drawing visualization techniques that is inspired by traditional illustration styles. In contrast to sparse lines that are placed at prominent features, hatching consists of a set of compound lines which cover larger parts of the surface to convey a spatial impression. The hatching approaches can generally be divided into three categories (see examples for these classes in Figure 7): *image-space*, *texture-space*, and *object-space* methods. These three classes have in common that, for most of them, the hatching strokes are placed along the principal curvature directions (PCDs). Early work has provided several arguments for this choice [IFP96; GIHL00; SW04], in particular in the case of visualization. They differ, however, in the specific domain/space in which the hatching strokes are placed, with implications for their use in illustrative visualization.

3.3.3.1 Image-space Hatching

As the name suggests, image-space hatching approaches operate on the projection of the surface mesh. This is a natural approach to hatching as the traditional technique was created by manually drawing lines on a 2D image (e.g., see Dürer’s artwork). Image-space approaches can again be roughly categorized into three classes based on how they compute the hatching strokes: *texture projection*, *streamline calculation*, and *line integral convolution*.

TEXTURE PROJECTION This first class uses dedicated hatching textures for rendering the shapes. These textures are defined in the 2D image plane and are either aligned with a fixed direction depending on the view [LMHBoo] or are aligned with the principle curvature directions [LKLo6; KYLo8]. This class of hatching methods thus relies on an image-plane mask that ensures that the area of the surface mesh is filled with the hatching texture. While this general approach facilitates an easy implementation and is fast to compute (even for animated shapes in a frame-coherent manner [KYLo8]), it can lead to the shower door effect. This effect results in strokes being perceived as if they were connected to the image plane and not to the object itself, which leads to distraction which is not suitable for most interactive visualization applications.

STREAMLINE CALCULATION The approaches in this class generate hatching strokes by tracing streamlines along the PCD (or other directional properties such as gradients of properties such as illumination or depth [LUS13; LUS14]) in the image plane. Typically, these fields are first computed in object space and then projected into image space [HZoo; RKoo]. In principle, the method starts by defining seed positions and then applying the streamline calculation using a Euler scheme along the PCD to derive the hatching strokes. To deal with the natural noise in the PCDs, this approach can use a segmentation of the PCD buffer into regions with homogeneous principal direction fields as done by Rössl and Kobbelt [RKoo] or by smoothing the direction field before the integration of strokes [HZoo]. The resulting strokes can then be stylized, for instance based on illumination, while cross-hatching can be used for particularly dark regions (e. g., [HZoo]). The seed point placement is essential for this class of methods [LUS13] as each stroke is derived independent from the ones already placed. This is important, in particular, for animations to avoid the shower door effect by ensuring that the seed points stay locked with respect to their initial positions on the 3D surface [LUS14]. Because each stroke is independent, they can still come close or even overlap, even for a good set of seed points. This issue can, however, be addressed by ensuring that strokes only have limited influence, for example using Lichtenberg et al.'s [LSHL16] stroke regions that are maintained for each seed. Alternatively, the seeds are placed iteratively such that a new seed is placed at a given distance from a previous hatching stroke and the stroke is stopped when it comes too close to an already existing line [HZoo]. The result of this approach is typically visually pleasing (e. g., Figure 2a), partially because it resembles the manual hatching process.

LINE INTEGRAL CONVOLUTION Another approach to address the issue of close lines is to use line integral convolution (LIC) which was

initially conceived for 2D vector field visualization [CL93; SH95]. In contrast to streamline computation, the LIC method uses an underlying noise field to generate lines in form of shaded pixels. Different noise fields can be used, for example based on shading, color, or features [KYM12]. This method is also not reduced to using the PCDs as the primary directions for the strokes but can also use other such as tangent directions of isocurves of view-dependent features, tangent directions of isophote curves [Min13; Min15], or light gradients [LKEP14]. While the LIC approach produces generally hatching-like images that can convey shape, the resulting images only exist in form of pixels and are not as similar to the traditional illustration style as the results from the streamline-based approach.

3.3.3.2 *Texture-space Hatching*

Texture-space hatching approaches rely on special textures which encode the hatching style. These textures are attached to the 3D surface mesh of the visualized object and then rendered using the traditional graphics pipeline. Brightness variations and the use of single and cross hatching is possible. This approach has the advantage that no shower door effect is created and it results in the impression of an illuminated surface illustrated by hatching. Two different techniques exist to parametrize the texture mapping, either *local* or *global*, as explained next.

LOCAL PARAMETRIZATION A local parametrization maps a region on the surface to a 2D coordinate system, and vice versa. The local parametrization may include a region consisting of several triangles, where the number of neighbored triangles depends on the variation of the curvature [PHWF01] or it may just include a fixed number of neighbors [SBB13]. Praun et al. [PHWF01] presented an interactive hatching approach in which they interactively control the viewing and lighting conditions in a frame-coherent manner. In a pre-processing step, they create hatching textures by starting with a stroke and create several instances of it. The creation of these *Tonal Art Maps* comprises textures that vary in both tone and resolution. They impose a nesting property, i. e., hatching textures in one image appear in all the finer resolution images, and all the darker tone images. At run-time, a subset of the textures are selected and blended and are applied to the mesh surface. The projection is based on the work by Praun et al. [PFH00] which cuts patches from an example image. Then, each patch is repeatedly pasted in random places onto the surface until it is completely covered. A direction field is needed to cover the patches onto the surface mesh because the hatching strokes have a certain direction. This direction field consists of the principal curvature direction with the maximal absolute curvature value. Webb et al. [WPFH02] extended Praun et al.'s [PHWF01] work and

introduced *Volume Tonal Art Maps* to enhance the control of tone and the use of colors. Instead of using a set of textures to control the brightness, they used a volume texture that allows a finer transition in case of a change of the illumination. Another extension of Praun et al.'s [PHWF01] work was introduced by Gasteiger et al. [GTBP08] who used a hatching method for anatomical meshes which were derived from clinical volume datasets. Their main contribution is to add the model-based preferential directions of the underlying anatomical structures. A real-time hatching on large scenes was proposed by Suarez et al. [SBB13]. In their approach the stroke direction depends on the curvature direction as well as on the light direction. The triangle adjacency is employed to guarantee coherent strokes of the surface mesh. Later, they extended their approach [SBB16] and presented a hatching scheme that can be applied on surface meshes with an associated texture. Different hatching patterns are generated based on the surface mesh's texture, which also vary in their tone. The resulting textures are stored in a multi-resolution tonal art map. In addition, different shading types are possible including regular shadows, soft/cast shadows, and self-shadowing. Moreover, their algorithm works on static, as well as animated surface meshes.

GLOBAL PARAMETRIZATION A global parametrization assumes a bijective map from every vertex on the surface mesh to a 2D coordinate system. The calculation of a global parametrization, however, is a challenging task as a *good* parametrization should preserve angles and areas. Such a parametrization does not exist in general, thus several methods exist that try to find an approximate solution [FH05]. For example, Ritter et al. [RHDKV+07] presented a hatching technique for vascular surfaces, see Figure 2b. They first acquire the surface meshes including a segmentation of the CT/MR volume data with the underlying skeleton structures. Based on the skeleton and the distances from the centerline to the surface, they then create texture coordinates. These texture coordinates are then used for the hatching, in a fragment shader, identifying the coordinate for which the fragment needs to be black. Their hatching scheme was used to illustrate the distance of the camera to the surface, the distance of a lesion to the surface, and to encode distance-encoded shadows. In contrast, Szécsi et al. [SSK16] provided a hybrid hatching solution by introducing *Tonal Art Maps with Image Space Strokes (TAMISS)*. First, they assign every stroke of a tonal art map a unique ID. Then, they fit a curve on each fragment that shares the same ID. Finally, these curves are extruded to image-space stylized strokes.

3.3.3.3 Object-space Hatching

Object-space hatching relies on explicit line primitives on the 3D surface. The lines are seeded on the surface mesh and are then traced

along a direction field, typically the PCDs. Because this processing is independent of the later projection into image space, it generally cannot guarantee a minimum distance between two strokes and have to use dedicated treatment to address cases where lines come too close. The advantage of object-space hatching is, however, that it inherently supports the interactive exploration of the depicted objects. These approaches thus avoid a re-computation of the hatching lines at run-time and do not exhibit any visual animation artifacts that arise in image-space approaches from the 2D character of lines.

Elber [Elb99] first presented an object-based approach for freeform surfaces. He uniformly distributed points on the surface, integrated lines along inherent direction fields such as the parametric directions or isophote directions, and parameterized the lines according to illumination and surface properties. Elber and Cohen [EC06] later also explored PCDs and the parametrization according to a contour-based notion of visual importance. This general approach has also been applied to surface meshes [RKSoo]. All these techniques, however, lead to rather random stroke distances to the more or less random placement of seed points. To address this issue, Deussen et al. [DHRS+99] and Medeiros et al. [MSVF09] thus explored a hatching approach that is based on computing intersections of the 3D geometry with evenly placed plane objects. The planes are kept more or less parallel to each other, oriented according to some skeleton or spine. The resulting illustrations resemble some hand-drawn techniques well, but work best only for more or less cylinder-like structures. To also be able to deal with general shapes, authors thus turned to careful seeding and line integration, combined with the use of direction fields on the surface. Singh et al. [SS10], for example, used illumination-based directions, orienting their hatching lines similar to view-dependent feature lines such as suggestive contours and apparent ridges. Zander et al. [ZISS04], instead, used direction fields based on the PCDs and described an elaborate line shading technique that can be controlled and changed at run-time. Moreover, they addressed the issue of object-space hatching lines coming too close to each other at regions perpendicular to the view direction by locally changing the line stylization to become less visually prominent. Lawonn et al. [LMP13] then described a GPU-based approach for the line integration on the 3D surface to make it possible to use the techniques in interactive scenarios and for complex meshes. All these techniques, however, relied on purely mathematical concepts to guide the line placement and parametrization. Gerl and Isenberg [GI13] thus use an example-based approach (Figure 2c) that first learns line placement and parametrization from a handmade example illustration and associated model, to later be able to automatically or interactively create new illustrations for other shapes, but using the initially captured hatching style.

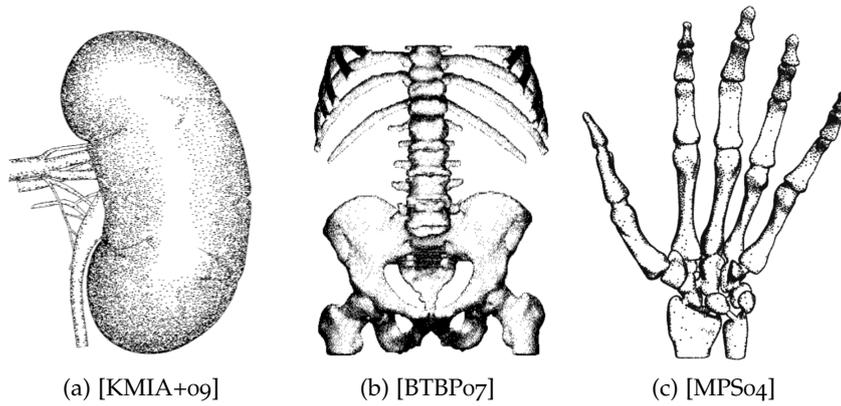


Figure 3: Three stippling examples, (a) as an image-based hatching, (b) as a texture-based hatching, and (c) as an object-based stippling.

Various application scenarios for such object-based hatching have been discussed. The authors of the cited papers have applied their techniques, for example, to the illustrative visualization of medical objects [DHRS+99; ECo6; GI13; LMP13; MSVF09; SS10], of mathematical shapes [Elb99; SS10; ZISS04], and of botanical models [ZISS04]. Generally, the line integration techniques are useful for shapes with well-defined direction fields (e. g., mathematically defined objects), while the surface intersection techniques are more robust to noise and can also be useful for ‘less clean’ objects such as those resulting from medical scans [SEI10].

3.3.4 Stippling

Similar to computer-based hatching, computer-based stippling is inspired by a traditional illustration style and uses primitives distributed over the surface of the depicted objects to illustrate its shape and properties. In contrast to hatching, however, stippling does not use directional primitives and restricts itself to dots. The placement of the dots and their size and potential overlapping is thus used to indicate, for example form, illumination, and materials. Furthermore, unlike feature lines and hatching approaches, stippling was not motivated by perceptual studies. Several computer-based stippling approaches exist [DI13; KCWI13]. Like the hatching approaches, the stippling approaches are roughly categorized into *image-space*, *texture-space*, and *object-space* methods, see Figure 3 for examples and Figure 7 for an overview.

3.3.4.1 Image-space Stippling

Generally it is possible to use virtually any traditional image-space stippling technique also for the illustrative visualization of surface

meshes—simply by rendering the surface into an image and then applying the stippling (for an overview see the existing surveys [DI13; KCWI13]). Often dot distributions are computed based on relaxation (e.g., [DHOS00; Seco2; BSD09]); or using dedicated distribution functions (e.g., [KCODL06; VBTS07]), example-based dot distributions (e.g., [KMIA+09], see Figure 3a), or scale-dependent schemes (e.g., [MALI11]). These approaches are well suited for producing illustrative visualizations see, for example, for medical illustration [KMIA+09; INCS+06; SHS02], biological objects [DHOS00; HHD03; KMIA+09], technical models [DHOS00], and even archeological artifacts or sites [INCS+06; MALI11].

3.3.4.2 *Texture-space Stippling*

Similar to the equivalent hatching approach, texture-space stippling techniques rely on dedicated stippling textures that are projected onto the surface mesh. Baer et al. [BTBP07], for instance, used the basic idea of tonal art maps [PHWF01] and created stipple illustrations for medical applications. To minimize distortions of the circular points, they employed a polycube representation (e.g., Figure 3b). Krüger and Westermann [KW07], instead, use 3D volumetric noise textures to assign a noise value to each surface fragment, which is then compared to the fragment's illumination to either render it as a stipple dot or not. This approach works for both surface and volume models and the authors demonstrated its use in medical visualization.

3.3.4.3 *Object-space Stippling*

Object-space stippling, finally, places the stipple dots directly onto the surface of the object to be illustrated. Like in object-space hatching, this process has the major advantage that animation or interaction artifacts are avoided such as the shower door effect. Moreover, this process facilitates the use of animated models, such as animations in biological education (e.g., [MPS04], see Figure 3c). For example, Meruvia Pastor et al. [MPS02; MPFS03] assign a dot to each vertex of the mesh, while Costa Sousa et al. [CSFWS03] assign marks to each mesh edge. Meruvia Pastor et al. then randomly perturb the dot positions to avoid artifacts from regular meshes and, to adjust the density of the dots to the illumination, use mesh subdivision and progressive meshes. Costa Sousa et al., in contrast, rely on dense meshes and adjust each mark according to the illumination and other properties. In an alternative technique, Meruvia Pastor and Strothotte [MPS04] use point hierarchies to adjust the stipple density. Similarly, Lu et al. [LTHE+02; LMTE+03] place several dots per polygon and only render some based on the illumination and size on the screen, the latter to take scale issue into account. In a fundamentally different approach, Yuan et al. [YNZC05] used a conformal parametrization to ob-

tain the stippling results, effectively operating in the geometry-image domain. This approach combines the benefits from image-space approaches (e.g., use of edge detection) with benefits from the purely 3D techniques (e.g., frame-coherence in animation). Purely object-space techniques, on the other hand, have been demonstrated to also be extendable to volumetric models [LMTE+03]. Examples for the use of object-space stippling for visualization purposes include the application to medical models [LTHE+02; LMTE+03; MPFS03; MPS04; CSFWS03], technical models [LMTE+03], terrain models [CSFWS03], and archeological artifacts [MPS04; CSFWS03].

REMARKS ON STIPPLING In contrast to curvature-based hatching, where studies indicate a clear advantage for shape perception (see 3.4), no rigorous study shows an advantage of stippling over conventional shading. Thus, it remains open whether it is merely an aesthetically pleasing approach.

3.3.5 *Illustrative Shading*

In contrast to the former techniques, where the model was illustrated with primitives such as lines or points, this section explains various techniques that illustrate the model based on lighting conditions. The term *illustrative shading* refers to those techniques that perform shading that is neither physically-based simulation of light propagation, nor it attempts to mimick such a simulation process. The central property of illustrative shading is that the shading emphasizes specific structural aspects and communicates them more effectively than a result of a photorealistic shading.

There is however not a clear separation between line-drawing techniques and illustrative shading techniques, as some line drawing generators in fact use the shading in the line generation algorithm. The work of Lee et al. [LMLH07] discussed in Section 3.3.2 presents one example centered around the idea that line drawing can be considered as an abstraction of shading. Their work is presented in Figure 5b, where the lines are combined with a toon shading, discussed below. The PEL-generator produces slightly different results, but is based on the same claim that illumination should be an important factor in the line drawing design [XHTS+07]. While these two techniques have exemplified how shading can serve as a basis for a line generator, the relation between shading and line generation can be also the other way around, shading can be expressed by means of lines [DR07] as is the case in the above described highlight lines. The entire category of hatching techniques uses lines to convey shading. So the first category of shading techniques can be categorized as both lines as well as illustrative shading techniques.

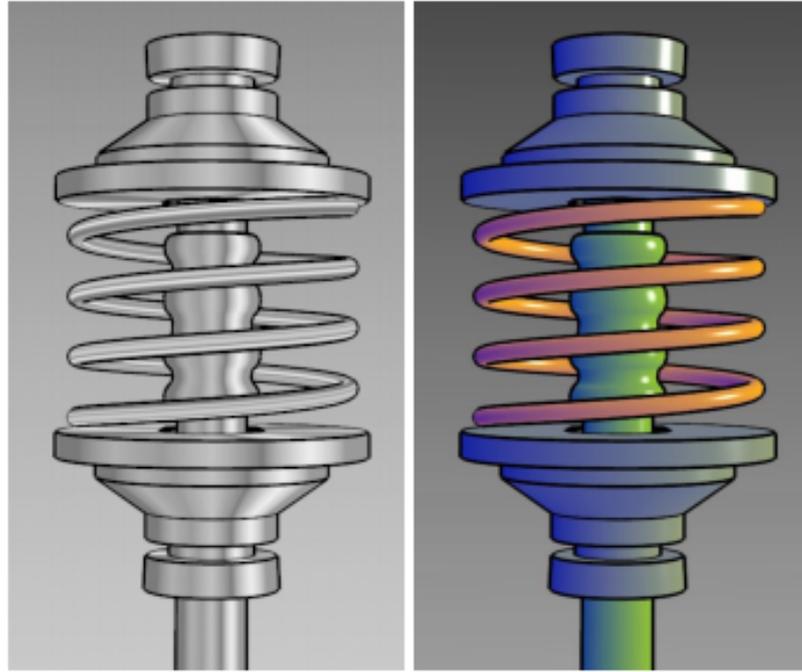


Figure 4: One of the first illustrative shading techniques by Gooch et al. [GGSC98]. Strong edges are illustrated in black color, highlights with white, and the surface shading is based on the luminance.

The recent work by Liu et al. [LML16] has extended the line drawing algorithms by the notion of a *global tone*, an illustration method based on the visibility of a region or a point on the surface, related to ambient occlusion. The mild toning adds to the line drawing additional shape cues and the entire illustration looks closer to a hand-crafted depiction when charcoal is used as the depositing material.

TOON SHADING The most classical approach to illustrative shading originates from the cartoon industry. There, because of lack of colors, one can quantize shading into few bands of color tones only, instead of a continuous threshold. This creates a distinct appealing effect. Such shading style is known as *toon* or *cel* shading. There are many variants of toon shading and it is hard to date its first computerized use. The most straightforward implementation of toon shading first defines a desired amount of luminance bands. Then the shaded intensity coming from a Lambertian shading would be quantized to a central intensity of a corresponding band. Toon shading can be used in illustrative rendering for encoding quantity intervals, or it can be used for shading the context in less of detail, so that the focus area is differentiated by a richer shading scheme.

Barla et al. [BTM06] extended the idea of toon shading and added an eye-space depth as an additional parameter in their *X-Toon Shader* technique. The dot product between the surface normal and the light direction serves as an argument together with the depth value into

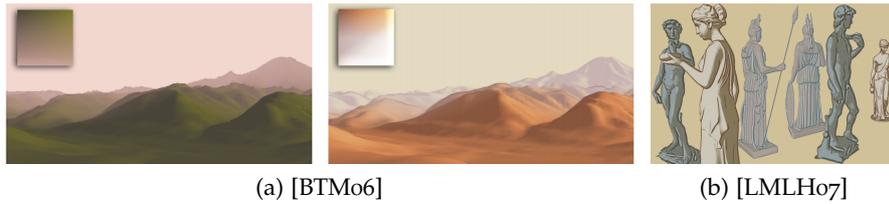


Figure 5: In (a) the results of the X-Toon Shader by Barla et al. [BTM06] is shown, where the 2D look up table is presented as an inlet. In (b) the line drawing by abstracted shading combined with toon shading is depicted by Lee et al. [LMLH07] is illustrated.

a two-dimensional lookup table with pre-computed shading values. Such a pre-computed lookup table basically represents a distinct rendering style that incorporates an appealing haze effect for more distant objects. The example of X-Toon Shader results are shown in Figure 5a. This principle has been extended by Hao et al. [HCZW10] who used a 2D texture to assign the fragment of a surface a specific color. The x -component is used as the radial curvature, similar to the curvature in suggestive contours. The y -component of the texture encodes the surface shading, i.e., the positively clamped dot product of the surface normal with a light vector.

The idea of capturing the shading behavior in a texture has been taken further in the work that proposes the *Apparent Relief Descriptor* [VBGS08]. This is essentially a texture that defines how the surface with particular curvature properties will be shaded, a concept similar to *Curvature-based Transfer Functions* for volume rendering [HKG00]. The shape descriptor is characterized by the principal curvatures. A lookup texture is used to encode the shading (or coloring) based on the combination of these two curvature parameters. This way ridge and valley areas can be made visually distinct, or any other features that have distinct curvature properties.

COOL-TO-WARM SHADING Another very well known example of illustrative shading is the technique proposed by Gooch et al. [GGSC98]. The key idea comes from observation of fine arts, where painters rarely use black paint for shaded objects or objects in shadow. Instead they are using a more *cool* color, typically darker blue, so that the light propagation is encoded, but the areas in shadow are still perceivable. The proposed shading scheme therefore does not only change the luminance of an object at a particular location on its surface, but the color changes simultaneously in hue. This technique is known as *Cool-to-Warm* shading. Highlighted areas become more yellowish, while strongly shaded areas become more blue. An example of this illustrative technique is depicted in Figure 4.

SHADING BY TRANSFER The X-Toon Shader [BTM06] discussed above uses a pre-computed texture for fetching the values during the shading process, based on light vector, normal vector, and the eye-space distance, as input parameters. The idea of pre-computing shading and fetching it during the image synthesis has been around many years already. The famous example on how to transfer hand-drawn shading using the *pre-computed* approach is known as the *Lit-Sphere* concept [SMGG01]. The artist can provide a large amount of various visual styles by drawing a 3D sphere. Such sphere is used as a lookup texture using the eye-space normal as the input parameter. On the particular location on the target geometry, the shading value is determined by computing the eye-space normal vector at this location that is input into the lit-sphere lookup texture. The point on the sphere which has the same normal vector will be taken as an output of the lit-sphere lookup and will be used for shading the target geometry.

Very recently, the lit-sphere concept has been extended by a *sphere on the table* in the StyLit approach [FJLS+16]. This way, not only the local shading can be captured, but also the global-illumination effects such as shadows and color bleeding. Previous techniques for style transfer such as *Image Analogies* [HJOC+01] or *Neural Stylization* [GEB16] use image statistics or convolutional neural networks for transferring the style. These turn out to be not satisfactory for a given purpose of the illumination transfer. However it is possible to decompose the light into *light-path expressions* [Hec90], namely into the direct diffuse, direct specular, indirect light and 1st and 2nd bounce components. If image statistics for each light path expression is computed independently, then a notably better result could be achieved. This result can be furthermore improved by using a texture-synthesis that controls the level of usage of particular image patches to avoid image patch overuse.

Prior to the above work Tietjen et al. [TPBG+08] proposed the so-called *shading maps* where each map contains a particular light type or geometric property encoding. These maps are then combined with various weights in the compositing stage to obtain a final image. First, the surface mesh is rendered in different ways, e.g., with shading of different light positions, the depth, contours, feature lines, curvature, etc. The results are stored in the framebuffer. The final image is generated as a weighted sum of all previously calculated results. Therefore, the final image slightly attenuates features that would not be shown with standard shading.

DYNAMICALLY COHERENT STROKE SHADING Above examples of the stylized shading in principle work for one image only. If creating an animation sequence, each image is created independently and naturally leads to temporal visual artifacts. This can be addressed by the work of Breslav et al. [BSMB+07] that used a coherent approach to

fix 2D texture patterns on the surface mesh. As a first step, they distributed 3D sample points over the surface, during the run time the positions on the view plane is traced. Thus, the position of the previous frame and the current frame is stored. With this observed motion a transform is applied to the pattern such that the visualization is frame-coherent during interaction.

SCALED SHADING A prominent category of illustrative shading consists of techniques that put a special emphasis on strongly conveying the surface details by modifying the standard shading equations. Rusinkiewicz et al. [RBD06] with their *Exaggerated Shading* technique enhanced details on the surface mesh, motivated by the rules for cartography heightmap design. The standard diffuse Lambertian shading is therefore modified to support levels of exaggeration of slope of changes based on the parameter α :

$$e = \frac{1}{2} + \frac{1}{2} \text{clamp}_{[-1,1]} \alpha \langle \mathbf{n}, \mathbf{l} \rangle \quad (41)$$

This exaggeration further utilizes a spectral approach, where the surface is decomposed into several bands of surface details. These levels can then be exaggerated individually and the final appearance is composited as a weighted sum of all surface-detail frequencies. To achieve various levels of detail, the surface normals are smoothed in multiple iterations and each iteration represents one scale.

A work that is visually similar to the exaggerated shading is based on light warping driven by view-dependent curvature [VPBG+09]. The perceptual reasoning why this technique conveys shape effectively, stands on visual perception research that describes the way how the perception of curvature depends on the compressions of reflected light patterns on the surface. Interesting surface details are characterized by non-zero curvature values. The technique essentially exaggerates the reflection vector based on view-dependent curvature so that it enriches the light pattern that is compressed onto the interesting surface area.

While these above techniques clearly emphasize surface features, a negative side-effect is that the modification of the shading changes also the perception of the material. A glossy material might easily change to a brushed appearance for example. In a follow-up work, the same authors have therefore presented a technique that both enhances the surface depiction and preserves the material appearance. The technique is referred as *Radiance Scaling* method [VPBG+11]. The key idea is to modify reflected light intensities based on view-dependent surface curvature as well as material characteristics. In the case of Phong illumination model, each light component, i.e., ambient, diffuse, and specular, has a separate scaling function. An example of surface landscape has been emphasized by radiance scaling in Figure 6b.

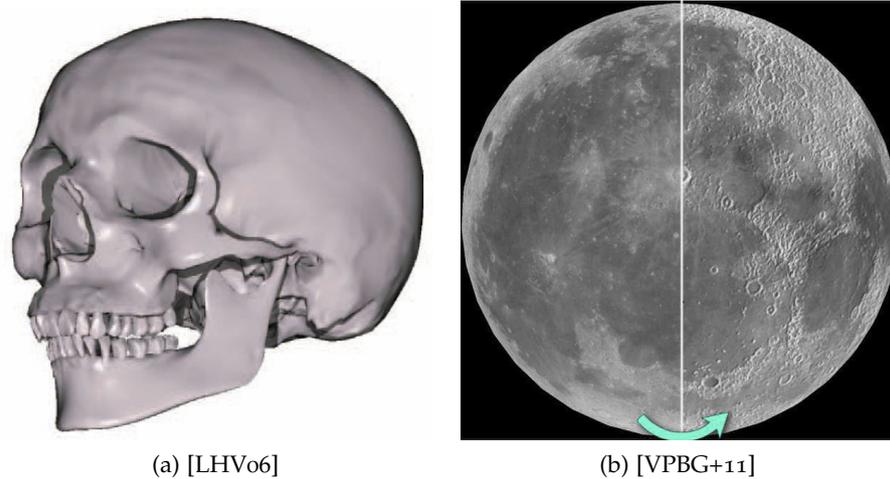


Figure 6: In (a) the Light Collages technique by Lee et al. [LHV06] is shown, where several light sources contribute to the object illumination. In (b) the method by Vergne et al. [VPBG+11] shows how radiance scaling can significantly improve the perception of surface features, without affecting the perception of material.

LIGHT COLLAGES Several techniques above decompose illumination or geometry into number of aspects that can be individually controlled. In the *Radiance Scaling* or the *StyLit* techniques it was the different light expressions, in the case of *Exaggerated Shading* it was the surface detail frequency. The technique called *Light Collages* [LHV06] also decomposes the geometry of the surface into several patches separated by high curvature borders. For each patch an optimal light source is calculated, so that the diffuse and the specular lights convey the surface details to the viewer. Once these are calculated, the light corresponding to a particular patch only illuminates that patch and not the rest of the surface geometry. At the patch borders the illumination originating from different light sources is mixed together to create a continuous illumination effect that is locally optimal but globally inconsistent. An example of such expressive physically inconsistent lighting is shown in Figure 6a.

3.4 EVALUATION

Just like for visualization in general, evaluation is also essential for illustrative visualization. The techniques employed for this specific sub-field of visualization are at the intersection between evaluation for visualization in general [IICS+13] and evaluation for non-photorealistic/expressive graphics [Ise13]. Below several specific examples are given and discussed how they are shedding light on the generation and use of illustrative visualizations.

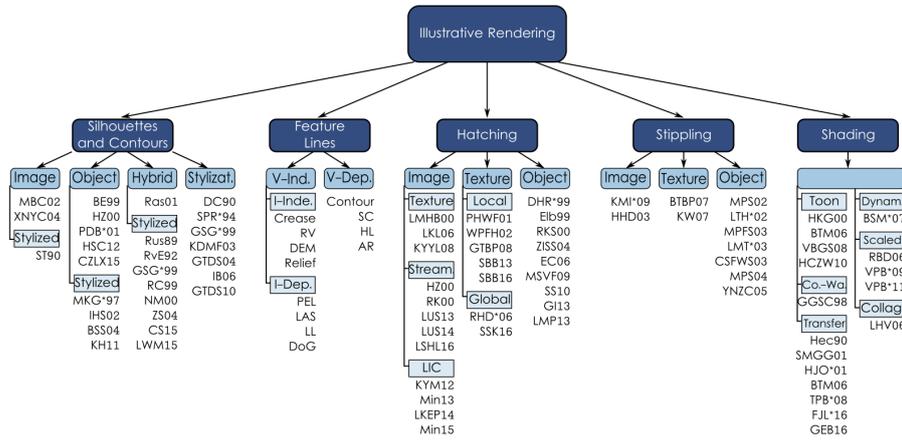


Figure 7: An overview of the discussed illustrative visualization techniques. Every field is divided in different categories, which reflects the subsections of the corresponding section.

Almost all techniques that were considered are actually meant for supporting interactive exploration; thus external validity is limited if only images are used.

Since illustrative visualization both derives its inspiration from traditional illustration and uses actual data to create *case-specific* (i. e., not generalized) *visual representations*, evaluation techniques can both be used to assess the quality of the specific representation or to understand the underlying *illustrative principles in general*. Moreover, illustrative visualizations are often either evaluated *qualitatively* to assess people's preferences and the suitability of a given visualization to a specific task or *quantitatively* to assess depth and shape perception. Perceptual studies (e. g., see [PBCI+16]), in particular, are often used to get an understanding of perception aspects.

As an example for an evaluation of a specific technique, Tietjen et al. [TIP05] used different rendering styles in combination on medical CT datasets. Based on an existing segmentation, their technique illustrated the anatomy in a combination of contours and shading. 33 participants (8 surgeons) filled out a questionnaire to assess the suitability of the generated illustrative visualizations. For each set of questions, two visualization results of the same anatomy were shown that used a different combination of visualization techniques. The questionnaire then asked about a preferred visualization and posed further questions about the images. The analysis of the questionnaires showed that, in the used example of thorax and liver visualization, the use of contours are appropriate for surgical planning. Additional information such as colored contours and transparent surfaces, however, were rated as useful. Based on their observations, the authors express their belief that their hybrid visualizations are useful in general for medical applications.

Later, Ritter et al. [RHDKV+07] evaluated illustrative visualizations of vascular structures. Their real-time technique relied on perceptual principles and in their evaluation they specifically concentrated on the used techniques of stroke hatching, distance-encoded surfaces, and shadows. They conducted a web-based perceptual study with a large pool of participants (160 people), with almost a quarter being medical professionals. Because their techniques were designed to improve depth perception, they based their evaluation on precise depth judgments. Their results demonstrated that the explicit encoding of distance by illustrative means leads to a more accurate perception of relative distance to the viewers compared to traditional shading. Moreover, the shape of the vessels was equally well communicated by illustrative means compared to traditional shading. Finally, their illustrative distance-encoded shadows were also effective—they led to more accurate depth judgments compared to visualizations without such shadows.

In a subsequent evaluation of an extended visualization system by the same group [HWRR+10], the authors conducted a lab study that also relied on precise judgments of the perceived distance of parts of the visualization. This evaluation looked at three different application scenarios, specifically designed for an augmented reality applications for liver surgery. They also added the illustrative visualization of resection surfaces compared to the original technique. Since the evaluation was conducted under controlled lab conditions, the authors were also able to capture the participants' reasoning by applying a think-aloud protocol. Therefore, the authors were able to not only confirm the effectiveness of the illustrative techniques but also captured valuable qualitative feedback for the further refinement of the application of illustrative visualization in practical applications.

Hummel et al. [HGHH+10] describe an informal and brief evaluation of illustrative visualization in a different application domain: the illustration of integral surfaces. They collected feedback from domain experts and collaborators on the effectiveness of the different illustrative techniques they used. In line with the results reported by Tietjen et al. [TIP05], they reported that the use of silhouettes/contours was able to improve the visualizations and that purely abstract (lines only) illustrations were not considered to be the best choice. Moreover, respondents confirmed the effectiveness of their adaptive transparency technique and preferred techniques that also made use of high-quality illumination.

In a final example by Lawonn et al. [LLPH15], the authors again evaluated illustrative visualizations of vascular structures. They first conducted a qualitative evaluation with 15 expert participants, a third of whom were surgeons. The discussion with the experts lead to a slight change of the visualization technique. Then they conducted a subsequent quantitative evaluation in form of a web-based question-

naire with 50 participants, 16% of whom were physicians. The results showed that their technique improved depth perception in comparison with other well-known methods.

While these evaluation results of specific visualization examples can likely be generalized beyond their very specific application scenarios, there are also more fundamental questions that relate to the use of illustrative techniques for visualization. In particular, the general utility of certain illustration principles is one aspect that can be established in general such that they can be applied to illustrative visualizations of 3D shapes. Also, due to the fundamental motivation of the field based on the tradition of hand-made illustrations it is important to question the relationship of computer-supported techniques to hand-made examples.

In particular to explore this second point, Isenberg et al. [INCS+06] conducted an observational study for non-photorealistic rendering (NPR) techniques compared with traditional hand drawings—using specifically illustrations and illustrative renderings of 3D shapes. They used three surface models (a human torso, a plant part, and a scan of an archeological artifact) and, with 5 different NPR techniques (primarily hatching and stippling, combined with contours and feature lines), created illustrative visualizations. In addition, they also asked five professional illustrators to produce illustrations of the same models. The resulting 30 images were then printed and used in a pile-sorting study: The study participants were asked to sort the illustrations into groups according to their own preferences, with virtually no constraints, and then to discuss these groups. Only after this discussion the experimenters also conducted a semi-structured interview, also asking specific questions. The results of the study included a better understanding of the differences between computer-generated and hand-made illustrations, including insights on what people appreciate in either group, what makes an illustration look hand-made, and how recommendations on what to improve for illustrative rendering.

To better understand the fundamental differences between hand-drawn and computer-generated line drawings, Cole et al. [CGLB+08] compared algorithmically generated lines (silhouettes/contours and view-dependent feature lines) with hand-drawings of the same shape. These drawings were produced by 20 artists for 12 different shapes using a clever study design that asked the artists to first freely draw and then to copy the drawn lines onto a low-contrast image to make them comparable with each other. The comparison of the scanned hand drawings and the computer-generated illustrations showed that most algorithmic lines indeed match with hand-drawn ones, for both object-space and image-space line extraction. There are, however, also some hand-drawn lines that cannot be explained by the analysis of local properties—artists sometimes choose to place lines in regions of

weak ridges or valleys, sometimes depending on the semantics of the subject matter.

In a follow-up study Cole et al. [CSDF+09] then raised the question of how well line drawings can depict the underlying shape. They thus again used twelve 3D models and created visualizations using shading and different line-based techniques. They then asked 560 participants (on Mechanical Turk) to orient randomly placed surface gauges such that they marked the local surface orientation, resulting in a total of 275,000 gauge placements. The analysis of this data showed that, for half of the examined 3D models, the line drawings could depict shape nearly as well as a shaded image. For other shapes the participants had more trouble interpreting the surface shape solely based on the used sparse line renderings. In all but one case, however, the best computer-generated drawing led to a slightly lower error than the artist's drawing, but the specific algorithm depended on the specific 3D shape.

A comparative evaluation to assess feature lines was conducted by Lawonn et al. [LBSP14]. They asked 149 participants to order six computer-generated line illustrations generated using ridge and valley lines, suggestive contours, apparent ridges, PELs, demarcating curves, and Laplacian lines. The participants were asked to rank based on realistic assessment, aesthetic depiction, and general preference. Each of the six line illustration used a different type of feature line, and the study was carried out for two different 3D shapes. Based on the analysis of these rankings, the authors conclude that suggestive contours, apparent ridges, and Laplacian lines were the preferred techniques. Overall, it was recommend to use suggestive contours because they do not use third-order derivatives. A later analysis by Baer et al. [BLS15] re-tested these results and confirmed their statistical significance.

3.5 APPLICATIONS

As already pointed out for the individual approaches, illustrative visualization techniques can be applied to numerous domains. Below we name a selection of these domains and briefly discuss some examples of how illustrative techniques have been used.

ARCHAEOLOGY Archeology has traditionally been relying on traditional illustrations, both for the documentation of excavations and artifacts as well as for the illustration of reconstructions. For many of these tasks it is possible to also use illustrative techniques, based on 3D models or 3D scans. For example, virtually reconstructed building [SMI99; MS98] were visualized with sparse line rendering, using line stylization to encode semantics of the excavation or reconstruction. In illustrative renderings of scanned ancient objects [KST08;

LLZ11; LTPH17] the visualization techniques can enhance, for example, the shape of the depicted objects. Especially relevant becomes illustrative visualization for archeology and for other fields when a certain information is combined with information that is uncertain or even fully speculative. In excavations one can find only few pieces and speculate how the entire object looked like. In such scenarios illustrative rendering techniques based on line or point elements may communicate the uncertain part very effectively, in contrast to the certain information which is depicted with a very detailed or even photorealistic rendering technique.

MOLECULAR SCIENCES In the context of molecular sciences, illustrative visualization techniques can be applied to molecules to explain the continuous character of the different abstraction schemes [ZLBI11], make additional visual variables available to visualize surface information [CG07; CJG09; LKEP14], or illustrate complex temporal aspects and reactions [MPSV14; MWPV15].

MEDICINE The applications of illustrative visualization in medicine are manifold since medical illustrations has traditionally played an important role in the field. For instance, vascular structures were displayed with distance encoding to a tumor or to the user [RHDKV+07; LLPH15], illustrative augmented reality has been used for liver surgery [HWRR+10], or virtual endoscopy has been augmented with illustrative methods [LSP14]. Illustrative visualization has also been used for medical training [SESS05]. Even fluoroscopic images with extracted 3D surface were combined with contours, stippling, or hatching illustration [RSBS+13]. Finally, some approaches have combined various illustrative techniques to convey semantic information [TIP05; LSHL16].

GEOSCIENCES For geovisualization the use of illustrative rendering has a great potential to communicate complex and spatial information [Dö7]. For 3D city scenes, for example, visualization techniques can convey level-of-abstraction transitions during the interaction [STKD12]. Furthermore, illustrative visualization can emphasize the salience of user-specified regions of interest to depict the whole scene with a focus on such areas [PCGC+11].

FLOW VISUALIZATION Another important application area is flow visualization [BCPV+12]. Major trends are flow pattern enhancement [BMGS13], the support of depth perception, e. g., by means of halos [EBRI09], and the illustration of blood flow [BPMS12].

PHYSICS Illustrative techniques have been applied in relativity and astrophysics [WBEF+05; WBEF+06]. The main goal in this context

was to convey various aspects of the theories of special and general relativity and of related fields of cosmology and astrophysics.

MISCELLANEOUS Illustrative visualization approaches are applied in many more fields such as biology [DHOS00], technical illustrations [HBPE+07; ND05], and mathematics [SS10]. Interaction frameworks exist [KMMK+02; GI13] to support both the automatic generation and manual adjustments. Another interesting application is the retrieval of a 3D model from a simple line drawing. If the user draws a model based on an abstract depiction, a matching 3D surface can be obtained [ERBH+12]. Finally, illustrative visualization techniques can also be used to produce physical visualization objects such as 3D laser-engraved glass [HCW10].

3.6 DISCUSSION & FUTURE WORK

This overview has shown, illustrative visualization has a lot of potential for visualization in general. As discussed in Section 3.1 and pointed out throughout the discussion, it allows visualization designers to convey salient information, encode (additional) information, improve perception, and guide attention through abstraction and emphasis. This chapter largely focused on surface-based models, but many techniques can also be applied to other types of data such as volumetric models—or dedicated algorithms exist for such data. Moreover, low-level techniques are more universal and applicable to multiple domains, while higher-level techniques tend to be more domain-specific (even though some more general approaches exist [RBGV08]). Yet, even for the discussed low-level techniques there are a number of constraints and limitations which are discussed next.

The field of contour detection has extensively been addressed in the literature as shown in Section 3.3.1. In fact, today it is easy and effective with modern GPUs. It is still challenging, however, to stylize the resulting strokes consistently and continuously using a parametrized curve. If a contour is closed with length L , for example, it would be desirable to have the position on the curve by a value in $[0, L]$. This setup would allow us to compute view-aligned quads along the contour with a parametrization such that the quad could be stylized with arbitrary patterns. The challenge, however, still is to change the parametrization consistently during the interaction and without sudden or large changes to maintain frame coherence. First approaches tackle this problem by searching nearby contour triangles to assign the new value after rotation [LWM15].

As discussed in Section 3.3.2, a whole zoo of feature lines exists, each with their different advantages and challenges. For most the computation is based on local criteria, so the exploration of more global approaches would be interesting. For example, there is poten-

tial to include information theory approaches to detect features in the surface mesh. It may be possible to analyze lines with respect to the information they would add if they would be drawn. Large lines that are distant from each other may add more information to the mesh than lines that are close to each other. Another interesting idea is to analyze smoothing approaches that remove noise without removing essential features from the surface [JDD03; KCL09; HS13; YWQW+14; WYPW+15]. These approaches could be used to generate new feature line techniques or to apply these approaches using existing feature line extraction techniques, without having to assume a perfectly smooth surface.

Some feature line techniques use ambient lighting to determine the placement of lines. Here, it might be interesting to analyze different lighting methods before the feature line technique is applied. Especially exaggerated methods [RBD06] may be promising.

The evaluation of novel feature line techniques should be an interesting point for future work. Cole et al.'s [CGLB+08] study compared hand-drawn images with feature line techniques. It may be that an artist would draw the features of an object differently if the object is well-known. For instance, if the artist should draw a cow model, he or she would likely use fewer lines because the shape of a cow can be recognized from a few contours and features. In contrast, it may be that the same artist would draw more lines when confronted with an unknown object [SNEAF+16]. To better study such questions, eye trackers could be used in observational studies to analyze how the eyes scan an illustrative visualization.

The actual use of illustrative visualization in interactive systems, e. g., for education, engineering, therapy planning, deserves more attention. In this same context, more work is needed to study the illustration of animated surfaces as many existing approaches for line extraction cannot easily deal with them at the interactive frame rates needed for practical applications.

Hatching methods have also been explored in detail. A big challenge is still to apply hatching methods to animated surfaces, in particular when trying to ensure a minimal distance between the hatching strokes. One method was introduced by Lichtenberg et al. [LSHL16] where an image-based approach was employed that used LIC with seed points determined with a *contact region* instead of the normal noise texture. This method ensures a small distance between neighbored hatching strokes, but a minimal user-defined distance cannot be achieved. Another idea to apply hatching strokes on a surface mesh is to determine a global parametrization, cf. Section 3.3.3.3. On the 2D map of the surface, the method by Jobard and Lefer [JL97] could be applied to ensure a minimal distance. A reasonable parametrization might be the *least squares conformal map* by Lévy et al. [LPRM02]. This parametrization is generally not rigid, i. e., not

distance-preserving, however, so one should take care that the distances are encoded in the 2D map.

Stippling methods were also extensively investigated. An interesting extension would be to encode information on the radius of the stippling points, e. g., the distance to a region of interest—this idea was sketched in the work by Ritter et al. [RHDKV+07]. This general approach can also be extended to the drawing of glyphs on the stippling circles to encode information.

As the field of illustrative shading is quite large and the presented methods did not follow a specific common goal. It is thus difficult to judge what is missing in this field. Thus, some ideas of what could be of interest and what could thus be considered for further improvements are mentioned. An interesting illustrative approach by Tietjen et al. [TPBG+08] combined various shading maps, cf. Section 3.3.5—shading by transfer. These shading maps could be arbitrarily extended by other maps. It might be interesting, for example, to combine these techniques with, e. g., *exaggerated shading* [RBD06]. The combination could not only be determined by the user but also automatically. Objects that might be of interest should be illustrated with maps that enforce an emphasis on structures, while surrounded context objects could be illustrated with low-information maps. This idea could also be applied to single objects with an underlying scalar field of interest, e. g., a distance map. In the case of a close distance, the shading map should then be chosen differently from region with a high/distant value.

Lit spheres [SMGG01; FJLS+16] may be extended to take the volume into consideration. Instead of mapping the normals to the sphere and applying the color scheme to the model, one could add information inside the sphere such as a volumetric sphere that contains information. Then, the shading could be defined that first assigns the normals to the sphere and then uses a second scalar field that yields the radius. Overall this approach would provide a volumetric lookup texture that encodes two types of information.

3.7 CONCLUDING REMARKS

Illustrative visualization techniques show potential to convey information and to abstract surface objects. They are suited for simplifying structures, but it may also support depth and shape perception. Various studies confirmed the potential use of illustrative visualization methods in comparison with standard visualization techniques, e. g., Phong shading. Furthermore, different applications exploit illustrative visualization methods to gain insight or to transfer knowledge. The manifold application areas confirm the potential of illustrative visualization techniques. However, they are still rarely used in com-

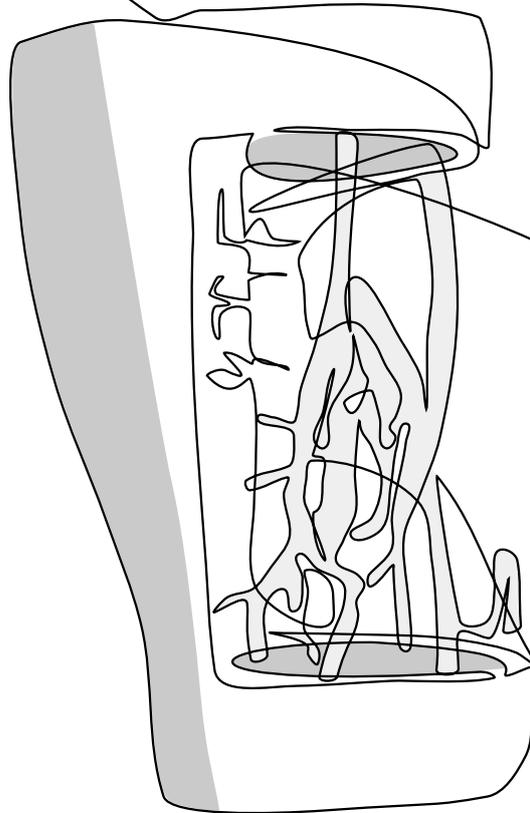
mercially available software. Thus, more exchange between academic research and developers is desirable.

Contours and feature lines give a first impression on the surface mesh and may be a good alternative to other visualization methods. They can be applied if the perception of the spatialness is not paramount. Studies have shown that feature line techniques show also potential to depict the shape, but it strongly depends on the underlying surface mesh. For an improved shape perception hatching strokes, stippling points, or even illustrative shading methods should be employed. Thus, the order of the presented illustrative visualization techniques may reflect the shape perception, but the more certain it is, the more information are used to give that impression. Here, hatching and stippling is regarded as comparable and similar. The question which visualization technique should be applied to a scene with various objects depends on underlying problem. The more information, e.g., color, primitives are used to illustrate the objects the more attention it will receive.

Most of the presented methods were used as an alternative depiction for surface meshes. Alternatively, these methods can be extended to not only improve the spatial or depth perception, but also to encode additional information by the style of the used primitives.

4

MULTI-VOLUME RENDERING



This chapter is partly based on:

Kai Lawonn, Noeska Smit, Katja Bühler, and Bernhard Preim.

“A Survey on Multimodal Medical Data Visualization.”

In: Computer Graphics Forum (2017), submitted.

MULTIMODAL medical datasets consist of multiple scans of the same subject using various acquisition methods. In this way, it is possible to image different tissue characteristics. Several modalities can also be integrated directly into the hardware of a single scanner. Of these hybrid medical imaging techniques, PET/CT is currently the most widespread. More recently, PET/MR became clinically feasible and additionally SPECT/CT is also used in clinical practice. The common denominator in the aforementioned hybrid imaging combinations is that often the radiologic modality (CT or MRI) depicts the anatomy of the patient in high spatial resolution, whereas the nuclear medicine modality (PET or SPECT) depicts functional processes, such as metabolism, in a lower resolution. Additionally, it is possible to combine structural information from a scan with functional information from the same scanner, such as combining an anatomical MRI acquisition with functional MRI information. CT and MRI, however, are able to visualize anatomical information at a much higher resolution, but lack metabolic information. Thus, a combination of both imaging modalities allows visualization and localization of metabolic information in anatomical context. The overall visualization goal in this case is to accurately localize regions featuring abnormal functional values based on their relation to structural anatomical information.

Typically, radiologists or specialists in nuclear medicine examine data in a slice-based fashion consisting of superimposed 2D images with a combined visualization of CT and PET image data, as is shown in Figure 8. Here, physicians browse through the slices and set window and level parameters to adjust the brightness and contrast, and examine structures of interest. Unfortunately, presenting these 2D images in such a way can hinder full and quick analysis of the data. First, the physician needs to check every slice, which may be time-consuming as the number of slices increases. Furthermore, physicians need to mentally fuse all information from these slices to form a correct diagnosis. Therefore, visualization techniques that provide a 3D overview are helpful to see potential abnormalities at a glance. This enables the experts to navigate to these suspicious regions for a more detailed exploration in the slice views.

In recent years, the amount of work focusing on multimodal medical data visualization increased. In 2010, there was an IEEE VIS contest on multimodal visualization for neurosurgical planning [DPLB+11], further highlighting the interest in and importance of multimodal

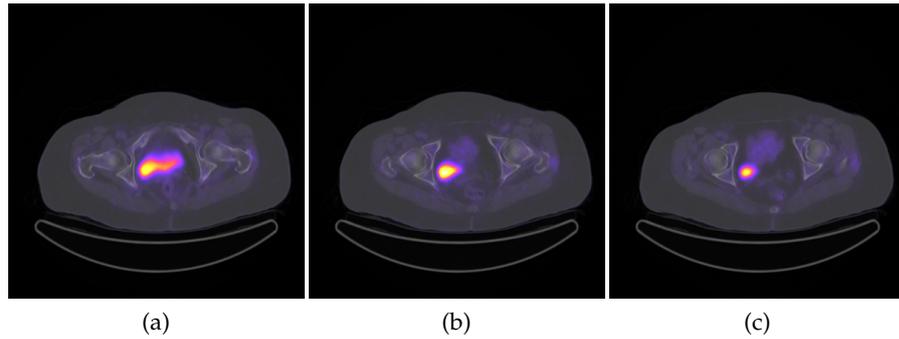


Figure 8: 2D images with superimposed CT and PET data, represented in gray scale and color respectively.

visualization. With this overview a survey of this literature is provided. A number of medical visualization techniques are described in a survey style in the book by Preim and Botha [PB13]. Multimodal visualization, however, is only slightly touched upon as an add-on to registration and is also not considered in other survey articles up until now. In general, concepts used for visualizing medical data can be extended to other multimodal visualization application domains, but there are some specific challenges in dealing with medical multimodal data. Thus, the presented works are of interest to any visualization researcher dealing with multimodal data.

The contributions of this overview are the following:

- An overview of multimodal acquisition techniques and relate this to requirements and challenges.
- A taxonomy of multimodal medical data visualization applications.
- An outlook on open problems in multimodal medical visualization and a perspective on future research directions.

Scope. There have been several surveys on multi-field data visualizations [STSo6; FH09]. In contrast to general multi-field data, measured medical image data is not as ‘clean’ as simulation data, due to the acquisition process. Living tissue is imaged with scanning parameters that favor the patients’ safety over image quality, unlike for instance applications in material testing. Furthermore, several tasks such as searching for metastases, or assessment of infiltration, are unique and therefore also evaluation of multimodal medical visualization that needs to consider such or similar relevant tasks. While multimodal medical image data are available at both microscopic and macroscopic level, this chapter focusses on the macroscopic level, e.g. radiological image data. Interaction is also considered with multimodal data. Due to the complexity of these data, interaction is even more important than in the exploration of single data.

Please note that in this overview the focus is on combining imaging modes acquired by different scanning techniques, and less on combining imaging modes of a single scanner, e.g. MRI T₁ and T₂. A survey focusing on DTI visualization techniques can be found in the recent work by Isenberg [Ise15]. DTI can additionally be used in multimodal applications when it is combined with fMRI or additional imaging modalities. Multimodal data is often acquired from separate scanners, which creates a need for software-based registration to align the volumes. Due to the time between the scans and patient pose differences it is not straightforward to register multiple volumes accurately. This in turn generates additional challenges in dealing with uncertainty in the form of processing errors introduced by the registration process. For an overview of medical image registration the reader is referred to the survey by Maintz and Viergever [MV98] and more recent work on mutual-information-based registration methods by Pluim et al. [PMV03]. In essence, there are many registration techniques that can be employed to align multimodal datasets, such as maximizing an information-theoretic measure, e.g. normalized cross-correlation and normalized mutual information. Furthermore, landmark-based registration can be employed, which requires appropriate crisp anatomical landmarks, e.g., vessel bifurcation points or brain folds. These optimization strategies can be costly, in terms of processing time required, in which case GPU support is necessary [FVWK+11].

While blood flow measurements can be considered functional information, a review of these techniques is out of the scope of this survey (see the recent survey on Cardiac 4D PC-MRI [KBPP+15]).

This overview focusses on three core applications areas, namely visualizations aimed at research, diagnosis and treatment planning. Within these areas, different application subdomains are distinguished. The applications aimed at diagnosis is divided in oncologic and cardiologic visualizations, and the applications aimed at treatment planning is divided in neurosurgery and radiotherapy planning.

Organization. This chapter is structured as follows. In SECTION 4.1 an overview of (hybrid) imaging acquisition the characteristics, advantages and disadvantages of each of these modalities are described. In SECTION 4.2, the typical workflow of a specialist in radiology and nuclear medicine for exploration and analysis of the data is explained. Based on this workflow, requirements for 3D visualizations are derived. In SECTION 4.3 a brief overview of relevant visualization techniques, such as smart visibility techniques and focus-and-context techniques, that can be employed for multimodal visualization is provided. SECTION 4.4 continues with multimodal rendering (Section 4.4.1) and interaction (Section 4.4.2) techniques. In SECTION 4.5, the focus lies on applications of multimodal visualization techniques to real-world clinical data. This section is divided into the application areas these visualizations are developed for:

- Research (Section 4.5.1),
- Diagnosis (Section 4.5.2),
- Treatment planning and guidance (Section 4.5.3).

SECTION 4.6 concludes the survey and outlines unsolved problems and challenges for future research. Finally, SECTION 4.7 concludes with a brief summary.

4.1 MEDICAL BACKGROUND

In this section, an overview of multimodal and hybrid imaging, as well as their applications in clinical practice is provided. The characteristics of each of the modalities involved are summarized, and their advantages and disadvantages and associated visualization challenges are discussed.

4.1.1 *Computed Tomography (CT)*

CT is an x-ray based tomographic imaging technique that creates stacks of 2D cross sectional images that can be combined to 3D volumes. It is especially suitable to distinguish tissues such as bone, water, fat, and the air in the lungs. A contrast agent can be applied to enhance vascular structures. Recently, hybrid scanners such as dual source, or dual energy CT scanners are available delivering a final image, which fuses information from high and low voltage image acquisition performed at the same time (see Figure 9) [KSF11]. Depending on the chosen imaging protocol this technique allows for a differentiation of structures like bone and contrast-enhanced blood vessels.

CT data is especially suited for high quality direct volume rendering, due to its high resolution (in general 512×512 in slice resolution and 0.3-2 mm slice distance), high signal-to-noise ratio and standardized intensity values (Hounsfield Units), allowing the definition of re-usable and task specific transfer functions.

4.1.2 *Magnetic Resonance Imaging (MRI)*

In Magnetic Resonance Imaging (MRI), a scan is made using a powerful magnetic field. In contrast to CT, MRI scanners are highly configurable and provide a large variety of imaging protocols allowing to capture structural as well as functional information. In general, several different MRI sequences, such as T1 and T2-weighted scans, are acquired at the same time, leading to (more or less) co-registered images. Since signal strength relies on the presence of Hydrogen atoms, MRI sequences are emphasizing image contrast in tissue containing

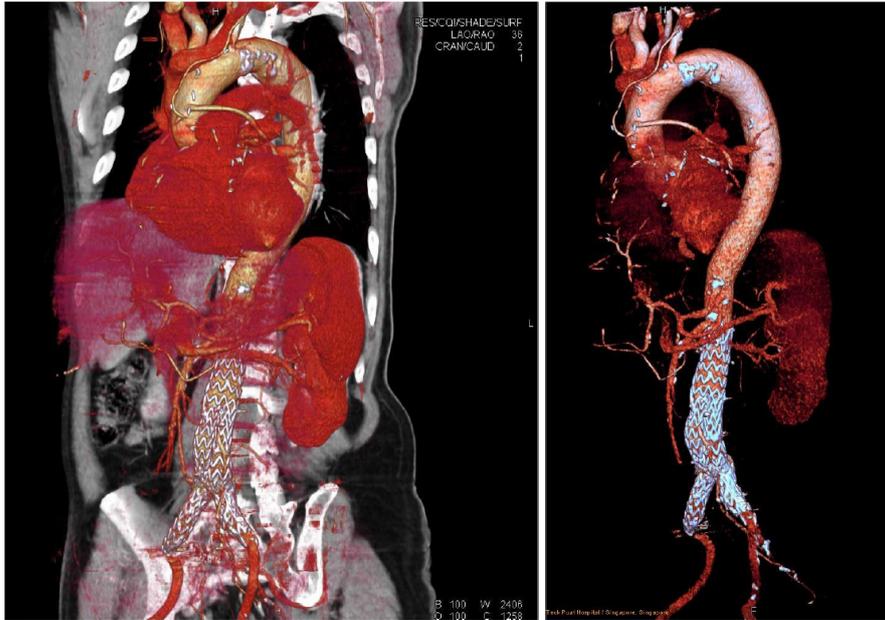


Figure 9: Volume renderings using dual energy CT scan of an aortic stent. On the left an overview shows the anatomical context, and on the right only the aorta and stent are shown. (From [KSF11]).

water. Intensities in MR images are not standardized. MRI data often exhibit an inhomogeneous gray level distribution, requiring careful preprocessing of the data, and intensity values vary depending on scanner vendor and clinic. Therefore MRI data is challenging to visualize.

MRI data can not be easily classified, i.e. a simple 1-D transfer function does often not deliver a clear separation of structures due to the inhomogeneities in the magnetic field. Furthermore, due to the unpredictability of the intensity values, transfer functions are not directly applicable across several datasets without dynamic adaptation [RSHSGoo]. Finally, MR images have generally a lower resolution and lower signal-to-noise ratio than CT images.

Besides, the standard MRI scanning protocols, there are specific MRI sequences and protocols, such as MRSI, DCE-MRI and DTI. In MRSI, spatially localized metabolites in body tissues are measured. DCE-MRI (Dynamic Contrast-Enhanced MRI) is a perfusion imaging technique that measures the perfusion of tissues by blood indicating regions damaged by stroke or infarction as well as characterizing the vascularization of tumors, helping to assess whether they are benign [TBBE+99]. DTI is an extension of Diffusion Weighted Imaging (DWI), that detects the direction of white matter tracts in the brain, which represent connectivity between different areas of gray matter. DTI is used in clinical practice to assess the deformation of white matter by tumors, neurosurgical planning and for (early) diagnosis of brain pathologies such as Alzheimer disease, schizophrenia and

multiple sclerosis [LBMPC+01]. DTI data is often visualized as a scalar field representing the Fractional Anisotropy (FA), using glyphs or fiber tracking [HS15].

Functional MRI (fMRI): Besides using MRI to acquire structural information, it is possible to detect neural activity. fMRI records subtle changes in blood flow in response to stimuli or actions and uses this information to visualize cortical activity. The most frequently employed technique is blood oxygenation level dependent (BOLD) fMRI. By having the subject perform tasks categorized into visual, motor, speech or memory tasks, different functional areas of the brain 'light up' and can be associated to the tasks performed. Additionally, fMRI is used in a research context to improve the understanding of neural networks in the brain even when the user has no task, as is the case in resting state fMRI [VDHP10].

4.1.3 *Ultrasound*

In medical ultrasound, high-frequency sound waves are employed to characterize tissue. Ultrasound can be used both for diagnosis and to guide interventional therapeutic procedures. Due to the nature of the modality, ultrasound is suitable to image soft tissues, such as tendons, vessels and organs, but cannot visualize bone and air, or structures lying underneath these tissue types. Based on the Doppler effect, blood flow in the heart and blood vessels can be detected. Advantages of ultrasound compared to other modalities are that it is cheap, safe, portable and real-time. However, ultrasound is difficult to interpret, due to the low signal-to-noise ratio, artifacts, and the limited field of view. Recent advances in ultrasound technology include 3D ultrasound [VBSH+13], elastography and contrast-enhanced ultrasound using microbubbles.

4.1.4 *Modalities from Nuclear Medicine*

Positron Emission Tomography (PET) relies on the indirect detection of gamma rays emitted by the patient after a positron-emitting radionuclide (tracer) is introduced. Using this technique, metabolic uptake processes can be visualized. While CT and MRI scans can provide detailed anatomical data, PET scans are able to reveal functional information, such as metabolism. A common application of PET scans is to search for metastases and for this, the radioactive substance fluorodeoxyglucose (FDG) tracer is used, a substance similar to glucose. The metastases have higher glucose uptake than normal and specific abnormal metabolic activity can be captured in this way. Besides oncological applications, PET is also used for neurological and cardiological diagnostic purposes. While commonly used, FDG is not the only

available tracer for PET and different tracers may be better suited for specific applications [TCVP+13].

The PET data needs to be attenuation-corrected before visualization. Visualizing PET data in 3D is challenging, since normal metabolic information is also contained in PET, the highest activity measures are not always the most interesting.

Single-photon emission computed tomography (SPECT) is a nuclear medicine tomographic imaging technique that uses radioactive tracer material to detect gamma rays. In this way, it is similar to PET, however in contrast to PET, the gamma radiation is measured directly. SPECT can be used for oncological diagnosis, but also for infection, thyroid or bone imaging. Besides these applications, SPECT can also provide localized function within organs for functional cardiac or brain imaging. SPECT suffers from a lower spatial resolution and contrast than PET [BW13]. Similar to PET, SPECT data is not straight-forward to render in 3D.

4.1.5 Hybrid Scanners

When combining acquisitions from multiple modalities from different scanners, registration problems arise. To simplify the registration process, efforts have been made to integrate multiple modalities into hybrid scanners that combine the best of both worlds in structural and functional imaging. Hybrid imaging scanners combining PET or SPECT with CT are already commonplace in clinical practice, while combining PET, SPECT with MRI is more recent development [Cheo9] for which the first prototypes have been made [PKNS10]. In a preliminary study comparing clinical impact of PET/CT and PET/MRI, PET/MRI imaging outperformed PET/CT and more frequently affected patient management [CRSH+13]. Furthermore, integrated whole-body PET/MRI was found feasible in a clinical setting with comparable reliability to PET/CT for this purpose [DSEB+12]. Since 2014 the very first devices are legally allowed and used in clinical practice, focused on head and neck imaging initially.

PET/CT: While both PET, CT, and MRI provide valuable information by themselves, they can be combined to provide more insight into the exact localization of suspicious metabolic activity. By combining PET and CT/MRI scans, physicians are able to detect (abnormal) metabolic activity using the PET scan and to localize this activity using the CT or MRI scan. This can aid the user in distinguishing which activity is physiologically normal and which is pathological. Since the adaptation of PET/CT is more widespread clinically, this overview focusses on this hybrid imaging techniques in the rest of this section (see Figure¹ 10).

¹ <https://upload.wikimedia.org/wikipedia/commons/c/c4/16slicePETCT.jpg>



Figure 10: A PET/CT scanner. The frontal part of the gantry features a CT scanner, while the rear part contains a PET scanner.

The CT data is also used to perform noiseless attenuation correction on the PET data, thus eliminating the need for an additional PET transmission scan for this purpose. This can reduce the total scanning time by up to 40 percent [TCYH04]. A drawback of PET/CT is that the imaging is performed sequentially, which takes more time and eliminates temporal correlation between the modalities [JWNC+08].

Clinical oncological applications of the combined PET/CT scanner include diagnosing and staging primary tumors as well as localization of metastatic disease in almost any region of the body [BTBK+00]. Further applications include decision making with respect to surgical operability of the tumor or treatment selection. PET/CT is additionally used to determine if cancer has recurred or to determine the difference between scar tissue and active cancer tissue. Especially when monitoring therapy, PET can detect changes in tumors earlier than CT, because metabolic changes occur sooner than anatomical size changes.

For a more elaborate description of PET/CT acquisition, the reader is referred to Townsend et al. [TCYH04].

SPECT/CT: The benefits of combining PET and CT also extend to combining SPECT and CT. As with PET/CT, the CT can provide structural anatomical information for localization, while the SPECT can provide metabolic information. Similarly, the CT can be used to perform attenuation correction for the SPECT. SPECT is less expensive than PET, but suffers from lower contrast and spatial resolution [HH12]. A general description of the technical developments and

future directions of SPECT/CT can be found in the review by Buck et al. [BNZB+08]. A more elaborate discussion on the clinical uses of SPECT/CT can be found in Mariani et al. [MBKK+10].

In essence, CT and PET or SPECT provide complimentary information that needs to be integrated and thus fused visualizations are desirable.

4.2 CLINICAL WORKFLOW AND REQUIREMENTS

In this section the clinical workflow for the analysis of multimodal medical imaging data is examined. Furthermore, an overview of visualization requirements is provided that should be considered when developing a novel visualization technique for such data.

4.2.1 *Clinical Workflow*

Multimodal medical imaging acquisitions are interpreted by radiologists and/or nuclear medicine physicians. Since each of the two disciplines have their own specific skill-set, ideally, two imaging specialists (one from each discipline) would interpret multimodal images together. However, due to availability and workload this is often not possible [BOBMR+07]. Additionally, there are now more and more physicians specialized in both radiology and nuclear medicine in joint departments for radiology and nuclear medicine. While specific protocols may vary per hospital, in general the images are analyzed similarly in radiology workstations of major vendors. The modalities can be separately analyzed in the views that are common for the individual modalities, such as slice-based viewing in the axial, coronal and sagittal direction. For nuclear imaging modalities, Maximum Intensity Projections (MIP) are commonplace in clinical practice. Often these MIP renderings are viewed from the front or the side of the patient. Depending on the clinical indication, more elaborate visualizations can be used, such as polar maps for cardiac SPECT data [LHGT+06].

To visualize both modalities simultaneously, the typical approach in for instance PET/CT data, is to examine the CT images in grayscale with the PET images superimposed using a colormap in a 2D fused view (Figure 8). The combination of a colored overlay of functional information with gray-valued anatomical information is quite effective since brightness and color are different perceptual channels and thus can be processed simultaneously [TG80]. Interestingly, the choice of colors in the functional color map is not standardized and varies per vendor. While most systems map the highest activity region to red, for instance General Electric (GE) workstations map it to blue. Common clinically used visualization combinations of PET/CT data can be seen in the work by Griffeth [Gri05]. The cases presented here are

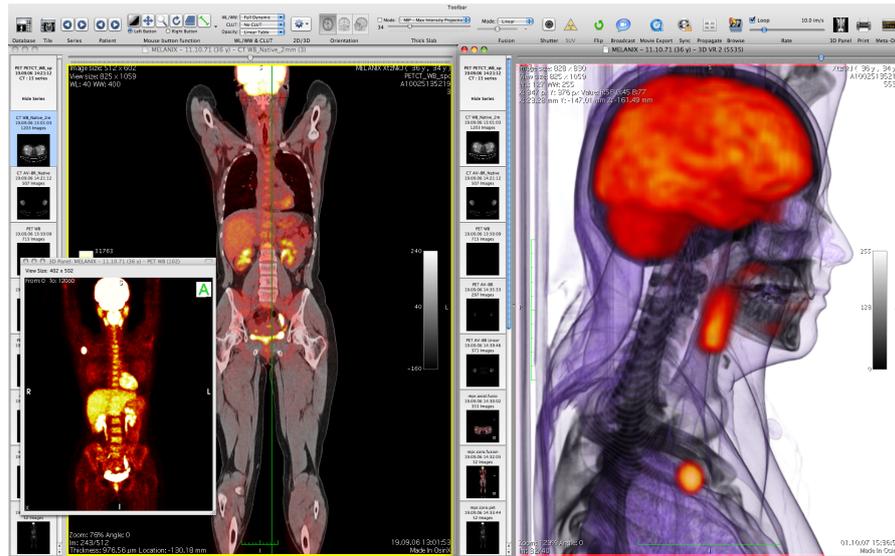


Figure 11: 2D and 3D fused PET/CT images with superimposed CT and PET data rendered in the OsiriX software [RSPR06].

often shown with a frontal MIP of the PET, axial slices of both the PET and CT separately, and finally a fused superimposed slice of the PET and CT combined. The exact configuration of the available views may vary, but typically there will be a central view with supporting linked views or a display featuring linked equally-sized views of the different modalities and slice directions. For oncological diagnosis and treatment planning, also the combination of CT and MRI data may be important [BHWB07]. Since these data are both high-resolution data, their combination is more challenging. The simultaneous exploration of CT and MRI, with either PET or SPECT, is currently not feasible with clinically available software.

While more advanced visualization techniques could be employed to visualize multimodal images, these are not yet broadly available in clinical practice. Examples of these techniques can be found in the OsiriX software, which was specifically designed for navigation and visualization of multimodal and multidimensional data [RSPR06]. Visualization options include Multiplanar Reformation (MPR), surface rendering and direct volume rendering for fused multimodal datasets (see Figure 11).

4.2.2 Requirement Analysis

From a medical visualization perspective, the goals of multimodal medical data visualization include:

- reducing complexity and thus cognitive load,
- enabling, improving, or accelerating decision making, and

- providing tailored visualizations for specific applications.

In diagnosis for instance, examining all 2D slices individually to examine anatomy, pathology and metabolic uptake can become time-consuming and cumbersome. Therefore, besides the traditional 2D approaches, 3D techniques can additionally be used to give an overview of the full datasets at a glance. Existing methods such as maximum intensity projection (MIP), which was first developed for use in Nuclear Medicine [WMLK89], can provide such an overview, but suffer from depth perception issues. In most cases, since the intensities are not comparable, a combined volume MIP can suffer from occlusion by the highest intensity regions.

3D visualization techniques that are suitable for displaying hybrid multimodal data would allow the users to get a quick overview of areas of interest and localize for instance foci of elevated metabolic activity in an anatomical context. When developing a novel visualization technique or application designed for use with multimodal medical data, there are some general requirements that should be kept in mind. The requirements for designing a 3D multimodal medical visualization technique are the following:

1. Visualization parameters should be easily adjustable to fit the needs of the user.
2. In order to be suitable for clinical use, the technique should be fast and interactive, with minimal or no pre-processing required.

When combining functional with structural information for diagnostic purposes, the following additional requirements should be fulfilled [LSPV15]:

1. The technique should show the combination of the two or more modalities in a fused view in which the functional activity of interest is always visible.
2. The technique should relate metabolic activity to nearby anatomical structures for accurate localization.

Such a technique can be used to guide the exploration of the datasets by bringing the attention to regions of interest, after which detailed inspection of these regions can be performed in 2D images. Furthermore, besides aiding diagnosis, 3D visualization techniques could be beneficial for research and treatment planning.

For research purposes, the visualization requirements are similar to the requirements for a diagnostic application, and further specific requirements are application-dependent and should be formulated based on the needs of the domain experts. Surgical treatment planning could benefit from multi-modal 3D patient-specific visualization

to support access planning. In oncologic neurosurgery for instance, an access path to the tumor needs to be planned, taking the location of risk structures, such as arteries, into account. In this neurosurgical context, a visualization application should [BHWB07]:

- Provide high-quality, interactive and flexible 3D visualization
- Offer multimodal visualization for modalities such as CT, MRI, fMRI, PET or DSA.
- Provide interactive manipulation such as simulated surgical procedures, endoscopic views or virtual cutting planes.

In radiotherapy treatment planning, the pathology, i.e., radiation target, should be visualized in the context of the healthy structures at risk to receive undesired radiation damage. The requirements for such an application include [SFNB14]:

- Support for 4D PET and CT data and fusion of these modalities in a 3D view
- Visualization of segmentation data
- Visualization of dose information
- Clipping and/or masking parts of the volume
- No pre-processing and interactive parameter adjustment

These requirements demand advanced visualization techniques to fulfill the presented requirements. In the next section, several visualization concepts are presented that have the potential to be applied in multimodal medical applications.

4.3 VISUALIZATION TECHNIQUES

Due to the nature of multimodal medical data visualization, there are always two or more volumes that need to be visualized. With overlapping extents, each of these volumes may be occluded by the others, and thus no clear insight can be given into the features of each of the modalities simultaneously. Visualization techniques supporting this visualization problem need to incorporate heuristics for assessing the importance of information as well as emphasis techniques to adjust the visualization to the derived importance. Therefore, smart visibility, focus-and-context techniques and other emphasis techniques are crucial when dealing with multimodal medical data. In this section, basic visualization techniques are described to cope with challenges such as dealing with occlusion, improving depth perception and presenting relevant information from multiple sources in an insightful way.

4.3.1 Basic Representation

For multi-modal visualization, often pre-processing is required in terms of registration or segmentation. When combining datasets from multiple modalities using separate scanners, registration needs to be performed to align the volumes. An example of this can be found in the work by Wein et al., who proposed a fully automatic registration method to align 3D ultrasound with CT scans [WKCK+07; Wei+07]. Segmentation is not performed routinely in all clinical applications, but for instance in radiation treatment planning, the target structure (tumor) and all organs at risk are segmented manually. Here, segmentation may be performed as a basis quantification for further analysis, e.g., volume measurements.

Both direct volume rendering (DVR) and indirect volume rendering (IVR) techniques can be applied to visualize multimodal medical imaging data. DVR visualizes the data directly, whereas IVR attempts to reconstruct a triangulated surface mesh representative of the original image data and as such is suitable to define crisp boundaries.

In IVR, to determine the triangulated surface mesh, usually the Marching cubes algorithm is applied to a segmented subset of the image data [LC87]. The method may lead to ambiguities such that correct results cannot be guaranteed. Using the resulting mesh, various shading techniques as well as different visualization methods can be applied. On such surface meshes, it is easier to perform illustrative visualization techniques such as stippling and hatching, than it would be in DVR, and these techniques have potential to be applied in multimodal medical visualization. To determine the surface, a segmentation of the medical image data needs to be performed. Segmentation can be performed by point-based, edge-based, region-based, or model-based approaches. Point-based approaches includes intensity-based methods, e.g., setting a threshold. Thus, there are various algorithms that determine an optimal threshold [Ots79; LCC+01; SKLB+10]. For edge-based approaches the LiveWire [MMBU92; USB92] technique can be applied. Region-based approaches include region-growing methods [SPSP02; BRLK+04] and watershed segmentation [Ser82; HGP05]. Model-based methods attempt to fit a predefined model to the intensities, gradients, curvature and/or texture of an image. Active contours or snakes are mostly used fit in a closed curve to a certain structure [KWT88]. Since discussing alternative methods for image segmentation is out of the scope of this section, the books by Bankman et al. [Bano8] and Setarehdan et al. [SS12] are referred for additional information on this topic.

In contrast to IVR, DVR is not susceptible to segmentation pre-processing. In this way, no pre-determined thresholds are needed and do not risk losing critical aspects of the original data, based on these choices. However, suitable transfer functions still need to be defined.

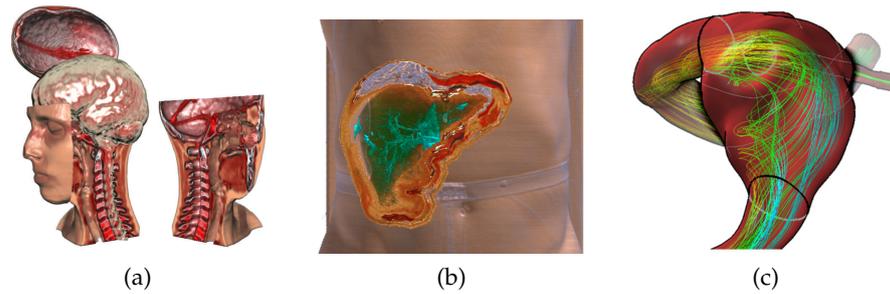


Figure 12: Examples of smart visibility techniques applied to medical data: exploded view of the human head based on CT data [BG06], a cutaway revealing the liver inside the abdomen based on CT data [VKG04], and blood flow showed inside an artery using a ghosted view based on a mesh and simulation data [GNKP10].

4.3.2 *Smart Visibility*

To display multiple structures from different modalities leads to challenging problems. Mostly, the user is interested in one specific structure, e.g., a tumor, or a class consisting of several important structures. These important structures may be surrounded by other objects, resulting in structures that may occlude or overlap each other. In the following, advanced visualization techniques that require the important structures to be segmented are described. This prerequisite is fulfilled in some clinical applications, e.g., in radiation treatment planning and complex neurosurgery planning. A primary concern is how to resolve the occlusion problem such that the most important structures are clearly visible, while still indicating surrounding structures. Hiding the surrounding structures is not appropriate, since they provide spatial context. Whenever a class of interesting objects are occluded, smart visibility techniques can be applied. For the smart visibility term, the reader is referred to Viola and Gröller [VG05] who gave the following definition “expressive visualization techniques that smartly uncover the most important information in order to maximize the visual information in the resulting images”. This visualization technique category comprises several sub-categories and different ways to remedy the occlusion problem. Here, the most prominent ways to resolve visibility problems comprising of cutaways, ghosted views and exploded views are introduced and preliminary guidelines are presented on how they could be applied to multimodal medical data.

4.3.2.1 *Cutaway Views*

Another way to depict the focus object and to separate it from the surrounding context structures is to use cutaway views. Cutaway views employ, for instance, a conic object placed around the focus object. The conic object is aligned towards the camera in such a way that the

camera looks inside the object from the base to the tip. Every non-focus structure that is inside the conic object is cut such that only the focus object inside the cone is visualized. The main advantage is the clear illustration of the main object without occlusion by the surrounding objects. Thus, the user gets a direct view on the focus object. Unfortunately, structures such as vessels around the object, as well as important regions between the focus object and the view point of the camera are lost. In general, cutaways vary from *predefined models* [LHV12] to *focus-oriented objects* [BF08; KTP10]. The algorithmic complexity differs strongly based on the object used.

Predefined models: With predefined models, view-dependent and view-independent approaches can be distinguished. Independently, the predefined model is loaded into the framework and subsequently all structures are tested to see if they are inside or outside the focus object. For instance, a lymph node can be approximated by a cylinder with a suitable size. Here, during the rendering an inside/outside test can be applied to the surrounding (context) objects with the predefined model. In case the context objects overlap with the predefined model, the fragments are discarded; otherwise the fragments are drawn. Based on the view-dependence, the predefined model is either aligned along the view direction of the camera or not. Mostly, view-dependent cutaway techniques are used, where the conic objects are oriented along the view-vector. Predefined models can be used if the focus object has a simple shape. Moreover, it is also only applicable if the model does not vary in time. In case of animated blood flow, predefined models cannot be employed, due to the change of the form and position of the blood flow [LGVP+15].

Focus-oriented objects: The second class of cutaway techniques uses focus-oriented objects. Here, the conic object is constructed based on the focus object and can be arbitrarily shaped. To achieve real time performance, the determination of the conic object and the associated cutaway needs to be efficient. Previous work, such as the work by Viola et al. [VKG04], uses a Chamfer distance transform approximation [Bor86] to the conic object in order to calculate the cutaway (see Figure 12b). Thus, a depth image of the conic surface is calculated, which is then used to cut away regions that are closer to the viewer than the conic object itself. If the depth of the surrounded objects is smaller than the depth of the conic object, the corresponding fragment is discarded, otherwise it is shown. Other approaches were developed to speed up this calculation [RT06]. Here, the cutaway surface function C is defined by:

$$C(p) = \max_{q \in R} (q_z - m \cdot \|q_{xy} - p_{xy}\|),$$

where R contains the focus object, m defines the slope of the cone by $m = \tan^{-1} \theta$, and p_{xy} is the pixel's current 2D location. The key idea is to determine the conic object in the view plane with a jump

flooding algorithm. Usually, the conic object is generated by applying the standard flood fill, an iterative process in which in every iteration a pixel passes its value to its direct neighbors. Contrarily, the jump flood algorithm uses the same approach, but the neighbors vary for each iteration. Thus, a pixel passes its value to its neighbors, which have a certain distance. For each iteration the step size is halved. Cutaway generation was used to visualize animated blood flow in such a way that a clear view on the blood flow within a vessel is always guaranteed [LGP14a].

Application to multimodal data: Cutaway views can be directly employed in multimodal medical data visualization applications without much adaptation. For instance, one modality could be displayed outside the cutaway region while another is displayed within the cutaway. Furthermore, cutaways can be used on both modalities to reveal structures of interest hidden within a larger anatomical structure.

For instance, a cutaway can be used to display hybrid imaging PET/CT data, where high PET activity needs to be depicted. By using a cutaway view in these regions, the high PET activity is always visible and it provides a clear view inside the anatomical CT data and can help guarantee a correct localization of the occurrence, whereas a simple maximum intensity projection cannot give reasonable depth cues. Such a view might be useful not only for diagnosis, but also for patient-doctor and interdisciplinary communication, as well as for treatment planning.

4.3.2.2 *Ghosted Views*

An easy way to illustrate the focus object and the surrounding objects is to use transparency. The main drawback of employing multiple transparent objects is that they may lead to visual clutter. Furthermore, it is hard to interpret the shape and ordering of overlapping regions. To remedy this, ghosted view techniques attempt to adjust the transparency depending on the underlying object. Here, a distinction can be made between: smart transparency techniques and interactive approaches. In smart transparency, the strength of transparency depends mostly on curvature or normal vector information. One approach is to set the transparency very high if the surface is nearly planar and to increase the opacity if the curvature varies. That helps to identify regions where the morphology changes strongly. This was successfully applied for instance to vessel visualization with embedded flow [GNKP10; GLHS+14], an example of which can be seen in Figure 12c. The transparency is set by $F = 1 - \langle V, N \rangle^r$, where V represents the view vector, N the surface normal, and r represents the shininess constant. The superiority of this approach over standard semi-transparent visualization related to cerebral blood flow was confirmed in the perceptual study by Baer et al. [BGCP11].

An additional transparency technique was inspired by suggestive contour lines, a line drawing technique that emphasizes characteristics of the surface. We used this approach to develop a shading approach, based on a suggestive contour scalar field [LGP14a]. Here, the scalar field is used to determine zero-crossings after which positive and negative values are used to shade the object in orange or cyan, respectively.

Application to multimodal data: Ghosted views can be applied to multimodal data visualization in order to remove emphasis of structures that are less important from one or more modalities. A ghosted view can be employed to show an anatomical context in order to localize functional data without occluding the view. For instance, a brain surface can be rendered as a ghosted view, while functional data is rendered in an opaque style. This yields a visualization where the brain surface can be recognized, without occlusion, whereas the functional brain data is directly depicted. The main challenge is to provide a technique that is not disturbing, but provides as much morphological context as possible, while not occluding underlying structures. Conventionally, a simple transparency setting is used which does not provide a strong shape impression.

4.3.2.3 Exploded Views

To illustrate a focus structure and representing the numbers of layers around the focus region, a carefully arranged exploded view, such as the one used in technical illustrations, may be suitable. Though it is not directly useful for clinical medicine, when this technique is employed by a skilled illustrator, it may result in valuable educational illustrations. To uncover a hidden focus object, either the visual representations can be changed or the spatial relations can be adjusted in order to provide more insight. The exploded view approach makes use of the latter to 'break' the object apart and order it differently in the region surrounding the focus object. This technique was first applied on technical illustrations [LAS04; LACS08], before it was adapted to visualize medical data sets [BGo6; RVBB+08] (see Figure 12a). When using exploded view techniques, the original spatial relations may be lost. Furthermore, it may happen that important structures around a focus object, e.g., vessels around the tumor, are relocated, or even worse, break apart.

Application to multimodal data: For reasons outlined above, exploded views are not directly suitable for multimodal medical data visualization. One could consider to have one modality rendered in a traditional style with the information from an additional modality arranged around it. However, since the spatial relations, which are crucial in the interpretation of medical data are lost, these techniques are not promising from this perspective.

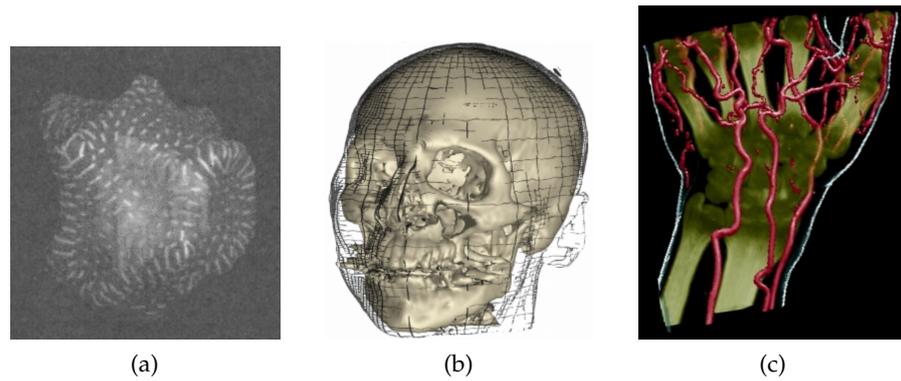


Figure 13: One of the first focus-and-context illustrations with hatching patterns used for the isosurface of radiation dose along with a tumor to be destroyed [IFP96] and a skull where the skin is illustrated using lines based on CT data [TCoo]. Furthermore, a focus-and-context visualization applied to a hand data set with arteries [HMIBGo1]

4.3.3 *Focus-and-context visualization*

For focus-and-context visualizations, it is assumed that at least the focus object is segmented. Even though a lot of techniques exist that provide insight into the focus object, the question is how to represent them in the context of the surrounding structures. A straightforward approach is to use the same shading technique on all structures, which may lead to a hindered perception of the focus object by the context objects. Thus, the main goal of the *focus-and-context* visualization is to illustrate the focus object so that it is perceivable, and simultaneously illustrating the surrounding context structures without distracting from the focus. Techniques vary from using unsaturated colors for shading, to different shading methods per object class, to even varying entire rendering concepts.

Line drawings. One of the first line drawing that appeared in combination with focus objects was presented by Interrante et al. [IFP96], see Section 3.3. They used a hatching technique for the surrounding objects which are represented by the isosurface of radiation. In this case, anatomical data is combined with a scalar field resulting from the simulation of dose distribution, a standard approach for radiation treatment planning. The lines were aligned along the principle curvatures directions for an improved shape perception, whereas the inside focus object was simply shaded (see Figure 13a). Another line drawing technique that appeared in combination with focus objects was presented by Treavett and Chen [TCoo]. They developed line drawings that can be applied on volume data sets and furthermore, they showed how it can be used to illustrate the surrounding context object, e.g., the skin, with the focus object, e.g., the skull (see

Figure 13b). Again in the context of volume rendering, Hauser et al. [HMIBGo1] used combined visualization techniques to illustrate different objects. For instance, bones are rendered with direct volume rendering, surface rendering is applied to the vessels, and contours indicate the skin (see Figure 13c). Lum and Ma [LMo2a] presented a way to render a dataset with non-photorealistic techniques efficiently. Several illustrative techniques could then be applied with interactive framerates (see Figure 14a). It was shown that line drawing concepts can successfully be used for focus-and-context visualization [TIP05]. The core idea here was to provide the flexibility to combine surface-, volume-data, and line drawing for a sparse visual representation of context objects. Here, not only line drawings, but illustrative visualization techniques in general were used to visualize context objects. The key idea of illustrative visualization is to provide a simplified and expressive depiction of a scene or problem [Law15]. Here, details are omitted purposefully in order to convey only the important information.

Indeed these different rendering styles may be used for multimodal visualization, e.g. contours representing boundaries of objects in one modality may be overlaid to the visualization of a second dataset, where no contours are used.

Tietjen et al. combined and tested [TIP05] various rendering techniques. While their technique was not shown on multimodal data, it is applicable for feature fusion. They confirmed that line drawings can be a meaningful extension to existing visualization techniques in an evaluation. Later, Tietjen et al. [TMSP+] used slice-based visualizations in combination with 3D illustration techniques (see Figure 14b) for an example of focus-and-context visualization in the medical context. In a similar way, a slice or a slab of CT/MRI data may be integrated in a SPECT/PET visualization or vice versa. The slice/slab may be even enhanced with contours from a segmented object. For PET/CT, Kim et al. [KEFo7] used different rendering techniques to illustrate the PET enrichment and the anatomical context by using the CT data. Various rendering techniques were also used in the context of simultaneously rendering anatomical and functional brain data [JBBS+08]. For an overview of several focus-and-context approaches, the reader is referred to Cohen [Coh06] and Bruckner et al. [BGMP+10]. Wan et al. [WH10] proposed an importance-based accumulated transparency modulation to visualize focus objects. Their approach is a generalization of the maximum intensity difference accumulation technique. Another importance-based direct volume rendering approach was presented by Pinto and Freitas [MPF10]. Furthermore, they combined standard volume rendering approaches with illustrative visualization techniques (see Figure 16). For focus-and-context visualization with PET/CT data, Bramon et al. [BRBB+13b] suggested visualization techniques to emphasize the PET activity. An-

other approach to explore and visualize interesting regions was proposed by Abellán et al. [ATGP13]. They showed their approach on MR/PET data and used illustrative visualization techniques to convey information.

Lens-based visualization. Focus-and-context visualization can also be applied interactively. One way to approach this, is to place a region on the surface that allows the user to look inside the data and reveals everything underneath the surface. This idea is called *magic lens* as it can be used to look through objects and to reveal interesting parts [BSPB+93]. The magic lens can be applied in many situations, ranging from whole body that uses a lens to reveal organs, to vessels where a lens depicts special properties of the blood flow [GNBP11]. A special application in medicine is to employ the vertical symmetry of the human body and propagate the movement of the lens in one part of the body to the other part, thus supporting the comparison between a suspicious region and the analogue region on the other part of the human body. Such lenses are employed, e.g., in the exploration of brain perfusion data, where one perfusion parameter is shown in the lens region and the other in the background. For 3D ultrasound, Schulte zu Berge et al. [BBKN14] presented a framework where the user can set various predicates such that interesting regions are highlighted. Increasing the slider for the vessels makes it more visible than the surrounding structures (see Figure 15a). In the context of vascular models and tumor visualization, Lawonn et al. [LLPH15] presented a way to depict the vessels with line drawings techniques whereas the tumor is visualized with diffuse shading (see Figure 15b). Furthermore, lenses can also be employed to control which source is displayed, e.g., in the lens region PET (or PET and CT combined) and outside only CT. Thus different layers may be shown interactively [KSWo6].

The distinction between smart visibility and focus-and-context visualization is not that clear. Focus-and-context visualization techniques are a subcategory within the broader category of smart visibility approaches. The main goal of smart visibility is to provide insights into interesting regions or objects and in this way, the same can be said about focus-and-context techniques. However, this goal can be also be achieved by adding 2D image planes in a 3D visualization representation [HE98; BHWV+07]. In this case, there is no focus object involved, although it is still clearly a smart visibility technique. Thus, focus-and-context visualization is considered as a subgroup of smart visibility techniques, which is a broader category including cutaways, ghosted views or exploded views, as well as techniques in which there is no focus object.

Application to multimodal data: Focus-and-context techniques may be applicable to multimodal medical data and can be employed in two main ways. First, both modalities may be combined and a region

of interest within the datasets can serve as a focus area, while the context is rendered in a different style. Secondly, one modality with a higher importance can be rendered as the focus dataset, while the other is rendered to provide context. The rendering style required strongly depends on the application. For instance, if the focus region consists of a tumor and the physician is interested in the distance to nearby vessels, which serves as the context object, then the morphological vessel structure may be not as important as an indicated position of the vessel structures. In this case a simple contour may be sufficient that shows only the outline and allows for a fast recognition of the spatial extent. But in case of a liver tumor, the spatial impression and the shape of the liver is important for the treatment planning. In this case, a more advanced approach needs to be applied to illustrate the tumor in a reasonable way while simultaneously depicting the liver surface with important spatial and morphological features.

4.3.4 *Summary*

There are different channels, with which information, e.g., color, or orientation of elongated objects are perceived. Thus, information can be perceived simultaneously if it is encoded in different channels. The major theory here is the Feature Integration Theory, developed by Anne Treisman [TG80]. Multimodal visualization applications may benefit from the use of different visualization techniques that can be easily perceived simultaneously, e.g., color is only used for depicting one dataset, and line drawing techniques for the other. In order to represent multiple modalities in a single rendering, smart visibility techniques are crucial to prevent visual clutter and occlusion problems. Specifically, multimodal medical visualization can benefit from deciding which dataset has priority, i.e., the "focus dataset", and which is the context dataset. In multimodal visualization design, one could consider indicators of importance of a modality, such as the size and resolution, which could for instance be indicative of which dataset is best used for anatomical detail. Then one could think about which techniques are primarily suitable to show the focus dataset, likely revealing more details, and which to show the context dataset. This approach forms a bridge between single data focus-and-context techniques and focus-and-context in the context of multimodal data. There may be also situations where both volumes have the same importance, e.g., CT and MRI, and in these situations, it might be more appropriate to consider a region within both datasets the focus area, while the region outside is rendered as context. In conclusion, visualization techniques that have been developed for single modality visualization can be adapted to multimodal data, but this requires careful

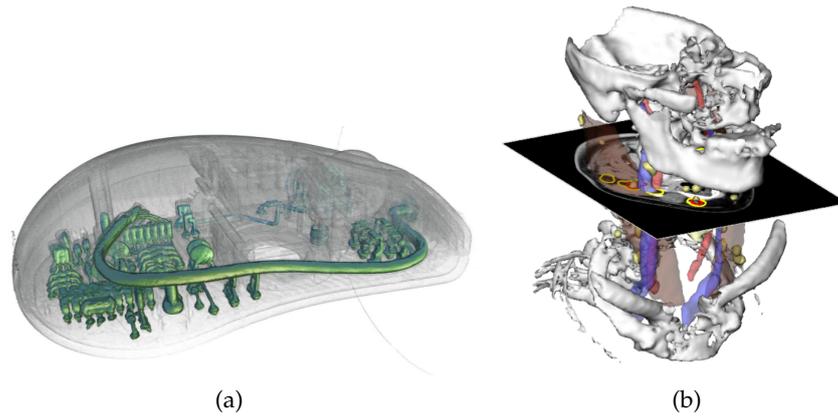


Figure 14: A focus-and-context visualization applied to a CT scan of a computer mouse to illustrate the insides [LMo2a] and a CT scan of the neck revealing lymph nodes and vascular structures with a slice of imaging data [TMSP+].

consideration of the specific requirements and in some cases extensions of existing techniques.

4.4 MULTIMODAL RENDERING AND INTERACTION TECHNIQUES

In this section, rendering and interaction techniques for multiple or fused volumes are discussed. First, recent contributions on rendering aiming at fusing multiple volumes together at different points in the rendering pipeline are described. After this, a discussion on interaction techniques that are designed for the use with multimodal medical data, including clipping and specialized transfer functions is introduced.

4.4.1 *Rendering Techniques*

In this subsection an overview of rendering techniques suitable for rendering two or more volumes simultaneously is provided. Although, Drebin [DCH88] and Levoy [Lev88] are often cited as the developers of the first volume rendering approach, actually Höhne [HB86] was the first who rendered a volume based on CT data. As a follow up, Höhne [HBTR88] were the first who developed a multimodal visualization technique visualizing MRI combined with MRA (see Figure 17). Cai and Sakas presented a seminal work that allows rendering of multiple volumes [CS99]. They presented a method to achieve a reasonable mix of opacity, color, and illumination. Furthermore, they defined three levels of volume intermixing:

- Illumination model level intermixing
- Accumulation level intermixing



Figure 15: Focus-and-context visualization of a 3D ultrasound dataset where the arteries are shown [BBKN14] and a visualization where vessels and a tumor, reconstructed from CT data, are illustrated with line drawings [LLPH15].



Figure 16: Focus-and-context visualization based on a CT scan of the hand and wrist. The bones are shown by applying a cutaway technique [MPF10].

- Image level intermixing

In illumination model level intermixing, the opacity and intensity at each sample point is calculated directly from a multi-volume illumination model, instead of mixing several opacity and intensity values. This type of volume intermixing is the most complex, but also the most realistic. In accumulation level intermixing, the intermixing is performed during ray sample accumulation by mixing opacities and intensities from different volumes. This method requires changes to the rendering pipeline and the time consumption can be high, but provides correct depth cues. The simplest intermixing method is image level intermixing, where the two rendered images are merged on the pixel level. For this, no changes need to be made to the rendering pipeline, but there is a lack of correct depth cueing. This approach

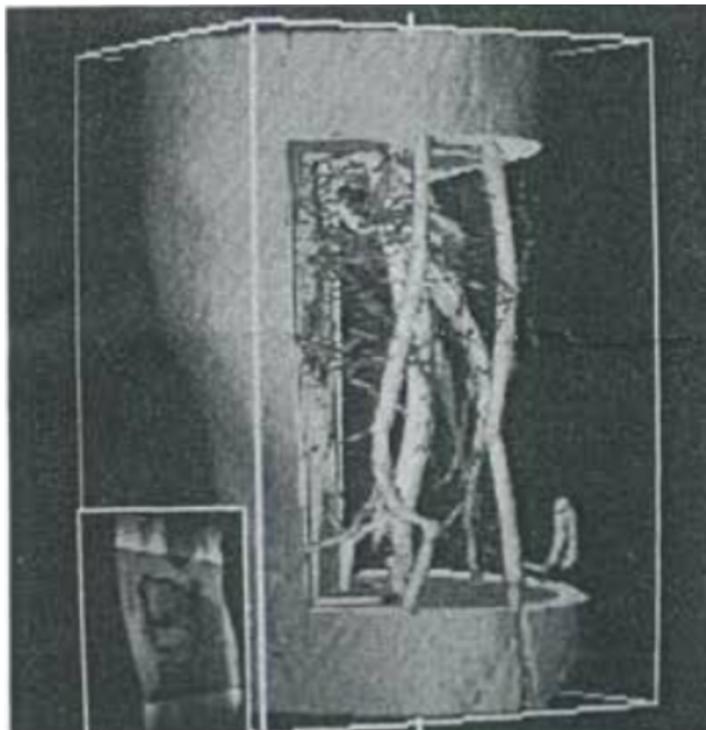


Figure 17: First multimodal volume rendering technique by Höhne et al. visualizing vascular MRA and anatomical MRI data simultaneously [HBTR88].

was successfully applied in the context of radiotherapy treatment planning, combining CT, dose and segmented object volumes.

Wilson et al. presented a visualization approach to depict different image modalities together at interactive framerates using different rendering styles [WLM02]. To render the result efficiently, they used the GPU and demonstrated their results with two case studies. One of the case studies featured a multimodal mouse data set combining PET and MRI acquisitions. Ferre et al. discussed strategies to visualize multimodal volume datasets using direct multimodal volume rendering (DMVR) [FPT04]. Furthermore, they stated requirements and proposed five rendering methods, that differ in the step of the rendering pipeline at which the fusion is performed: property -, property and gradient -, material -, shading - and color fusion. They used SPECT and MRI data to present a combined visualization and evaluate the suitability of the five methods for this data. Another technique to render multimodal volumes that additionally solves depth cuing issues was presented by Hong et al. [HBKS05]. They conducted experiments and could show that their method can distinguish the depth of overlapping regions and produces rendering results faster than conventional software using High Level Shader Language (HLSL) for efficient volume fusion. They show the results of their method on PET and CT images in a performance evalua-

tion. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively, focused on functional brain images [RTFE+06]. Subsequently, they introduced dynamic shader generation for multi-volume ray casting [RBE08]. In an accumulation-level volume rendering context, their system automatically generates different shaders per volume from the configuration of an interactive abstract render graph. Brecheisen et al. employed a framework to render different volumetric datasets together [BBPHRo8]. They used depth peeling to visualize overlapping volumes that can be intersected with an arbitrary number of geometric shapes. Kainz et al. proposed a framework for rendering multiple volumes on the GPU [KGBH+09]. Theirs was the first framework for multi-volume rendering that still provided interactive frame rates while rendering over 50 arbitrarily overlapping volumes. Recently, Sunden et al. introduced a volume illumination technique specifically designed for use with multimodal data [SKR15]. They proposed a new *light-space-based* volume rendering algorithm, that employs illumination-importance metrics to compress and transform multimodal data into an illumination-aware representation. Their method was applied to CT and MRI data, as well as microscopy and simulation data and evaluated based on the quality and performance. Lindholm et al. [LFSB+15] presented a rendering algorithm for hybrid volume-geometry data. They combined volume data and geometry-based data, i.e., a mesh. They applied their approach to proteins, vessels with blood flow, and DTI.

Additional related work can be found in the state-of-the-art report on the visualization of multivariate scientific data presented by Fuchs and Hauser [FH09]. Schubert and Scholl provide a performance and perceptual comparison of GPU-based multi-volume ray casting techniques [SS11]. Furthermore, they presented an overview of visualization techniques that use data intermixing approaches and direct volume rendering methods that use ray casting. They add Classification Level Intermixing to the three levels defined by Cai and Sakas [CS99], which mixes volumes by linear combination of the sample values at the beginning of the rendering pipeline. For a general overview of large-scale volume visualization techniques, the reader is referred to the state-of-the-art report by Beyer et al. [BHP15].

In the area of medical image fusion, Gupta et al. [GRBo8] presented a measure to evaluate PET/MRI image fusion. They could confirm the usefulness of their approaches in different experiments. For this, they used an entropy measure to combine the interesting parts in PET and MRI, but this approach also was applied to CT combined with MRI. Lindholm et al. introduced fused multi-volume direct volume rendering using a BSP, aimed at medical volume rendering of multimodal data [LLHY09]. Lindholm et al. provided a GPU-based ray casting approach to visualize intersecting volumes with regard

to an efficient depth sorting of the resulting fragments. Bramon et al. fused image modalities and applied their technique to CT, MRI and PET data [BBBR+12]. This information-theoretic framework automatically selects the most informative voxels from two volume data sets. They evaluated the potential of their technique with medical experts and concluded that it is potentially useful for planning radiotherapy, treatment monitoring, and planning brain surgery. In their evaluation, they proposed several information maps and fused data sets and had the experts vote on the quality. Their method performed well in differentiating between bone and cerebral tissue, as well as between morphological and functional data. More recently, Kim et al. introduced a slab-based intermixing method for fusion rendering of multiple medical objects [KKLS+16].

The reader is referred to James and Dasarathy [JD14] and Galande and Patil [GP13] for a more extensive discussion about image fusion.

4.4.2 *Interaction Techniques*

In this subsection several interaction techniques that were successfully used in the context of multimodal medical data are described. Interaction techniques primarily serve to adjust which portions of the dataset are visible. Visibility may be adjusted with clipping and cutting on a geometrical basis and with transfer functions on the basis of attribute values, e.g., intensity values, gradient magnitude or curvature.

Clipping and cutting. Hastreiter et al. employed clipping planes per volume to cut the separate modalities in different ways [HE98]. This can be used to look at the slices while simultaneously analyzing the 3D visualization. Furthermore, they combined CT and MR images in a single view. Fairfield et al. proposed a custom clipping solution for co-registered MRI and CT, which they refer to as ‘curtaining’n [FPJT+14]. This allows the user to define a clipping region in which the other modality will be shown. Manssour et al. propose a method to visualize inner structures in multimodal volume data employing cutting and data intermixing [MFO+02]. They applied it to the skull where cutting planes were used to analyze the brain. More advanced ways of clipping data were developed by Weiskopf et al. via per-fragment operations in texture-based volume visualization [WEE02]. While their method was applied to a single modality, their technique was successfully employed by Rößler et al. to clip the brain based on regions in an atlas of the human brain combined with fMRI data [RTFE+06]. More recently, a membrane clipping approach was proposed that avoids cutting through features in the data [BBBV12]. While it was not directly applied to multimodal datasets, it could be used in such a way. Furthermore, their technique

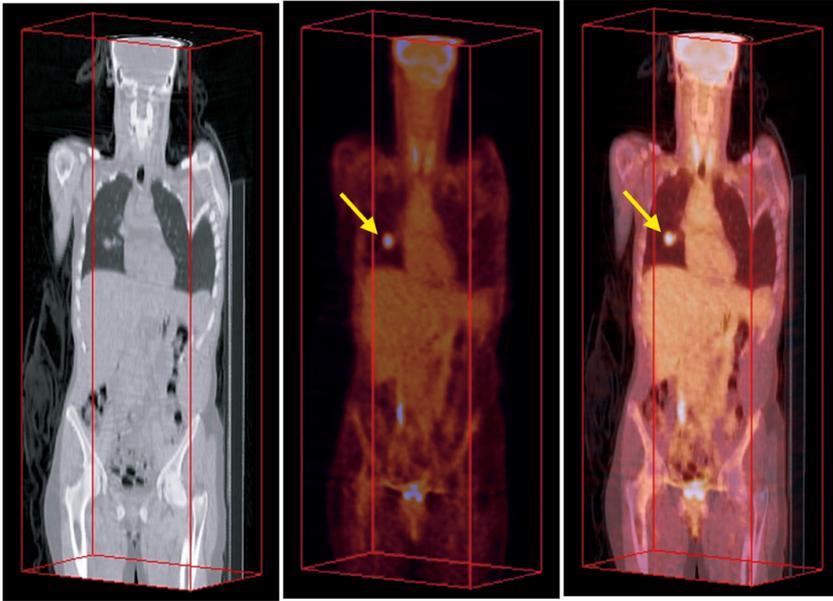


Figure 18: PET/CT data are mixed to visualize the volume in a single view [KEF07].

features slab rendering which is a trade-off between slice-based viewing and full 3D rendering.

Transfer functions. To tackle occlusion problems, transfer functions are essential and have been designed for multi-modal data. 2D transfer functions, introduced by Levoy in 1988 [Lev88], are an essential method for multimodal visualization, since they allow the user to emphasize boundaries and details within structures. Kniss et al. applied a multi-dimensional transfer function to multivariate weather data that allows experts to readily identify specific zones for their analysis task [KHGR02]. They presented a case study in which they explore the utility of multidimensional transfer functions for the visualization of multivariate fields. Kim et al. introduced a dual-lookup table for PET/CT data such that medical experts can set different transfer functions for every volume in a single view [KEF07]. They provided techniques such that the volumes can be merged properly with different transfer functions, see Fig. 18. A general information-based approach for transfer functions was introduced by Haidacher et al. [HBKGo8]. They provided a transfer function space to visualize certain tissues of multimodal data. The method was successfully applied to PET/CT data of the brain. Maciejewski et al. proposed a clustering technique to generate 2D transfer functions [MWCE09]. Their approach yields a clustering in the 2D histogram, which can then be adjusted by the user to get appropriate 3D visualization results. While this is an advanced and elegant solution, the necessary type of interaction limits users to visualization experts.

Furthermore, Bramon et al. introduced an information-theoretic observation channel for volume visualization, which quantifies the information transfer between the source data set and the rendered image [BRBB+13a]. With these quantities the authors provide the possibility to select informative viewpoints, to design transfer function, and to position the light. Joshi et al. introduced new interaction techniques to explore and visualize multimodal data [JSVS+08]. Their technique allow for a precise control of the shape of the region in the brain that can be used to gain insight into the data during exploration and surgery. For their approach they used SPECT/MRI and fMRI data. Additionally, Ropinski et al. proposed interaction techniques to show a close-up of an interesting region in the context of the rest of the data [RVBH+09]. They applied their technique to PET/CT data registered with MRI. More recently, Haidacher et al. proposed an approach for analyze a volume based on multimodal surface similarity [HBG11]. With this, the similarity space generated can be used for isosurface selection in applications like Dual Energy Computed Tomography (DECT). The similarity map can be used to set different isovalues for, e.g., DECT, such that inner structures can be set opaque whereas outer structures can be set as transparent. Correa and Ma [CM11] employed visibility-driven transfer functions to illustrate important structures in comparison to the surrounded regions. Their approach provides the user with graphical cues that inform about the contribution of particular scalar values to the final image. Furthermore, they presented a semi-automatic transfer function design that solves an energy minimization problem, such that the visibility of an initial opacity transfer function that provides the desired importance is maximized.

4.5 APPLICATIONS

In this section, an overview of recent multimodal medical visualization papers are provided that are application-oriented. Furthermore, this section is subdivided in applications developed in a research-oriented, diagnosis or treatment planning and guidance context. For diagnostic purposes, applications developed for cardiology and those aimed at oncology are distinguished. The treatment planning and guidance applications are further categorized into neurosurgery and radiotherapy planning.

4.5.1 *Medical Research*

Multimodal medical data visualization applications were developed in a research context related to vascular pathology development as well as neuroscience.

Table 2: Table of references sorted according to the application type (Type), area (App), and to the acquisition scanners (hybrid scanner, single scanner and multiple scanners). It is stated if any segmentation (Seg) is needed, if the approach was evaluated (qualitatively (Ql), quantitatively (Qt), with domain experts (D) or non-domain experts (N) or the performance (P), and what types of image modalities are used. The visualization techniques that are used are mentioned, subdivided into cutaways (CA), ghosted view (GV), focus-and-context (FC), and illustrative visualization (I-Vis).

References	Type	App	Scanner	Seg	Eval	Image Modality								Vis technique			
						CT	MRI	fMRI	DTI	PET	SPECT	US	CA	GV	FC	I-Vis	
[DHSK+13]	Res	Vasc	Hybr	✓	QID	✓	×	×	×	✓	×	×	×	×	×	×	×
[RHRS+09]	Res	Vasc	Hybr	×	QIN	✓	×	×	×	✓	×	×	×	×	×	×	×
[ESMH+05]	Res	Neuro	Mult	✓	×	×	✓	×	✓	×	×	×	✓	✓	✓	✓	×
[DMLB12]	Res	Neuro	Sing	×	QID	×	✓	✓	×	×	×	×	×	✓	✓	✓	×
[JNMS+09]	Res	Neuro	Sing	×	QIN	×	✓	✓	×	×	×	×	×	×	×	×	×
[NOEH+10]	Res	Neuro	Sing	×	QtN	×	✓	✓	×	×	×	×	×	×	×	✓	×
[RTFE+06]	Res	Neuro	Sing	×	QtP	×	✓	✓	×	×	×	×	✓	✓	✓	✓	×
[SZV01]	Res	Neuro	Sing	✓	QID	×	✓	✓	×	✓	✓	×	×	×	×	×	×
[SZPR+97]	Res	Neuro	Mult	✓	QID	×	✓	×	×	×	✓	×	×	×	×	×	×
[JKEF+13]	Diag	Onco	Hybr	✓	QtP	✓	×	×	×	✓	×	×	×	×	✓	✓	×
[KEFo7]	Diag	Onco	Hybr	×	×	✓	×	×	×	✓	×	×	×	×	✓	✓	×
[LSPV15]	Diag	Onco	Hybr	×	QID	✓	×	×	×	✓	×	×	✓	✓	✓	✓	✓
[VNØU+08]	Diag	Onco	Mult	○	QIN	✓	×	×	×	✓	×	×	✓	×	✓	✓	✓
[HSFM+08]	Diag	Cardio	Mult	✓	QtD	×	✓	×	×	×	×	×	×	×	×	×	×
[OKGH+06]	Diag	Cardio	Mult	✓	×	✓	✓	×	×	×	×	×	×	×	✓	×	×
[TBBV+07]	Diag	Cardio	Mult	✓	×	×	✓	×	×	×	×	×	×	×	×	×	×
[KGSAY+14]	Diag	Cardio	Mult	✓	QtD	✓	×	×	×	×	×	✓	×	×	✓	✓	×
[BHWB07]	Trea	Neuro	Mult	○	QID	✓	✓	✓	✓	✓	×	×	×	×	×	✓	×
[BBMN+07]	Trea	Neuro	Sing	✓	QID	×	✓	✓	×	×	×	×	×	×	✓	✓	○
[BLEL+13]	Trea	Neuro	Mult	✓	QID	✓	✓	×	×	×	×	×	✓	✓	✓	✓	×
[BJHS+09]	Trea	Neuro	Sing	✓	×	×	✓	✓	✓	×	×	×	✓	✓	×	✓	✓
[DPLB+11]	Trea	Neuro	Sing	✓	QID	✓	✓	✓	✓	×	×	×	✓	✓	✓	✓	×
[JBBS+08]	Trea	Neuro	Sing	✓	QID	×	✓	✓	×	×	×	×	✓	✓	×	✓	✓
[JFSG+00b]	Trea	Neuro	Sing	✓	QID	×	✓	✓	×	×	×	×	×	×	×	×	×
[JSVS+08]	Trea	Neuro	Mult	✓	QID	✓	✓	✓	×	×	✓	×	✓	×	✓	✓	×
[KNST+12]	Trea	Neuro	Mult	✓	QtD	✓	✓	×	×	×	×	×	×	×	✓	✓	×
[MSEN+06]	Trea	Neuro	Mult	×	QtP	×	✓	×	✓	×	×	×	×	×	✓	×	✓
[NTSW+10]	Trea	Neuro	Sing	✓	QID	×	✓	×	×	×	×	×	×	×	×	✓	×
[NMMW+04]	Trea	Neuro	Mult	✓	QID	✓	✓	×	×	×	×	×	×	×	×	×	×
[RRRPO8]	Trea	Neuro	Sing	✓	QID	×	✓	✓	✓	×	×	×	×	✓	✓	×	×
[RSHPo8]	Trea	Neuro	Mult	×	×	×	✓	×	×	×	×	×	×	×	×	✓	×
[SKGH+98]	Trea	Neuro	Mult	✓	QI/QtD	✓	✓	×	×	×	×	×	✓	×	○	×	×
[WRDW+11]	Trea	Neuro	Sing	✓	QID	×	✓	×	×	×	×	×	✓	×	✓	✓	×
[NRSK+14]	Trea	Radio	Sing	○	QI/QtD	×	✓	×	×	×	×	×	×	×	○	×	×
[SFB14]	Trea	Radio	Hybr	✓	×	✓	×	×	×	✓	×	×	×	✓	✓	✓	×

PET/CT scans are often requested in cases related to clinical oncology for initial cancer staging or a follow-up during or after treatment [BOBMR+07]. However, Ropinski et al. [RHRS+09] visualized mouse aorta PET/CT scans in a medical research-oriented application related to the formation of plaque. They propose a linked multi-view approach, using a specialized straightened multipath curved planar reformation combined with a multimodal vessel flattening technique (see Figure 19). In a follow-up work, Diepenbrock et al. additionally looked into PET/CT visualization of mice arteries [DHSK+13]. They use the vessel wall extracted from the CT scan to perform a normalized circular projection which allows the user to judge PET signal distribution in relation to the deformed vessel.

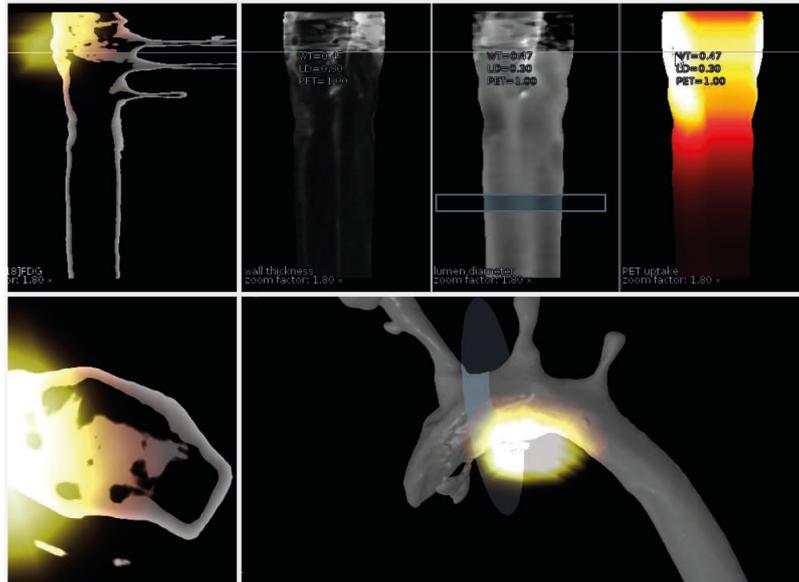


Figure 19: Multiple linked views allow for comparative visualization of vascular PET/CT in mice [RHRS+09].

Nguyen et al. used an approach to visualize and interact with real-time fMRI data in a neuroscience research context [NOEH+10]. They treat the fMRI signal as light emission and render it in the context of a patient-specific high resolution reference MRI scan. Using this technique, the brain glows and emits light from active functional regions with a 2 second delay of the measured fMRI signal. In the work by van Dixhoorn et al., the focus is on examining whole brain functional network connectivity at a voxel-level [DMLB12]. They visualize the correlation of the functional activity, where fMRI time-signals at each voxel are correlated with every other voxel in the brain to determine functional connectivity. Therefore they propose an application for the interactive visual analysis of this high resolution brain network data, both in a linked matrix representation as well as in its anatomical context based on MRI in a GPU raycasting framework.

4.5.2 *Diagnosis*

Visualization techniques aimed at improving diagnostic value of multimodal medical data have been primarily developed in the fields of oncology and cardiology.

4.5.2.1 *Oncology*

Among the oncological applications, Kim et al. introduced a dual-lookup table specifically designed for use with PET/CT data such that medical experts can set different transfer functions for every volume in a single view [KEF07]. Jung et al. employed a novel visual-

ization approach by integrating a visibility-driven transfer function specifically for PET/CT data [JKEF+13]. Furthermore, they provided an intuitive region of interest selection tool for further exploration. Lawonn et al. recently developed an illustrative technique for focus-and-context rendering of PET/CT data [LSPV15], see Chapter 5. In their approach, the functional information from the PET data is used as the focus, while the CT provides anatomical context using dynamic cut-away views. Viola et al. presented illustrative techniques to visualize liver ultrasound combined with CT information serving as anatomical context [VNØU+08]. They aim to shorten the long learning curve for ultrasound practitioners in training, as well as assist the interpretation of liver examinations.

4.5.2.2 *Cardiology*

Due to the severity and frequency of cardiac diseases this application area is of utmost importance. Essential diagnostic tasks include plaque assessment in the coronary arteries, detailed diagnosis of a heart infarction, e.g. the extent of infarct core, assessment of the heart valves, and assessment of abnormalities, such as congenital heart failures. Morphological information, extracted from CT and MRI, as well as functional information, e.g., wall motion extracted from ultrasound, are important.

An example of a multimodal visualization problem that was considered in visualization research is the use of perfusion data combined with other MR imaging modes of for the diagnosis of the coronary heart disease. Perfusion data of the heart, which indicates the blood perfusion in the myocardium and can be acquired from MRI or SPECT scanners, needs to be combined with anatomical data to provide morphology of the heart muscle and the coronary arteries. Similar to combining PET and CT data, the resolution of the perfusion data is much coarser. Moreover, the acquisition of cardiac perfusion data exhibits gaps, i.e., there are regions in the morphological data, where no corresponding perfusion data is available. Simply interpolating the missing information is not a satisfactory solution. Thus, the overlay of both information must convey where slices of the morphological data correspond to morphologic slices and where no such correspondence exists. Termeer et al. proposed a visualization to aid diagnosis of patients with Coronary Artery Disease using cardiac perfusion MRI data [TBBV+07]. They extend the traditional bull's eye plot to a continuous volumetric version that reveals transmural of scar tissue and link this to an anatomical view of the heart. Transmurality indicates whether the whole wall of the (left) ventricle is affected by an infarction—information that is essential for prognosis and treatment. Hennemuth et al. developed a comprehensive approach to the analysis of contrast-enhanced cardiac MR images [HSFM+08]. They propose a full workflow to align the differ-

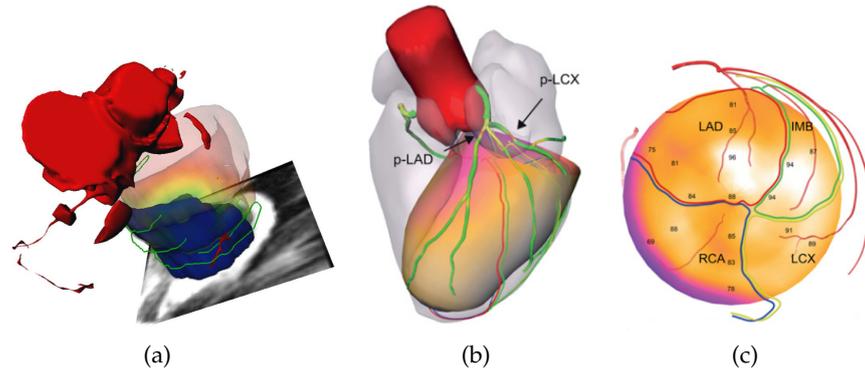


Figure 20: (a): Combined visualization revealing an infarcted area in dark blue, with hypoperfused regions represented by contours. The myocardium is shown as a transparently colored surface that encodes the distance to the infarction [HSFM+08]. (b),(c): Combined CT Angiography and SPECT myocardial perfusion imaging are visualized to illustrate coronary artery disease [KGSAY+14].

ent datasets, extract surfaces and analyze them in a combined way (see Figure 20). Oeltze et al. presented an integrated multi-modal visualization of morphologic and cardiac perfusion data for the analysis of coronary artery disease [OKGH+06]. Myocardial perfusion is measured using an MRI scanner and combined with CT Coronary Angiography (CTCA) which depicts the anatomy. In their visualization, colored icons, heightfields and lenses are used to visualize and explore the different parameters measured. Kirişli et al. combine CT Angiography (CTA) and SPECT myocardial perfusion imaging (MPI) in a visualization aimed at assessing coronary artery disease [KGSAY+14]. They evaluated the diagnostic value of a software-based image fusion system over conventional side-by-side analysis and found improved diagnostic performance when using their application (see Figure 20).

4.5.3 Treatment Planning

Multimodal medical data visualization applications have been developed mainly in the areas of neurosurgical planning and guidance. Some recent works have focused on radiotherapy treatment planning.

4.5.3.1 Neurosurgery

A primary focus in neurosurgical planning applications is identifying various surgically relevant structures that are at risk to be damaged, such as the functional areas in the gray matter and the white matter fiber tracts. This enables an understanding how they are related to the lesion that needs to be removed. The surgically relevant structures can be acquired using several modalities such as CT and MRI,

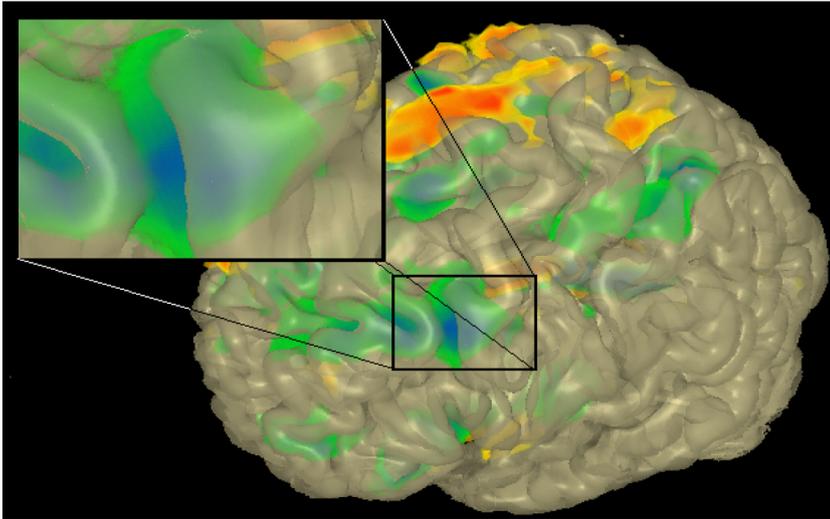


Figure 21: A non-polygonal isosurface revealing patient-specific functional information in anatomical context of a template brain [RTFE+06].

but also special MRI modes such as fMRI, DTI as well as multiple different scanning parameter combinations are often employed. Additionally, neurosurgeons need to plan the optimal or safe access path to such a lesion, with the least amount of damage to vital functional areas, derived from fMRI, and fiber tracts, derived from DTI. Visualization applications aimed at neurosurgical planning and guidance often focus on brain tumor resection, but also applications aimed at Deep Brain Stimulation (DBS) and Arteriovenous Malformations (AVM) surgery exist.

Jannin et al. fused several modalities and imaging modes from MRI for neurosurgery navigation [JFSG+00b]. They combined segmented structures and vessels with functional areas from magnetoencephalography (MEG) and fMRI data in a neuronavigation application. As one of the most notable early techniques where surfaces are used for combined information, Stokking et al. presented a visualization method that combines functional input data and a surface extracted from anatomical data [SZV01]. They map fMRI values onto a brain surface using the surface normal and evaluate their technique on both registered (S)PE(C)T/MRI and fMRI/MRI. This technique was an extension from earlier work on combining software-registered SPECT with a surface extracted from an MRI scan, by mapping functional values of the SPECT to the surface of the brain along the normal [SZPR+97]. While these techniques were applied to the brain, the idea of using a surface can be transferred to other organs, such as the heart muscle. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively (recall Section 5.1.1), and their technique was specifically focused on functional brain images [RTFE+06]. Their tool was aimed at supporting cognitive neuroscientists in experimen-

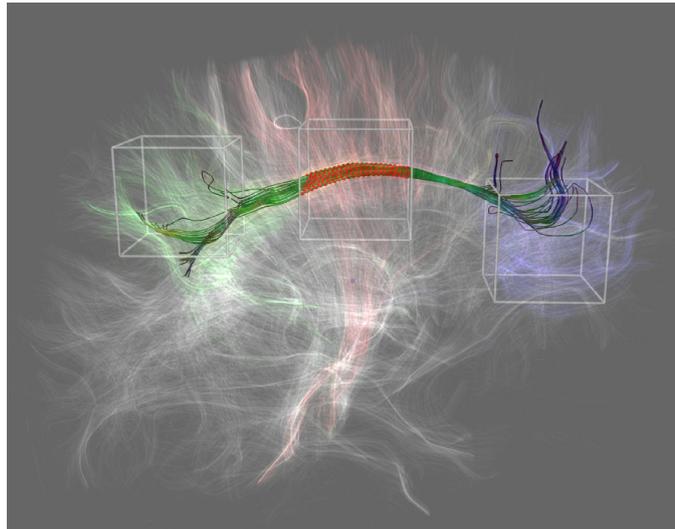


Figure 22: Fiber bundle filtering using multiple boxes that can be combined to select bundles using arbitrary logic combinations for selection of fibers to be shown in multimodal visualization [BBMN+07].

tal studies and to communicate results to non-experts and employs a template brain with patient-specific fMRI data (see Figure 21). Blaas et al. proposed an approach that fuses fMRI and DTI data for planning brain tumor resections [BBMN+07]. In this fused volume, users can extract fiber bundles that pass through a region around the tumor. These bundles can then be explored by filtering on distance to the tumor, or by selecting a specific functional area using arbitrary convex geometries as selection criteria (see Figure 22). Rieder et al. employed a combination of fMRI and DTI data for neurosurgical therapy planning [RRRPo8]. Their application visualizes pathologies using a distance-based transfer function, as used by Tappenbeck et al. [TPDo6], and only shows functional data in close proximity to the lesion. Furthermore, they increase depth perception by including a distance-ring, which visualizes how deep the lesion is situated in the brain from the current viewpoint. For the surgical planning itself, they provide access path visualization and rely on identification of superficial landmarks which can be translated to the pre-operative context (see Figure 23).

A visualization approach for combined MRI and fMRI brain data was presented by Jaineck et al. [JBBS+08]. Various rendering styles were used, e.g. ambient occlusion, and illustrative techniques were employed to enhance the visual output. Following up on this work, Born et al. extended it to include DTI tracts, which reveal reconstructed nerve fibers that connect functional areas [BJHS+09]. They also enhanced depth and shape perception by applying silhouettes and dithered half-toning (see Figure 24).

Janoos et al. [JNMS+09] proposed a method to visually analyze brain activity from fMRI data, with a special focus on temporal dependencies. They propose a methodology to analyze the time dimension

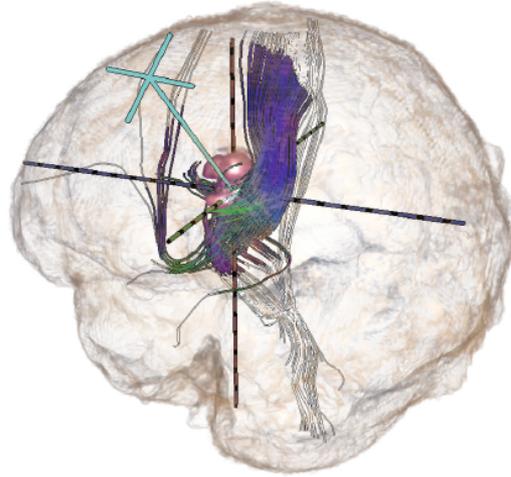


Figure 23: Path planning for oncologic neurosurgery. An entry point in the skull is searched to reach the target area in the tumor without destroying fiber tracts and functional areas based on MRI, fMRI and DTI data [RRRPo8].

through volumes-of-interest, of which the selection is guided by a hierarchical clustering algorithm in the wavelet domain. They visualize these volumes-of-interest overlaid onto MRI data of the brain and show the cluster time-series of selected clusters in a separate view.

Diepenbrock et al. were the winners of the IEEE VIS Visualization contest on multimodal visualization for neurosurgical planning [DPLB+11]. They provided a 3D view inside the brain featuring an interactive probe containing structures acquired from fMRI, DTI and various additional MR imaging modes. Uncertainty in the fMRI and DTI data is revealed using uncertainty borders and a different rendering style. Furthermore, they provide a cylindrical access path projection representing the distance to the structures at risk along the planned path. A close-up view of the tumor allows a more detailed view on the area that the surgeons are planning to resect.

Serra et al. developed a multimodal pre-operative neurosurgical planning system designed for use with a mirror-based Virtual Reality system workbench [SKGH+98]. They render CT, MRI and MRA scans using 3D textures and allow the user to interact with the volumes using a tracked 3D pointer. Neubauer et al. proposed a surgical simulation application for endonasal transsphenoidal pituitary surgery, which is a minimally invasive endoscopic procedure [NMWW+04]. They employ CT and MRI to simulate this endoscopic procedure for training and pre-operative planning. For this, they used CT for skull anatomy, MRI for the tumor, optical nerve and pituitary gland, and contrast-enhanced CT and MRI to highlight the internal carotid artery. They provide interactive threshold adjustment to adjust the surface visualizations.

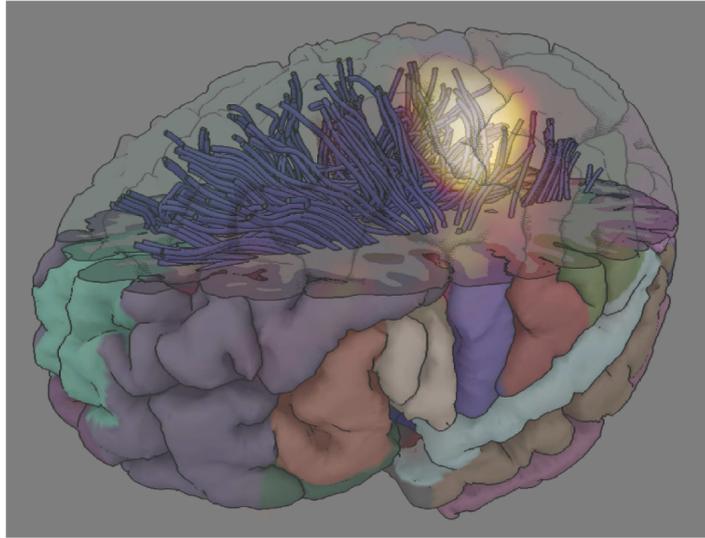


Figure 24: Visualization of cortical anatomy (MRI), brain activity (fMRI), and nervous pathways (DTI) [BJHS+09].

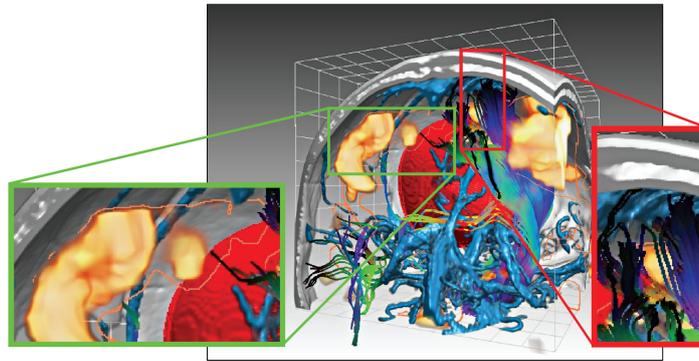


Figure 25: In a multimodal 3D view, neurosurgeons can explore and plan an initial access path by probing [DPLB+11].

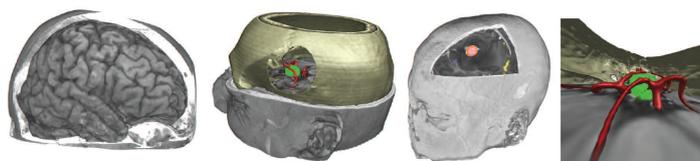


Figure 26: From left to right: skull peeling, Multi-volume rendering of segmented data (green: tumor (MRI), red: vessels (MRA), brown: skull (CT)), Multi-volume blending (black/white: brain (MR), red: metabolic active part of tumor (PET), yellow: brain areas active during speech - fMRI) and perspective volume rendering for simulating keyhole surgery [BHWB07].

For planning neurosurgical procedures, Beyer et al. introduced a framework that can visualize multimodal volumes such as CT, MRI, fMRI, PET and Digital Subtraction Angiography (DSA) [BHWB07].

They propose a skull peeling algorithm that can automatically remove an occluding bony area to reveal the underlying brain. Furthermore, they developed a masked-based approach to show multiple modalities concurrently and a rendering technique to render binary segmented objects with a smoothed appearance when available (see Figure 26). Within masked parts of the volume, different datasets are rendered, each with their own transfer function settings. Joshi et al. presented interaction techniques in a neurosurgical planning application that includes stereotactic navigation [JSVS+08]. They visualize MRI, fMRI, DTI and SPECT in a visualization that allows the user to crop volumes using an interactive line widget. Kin et al. developed a neurosurgical planning tool specifically focussing on brainstem malformations in which they fuse MRI, CT, and 3D rotational angiography together [KNST+12]. They mainly rely on combining multiple surface reconstructions of the individual modalities, which demands a significant amount of preprocessing time in manually segmenting important structures.

In the area of neurosurgical tumor planning, Rieder et al. proposed a multimodal visualization of intracerebral pathological tissue [RSHP08]. They use multiple MRI sequences (T₁, T₁ce (contrast enhanced), T₂, FLAIR) as the input for their visualization application and perform clustering to determine pathologic regions. Next, they blend these pathologic regions with the anatomical context information using an automatically calculated transfer function. Furthermore, they propose an automatic cutting tool and brain peeling to reveal hidden structures of interest.

Bock et al. recently presented a tool for planning and guiding deep brain stimulation (DBS) interventions by fusing multimodal data that includes uncertainty regions from the acquisition process [BLEL+13]. In order to guide these procedures, CT and MRI are fused and visualized integrated with results from Microelectrode Recordings (MER), which measure the electric field in the brain intra-operatively. Their tool features a planning, recording and placement phase in which the corresponding steps of the intervention can be performed (see Figure 27).

Weiler et al. presented work on neurosurgical planning for treatment of Arteriovenous Malformations (AVMs) in the brain [WRDW+11], an example of which can be seen in Figure 28. In these procedures, a precise identification of the arteries and veins is crucial to understand the complex inflow and outflow in these vascular pathologies. They combine several differently weighted MRI scans, such as T₁-weighted images with and without contrast agent, arterial time-of-flight (TOF), and MR-Venographies (MRV), into a single multi-volume visualization to facilitate understanding of the lesions angio-architecture. To prevent clutter, they propose a focus of attention technique, based on the distance to a point of interest, that allows the user to atten-

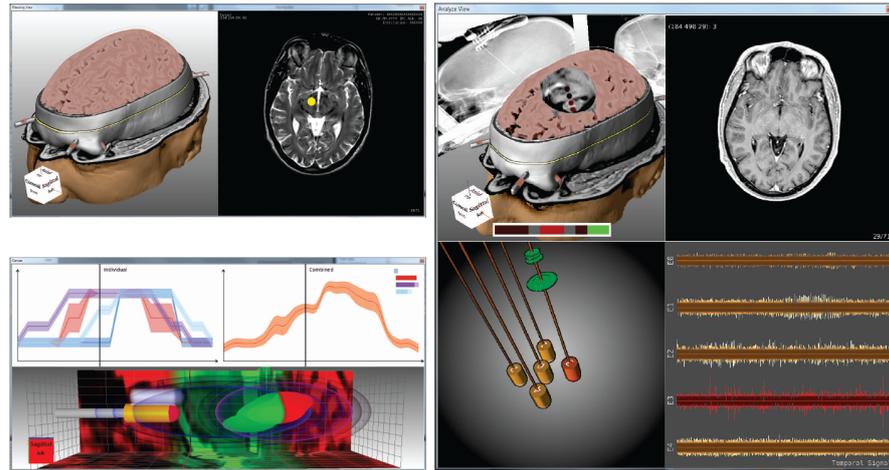


Figure 27: The DBS planning phase in the top left, recording phase on the right, and the placement phase in the lower left. [BLEL+13].

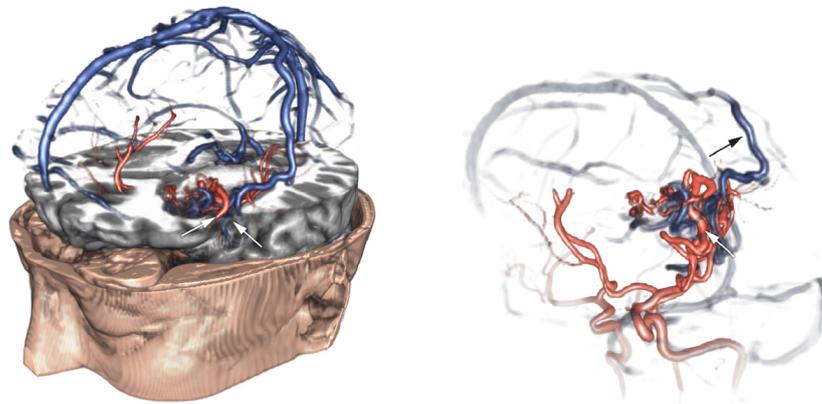


Figure 28: Visualization of neurovascular anatomy for treatment planning of AVM surgery with context (left) and focused on vascular structures only (right), based on MR-Venographies and T₁-weighted MRI [WRDW+11].

uate importance using transparency and saturation manipulation for structures outside the region of interest. Navkar et al. developed visualization tools for planning neurosurgical interventions with straight access, such as biopsies, deep brain stimulation and ablation of brain lesions [NTSW+10]. For this, they generate various access maps based on vascular structures on the surface of the skin of the patients head for guidance in selecting a safe entrance point.

Much research has been done on the visualization of fiber tracts from DTI scans, often using anatomical information from an MRI scan as context. Enders et al. proposed to visualize white matter tracts reconstructed from DTI with wrapped streamlines [ESMH+05]. They generate surfaces that wrap around the convex hull of fiber bundles for a more intuitive representation of tracts and combine this with anatomical information from a T₁-weighted scan. Merhof et al. proposed a new visualization technique for white matter tracts using

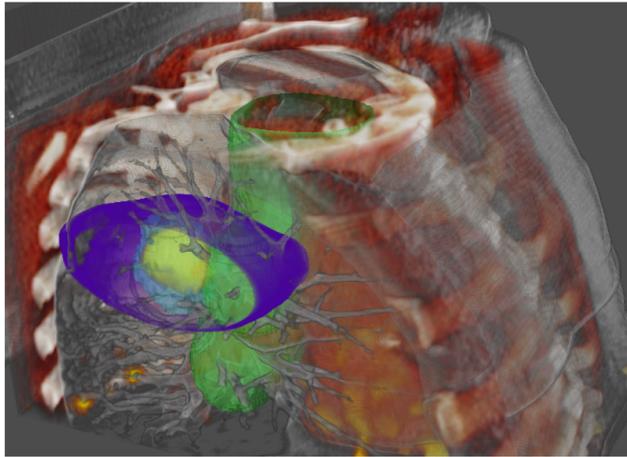


Figure 29: Fusion visualization of 4D PET/CT with segmentation and dose information combined with clipping [SFNB14].

triangle strips and point sprites [MSEN+06]. Their novel DTI hybrid rendering approach speeds up the rendering process and can then be combined with DVR of anatomical data to provide additional information for surgical planning. Since a full overview of these visualizations is outside the scope of this chapter, the reader is referred to a recent survey written by Tobias Isenberg on illustrative visualization techniques for diffusion-weighted MRI tractography [Ise15].

4.5.3.2 Radiotherapy

Radiotherapy planning is generally based on CT, MRI and PET/CT scans. These volumes are used to define target areas for radiation but, also organs at risk that should not receive high radiation. Especially for target areas such as chest and upper abdomen, that move strongly due to breathing, the recent possibility to acquire 4D PET/CT data opens the possibility to capture and integrate the movement of tumors into the radiation target volume definition, building the basis for the dose calculation. In contrast to many other diagnostic and treatment planning procedures that occur under severe time pressure, radiation treatment planning is advanced and time-consuming. Target structures and structures at risk are segmented regularly to assess how they would be affected by a certain dose distribution. Thus, advanced multimodal visualization techniques would be applicable.

Schlachter et al. [SFNB14] showed that 3D and 4D visualization of images, combined with delineated regions, and the calculated dose (also available as volume dataset), complements the common slice views. This provides a fast overview over the spatio-temporal configuration of all delineated areas and the related dose distribution resulting finally in a faster quality checks of the radiation plan.

Magnetic Resonance Spectroscopy Imaging (MRSI) data provides metabolic information that quantifies the concentrations of multiple

brain metabolites, such as Choline and Creatine, per voxel. Currently MRSI is only performed in a clinical research context, but not yet employed in clinical practice. In the work by Nunes et al., MRSI data is fused with multimodal radiology imaging in an integrated visual analysis system aimed at radiotherapy treatment [NRSK+14]. They linked the medical imaging framework MITK and general purpose data exploration tool ComVis to analyze, relate and visualize MRSI data together with multimodal images.

4.5.4 Discussion

An overview of all papers according to the application type and application area can be found in Table 2. The table states if any segmentation is needed for processing the data in order to visualize it, which is important to estimate the amount of preprocessing an application requires. Furthermore, the type of evaluation (if any) was performed to verify the utility of the applications is mentioned. For this, the evaluations are distinguished between quantitative (Qt) and qualitative (Ql) studies with domain experts (D), non-domain experts (N), or if the performance was tested (P). Furthermore, the image modalities visualized is mentioned and the visualization techniques is described that are applied for the visualization of the multimodal medical data. Additionally, a Venn-Edwards diagram of all application papers categorized according to the visualization techniques used and the application type is illustrated in Figure 30.

The table shows that a fair amount of work has already been done on neurosurgical treatment planning, as well as work on guidance. Not much work has been done yet on radiotherapy planning, while this field especially uses multimodal data and is in need of suitable visualization methods. In radiotherapy, relevant structures are regularly segmented, making all described visualization techniques in Section 4 practically applicable. There has only been one visualization application paper combining ultrasound with another modality, while ultrasound has great potential and benefits that can be complementary to other modalities, such as the ability to provide real-time information. The limited field of view and noise in ultrasound data could be alleviated by combining it with additional modalities, and could be of value also in for instance biopsy guidance. SPECT is involved only in cardiac diagnostic and neuro-science and -surgical planning applications. A popular modality combination for both vascular research and oncologic diagnosis is PET/CT, which is used in all the listed application papers. Due to the nature of MRI, fMRI and DTI data, these modalities are often involved in both neurosurgical treatment planning and research. For obvious reasons, fMRI is completely neuro-specific, but DTI is also clinically applied so far mainly in neurological contexts.

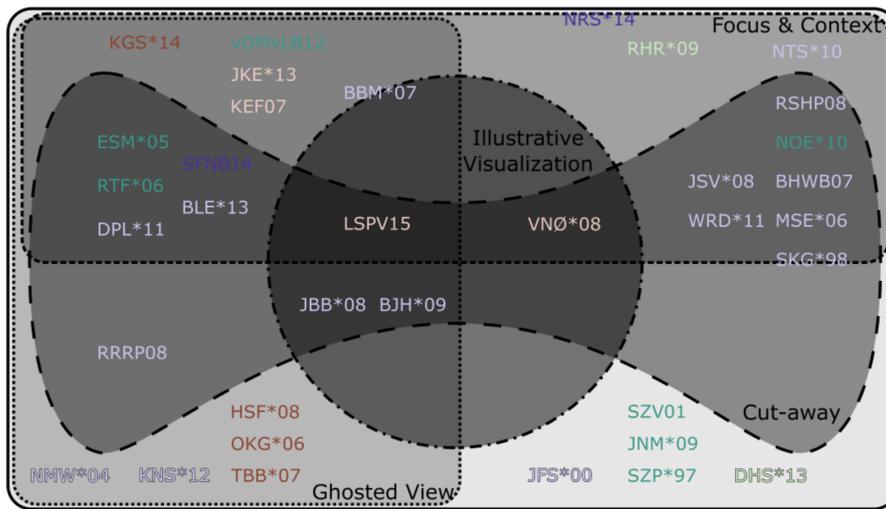


Figure 30: A Venn-Edwards diagram of the references categorized according to the applied visualization techniques. The color of the references corresponds to the application type colors: vascular research, neuroscience, diagnostic oncology, diagnostic cardiology, neurosurgical treatment planning and radiotherapy planning.

While the visualization techniques used per application area vary too much to make strong conclusions based on the limited number of samples, it is clear that ghosted views are often employed for diagnostic purposes (see Figure 30). Furthermore, cutaways are most successfully applied to neurosurgical planning, which makes sense due to the nested anatomical and pathological structures involved and the need for exact access path planning through the skull. There are five papers not employing smart visibility or illustrative techniques, mainly in a research context.

In most works qualitative evaluations with domain experts are presented, but quantitative evaluations as well as more elaborate clinical studies are still often needed. Only a few works applied illustrative visualization techniques [LSPV15; BJHS+09; JBBS+08; VNØU+08]. However, the application of illustrative visualization technique seems to be promising [TIP05; Law15] and it shows a potential to use these techniques for further research on multimodal medical visualization. In multimodal medical visualizations, there are often many overlapping and nested structures, which causes occlusion problems that the abstraction in illustrative techniques can help alleviate. However, the segmentation required often limits the clinical uptake. It might also be helpful to employ novel focus-and-context techniques, by defining one modality as the focus and the other(s) as the context.

4.6 CONCLUSION AND FUTURE CHALLENGES

For many techniques described in this survey, it is not clear how they fit in clinical workflows, how much additional and relevant information they provide and if so, whether this justifies a potentially larger effort, e.g., due to the necessity to segment structures or adjust 2D transfer functions.

Dealing with medical imaging data involves uncertainty in the form of imaging errors, for instance resulting from noise, field bias, patient motion, or imaging artifacts. Additionally, processing errors may occur due to errors in the segmentation or registration process. When visualizing multi-modal datasets resulting from combining acquisitions of multiple scanners, additional care has to be taken to visualize the uncertainty resulting from these processing errors in the registration process. This is a challenging and so far unsolved research challenge [Hah05].

The existing multimodal visualization techniques need to be better evaluated and compared. On the one hand, evaluations that consider effectiveness in specific medical tasks are needed. On the other hand, perception-based evaluations are essential, e.g., to understand how effective certain emphasis techniques actually are, how well users could discriminate values with certain color scales (used as overlays) and how accurate they can locate functional abnormalities.

The existing techniques are spread over many research prototypes and commercial solutions. It is thus very difficult to compare them. A framework that integrates at least the most commonly used techniques and flexible parameterization options would be very valuable. A comparison also benefits from benchmark tasks and data. A first step in this direction was the IEEE Vis. 2010 contest on multimodal brain data for neurosurgery planning. A few more such tasks, e.g., in cardiology or radiation treatment, would be beneficial.

A multitude of visualization techniques were shown. Interestingly, these are exclusively techniques that were developed first for displaying single datasets and then adapted and refined. An open question is whether truly multimodal techniques can be developed. Multimodal visualization obviously benefits from multiple views. More research is necessary to understand which views are essential, how flexible viewing configurations should be and which synchronization/coordination techniques are needed.

The visualization of single medical datasets may benefit from automatic viewpoint selection to support the user, i.e., choosing a good initial view on a dataset as starting point for further analysis and exploration [VFSH01; BS05; MNTP07; KBKK07; KBKG08; RCL16]. Measures of visibility, size of objects and viewpoint stability are employed for the viewpoint selection. Automatic viewpoint selection is also promising for multimodal data visualization and requires careful adap-

tation of measures and application to selected case studies. Also the user interface must be considered to enable a physician to specify what is important for her as input for an algorithm.

Neurosurgery is among the most advanced user communities in this field, while other disciplines like radiotherapy are not yet exploiting the full potential of visualization techniques available today to support therapy planning and monitoring. There are many applications in neurosurgery because of the relatively easy and reliable registration process, which can be automated, as well as specific application demands. First, CT and MRI contribute a significant amount of relevant and complementary information, while in lung surgery for example you can completely rely on CT. Second, neurosurgery has much higher accuracy demands compared to any type of abdominal surgery. A few millimeters in abdominal surgery usually do not matter, while in neurosurgery 2mm may be a lot.

Challenges also appear in case when the input does not consist of surface meshes. For volumetric data, some visualization techniques are more challenging than when applying them to surfaces. Feature lines are mostly defined on surfaces, but some methods can also be applied on volumetric data. Hladůvka et al. [HKG00] visualized ridges and valleys depending on curvature values. With an adapted transfer function, these lines could then be highlighted, which is similar to the approach by Lawonn et al. [LSPV15]. In contrast to feature lines, hatching approaches were also successfully applied to volume datasets [DCLK03; YC04; CD05; PVW08; BVW08]. Mostly particles or points are placed in the dataset on a specific isovalue and then the points are traced along the principle curvature directions. Due to the challenging character of feature lines, these techniques were not commonly applied to multimodal data visualization so far. Although its potential was shown to illustrate surfaces [LP16] more work needs to be done in this field.

Developing efficient and effective visualization methods with real impact on daily clinical or research routine of potential users requires a highly interdisciplinary effort fusing knowledge from physicians, visual computing specialists, HCI, imaging and image analysis specialists. This becomes even more important with the increasing availability of new imaging modalities, that can be fused with imaging pipelines delivering sets of highly heterogeneous data at different scales. A deep understanding of this data and its application context is necessary to generate useful results, on the other hand technical skills are getting more and more important, as the amount of data to be processed increases tremendously. As a consequence, multimodal visualization will be even more interdisciplinary than it has been ever before.

4.7 SUMMARY

An overview of the current state of the art in medical multimodal visualization was presented. Relevant medical imaging modalities and acquisition techniques, i.e., CT, MRI, fMRI, DTI, PET, and SPECT, were introduced and their properties as well as the associated visualization challenges were described. Afterwards, the current clinical workflow for the exploration and analysis of multimodal medical modalities were examined. This chapter summarized the requirements in designing a visualization technique to maximize the insights into relevant details of multimodal medical data.

Subsequently, the most common visualization techniques that solve this visualization problem were highlighted. These techniques need to incorporate heuristics for assessing the importance of information and emphasis techniques to adapt to the importance values. For this reason, smart visibility approaches, including focus-and-context techniques, ghosted views and cutaways are highly relevant for visualizing multiple volumes. While these techniques are frequently used in visualization research, they are not part of any commercial solution or available in radiology workstations and the advantages are not verified by appropriate studies. This is probably because these techniques require time-consuming preprocessing and are quite complex, as well as unfamiliar to physicians.

Rendering multiple volumes is an associated research challenge for which many approaches have been developed. Furthermore, special interaction techniques have been designed for exploring multimodal datasets, which can be as important in providing insight in multimodal data as rendering refinements.

In the main part of this chapter, an overview of 35 visualization application papers was provided. The scope of the chapter was defined by focusing on an overview on multimodal medical visualization applications associated to research, diagnosis and treatment planning or guidance. Within these categories, the main application areas were related to oncology, cardiology, radiotherapy and neurology. The findings were summarized in a table featuring the three aforementioned categories, a further subdivision according to the medical application domain, whether preprocessing is required in terms of segmentation, image modalities used, and visualization techniques employed. From the table, conclusions were drawn which application areas are active, upcoming or sufficiently researched. While a fair amount of work has been done on neurosurgical planning, radiotherapy planning seems to be a rising field with many opportunities for interesting research, and the same holds for combining ultrasound with other modalities.

Finally, open problems in multimodal medical visualization were mentioned that serve as a starting point for future research in this area. Evaluation of medical task performance and in the form of

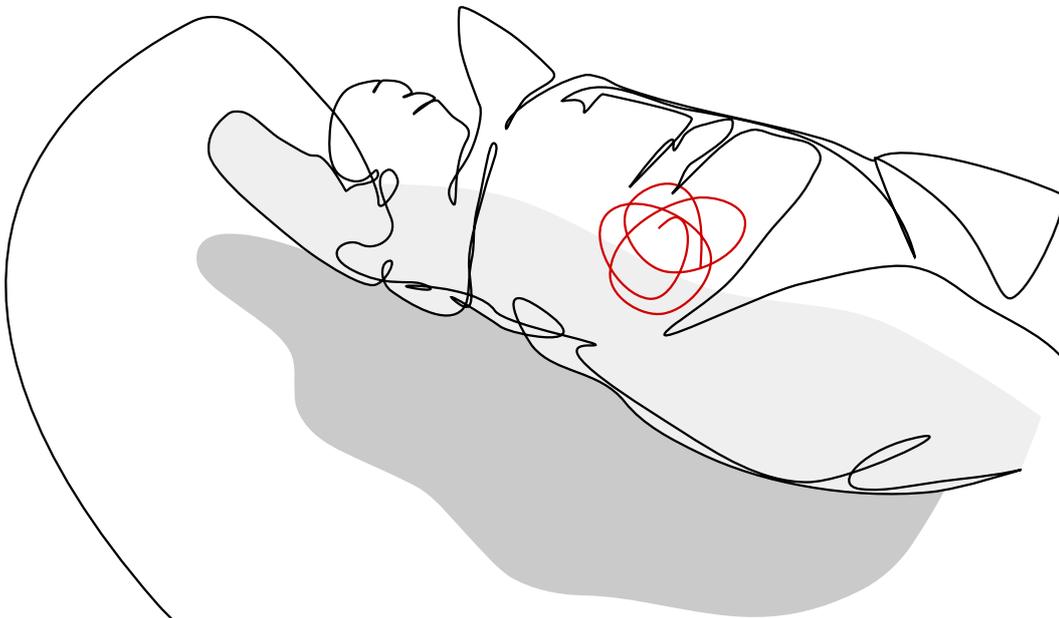
perceptual studies are needed to validate the effectiveness of visualization application designs. Furthermore, most developed techniques were extended from existing single modality techniques, while truly multimodal techniques have yet to be developed. With the increase in the amount of data acquired and new modalities being brought into clinical practice, multimodal medical visualization remains a promising research area for future developments.

Part IV

SELECTED APPLICATIONS

5

ILLUSTRATIVE PET/CT RENDERING



This chapter is partly based on:

Kai Lawonn, Noeska Smit, Bernhard Preim, and Anna Vilanova.

“Illustrative Multi-volume Rendering for PET/CT Scans.”

In: Proc. of Eurographics Workshop on Visual Computing for Biology and
Medicine (VCBM), 2015, pp. 103–112.

ILLUSTRATIVE MULTI-VOLUME RENDERING FOR
PET/CT SCANS

POSITRON emission tomography (PET) is a medical imaging modality that is used to detect metabolic activity and is able to highlight functional pathologies. However, exact localization of these abnormalities is not possible since PET is only able to display metabolic information and not able to distinguish different tissue types. X-ray-based computed tomography (CT) scanners provide structural anatomical information at a much higher spatial resolution, but lack the metabolic information that PET can provide. Therefore, combining the PET information with CT scans can depict metabolic abnormalities enriched with the spatial information of the anatomical imaging provided by CT.

Despite research projects aiming at fused 3D visualizations for clinical practice, radiologists currently explore the stack of two superimposed 2D images to analyze abnormalities in a slice-based approach for PET/CT. They are able to browse through the stack of slices and specify different image properties, e.g., window and level to adjust the contrast and examine structures of interest. The major problem with this approach is that physicians have to mentally fuse all information. As the number of slices per CT scan increases, browsing through all of these slices manually becomes more involved and time-consuming. An overview could potentially help to guide the user towards areas that warrant closer examination at a glance. Furthermore, patients and other specialists might be less accustomed to the slice-based representation and might benefit from a more familiar representation.

To help alleviate these problems, this chapter provides a focus-and-context visualization for PET/CT images, see Figure 31. By combining direct volume rendering with illustrative visualization techniques, an improved shape perception and minimizing occlusion technique is achieved. CT scans are visualized as anatomical context using line drawing techniques, while the PET scans are visualized as the focus object, drawing the attention of the viewer to the metabolic activity. The new visualization technique allows for a fast overview visualization that can be easily applied to PET/CT data and extended to other focus-and-context applications. To validate the utility of the visualization technique and its application to PET/CT data, a qualitative evaluation with medical experts was conducted. With this, the contributions of the novel visualization technique are as follows:

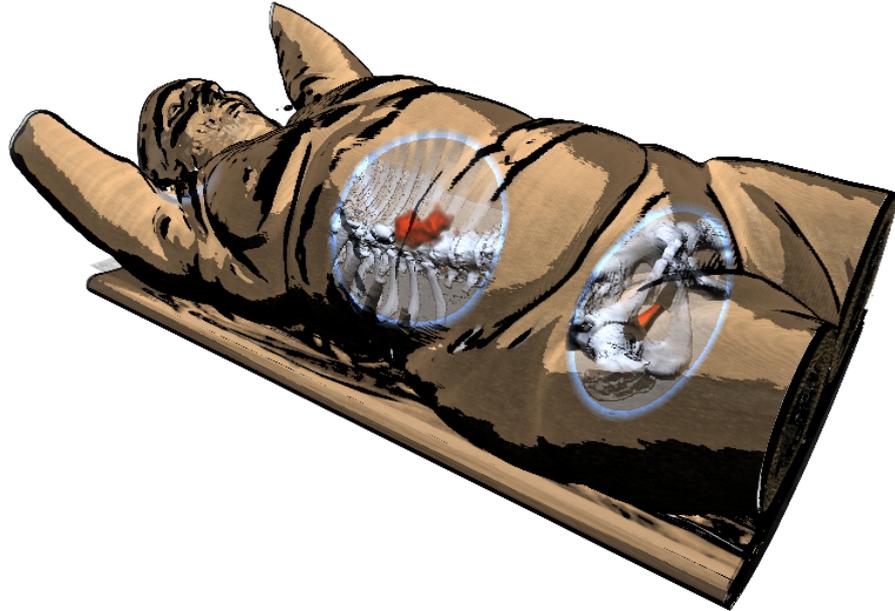


Figure 31: The novel illustrative visualization method encodes PET as a focus object, while the CT allows surrounding bone structures and skin to be visualized as the context.

- An illustrative focus-and-context visualization for multi-volume rendering applied to PET/CT data.
- An application that demonstrates the utility of the new visualization technique.
- An evaluation with domain experts that confirms the advantages of the presented work.

5.1 MEDICAL BACKGROUND

An overview on CT, PET and PET/CT acquisition and its applications can be found in Chapter 4. For a more elaborate description, the reader is also referred to Townsend et al. [TCYH04].

5.1.1 Requirement analysis

There are several options for visualizing combined PET/CT images. In the current clinical workflow, the typical approach is to examine the CT images in grayscale with the PET images superimposed using a colormap. Browsing through all the 2D slices individually to examine metabolic uptake can become time-consuming and cumbersome as the number of slices increases. Besides the 2D techniques, 3D techniques can additionally be used to give an overview of the full datasets at a glance. Existing methods such as maximum intensity projection (MIP), can provide such an overview, but suffer from

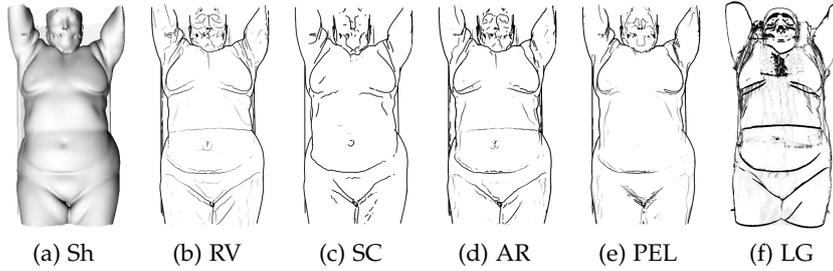


Figure 32: The shaded model (Sh) with different feature line techniques: *Ridges & Valleys* (RV), *Suggestive Contours* (SC), *Apparent Ridges* (AR), *Photic Extremum Lines* (PEL), and the presented approach LG. Compared to the feature line techniques, the presented approach enhances salient regions whereas the feature line methods detect lines based on different feature measures.

depth perception issues. In case of PET/CT data, since the intensities are not comparable, a combined volume MIP can suffer from occlusion by the highest intensity regions. Ideally, a combined volume 3D technique would allow the users to get a quick overview of areas of interest and to localize foci of suspicious metabolic activity in an anatomical context. The requirements for such a technique are the following:

Req. 1: COMBINED VISUALIZATION. It should show the combination of the two modalities in a single view in which the PET activity is always visible.

Req. 2: LOCALIZATION. It should relate the metabolic activity to nearby anatomical structures for accurate localization.

Req. 3: USER-FRIENDLY. Visualization options should be easily adjustable to fit the needs of the pathology being examined.

Such a technique can be used to guide the exploration of the datasets, bringing the attention to regions of interest, after which detailed inspection of these regions can be performed in 2D images. Furthermore, it could be beneficial for treatment planning, doctor-patient communication and interdisciplinary communication.

5.2 METHOD

In this section, different illustrative visualization techniques are described that are designed to depict the combined PET/CT datasets. Based on the requirements specified in the previous section, the visualization tasks are divided into three areas. The PET activity should

be the focus area and always visible (Req. 1). The CT scan is used as an anatomical context, and can provide structural information revealing the skeletal structures near the focus (Req. 2). Further away from the focus region, the structural information can be abstracted to show only the skin to aid shape perception and orientation. The presented technique then consists of a focus-and-context visualization and can be subdivided in three steps:

1. Illustrating the context, e.g., the skin.
2. Illustrating the focus, e.g., the PET activity.
3. Visualizing the region closely surrounding the focus area.

As a reminder for the reader the following notations will be used. The Hounsfield values of the voxels in the CT scan data are given by the function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ and represent the tissue density. For further calculations, curvature measurements are needed, which are based on the work by Kindlmann et al. [KWTM03]. The gradient ∇f of the CT data is determined using central differences, which gives us the normal vectors: $\mathbf{n} = -\nabla f / \|\nabla f\|$ of the dataset. Afterwards, the Hessian $\mathbf{H}(f)$ of f is calculated by central differences again. Then, $\mathbf{P} = \mathbf{I} - \mathbf{n}\mathbf{n}^\top$ is set and serves as an orthogonal projection operator that maps every point onto the plane defined by \mathbf{n} , then the operator $\mathbf{G} = -\mathbf{P}\mathbf{H}\mathbf{P}$ is determined. This gives the curvatures κ_1, κ_2 by determining the trace: $\text{tr}(\mathbf{G}) = \kappa_1 + \kappa_2$ and the Frobenius norm: $\|\mathbf{G}\|_F = \sqrt{\kappa_1^2 + \kappa_2^2}$. Using the quadratic formula, the curvatures can be computed. Principle curvatures directions are then given by the eigenvectors of \mathbf{G} .

5.2.1 Context Visualization

A combination of shading styles is used for the context object, e.g., the skin given by the CT dataset. For this, a given transfer function is used to determine the structure of interest. Afterwards, this structure is illustrated using toon shading. The light source is placed top left and it follows the camera, as in medical atlas illustrations normally no headlights are used [Hod03]. To provide an enhanced contextual visualization, additional features lines are used. Most feature line techniques are defined on triangulated surface meshes, where a feature scalar-field is determined. The isolines of this scalar-field then yield the feature lines. Working with volumetric datasets, allows to determine feature regions instead of lines. For this, the light gradient is calculated, which is obtained by determining the shading values per voxel $l = \langle \mathbf{n}, \mathbf{v} \rangle$, where \mathbf{v} is the normalized view vector. Afterwards, $\|\nabla l\|$ (LG) is determined and whenever this value exceeds a certain user-defined threshold the corresponding regions are illustrated in black. As a remark, sometimes a CT dataset features too much noise to get reasonable curvature and light gradient values. Therefore, an

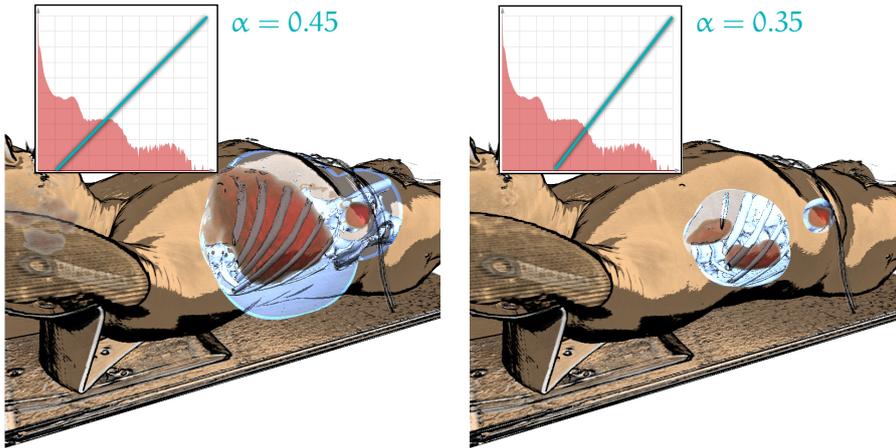


Figure 33: The initial TF for the PET dataset is defined using only a single value. After defining the α value, the TF can further adjusted if needed.

anisotropic diffusion filter [GKKJ92] is applied on these datasets for the purpose of determining the feature regions only. Figure 32 shows a comparison of the *PETCETIX* [Osi] dataset rendered with different feature lines techniques. Here, the shaded model (Sh) is illustrated with the *Suggestive Contours*, *Ridges & Valleys*, *Apparent Ridges*, *Photoc Extremum Lines*, and the novel approach *LG*. For the comparison, surface of the volume dataset is extracted and the feature line techniques are applied to it. Unlike the previous approaches, only the light gradient was taken into account. The results of the novel approach are comparable to the presented feature line techniques. Furthermore, valley lines in white and thickness-controlled contours as suggested by Kindlmann et al. [KWTM03] were added. For the contour, the curvature along the view direction κ_v was used and determined if $|\langle \mathbf{n}, \mathbf{v} \rangle| < \sqrt{T\kappa_v(2 - T\kappa_v)}$ holds, where $T = 3$ is a user-defined threshold.

5.2.2 Focus Visualization

Since the focus region should draw the attention of the viewer immediately, the PET activity visualization is rendered by using standard Phong shading in red, see Req. 1. The contour region of the PET activity is illustrated in white, similar to Fresnel shading. The red color is used to ensure that the experts' attention is drawn to these regions. Since this hue immediately stands out among the other colors, red is in this case used as a pre-attentive visual feature. For the transfer function (TF) used in the PET visualization, a single value $\alpha \in [0, 1]$ is used to set an initial TF. It starts with a vertical line when $\alpha = 0$ and as the alpha increases, the slope linearly decreases, as shown in Figure 33. This simplified transfer functions setting fulfills Req. 3.

5.2.3 Contextual Visualization of the Surrounding Focus

To illustrate the surroundings of the PET enhancement, the lens concept was used as it was introduced by Bier et al. [BSPB+93]. For a survey on interactive lens visualizations, the reader is referred to Tominski et al. [TGKD+14]. Instead of using interactive lenses, an automatic approach is used to set the position of the lens. For this, the PET scan is used and Gaussian blurring is applied. This results in a severely blurred version of the PET foci. Afterwards, a similar TF as used for the original PET on the smoothed PET volume is applied. Then a blurred result is used to create an inside view with a similar effect as cut-away views. During raycasting the ray is detected if it hits the surrounding region. If this is the case, a different shading technique as well as a different TF on the structures in this region is applied. For these regions, Phong shading and a standard TF that illustrates bones as anatomical context is employed. Thus, whenever the ray hits the surroundings a different visualization technique is applied. This gives us the illusion of having a window that allows the user to look inside the body to reveal local context structures.

For further emphasis on the lens concept, an additional illustration on the skin inspired by Lerotic et al. [LCMY07] is used. To create a smooth transition between the context structure and the surrounding focus, the intensity of the skin relative to the distance to the PET activity is decreased. By doing this, the skin slightly disappears for enhancement of the surrounding structures, e.g., the bones. In this way, bones close to the border are more faded than structures closer to PET activity. Furthermore, a blueish color is used to emphasize the active lens area and a blue border around the lens is added. However, the feature lines are always drawn on top of the lens to represent the skin surface. Finally, the PET data is illustrated and mixed with the skin illustration. In order to do this, the distance R from the skin along the view direction to the PET data is determined. The distance is then used as a weighting parameter for the transition of the skin with the PET. The closer the PET is to the skin, the brighter it appears. Here, the idea of the Rayleigh scattering is used and the image data is mixed with $1/R^2$ as a weighted coefficient. In case the PET activity is occluded by the bone structure, a slight change of the brightness of these regions is applied such that the PET regions are still visible.

5.3 IMPLEMENTATION

The implementation of the presented algorithm consists of several steps, see Figure 34. For fast rendering, all computations are performed using the OpenGL shader framework in order to be independent of graphics card vendors. In the first step, an anisotropic diffusion smoothing is applied on the CT dataset. An anisotropic method

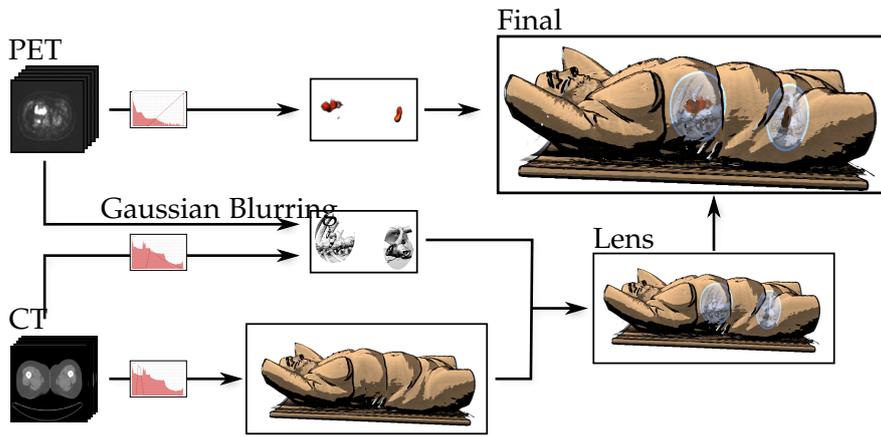


Figure 34: Illustrative multivolume visualization pipeline: The PET data is used to produce a Phong shaded result. We use the blurred PET combined with the CT data to create a visualization of the surroundings of the PET and we illustrate the skin of the CT data. Afterwards, we combine them and create a lens effect. Finally, we add the PET and create the final image.

is used to ensure that edges are mostly retained during the smoothing process. The smoothed result is more appropriate for the gradient calculation and for the illustrations of the feature lines. For the PET data, a simple Gaussian smoothing is used to define the surrounding region. In the next steps, a standard direct volume rendering is applied to composite the intensities along a ray. For the sampling of the ray, equally-spaced distances are used to calculate the intensities. The interpolation is performed with using trilinear interpolation at the sampling points.

5.4 RESULTS

To confirm the utility of the method, it was applied on different clinical and research-oriented datasets. The DICOM sample image sets that are provided by the publicly available *Osirix* [Osi] database were used. First, the *PROSTATIX* dataset, consisting of a PET/CT study with F18-fluorocholine in a male patient with prostate cancer was examined. The results of the technique on this study can be seen in Figure 33. The enhancement shows high-activity regions in the pelvis and abdomen. Using the surrounding illustration, the location of these regions can be identified. Combined with the structural information from the CT scan, normal metabolic uptake can be differentiated from suspicious regions. Next, the novel visualization technique was applied to the *COMUNIX* dataset, which consists of a PET/CT study of a patient with a neck tumor, as can be seen in Figure 35. There are two regions of PET activity, one in the neck and one in the brain. By adjusting the transfer function for the CT scan, the method can either emphasize the skin or the bones as CT context information

to determine the proximity of the PET activity with respect to the skin or bones. As a third dataset, the *MELANIX* dataset was selected, consisting of a PET/CT study of a patient with melanoma (see Figure 39c). Here, the PET activity is high in several regions, such as the armpit, heart, abdomen and pelvis. The CT-based context visualization allows us to localize the PET activity in the armpit as right in front of the shoulder joint. Lacking further information about the patient, it could be either activity in the primary tumor or a metastasis to a lymph node. In Figure 36a the visualization technique is applied to the *PETCETIX* dataset once more. Since this study also features contrast-enhanced CT scans, it is also possible to visualize vascular information as the CT context instead of the bones. In this way, it might be possible to see if there is any tumor angiogenesis.

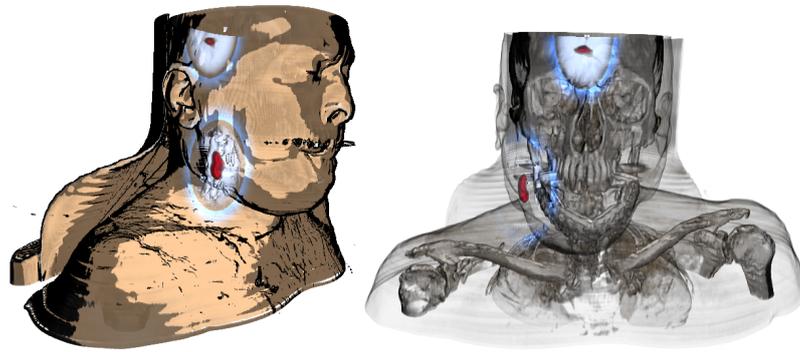


Figure 35: Different transfer functions applied to the *COMUNIX* dataset. The first image shows the skin as the context. In the second illustration, the transparency of the skin was increased to reveal the skeletal structures.

To test the flexibility of the presented technique in another context, it was applied to the *PANORAMIX* CT dataset that contains contrast-enhanced vascular structures. This dataset contains CT scans of a patient with an abdominal aortic aneurysm (AAA). Instead of using the PET region as the focus object, the aorta was segmented and used as the main object of interest, whereas the rest of the CT scan were used as the context that visualizes the surrounding bones and skin (see Figure 36b). The location of the aneurysm is visible using the novel technique, and the context information provided by the bones and skin, can help to localize the pathology. This particular dataset shows that the technique can be applied to other cases, such as highlighting atlas information in CT scans.

In addition to clinical applications, the technique can also be used in a research setting. The presented technique was applied to the PET/CT data in the publicly available digimouse project [SCSL+02; DSCLo7], a full digital mouse atlas consisting of labels, CT and PET data. This data was not acquired in a PET/CT scanner, but registered afterwards using markers. Another difference between this and the

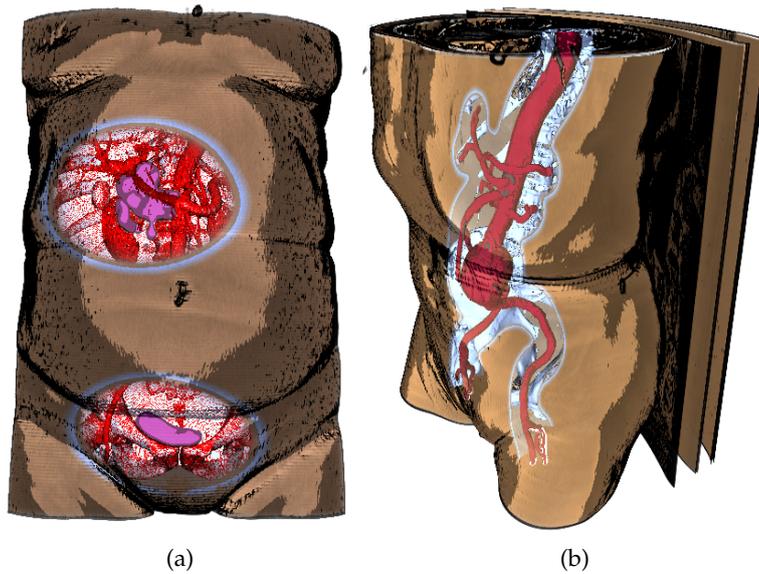


Figure 36: (a): The PETCETIX dataset with the arteries as well as the bone structures as a surrounding context object for the PET region. (b): the PANORAMIX CT dataset with the aorta as the focus object, the bones as immediate context and the skin as the context.

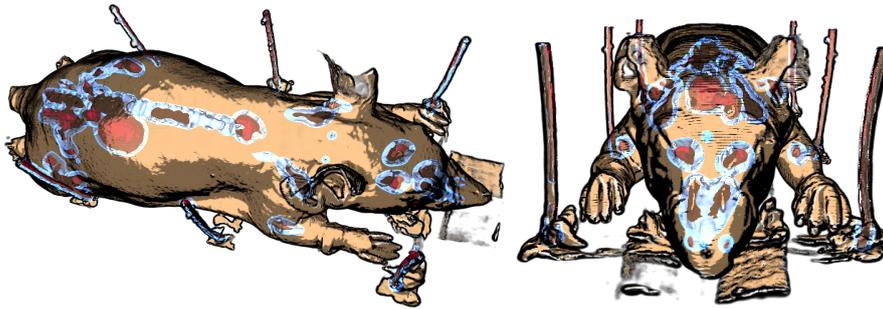


Figure 37: The technique applied to the digimouse set.

previously mentioned results is that the CT scan was acquired in a micro-CT scanner. The results of the presented method applied to this dataset can be seen in Figure 37. Normal metabolic PET activity is shown in several organs and the skin and bones from the CT are used as anatomical context.

5.5 EVALUATION

To assess the clinical applicability of the approach, two informal evaluations with three physicians in the initial evaluation and eleven physicians in the second phase were conducted. The setup for the experiments was as follows:

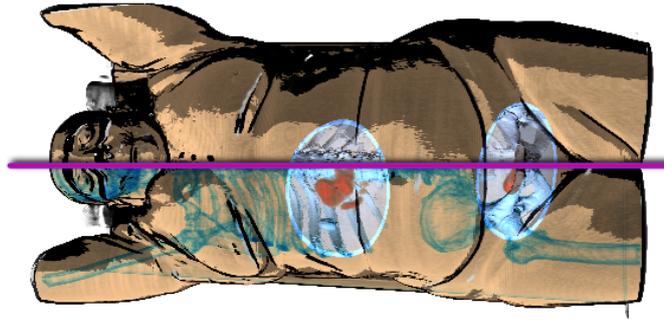


Figure 38: Optional bone context illustration.

PREPARATION: The participants were provided with a short introduction to PET/CT imaging data and the visual encoding chosen. First, the types of data used in the evaluation were introduced and what they represent. A 2D slice image was displayed showing the fused PET and CT data and the colormap used for the PET enhancement. Afterwards, a 3D rendering example was presented and how the PET activity and CT information is shown in these was demonstrated. This helps them to understand and rate the usefulness and the usability of the different rendering techniques presented in the video.

PRESENTATION: In order to gather as many participants as possible a video was presented online that shows the results. The video features three datasets represented in 2D slices as well as several 3D rendering methods: maximum intensity projection (MIP), standard raycasting with pre-defined transfer functions (Multi-Vol), and the new approach (Illustra-Vis). The goal of this phase was to enable participants to assess the quality of Illustra-Vis compared with other methods. All datasets were first displayed in the traditional 2D slice-based views, by presenting all the axial images from bottom to top and back. After this, they were shown using three 3D rendering techniques by rotating the volume around the z-axis. First, a maximum intensity projection (MIP) was shown, as it is currently used in medical practice for analyzing scans. MIP visualizes only the structures in the combined dataset that have the highest intensity along the ray. The second technique is standard multi-volume (Multi-Vol) rendering using a separate transfer function for both modalities. The Multi-Vol approach accumulates the values along the ray with a given transfer function (TF) for both volumes. To make the comparison fair, a pre-defined transfer function for the CT scan with a similar appearance to the novel visualization technique was used. The participants were shown the results of all techniques on all three datasets (see Figure 39).

QUESTIONS: To conclude the evaluation, the participants were inquired to give feedback on the presented techniques via a series of multiple choice and open questions. These questions allowed us to get their opinion on preferences regarding 2D slices versus 3D techniques as well as to compare the 3D techniques with each other. Additionally, several questions regarding the new technique were asked to evaluate the benefits and limitations. Finally, the users were asked open questions on potential applications and suggestions.

5.5.1 *First evaluation*

For the first evaluation session, three physicians/ medical students were asked to provide feedback on the new illustrative visualization technique. All three women had different specializations in the field of medicine and various levels of experience. The first participant P1 (45 years) is a radiologist with 20 years of experience mainly focused on CT. The second participant P2 (29 years) has a surgical background with 8 years of experience. The third participant P3 (22 years) is a medical student with 4 years of experience. P1 stated that to get a good orientation in 3D, the bones are the most important part for her. Therefore, the Multi-Vol visualization was rated less useful, as other organs occluded it as well as the PET enhancement under different viewing angles. Furthermore, she stated that *Illustra-Vis* is applicable for patient-doctor communication and education, but for an overview the lens is too small to give a good insight in the full skeletal context. Hence, she preferred the MIP technique as an additional technique to use besides the 2D slices she is accustomed to. P2 rated *Illustra-Vis* as the best 3D techniques for an overview. She argued that it provides a nice visualization of the skin that facilitates the depth and shape perception. The lens helps to focus on the PET activity and avoids occlusion, which is an issue in Multi-Vol rendering. Furthermore, she stated that *Illustra-Vis* could be applicable to surgical planning, if surrounding organs and vessels were made visible. She also saw potential in using the technique for interdisciplinary communication (in so-called Tumorboards) and in education for courses on radiology and surgery. P3 argued similar to P1, she stated that the full skeletal information is important to localize the PET in terms of the whole body. *Illustra-Vis* can do that as well, but then the lens should be enlarged so far that the shape perception from the skin is reduced. Based on the feedback of this initial user study, the new technique was adjusted, taking their feedback into account. As all physicians argued that a full skeletal visualization is important for the orientation and localization, the bones in the context visualization were optionally included. For this, the contour of the bone structure as well as the bones themselves were computed and added with a blue color on the skin (see Figure 38). Here, the contours are visualized more promi-

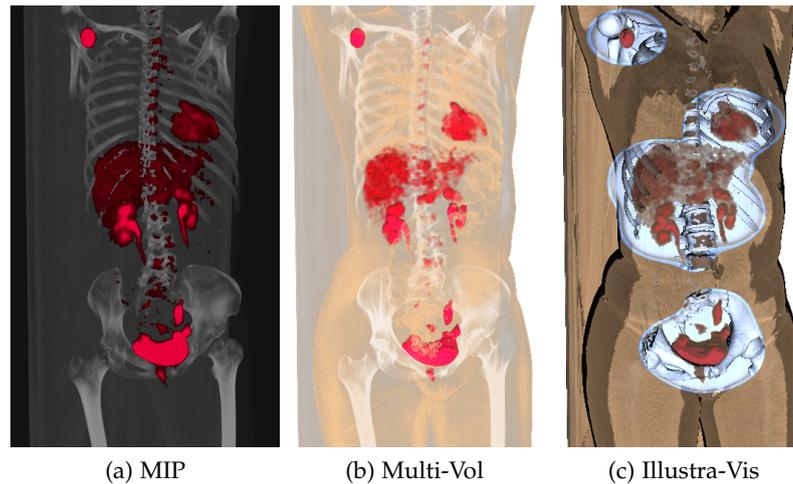


Figure 39: Comparison of different volume rendering techniques: MIP, Multi-volume, and Illustrate-Vis.

nently than the complete bone structure. Drawing the full bones on the skin could be either too subtle or too distracting. Therefore, as an optional feature besides only illustrating the skin, the bones were slightly highlight and the contours were added as the stronger indication for the bones.

5.5.2 Final evaluation

After the adjustment of the visualization, a second evaluation session was conducted to gather more feedback from the three participants. P2 already liked Illustrate-Vis the most, but appreciated the changes made and argued that it adds a better context to the skin. P1 and P3 argued that Illustrate-Vis improved significantly and that it combines the advantages of MIP and Multi-Vol. Additionally, the number of users were increased for the second evaluation in order to gather more feedback. A group of eleven participants (one male, ten female) ranging from ages 22 to 45 (average: 29, median: 28) with diverse backgrounds and years of experience ranging from 1 to 20 years (average: 6.59 years, median: 6 years) were gathered. Using a questionnaire, several specific questions were asked where the participants had to answer what technique they would choose (2D slices, MIP, Multi-Vol, Illustrate-Vis) in four scenarios:

- Showing the PET activity to a patient to explain diagnosis or treatment planning
- Discussing the results as a radiologist with another medical specialist
- Approximately localizing the PET in relation to the anatomy

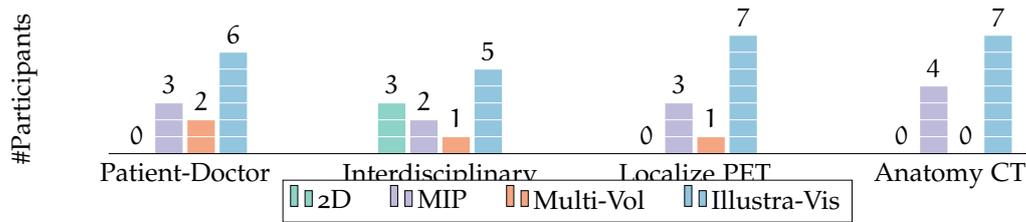


Figure 40: Answers of the eleven participants regarding which of the four techniques they would prefer in four scenarios.

- Visualizing the structural CT anatomical information

In all scenarios, Illustra-Vis was chosen most often, followed by MIP and Multi-Vol (see Figure 40). Finally, several questions on the practical use of Illustra-Vis were asked. It was asked if they would use the Illustra-Vis technique in: diagnosis (Q₁), treatment planning (Q₂), patient-doctor communication (Q₃), interdisciplinary communication (Q₄). Most participants would use Illustra-Vis in all cases (see Figure 41). Additionally, the users were asked whether the skin illustration supports the form/shape impression. Here three stated they strongly agree and the other seven agreed. Furthermore, regarding the depth/spatial impression, three stated they strongly agree that the skin illustration is helpful, while six agreed and two neither agreed or disagreed. In the open questions, one user commented that he would use Illustra-Vis specifically for patient-doctor communication. He mentioned: “Patients with a lack of knowledge probably find it easier to identify themselves when using the Illustra-Vis image to explain diagnosis or treatment. It shows the contours of the body, making it easier to imagine where the area of interest is roughly located.” For this purpose, he would use the illustration technique showing just the skin in the context, as the bones are less relevant to patients. He states that MIP reveals the PET activity more clearly, but Illustra-Vis and specifically the skin illustration supports the depth and shape perception. A participant who chose MIP for patient-doctor communication argued that the other 3D techniques show too much information. Another participant discusses the main advantage of Illustra-Vis: “The main advantage is you have a complete image of the patient’s body and you could do some prediction about what is normal or not from the image. For example: when you have an obese patient you could overestimate the shape or size of an organ and make a wrong diagnosis.” Additionally, a participant argued: “Illustra-Vis focuses on relevant areas (PET) and is not occluded by tissues as in Multi-Vol. Furthermore, it gives a fast overview of the PET activity as well as a good localization as the surrounding bones are illustrated.” Further comments were: “Illustra-Vis is more easily accessible for non-professionals.” and “It is an illustrative way to show patients their diagnosis and possible treatment planning.”

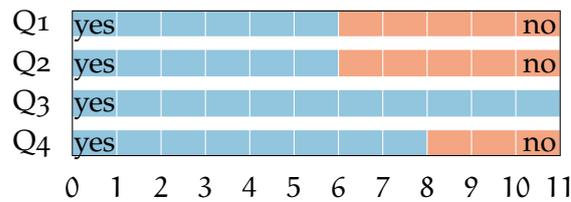


Figure 41: Would you use the Illustra-Vis technique in: diagnosis (Q1), treatment planning (Q2), patient-doctor communication (Q3), interdisciplinary communication (Q4).

5.6 DISCUSSION

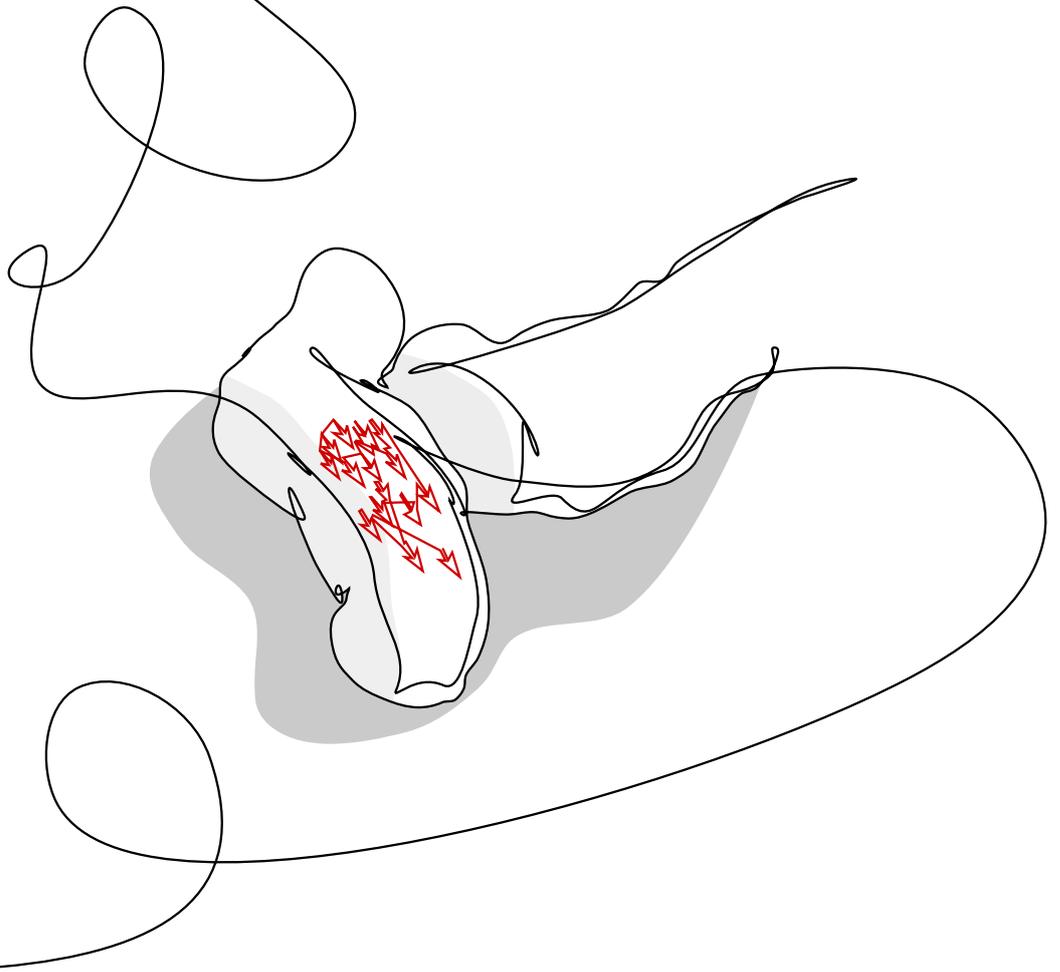
In the evaluation, the new technique was compared to the standard maximum intensity projection (MIP) approach and the multi-volume rendering (Multi-Vol) approach. Illustrating PET/CT with the MIP approach results in a dominant representation of the PET such that the PET region of high intensities are always drawn on top of the bones. MIP is a good technique to get a fast and simple overview of the data and many medical experts already have previous experience with MIP. Mostly medical experts use this technique to identify areas of interest and afterwards, they use the 2D slice view to analyze the affected regions. Nevertheless, this approach does not provide depth cues and therefore the slice view is needed in addition to accurately locating the suspicious activity correctly. Multi-Vol rendering, unlike MIP, can ensure that the correct drawing order of the structures is shown and has thus improved depth perception. This technique however, can result in occluded regions such that the important PET activity is not visible anymore under all viewing angles. The new approach illustrates the surrounding regions of the PET activity without causing occlusion, while providing more shape cues than MIP. This can be especially useful when the PET activity of interest has a lower intensity than the surrounding structures. With Illustra-Vis it is possible to have a contextual overview by means of visualizing the skin and bones as well as identifying the PET foci in a single view. Using a different rendering technique around suspicious metabolic activity, allows the users to identify the regions more quickly and precisely. In case the PET activity is occluded by bone, the color of the bones were attenuated such that the PET is always perceivable. Theoretically, Illustra-Vis combines the advantages of MIP (no occlusion of the PET activity) and Multi-Vol (improved depth perception) in a single comprehensive solution. During the evaluation, while most experts preferred Illustra-Vis and noted the advantages, several medical experts still preferred the MIP. These experts state that they appreciate the simple view a MIP provides and are accustomed to interpreting this representation, while they are not used to Illustra-Vis.

5.7 CONCLUSION

In this chapter, an illustrative visualization technique for PET/CT datasets was presented. Different illustrative techniques were used to represent the functional PET activity as a focus area combined with the anatomical CT information as a context object. First, feature line techniques were applied to the CT dataset and demonstrated that the results of the presented approach are comparable to other state-of-the-art feature line techniques. Afterwards, a halo around the PET activity was created by smoothing as a simple way to create a region of interest in the dataset. This region was used as the basis for a lens metaphor that allows the user to see the inside of the body surrounding the PET activity. Here, different shading techniques were used to illustrate the surroundings of the PET to give a better indication of the spatial relationships to the anatomy and thus a better way to localize the suspicious PET activity in the data. Furthermore, the presented illustrative multi-volume rendering was applied to different datasets to show its flexibility. Besides freely available clinical data, the technique was also applied to a research dataset, acquired using a Micro-CT scanner. Additionally, the method was applied to a single CT dataset, using contrast-enhanced vascular information as the focus object to demonstrate the flexibility of the method. A qualitative user evaluation showed the utility of the presented approach for treatment planning, doctor-patient communication, interdisciplinary communication, and diagnosis (when combined with 2D slice information). The evaluation revealed that the presented approach gives a fast and clear summary and can be used as a basis for further analysis. In summary, a novel illustrative visualization technique was developed, designed with PET/CT data in mind, which is potentially applicable for general (multi-)volume focus and context applications.

6

BLOOD FLOW VISUALIZATION



This chapter is partly based on:

Kai Lawonn, Sylvia Glaßer, Anna Vilanova, Bernhard Preim, and Tobias Isenberg.

“Occlusion-free Blood Flow Animation with Wall Thickness Visualization.”

In: IEEE Transactions on Visualization and Computer Graphics 22 (1).1 (2015), pp. 728–737.

An aneurysm is a blood-filled expanded bulge in the wall of a blood vessel that may occur everywhere in the arterial system, predominantly at the abdominal aorta and at the cerebral arteries. Inherent to aneurysms is the risk of rupture which may result in fatal consequences for the patient. This chapter focuses on cerebral aneurysms, i.e., those that occur in the brain. Cerebral aneurysms result from a pathological weakness in the vessel wall, causing a dilation of the vessel structure. The initiation and outcome of such a cerebrovascular disease comprises different factors such as hemodynamic information [CMWP_{11b}; MFKH+₁₂] and the thickness of the vessel wall [KDM₁₃]. Treatment options have to be evaluated compared to the aneurysm's rupture risk because treatment may cause a rupture as well. Current research in medical imaging has made great advances, in particular for intravascular data for (the larger) cardiac vessels. Moreover, wall thickness assessment will likely be possible for human cerebral arteries in the near future. Visualization techniques and tools are thus essential that are capable of presenting all this information in an integrated view. Furthermore, it is vital to incorporate the behavior of the time-dependent flow as well. The relation between the vessel morphology and the near-wall flow is particularly important. Near-wall flow may lead to a pathological vessel dilation, which may increase the risk of a new rupture at locations where a former rupture occurred. The main challenge is, therefore, to combine visualization techniques that simultaneously illustrate both the outer wall (i. e., the border between vessel wall and surrounding tissue) and the inner wall (i. e., the lumen and vessel wall border) to allow physicians to get an impression of the wall thickness.

Current visualization solutions can either illustrate the blood flow with the surrounding vessel or depict the wall thickness without integrated blood flow, but a combined visualization was not possible so far before. Therefore, the first visualization approach tailored to the occlusion challenges by presenting wall thickness with blood flow is presented. A visualization technique based on cutaway surfaces, such that occlusion problems with respect to the internal blood flow are resolved, is introduced. This technique is integrated into a framework that was developed in collaboration with experts from neuroradiology, computational fluid dynamics, and medical illustration. Furthermore, the focus lies on illustrating the blood flow by simultaneously depicting it with the vessel wall as shown in Figure 42. The visualization technique and the framework were developed based on hand-

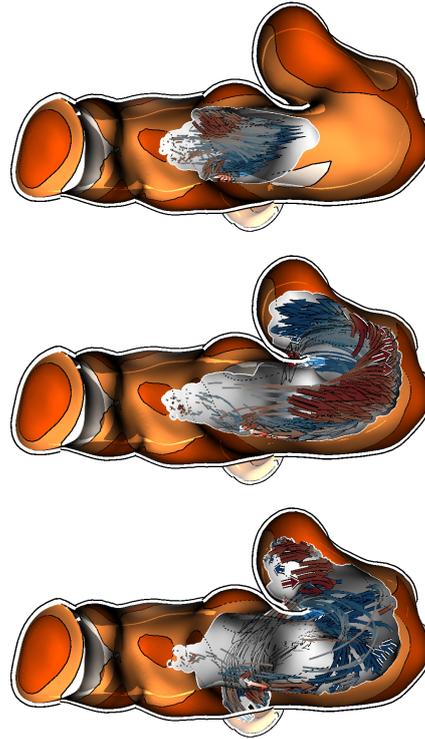


Figure 42: Three subsequent time steps of the animated pathlines in a vessel. The cutaway technique facilitates insight into the vessel during the animation of the blood flow. Furthermore, arrow glyphs represent the flow's pathlines.

drawn images provided by medical researchers. These images indicate how the experts would envision the tool, and how they would illustrate the aneurysms with all the mentioned structures. In summary, the main contributions are:

- A design study for the visualization and exploration of internal blood flow and wall thickness that is based on fruitful discussions with medical experts and their hand-drawn images.
- A novel illustrative visualization techniques for animated pathlines and the surrounding vessel wall thickness in a single, integrated, and interactive view.
- Collaborative tools to improve the exchange of information.
- Results of a qualitative evaluation that revealed the advantages of the developed techniques, which was compared with state-of-the-art methods.

6.1 RELATED WORK

Related publications comprise, in particular, the visualization of embedded structures, the visualization of blood flow, and the visualization and exploration of aneurysms.

VISUALIZATION OF EMBEDDED STRUCTURES. Diepstraten et al. [DWE03] proposed different cutaway approaches to achieve interactive frame rates and Viola et al. [VKG04] presented a view-dependent cutaway surface generation. In these methods, each part of the volumetric data has an object importance. Whenever unimportant structures occlude more important ones, an insight (i. e., a cutaway) is generated to ensure the visibility of important structures. Later, Liang et al. [LCKM05] introduced an approach to interactively cut a hole in a surface to reveal important structures behind it. As the use animated blood flow is employed, an automatic approach to cut away a surface is needed. Li et al. [LRAC+07] also proposed exploration techniques by interactively defining target structures which their system uses to generate a cutaway illustration. Burns and Finkelstein [BF08] introduced an adaptive cutaway approach to reach real-time performance which is also applicable to dynamic objects—useful for the goals of visualizing blood flow. An interactive framework to create cutaways based on sketching input was introduced by Knödel et al. [KHG09]. This framework, however, requires elaborate user interaction, which makes it too complex for medical researchers. Sigg et al. [SFCP12] proposed an approach to optimize the visibility of specific target objects. They use a view-dependent objective function to approximate the optimal solution of positioning cutaway objects. Lidal et al. [LHV12] used predefined objects to generate a cutaway surface for encoding context information on it to distinguish layers of geological models. For this medical scenario, however, predefined objects may be too large such that the vessel morphology is cut away. Pindat et al. [PPCP13] introduced *Gimlenses*, a visualization technique for the interactive exploration and navigation through complex 3D models.

VISUALIZATION OF BLOOD FLOW. Mattausch et al. [MTHG03] presented different approaches for the interactive exploration of flow. With this approach, flow features can interactively be analyzed by mapping them to different visual attributes, e. g., width and opacity. Gasteiger et al. [GNBP11] introduced the *FlowLens*, an interactive visualization tool for the exploration of data on and inside a vessel. It can be used to explore the interior blood by showing the animated pathlines. However, with this approach only blood flow in a certain area—defined by the lens—can be observed. Thus, it is difficult to follow the flow outside the lens. Neugebauer et al. [NLBJ+13] proposed an approach for the qualitative exploration of near-wall blood flow. They used several 2D widgets to simplify the underlying streamlines at different positions on the surface. Van Pelt et al. [PBBC+10] introduced a framework for the interactive exploration of 4D phase-contrast magnetic resonance imaging (PC MRI) data that represents cardiac blood flow. In this framework, cross-sections can interactively be translated along the vessel and serve as seeding points for the pathlines. For the interactive virtual probing of 4D blood flow, van Pelt et

al. [PBBC+11] introduced another framework to position a probe semi-automatically. This tool focuses on the blood flow and does not visualize the vessel morphology. Born et al. [BMGS13] used illustrative visualization methods to provide a simplified blood flow representation. The main aspect of their approach is to simplify the blood flow without representing the vessel morphology. An overview on the visualization of simulated and measured flow data can be found in the summaries by Preim and Botha [PB13] and Vilanova et al. [VPPG+14].

VISUALIZATION AND EXPLORATION OF ANEURYSMS. Based on contrast-enhanced computed tomography (CT) data, Higuera et al. [HSTN+04] introduced the automatic adjustment of bidimensional transfer functions, using measured intensities and gradient magnitudes to visualize aneurysms. Neugebauer et al. [NGBD+09] employed a 2D map display for the exploration of the cerebral aneurysm. They used a scalar field to decode information on the surface. They then visualized the surface as a 3D model as well as a flattened map surrounding the surface. While this approach enables users to explore the entire scalar field of the surface from a single viewpoint, no blood flow information is provided. Gasteiger et al. [GNKP10] visualize aneurysm surfaces based on Fresnel shading, which allows a simultaneous display of vessel morphology and blood flow. However, they shade the vessel parts which face away from the view direction more opaquely, thus occluding the interior blood flow information. Lawonn et al. [LGP14a] proposed a novel surface visualization based on suggestive contours. This shading technique distinguishes between convex and concave regions according to the viewpoint. Glaßer et al. [GLHS+14] presented a framework for the visual exploration of vessels with wall thickness and hemodynamic information. Here, they visualized the wall thickness but a visual representation of the interior blood flow was missing.

None of these approaches, however, provide a combined visualization of embedded structures and the interior dynamic blood flow. Thus, different visualization approaches need to be integrated into a framework, combining the advantages of this prior work for vessel visualization. Moreover, this work adds a dynamic view of the internal blood flow.

6.2 MEDICAL BACKGROUND

A brief summary with the medical background and the image data acquisition for evaluating cerebral aneurysms is presented. The rupture of a cerebral aneurysm can lead to a subarachnoid hemorrhage where blood flows in the ventricular system. This results in a high mortality rate of 45%–75% [RRWG+08]. Therefore, research is essential to improve treatment options. The evolution and initiation of the disease comprise different factors such as genetics, hemodynamics, ves-

sel morphology, inflammation, and further factors that are not well understood yet [AML^P+08; FPPK+04]. According to recent studies [CMWP_{11b}; MFKH+12], domain experts reported certain flow patterns to be important indicators for the initiation and outcome of a cerebrovascular disease. In particular, they observed a correlation between complex and instable blood flow patterns and the aneurysm's rupture risk. An effective visualization is thus necessary to allow medical researchers to assess the vessel hemodynamics. The needed multi-dimensional data can be derived from computational fluid dynamic (CFD) simulations or measured using 4D PC MRI with ultra-high field scanners (7 Tesla and beyond).

For hemodynamic analysis, neurosurgeons reported that an aneurysm's rupture risk depends on its wall thickness [KDM₁₃]. While a non-invasive in-vivo imaging technique to acquire the cerebral aneurysms' vessel wall is still missing, intravascular imaging such as intravascular ultrasound (IVUS) is currently employed in cardiology to depict the vessel lumen and wall morphology. Because it is not yet certified for use in the cerebral vasculature, current clinical research focuses on transferring intravascular imaging from cardiology to neuroradiology.

In this section, a real-world IVUS dataset with measured wall thickness and simulated blood flow (hemodynamic dataset) is used. In addition, several aneurysm datasets are used. For this, the internal vessel geometry was measured and the blood flow was simulated in this geometry. Afterwards, the wall thickness was artificially generated to approve the functionality of the presented approach on several datasets. The framework allows for both medical analyses: the evaluation of internal blood flow as well as the wall thickness.

An **IVUS dataset** was employed that was generated by measuring the wall thicknesses in an artificially created aneurysm in a porcine renal artery using IVUS (see Glaßer et al. [GLHS+14] for more details). Hence, a branching artery was dissected and the shorter branch was artificially closed to provide a saccular aneurysm. Then, the IVUS dataset was acquired by probing along the vessel with the ultrasound probe at a constant velocity. The resulting stack of 2D grayscale images was used to generate a 3D surface and to extract the wall thickness.

The geometry of **aneurysm datasets** was derived from computed tomography angiography (CTA), magnetic resonance angiography (MRA), and 3D rotational angiography. The high vessel-to-tissue contrast in all these modalities facilitates a vessel segmentation using a threshold-based segmentation followed by a connected component analysis. By applying marching cubes to the segmented image data, a 3D surface was generated. In the preparation of a geometric model for CFD simulation, it is necessary to manually correct artifacts because these can lead to erroneously merging branches (details can be found

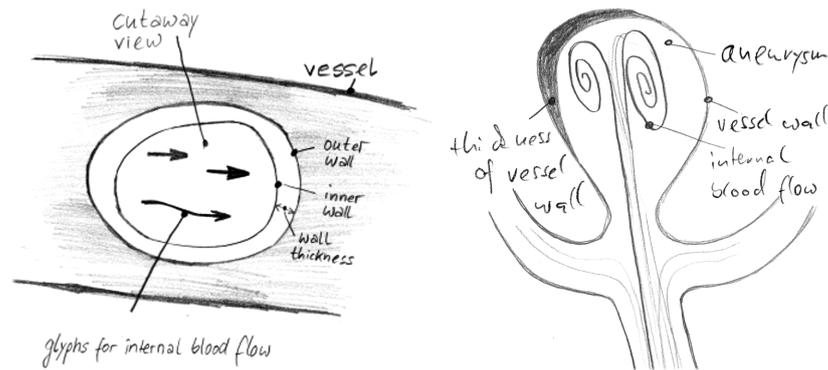


Figure 43: Hand-drawn example sketches by domain experts who were asked for a cutaway view providing internal blood flow as well as a simultaneous visualization of the wall thickness.

in [Gas14; MNP11]). To use these datasets, the meshes were manually smoothed by applying the technique of Mönch et al. [MLKW+13] and then the inner vessel wall was extruded to the outside to be able to visualize the wall thickness.

For both of the previously described dataset types, a separate simulated **hemodynamic dataset** is necessary to represent the blood flow. Using the generated or obtained 3D models, a tetrahedral grid was determined using *Netgen* (see Schöberl [Sch97]). Based on this grid, a CFD simulation was applied to approximate the interior blood flow. Blood is considered to be a Newtonian fluid, so the Navier-Stokes equation was used in the simulation (for details see Cebral et al. [CCAP+05]). Based on this flow dataset, the pathlines were extracted using an adaptive Runge-Kutta scheme [TGE97]. The reader is referred to the thesis of Gasteiger [Gas14] for a comprehensive overview including visualization and exploration techniques.

6.3 REQUIREMENT ANALYSIS

For developing a visualization framework that is able to depict the wall thickness with animated blood flow several requirements need to be fulfilled. The approach is based on discussions with four domain experts: two physicians and two CFD engineers who work on hemodynamic analysis. The participants were asked to produce hand-drawn images, indicating how they would imagine a visualization that illustrates both the wall thickness and the internal blood flow, see the examples in Figure 43. The drawings were discussed and the participants were asked to state the most important quantities that need to be visualized. As the surface morphology and the flow behavior influence each other, the experts stated that they are interested in an integrated view such that both modalities are shown. They also commented that strong local changes of the vessel wall thickness are important. Moreover, they stated that the vessel wall at irregu-

lar shapes, e. g., blebs at daughter aneurysms, is essential. The vessel wall's changes indicate remodeling due to plaques and are thus essential to assess the risk of rupture. Furthermore, they mentioned that the flow behavior might be an important aspect as it may initiate or evolve a cerebrovascular disease. Therefore, the vessel morphology is an important requirement for the visualization (see also [GNKP10]). Furthermore, the experts commented that they need to know the behavior of the flow near the wall. They stated that near-wall flow may lead to pathological vessel dilation and, therefore, needs to be depicted for further analysis (see also [GJ12; GTBP+10]). Based on these discussions, drawings, and the analysis of the state of the art, the main requirements for a visual exploration of vessel characteristics (i. e., wall thickness with internal blood flow) were summarized as follows:

Req. 1: **VESSEL VIEW WITH WALL THICKNESS.** An overview of the aneurysm is needed, comprising the vessel morphology as well as wall thickness information. For the thickness, the domain experts want to be able to set up different thresholds during exploration. The experts stated that they prefer discrete colors for the wall thickness on the vessel, in contrast to using a continuous color encoding. This facilitates a semi-quantification for the detection of dangerous regions, i. e., at thin vessel regions.

Req. 2: **VISIBILITY OF FLOW.** The blood flow needs to be visible during the whole animation. Experts can then trace the flow without encountering occlusion problems, and without the need to interact.

Req. 3: **REVEALING NEAR-WALL FLOW.** During the exploration, the experts also need to understand the flow's distance to the vessel to distinguish between near- from far-wall flow: Near-wall flow may lead to a pathological vessel dilation, which increases the risk of rupture. To ensure that these flow regions are passed by the pathlines, the seeding of the flow is also very important—but this is not the focus of this chapter.

Req. 4: **FLOW ANALYSIS AT THIN REGIONS.** The analysis of the blood flow behavior is essential for thin regions. Experts thus state that they first analyze the flow at thin regions before observing the pathlines in general. Further analysis tools are hence required that permit them to focus on the flow at different regions while also showing a contextual overview of the dataset.

Req. 5: **UNDERSTANDING THE VESSEL MORPHOLOGY.** For distant or overlapping vessel parts it is important to provide depth cues.

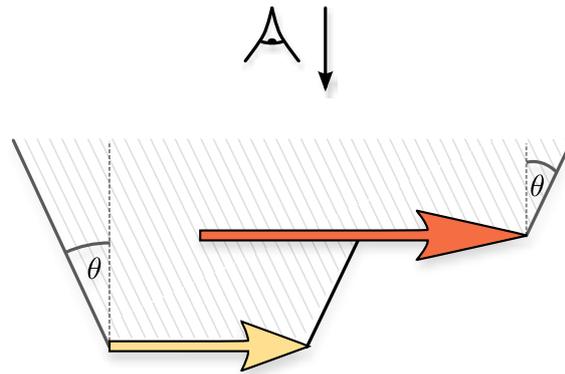


Figure 44: Arrows of the flow seen from the side (i. e., the viewpoint is on the top). The opening angle θ defines the (gray) cutaway surface.

These support users in understanding the vessel morphology and improve the perception of depth and the spatial layout of the vessel.

Req. 6: COLLABORATION TOOLS. For collaboration purposes it should also be possible to set landmarks with different animation steps. An external researcher should be able to analyze the regions at the marked landmarks and discuss them with other colleagues.

6.4 BASIC VISUALIZATION ALGORITHMS

Based on these guidelines an integrated representation of both the blood flow as focus and the vessel geometry as context is presented. To meet Req. 2, an algorithm is needed that ensures the visibility of the flow. Thus a real-time cutaway surface is applied on the pathlines that guarantee an insight into the vessel to observe the animated blood flow is required. Therefore, the cutaway technique by Burns and Finkelstein [BFo8] is employed, which uses the algorithm by Rong and Tan [RTo6]. Therefore, the concepts of both approaches are summarized next and how they are applied to animated pathlines.

6.4.1 Cutaway Technique

A technique introduced by Burns & Finkelstein [BFo8] is employed, which is based on the depth image of the current scene and Rong and Tan's [RTo6] *jump flooding* algorithm. Thus, only the depth image of the objects of interest — in this case the pathlines — is used to create the cutaway. Initially, the depth image is set to the following: if a pixel p contains the object of interest R , its depth value is set to the pixel's computed z -value. Otherwise, it is assigned to 0: Next a conic object



Figure 45: The *jump flood* algorithm on a single arrow glyph with $i = \{0, 2, 4, 6, 8\}$ iterations. First, only the depth image of the arrow is generated. Then the jump flood algorithm is applied and a depth image is generated that can be used for the cutaway technique.

is created around the objects of interest R with an opening angle θ , see Figure 44. A cutaway surface function C is used:

$$C(p) = \max_{q \in R} (q_z - m \cdot \|q_{xy} - p_{xy}\|) , \quad (42)$$

where R contains the objects of interest, m encodes the slope with $m = \tan^{-1} \theta$, and p_{xy} is the pixel's current 2D location.

6.4.2 Jump Flooding Algorithm

The *jump flooding* algorithm is used to determine the cutaway surface function C , recall Equation 42, with real-time performance. Given an image of size $n \times n$, a standard flood fill algorithm would pass the content of a pixel to its 8 neighbor points. This algorithm needs at most n iterations to pass the content through the whole image. Contrarily, the *jump flooding* algorithm needs at most $\log n$ iterations. It changes the step size at every iteration i to $k = n/2^i$. At each iteration, a pixel at (x, y) thus passes its content to its eight neighbors at $(x + l_1, y + l_2)$ with $l_1, l_2 \in \{-k, 0, k\}$. Further refinements and details are described in Section 6.5.

6.4.3 Pathline Animation

Inspired by Everts et al.'s [EBRI11] illustrative line styles, arrow glyphs are used to represent the pathlines. In addition, line predicates, e.g., velocity or distance are encoded to the surface by color. The pathlines consist of points \mathbf{x} , each of which is assigned to two scalar fields: the integration time $t(\mathbf{x})$ and the distance to the surface $d(\mathbf{x})$. The integration time refers to the time at which a specific point is reached during the animation. Also $T_{\max} = \max_{\mathbf{x}} t(\mathbf{x})$ is set as the global maximum integration time of all pathlines. During the rendering process, the current time t_{cur} with $t_{\text{cur}} \in [0, T_{\max}]$ is known. The respective part of the pathline is drawn whenever the absolute value of the difference of the current time t_{cur} and the integration time $t(\mathbf{x})$ of a specific point \mathbf{x} is smaller than the user-defined threshold width. This approach visually emphasizes faster pathline parts compared to other (slower) parts: the faster a section of the pathline is, the wider

this part is drawn. Color assignment and glyph representation details are explained in Section 6.5.

6.4.4 Centerline Extraction for Slice View Analysis

A slice view is also provided to facilitate a detailed analysis of the wall thickness. Users can slide along the centerline of the vessel (determined using the approach by Antiga et al. [APBEI+08]) using a cutting plane, giving them two views of the wall geometry. First, the intersection is displayed on the cutting plane itself, providing an initial orientation in the context of the integrated visualization. Second, a separate 2D view shows an orthographic projection to give users an undistorted view of the vessel geometry and the wall thickness.

6.5 VISUALIZATION DESIGN FOR INTEGRATED DEPICTION OF VESSEL MORPHOLOGY AND BLOOD FLOW

In this section the implementation details are outlined and this section explains how the requirements derived in Section 7.3 were fulfilled. First, the pathline glyph visualization is addressed. Then, the realization of the *jump flooding* algorithm is described and finally the slice view is introduced.

6.5.1 Pathline Visualization

A novel glyph representation of the pathlines is used that emphasizes near-wall flow (Req. 3). The pathline visualization is based on view-aligned triangle strips. For these strips a width ω is defined to extend the lines to become quads. For every adjacent point pair $\mathbf{x}_i, \mathbf{x}_{i+1}$ on the pathline, the tangent vector \mathbf{t} is determined by subtracting the point and normalizing the resulting vector: $\hat{\mathbf{t}} := \frac{\mathbf{x}_{i+1} - \mathbf{x}_i}{\|\mathbf{x}_{i+1} - \mathbf{x}_i\|}$. The direction of the quad's extent is determined by computing the cross-product of the normalized view direction vector $\hat{\mathbf{v}}$ with $\hat{\mathbf{t}}$. With each incoming adjacent point pair $\mathbf{x}_i, \mathbf{x}_{i+1}$, four points are passed to the geometry shader:

$$\begin{aligned} \mathbf{x}_{\text{out}1} &= \mathbf{x}_i + \omega \cdot \hat{\mathbf{t}} \times \hat{\mathbf{v}}, & \mathbf{x}_{\text{out}2} &= \mathbf{x}_i - \omega \cdot \hat{\mathbf{t}} \times \hat{\mathbf{v}}, \\ \mathbf{x}_{\text{out}3} &= \mathbf{x}_{i+1} + \omega \cdot \hat{\mathbf{t}} \times \hat{\mathbf{v}}, & \mathbf{x}_{\text{out}4} &= \mathbf{x}_{i+1} - \omega \cdot \hat{\mathbf{t}} \times \hat{\mathbf{v}}. \end{aligned} \quad (43)$$

This scheme constructs a view-aligned quad. Furthermore, texture coordinates are passed to the fragment shader to draw glyphs along the pathline. Here, the width encoded as a parameter, ranging from $[-1, 1]$ is needed only. Then, the texture coordinates are used to create the arrow glyph. For this, the domain expert can set two different thresholds f_1, f_2 to distinguish between near-wall and far-wall flow. Here, the values are set such that $f_{\min} \leq f_1 \leq f_2 \leq f_{\max}$ holds,

where f_{\min} is the minimal distance from the pathline to the vessel and f_{\max} corresponds to the maximum distance. The expert can thus set different regions R1, R2, R3 based on f_1, f_2 . The distance f is in region R1 if $f \in [f_{\min}, f_1]$, in R2 if $f \in (f_1, f_2)$, and in R3 if $f \in [f_2, f_{\max}]$. Initially, the values are set as $f_1 := 0.1 \cdot (f_{\max} - f_{\min}) + f_{\min}$; i. e., 10% of the nearest flow is set for the first value and 50% is set for the second value: $f_2 := 0.5 \cdot (f_{\max} + f_{\min})$. During the animation, glyphs are drawn that are in the range $[t_{\text{cur}} - \frac{\text{width}}{2}, t_{\text{cur}} + \frac{\text{width}}{2}]$ and, thus, set $u \in [0, 1]$. The width is set by the user and defines the width of the glyph.

The arrow glyph shape for the different regions R1, R2, R3 is defined as follows. For R1 with $f \in [f_{\min}, f_1]$:

$$v(f, u) := \begin{cases} 0.25, & \text{if } u \in [0, 0.5) \\ 2 \cdot (-u + 0.9), & \text{if } u \in [0.5, 0.9) \\ 0, & \text{if } u \in [0.9, 1]. \end{cases} \quad (44)$$

For R2 with $f \in (f_1, f_2)$:

$$v(f, u) := \begin{cases} \frac{f-f_1}{f_2-f_1} \cdot (-0.125) + 0.25, & \text{if } u \in [0, 0.5) \\ 2 \cdot (-u + 0.9) \cdot \frac{f_2-f}{f_2-f_1}, & \text{if } u \in [0.5, 0.9) \\ 0, & \text{if } u \in [0.9, 1]. \end{cases} \quad (45)$$

For R3 with $f \in [f_2, f_{\max}]$:

$$v(f, u) := 0.125, \quad \text{for all } u. \quad (46)$$

Here, f is the distance of the pathline's fragment to the vessel. In R2, the width of the contour of the pathline is changed slightly from 0 to 0.1 with $0.1 \cdot \frac{f-f_2}{f_1-f_2}$ illustrated in black. In R1, the offset is slightly changed from 0.1 to 0.2, while a white contour is also generated around the arrow from 0 to 0.1. Figure 46 illustrates this change. Furthermore, a small depth offset is added for the halo to obtain results similar to the depth-dependent halos by Everts et al. [EBRI09]. Finally, the glyph color is set according to its second line predicate, e.g., the distance to the wall. This color is defined by determining the difference of the vessels' back and front face depth and the difference of the pathlines' depth with the front face depth. This creates an encoding that shows whether the pathline is near the current vessel in view direction or not—recall Figure 42 and see Figure 52. Here, a color scale ranging from dark red to dark blue with white in between is used. Thus, depending on the pathline glyph, the expert can differentiate between near-wall flow from far-wall flow, and furthermore he can distinguish if the flow is near the front face of the inner vessel or the back face of the inner vessel, see Figure 47. Therefore, the changing style of the pathline glyphs as well as the color coding allows to detect near-wall flow and hence fulfills Req. 3.

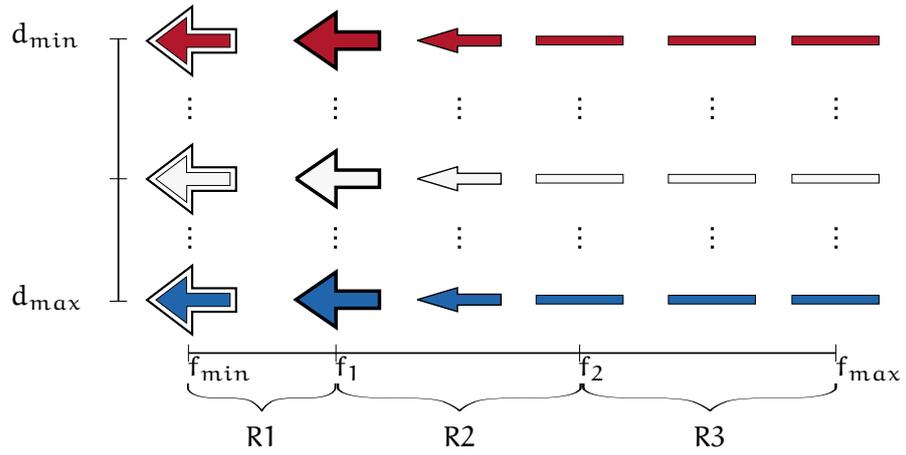


Figure 46: Illustration of the continuous change of the pathline glyphs from far-wall to near-flow flow. R1, R2, and R3 represent the different regions.

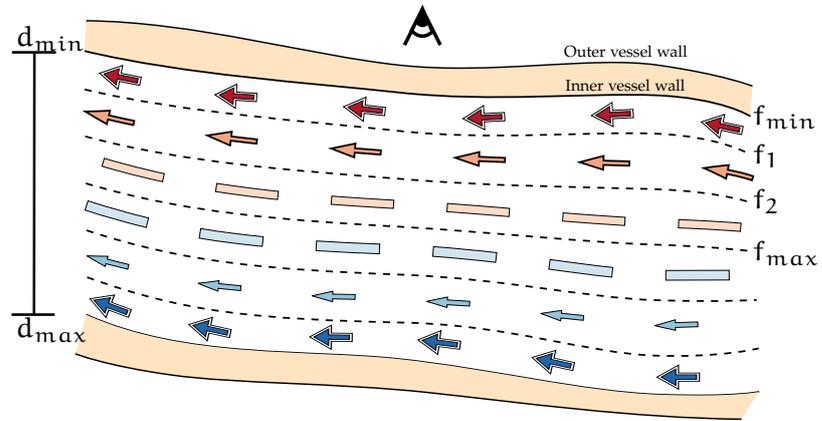


Figure 47: Illustration of the applied pathline glyphs with color coding. Near-wall flow is represented as arrows with a white contour. The flow near the camera appears dark red, whereas near-wall flow very far to the backside of the vessel is encoded as dark blue.

6.5.2 Jump Flooding Details

To determine the cutaway surface function C (see Equation 42) a *compute shader* and the OpenGL extension `SHADER_IMAGE_LOAD_STORE` is used. This technique is applied to the pathlines to increase the visibility of the flow (recall Req. 2). First, the depth image of the objects of interest is determined and is loaded into the `COMPUTE_SHADER` using frame buffer objects (FBOs). Then, the *jump flooding* algorithm is applied and the values are stored in the `RGBA` channels of a cutaway texture using the OpenGL function `imagestore()`. The `RG` channels store the 2D position (x, y) of the pixel that maximizes Equation 42. These coordinates change during the iterative process. Initially, these values are empty. The `B` channel carries the depth of the neighbor result-

ing from the algorithm described in Section 6.4.2, which is initialized with their own depth. The A channel stores the $C(p)$ value. In the first iteration ($i = 1$), the step length is $k = n/2$. Thus, for a given pixel p its (maximum) eight neighbors q^1, \dots, q^8 are considered. Then, the cutaway surface function is determined for p and its neighbors:

$$C(p) = \max_{j \in \{1, \dots, 8\}} (q_z^j - m \cdot \|q_{xy}^j - p_{xy}\|) . \quad (47)$$

Let us assume the pixel q^j maximizes $C(p)$. Afterwards, $C(p)$ is compared to p_w (which stores the depth of the best neighbor that maximizes Equation 42 from a previous iteration step). Its value is kept if the current depth p_w is greater than $C(p)$. Otherwise, $p_w := C(p)$, $p_{xy} := q_{xy}^j$, and $p_z := q_z^j$ is set, where $xyzw$ is encoded in the RGBA channel. Therefore, the pixel p stores the information of the depth of the cutaway surface, and furthermore it records the position and the depth of the neighbor that maximizes C . The algorithm terminates after $\log n$ iterations. The iteration process is shown in Figure 45.

6.5.3 Cutaway Illustration

The cutaway ensures the visibility of the flow (related to Req. 2). It is visualized using *frame buffer objects* (FBOs). The complete scene consists of images that are depth-sorted in a dedicated shader which uses the depth of the cutaway surface to determine if the outer or inner vessel is drawn. A thin white border is drawn around the cutaway to highlight the flow in the vessel; recall Figure 42. During the animation and interaction, the user is thus able to focus on the flow as it is emphasized.

6.5.4 Slice View

As described in Section 6.4.4, the centerline provides a natural choice for the slice view navigation, which is therefore related to Req. 1. The centerline consists of consecutive points $\mathbf{p}_i, \mathbf{p}_{i+1}$ and the tangent $\hat{\mathbf{t}}_i = \frac{\mathbf{p}_{i+1} - \mathbf{p}_i}{\|\mathbf{p}_{i+1} - \mathbf{p}_i\|}$ is used as the normal of the cutting plane. Having a user-defined parameter s for the slice view in the range $[0, 1]$, the exact position \mathbf{q} of the plane is first determined with:

$$\begin{aligned} j &= \lfloor n \cdot s \rfloor \\ \mathbf{q} &= \mathbf{p}_j + (n \cdot s - j) \cdot \hat{\mathbf{t}}_j. \end{aligned} \quad (48)$$

This plane divides the space into two half spaces. For every triangle of the vessel geometry it is tested if two of its three points lie in a different half space than the third one. In this case, the intersection points of the corresponding edges is determined with the plane and both points are connected. This geometry intersection is represented in the form of lines. Then the camera is positioned on the tangent vector

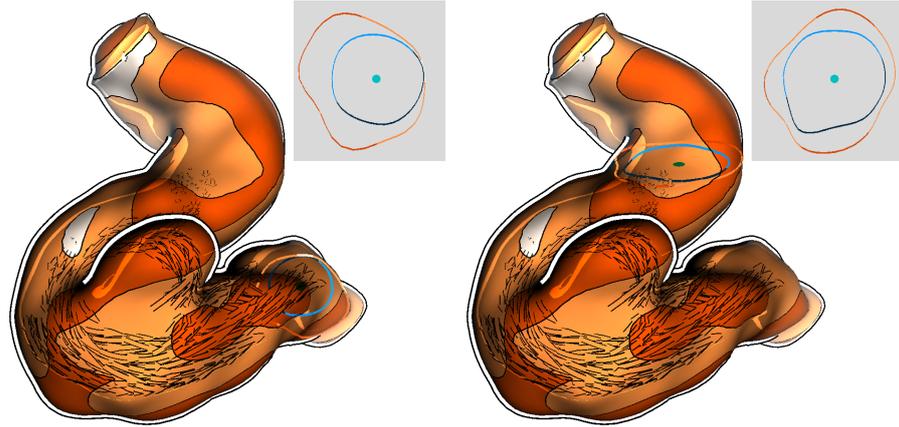


Figure 48: The slice view facilitates the thickness estimation. Here, two different positions of the slice view are represented.

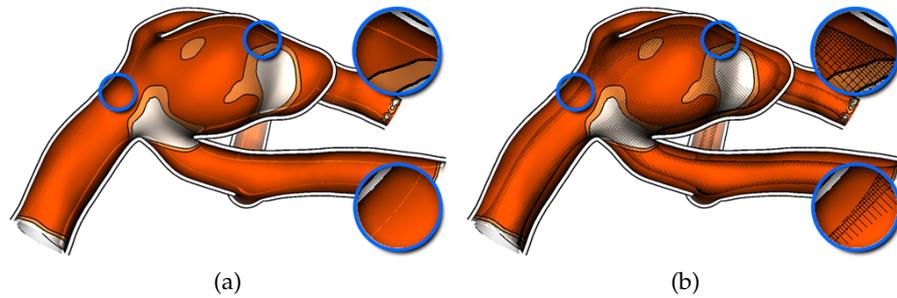


Figure 49: Different illustration techniques that visualize the inner vessel part. In (a) the contour of the vessel is illustrated, whereas in (b) it is illustrated using image-based hatch strokes.

and this view is added as an additional 2D widget in the framework, see Figure 48.

6.5.5 Vessel Illustration

In the following, the vessel illustration with the wall thickness visualization is described to satisfy Req. 1. The vessel visualization is divided in two parts: (a) the visualization of the outer vessel wall and (b) the visualization of the inner vessel wall. For the *inner vessel* wall, the user can choose between two different rendering styles. The first version is an illustrative margin around the contour (Figure 49a). This visual encoding provides contextual information of the inner structure. The alternative is the use of image-based hatch strokes, which are generated at the contour margin of the surface (Figure 49b). The hatching is parameterized by the dot product between the view direction and the normal, using more strokes for smaller dot products (i. e., closer to the actual silhouette). Figure 49 directly compares the two versions and Figure 54b shows another example of the hatching



Figure 50: Four colors representing the wall thickness from low to high. The user is free to adjust this number of discretized colors to his/her needs.

technique. The backfaces of the inner vessel are also drawn at regions only where the cutaway surface provides a view into the vessel. The *outer vessel* is first visualized with a scalar field representing the thickness, ranging from white to brown, see Figure 50. For the thickness computation per vertex \mathbf{v} of the outer vessel, the algorithm iterates over all vertices of the inner vessel wall and determines the Euclidean distance to \mathbf{v} . The smallest value is assigned to \mathbf{v} . The user can select the thickness values in order to discretize the thickness information. In addition, the opacity of the discretized values can be set as well. This supports the domain expert to investigate the flow at specific regions. Here, the domain expert can set the opacity at thin regions to 0 and can thus analyze the flow at these regions only. As default values, 0, 0.2, 0.4, and 0.8 are set with regard to the regions ranging from the thinnest to the thickest region. This allows to explore the near-wall flow in these regions and, therefore, satisfies Req. 4. In addition, screen-space ambient occlusion [Kaj09] is implemented to improve the perception of the vessel through depth cueing following Req. 5.

6.5.6 Depth Attenuation

To increase depth perception, an atmospheric attenuation was added to objects that are far away from the camera according to Req. 5. For this purpose, the depth of the current fragment z_f is determined. It is normalized with the minimum z_{near} and maximum z_{far} depth values in eye space: $z_n = \frac{z_f - z_{near}}{z_{far} - z_{near}}$. If z_n exceeds 0.45, $5 \cdot (z_n - 0.45) \cdot (1, 1, 1)$ is added as an RGB color to the current fragment to slightly increase its brightness. The threshold is set to 0.45 because a lower value would result in parts close to the camera becoming too bright, and a greater value would not affect the depth attenuation much.

6.5.7 Framework

This section describes the framework of the visualization tool. The GUI is divided in two parts: the standard GUI with all important settings and the playbar. The standard GUI provides pathline-related and vessel-related features. For the pathline settings, the user can adjust the animation speed, the pathline length, and the opening angle θ of the cutaway surface. The user can also choose between the novel glyphs and normal arrows as pathline representation. Furthermore,

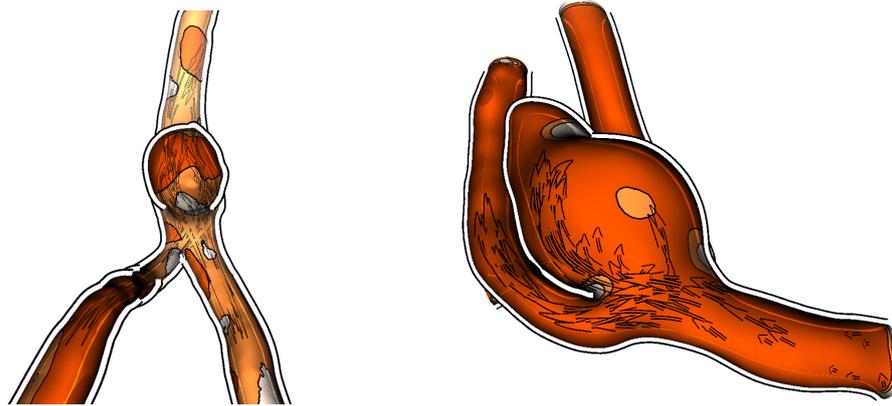


Figure 51: Different views to analyze the surface where pathlines serve as context information.

if the user wants to analyze the vessel wall surface, the pathline visualization can be turned off. The cutaway can also be deactivated to explore the surface with the pathlines as context objects, as shown in Figure 51. This might be important as the expert needs the pathlines only as context information to explore patterns inside the vessel at thin regions.

The settings for the visualization comprise facilities to adjust wall thickness regions, to adjust distance parameter for the pathline glyphs and highlighting of thin regions. Thin regions exhibit a higher risk of rupture, making these regions relevant for investigation as indicated in Req. 1 and Req. 3. Thus, the user can adjust a parameter such that those parts of the vessel are highlighted which fall below this thickness value. This approach facilitates a fast perception of potentially dangerous zones, even if the vessel includes many thin regions such that color coding may not be sufficient.

Req. 6 indicates that save and load facilities (for data, views, parameter settings, and animation) are needed to support the collaboration between medical experts. Furthermore, a feature that allows them to place landmarks on the surface is included. After the placement, a screenshot is taken and all important parameter settings and the camera properties are stored. Another domain expert can see the landmarks, represented as red spheres, and can select them. A short preview is then shown in the top right corner as depicted in Figure 52a. A click on it loads the scene as shown in Figure 52b. In this case, the domain expert made a very important finding. She discovered that the flow is close to the vessel wall, in a region where the wall is very thin.

For the visual exploration of the pathlines, an interface similar to typical movie players was created. Thus, standard buttons with their well-known functionality as well as a timeline to facilitate a fast navigation through the animation was provided. Users can click on the

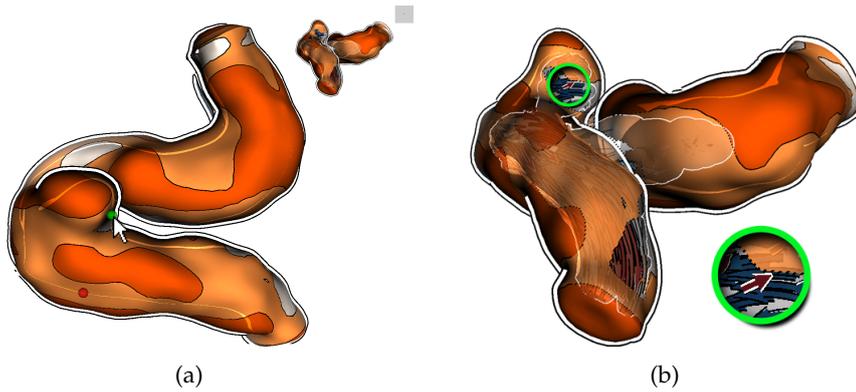


Figure 52: A domain expert created landmarks, represented as red spheres. Another expert can select the landmark, resulting in a small preview (a). A click on it loads the dataset with the first domain expert's settings (b).

timeline and the animation jumps to the corresponding keyframe. Furthermore, dragging the animation is possible for fast navigation through the simulation. Consecutive keyframes can also be shown by clicking on the *previous* or *forward* button.

6.6 COMPARISON WITH RELATED APPROACHES

Closest to the new visualization techniques is our own previous work [LGP14a] and [GLHS+14]. Lawonn et al. [LGP14a] presented a shading technique that enhances the surface while near-wall blood flow is emphasized by a change of the surface color. While this approach highlights the surface regions where the blood flow is near the vessel wall, it does not visualize the wall thickness. Glaßer et al. [GLHS+14], in contrast, focused on the wall thickness and used different visualization concepts to analyze dangerous regions based on the wall thickness, yet without showing the internal blood flow. Thus the novel visualization technique is compared to these other two concepts.

6.6.1 Near-Wall Visualization

The vessel visualization by Lawonn et al. [LGP14a] is based on a suggestive contour measure to determine the shading (see Fig. 53a). Based on the surface's local shape (convex vs. concave) as well as on the distance of the pathlines to the vessel, they determine the rendering parameters (stored in a 2D lookup table). During animation, the changing flow leads to a highlighting of those vessel parts that are close to the blood flow. This visual encoding facilitates an analysis of surface regions that are close to fast blood flow. However, it only supports a qualitative estimation of the near-wall flow, as demonstrated with a comparison of our previous work with the new approach in

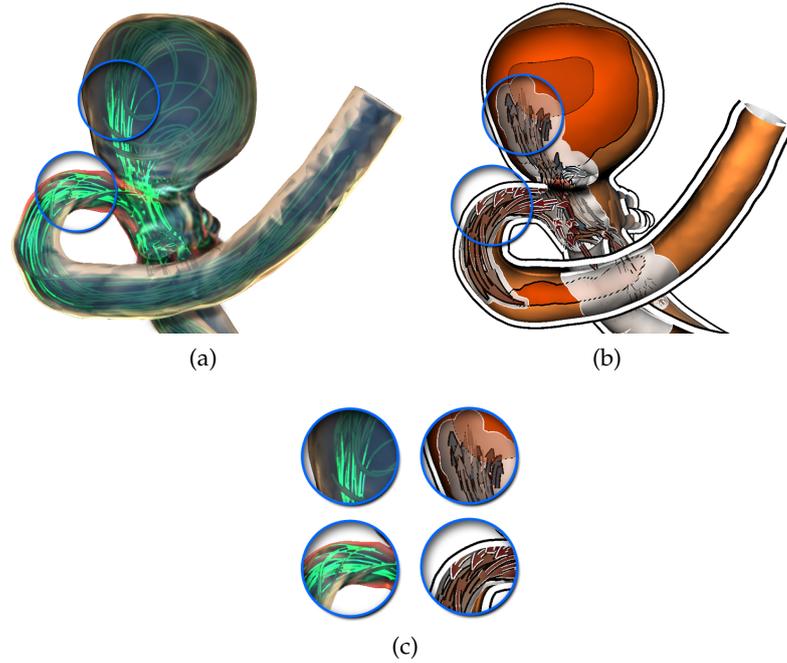


Figure 53: Close-ups (c) of the visualization technique [LGP14a] (a) and the presented new technique [LGVP+15] (b). The presented pathline glyphs support the illustration of the near-wall flow better since the arrow glyphs with the white contour indicate near-wall flow. Furthermore, it can be better distinguished if the pathlines are near the front faces or the back faces of the vessel wall, here colored in red and blue, respectively.

Figure 53. The first close-up in Figure 53c shows that the glyph-based approach reveals that the pathlines are close to the vessel wall — a fact that is invisible in our previous technique. The blue arrows in the novel visualization indicate that the pathlines are near the backside of the vessel. While the pathlines are also close to the wall in the visualization by our previous work, this fact cannot be perceived due to the region highlighting. Furthermore, unlike the new approach, important near-wall flow details may be overlooked in the visualization by our previous work, if the view does not show the right vessel region. This effect was demonstrated in the second close-up in Figure 53c: only the red color around the contour indicates the near-wall flow, but for surface regions directly facing the view direction this effect is hard to perceive. With the presented technique the white outline around the arrow shows that the pathlines are close to the vessel. Moreover, the red pathlines show that it is close to the facing vessel part.

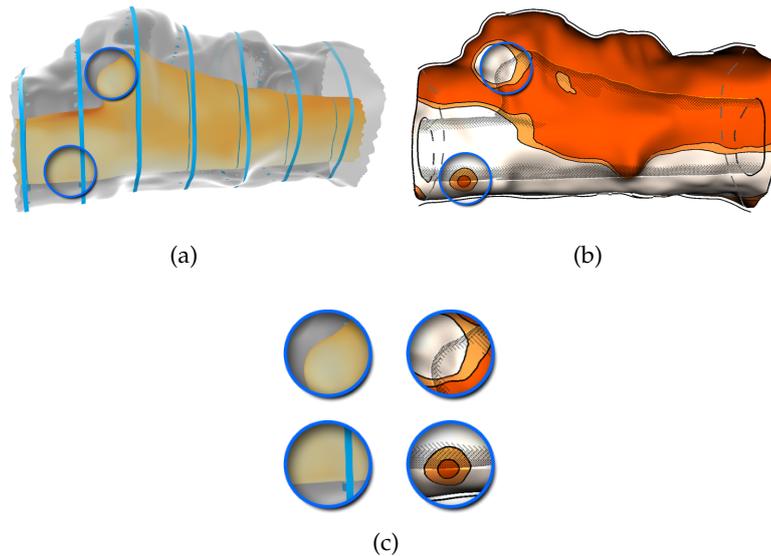


Figure 54: Close-ups of the visualization technique by Glaßer et al. [GLHS+14] (a) and the new technique [LGVP+15] (b) with the IVUS dataset. The novel wall thickness visualization allows to identify thin regions even if they are very small. Furthermore, the visualization of the inner vessel allows for a contextual visualization.

6.6.2 Wall Thickness Visualization

Glaßer et al. [GLHS+14] designed a framework for the analysis of the wall thickness. They used a slice view, distance ribbons, and an outer wall representation to visualize the thickness and to facilitate a qualitative analysis (Figure 54a). None of these techniques, however, allows physicians to measure the thickness quantitatively. Figure 54 directly compares the presented visualization technique with the method by Glaßer et al.; the two close-ups in Figure 54c demonstrate the advantages of the new wall thickness visualization. The first close-up illustrates a thin region, and both visualization techniques show the region well. However, with the presented discrete wall thickness representation a dangerously thin region can immediately be identified which can be overlooked in the continuous visualization by Glaßer et al. The related issue of identifying thick regions is shown in the second close-up in Figure 54c: the thick region can easily be missed with the approach by Glaßer et al.

6.6.3 Summary

The new technique allows physicians to analyze the wall thickness quantitatively while simultaneously representing the blood flow. By also depicting the pathlines with the glyph-based approach a quanti-

tative analysis of the flow near the vessel wall was facilitated. Therefore, near-wall flow can be perceived even if the flow is near the back-faces of the vessel (i. e., far away from the viewer). Furthermore, thin regions — independent of their size — can be perceived using the presented discrete visualization. In summary, compared to Lawonn et al., the presented technique is more appropriate to reveal near-wall flow (recall Req. 3). The method by Glaßer et al. allows a visual analysis, as brushing & linking is applicable, but in comparison to the presented technique, the new presented improves on the vessel wall thickness representation (recall Req. 1). The new visualization techniques combines the advantages of both approaches.

6.7 QUALITATIVE EVALUATION

The design process comprised several iterations with the domain experts. The first feedback cycle for the initial prototype is discussed that did not yet comprise the complete functionality presented so far.

Six unpaid participants (1 female, 5 male; 26–45 years old; median 35 years) took part in this study. Among them were the four domain experts that already helped us with defining the initial concept: two physicians (14 and 15 years of professional experience) and two CFD engineers involved in hemodynamic analysis (both 5 years of professional experience). To get additional independent feedback, two researchers with background in medical visualization were added.

The framework was briefly presented and its features to the participants were described. Next, the participants were encouraged to explore the tool on their own, while their aural comments and their interaction with the framework was recorded. Finally, they were asked to fill out a questionnaire.

6.7.1 Questionnaire

The questionnaire consisted mainly of questions about medical aspects, the presented visualization, the framework design, and the parameter choice. Depending on the question, the users were asked for yes/no responses, ratings on a five-point Likert scale (—, —, ○, +, ++), or open-ended comments.

Demographics/background. First, the typical participant demographics: gender, age, and medical and/or technical experience were collected. Also the participants' background in computer-assisted analysis of aneurysms and about their experience in analyzing simulated blood flow was inquired.

Medical questions. Next, the participants were asked about the importance of the wall thickness for medical evaluation and about the importance of pathlines. They were also asked about which elements

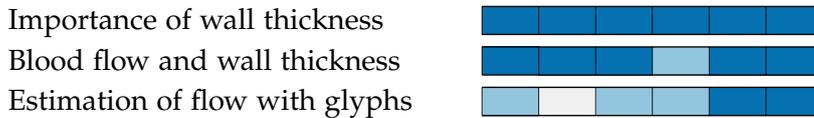
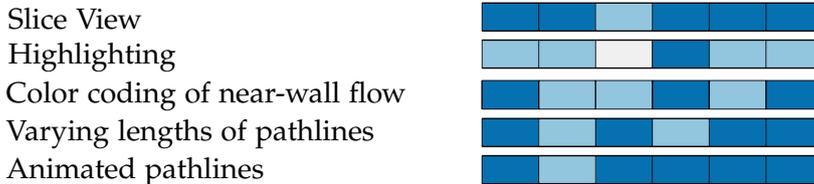
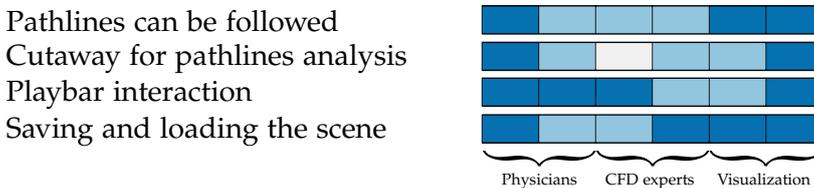
Individual colors for Likert score:**Medical questions:****Visualization questions:****Framework aspects:**

Figure 55: The results of the evaluation with color-encoded Likert scores. Each box represents the answer of one participant: the two left-most boxes depict the physicians, the middle two boxes show the CFD experts, and the two right-most boxes represent the medical visualization experts.

are important for a diagnosis of cerebral aneurysms and how this is currently visually supported.

Visualization questions. The participants were asked to comment on the different kinds of wall thickness visualizations that were provided. Specifically, the participants were asked about the slice view and about the highlighting of surface regions with walls thinner than a given threshold. Next, they were asked whether the color coding of the pathlines is useful to highlight near-wall flow. The users were also asked if the glyph-based representation with varying length allows to differentiate between slow and fast-moving flow. Then, they were inquired about the utility of the playbar for an effective temporal data exploration. Finally, they were asked if the cutaway visualization makes it possible to follow the animated pathlines and if it is helpful for the analysis.

Framework aspects. Here the users were asked if they had any problems when exploring the vessel with internal blood flow. Furthermore, the participants were inquired if the saving/loading options are useful for the analysis as well as for the collaboration with other medical experts. They were also asked whether any problems occurred during the exploration and about potential suggestions for

the future development of the technique. Then the users stated if they could imagine to use the framework for their research. Finally, additional remarks were asked.



Figure 56: Dynamic cutaway visualization of pathlines within the context of the vessels and aneurysm: Three steps within the cardiac cycle.

6.7.2 Results and Discussion

The overall result of the empirical evaluation is shown in Figure 55. No participant had any major problem with using the presented framework. After a short introduction, the participants could immediately use the tool. All stated that they would like to use the tool for the exploration of wall thickness and blood flow. In particular, the domain experts confirmed the high importance of analyzing the wall thickness for the rupture risk assessment of aneurysms. They also stated that the combination of wall thickness and blood flow visualization is important for further analysis of cerebrovascular diseases. Next, selected aspects and insights from the evaluation will be discussed.

Cutaway. All participants stated that the cutaway is useful to follow the animated blood flow and that it provides a natural view into the vessel. One physician stated that this animation is confusing due to the constant change of the cutaway. Hence, pausing of the animation was required to stop this change. Nevertheless, he appreciated the use of the cutaway for analyzing the pathlines. The usefulness of being able to follow the pathlines and to analyze the interior flow behavior was rated by the physicians with ‘++’ and ‘+’.

Glyph-based pathlines. Most participants described the color coding of the pathlines as essential to be able to distinguish between near-wall and far-wall flow. All participants appreciated the length encoding of the glyphs to represent the velocity. One of them stated that it is very helpful for still images because this representation conveys that the flow is slower in the aneurysm compared to the flow in the vessels. Moreover, one can directly see the flow direction.

Wall thickness. The results show that participants preferred the slice view over the 3D view to understand the wall thickness. This may be due to the fact that, in the first iteration, a continuous color coding was used to encode the thickness. The highlighting was also considered as essential to assess dangerous zones. Participants mostly

used the highlighting first; later they used the slice view for a more detailed analysis of the thickness.

Playbar. All participants found the playbar useful: they could explore all time steps and analyze the pathlines. Mostly, they first played the whole animation and then looked for specific points in time for further examination using the animation bar. Finally, the participants used the forward and backward buttons for further detailed analysis.

Application domains. The participants were also asked about application domains they could envision for the presented tool. Both physicians explained that the presented tool is useful for aneurysm risk assessment and for treatment planning. All participants also agreed that another important application field is research and education. Some participants mentioned that the presented technique may also be helpful for patient information and may be used to explore the predicted blood flow after different treatment options, such as coiling and stenting.

Performance. For all presented models (e. g., Figure 42, 51, 56, and 57), a real-time performance of approximately 41 fps was achieved. A laptop with an Intel Core i5 with 2.6 GHz, 8 GB RAM, and an nVidia GeForce 870M with 4 GB was used. This performance test was conducted with a resolution of 1920×1080 and the half of the resolution (approx. 960×540) was used for the cutaway surface.

Improvements. As already indicated, based on this evaluation and concrete suggestions several additional visualization techniques and encodings were added. One participant stated that it might be difficult to distinguish those vessels that are near to the camera from other vessels which are further away. Therefore, the depth attenuation described in Section 6.5.6 was added. Another participant wanted us to add features to only visualize the inner wall, to visualize the centerline, and to enlarge the slice view on a mouse-over. All these ideas were added and he liked the improvements. Then, the wall thickness representation was changed from a continuous color coding to a discrete color coding with parameters that can be set by the experts. Next, the new glyph-based representation of the pathlines was proposed to better distinguish between near-wall and far-wall flow. Furthermore, colors were added to make it clear if the pathlines are near the front or the back faces. As suggested by the medical experts, many additional collaboration tools were added such as the ability to add, store, and retrieve landmarks. Finally, the hatch-based illustration of the inner vessel wall was added since one expert had asked for a more contextual visualization.

6.8 EXPERT FEEDBACK

To evaluate the improvements, another evaluation similar to a focus group [LFH10] was conducted. Here, the participants had the oppor-

tunity to test the improved framework and to analyze its functionality. Then the tool was discussed with three researchers with background in CFD simulation (males; 27, 30, and 35 years old) and one physician with an emphasis on surgery (female; 29 years old). The CFD researchers are abbreviated with C₁, C₂, and C₃ and the physician with P₁. The participants were asked to make positive as well as negative comments. The interviews took approximately 40 minutes.

First, all participants stated that the analysis of the wall thickness with the animated blood flow is a great extension to state-of-the-art visualization frameworks. Then, there were asked to share their impressions of the framework. P₁ stated that it would be a nice feature to have animated camera paths to navigate around the vessel. C₂ had the idea that it might be an interesting feature to activate multiple linked views. Then, one could activate different parameter settings like wall shear stress and the user can analyze different parameters simultaneously.

Furthermore, the physician wished to have the original image data in the slice view. The CFD experts also wanted to see the original data to estimate the uncertainty of the vessel reconstruction. In fact, the experts stated that the presented technique is applicable to both the analysis of hemodynamic information according to near-wall flow as well as wall thickness. Moreover, they stated that the glyph-based approach is very helpful to observe where the flow is very near the wall. Setting own thresholds makes it easy to explore the dataset. P₁ explicitly emphasized the advantages of the slice view as certain regions can be observed. C₃ mentioned that it would be nice to have dashed contour lines whenever the flow is occluded by a vessel. In addition, C₃ stated that the color coding as well as the glyph representations with the white contour allows a fast perception of the near- and far-wall flow. Furthermore, he said that the color coding makes it easy to distinguished between the near-wall flow that is close to the front or the back faces. The participants also pointed out that a main benefit of the new approach is its placing of landmarks, which easily allows to collaborate with other experts. Here, C₃ stated that it would be nice to include text at the landmarks. With this, the mouse-over not only shows a small preview, but a text is also shown, which states the specific observation made by the expert. This was also stated by C₂ who said that the initials of the domain expert who placed the landmark should appear. This would be beneficial when several researchers work on the same data set. Furthermore, he recommended to blur the white line indicating the cutaway surface which would make the lines less obtrusive. P₁ stated that the color representation is a great benefit to easily distinguish between thin and thick regions. She also mentioned that the different glyphs make it easy to perceive near-wall flow and that the color coding helps to analyze the blood flow. In addition, P₁ said that seeing a blue glyph with a white con-

tour leads to a fast understanding where the pathline is located. In essence, P1 stated that dangerous regions can be easily identified and analyzed with the interior blood flow. Finally, the participants stated that a tool that includes the analysis of wall thickness and hemodynamics would be of great benefit to them.

6.9 CONCLUSION AND FUTURE WORK

In this chapter, a design study was presented comprising a visualization technique and an exploration tool for the analysis and exploration of vessels using an integrated wall thickness and internal blood flow representation. Some final results are shown in Figure 56 and Figure 57. Based on the requirements derived from consultations with medical experts and a literature analysis, a number of known rendering and visualization techniques with several novel technical concepts was combined. In particular, a technique for glyph-based visualization was contributed to represent data values along streamlines or pathlines. All these methods put together, it is—for the first time—possible to present an integrated illustrative visualization of all essential elements for aneurysm assessment—in particular including a representation of a vessel's wall thickness. The presented visualization framework thus allows users to analyze the pathlines using interactive animations and without occlusion issues. Furthermore, modern visualization methods were integrated to the surfaces and to the pathlines to depict further flow attributes such as stress forces and velocities. The presented technique makes use of illustrative elements such as color coded arrow glyphs to represent the flow direction, velocity, and pathline distance even on still images.

Thanks to the GPU implementation real-time rendering speeds were achieved even on mid-class computers. The visualization technique was integrated into a data exploration framework that facilitates the exploration of the vessel data. To assess the results, a qualitative evaluation with the collaborating experts was conducted. This evaluation provided evidence that the overall approach is valid, confirmed many of the design choices, and provided several additional paths for improvement—some of which were realized, presented, and re-evaluated with further domain experts. Thus, a comprehensive illustrative visualization approach for the analysis and exploration of aneurysms and their internal blood flow was provided that is applicable in medical research and other domains.

OUTLOOK Although the visualization technique was evaluated with positive feedback, some issues still need to be improved. First, while the slice view is helpful to get an impression of the wall's thickness at specific points, the medical experts are also interested in the original image data which are currently not displayed. Furthermore, to ana-

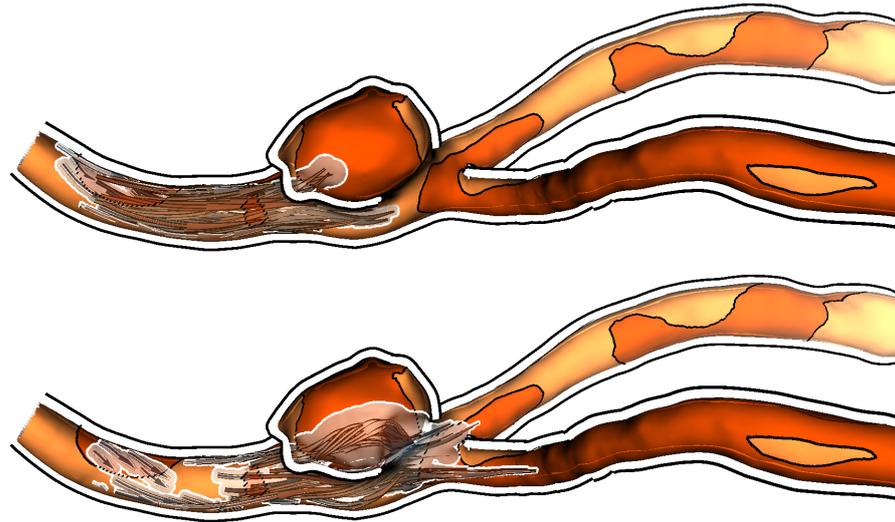


Figure 57: Two subsequent time steps of the animated pathlines in a vessel.

lyze the wall thickness it might be helpful to include a ruler to measure the thickness on specific parts of the vessel. More interaction can be integrated by selecting a region to obtain different quantities, e. g., the average wall thickness or the wall shear stress. As soon as modern imaging technology is available, additional information should be visualized on the inner vessel structure onto the cutaway surface such as calcified parts or plaque. This would allow medical experts to gain further insights into the vessel's structure. Finally, the CFD simulation should be improved to include biomechanic properties of the vessel such that dynamic changes of the vessels which could then also be included in the visualization tool.

Nonetheless, based on the described evaluations the presented method is considered as successful in filling the gap between visualizing wall thickness and illustrating hemodynamics with highlighted near-wall flow.

7

CARVING VISUALIZATION



This chapter is partly based on:

Kai Lawonn, Erik Trostmann, Bernhard Preim, and Klaus Hildebrandt.

“Visualization and Extraction of Carvings for Heritage Conservation.”

In: IEEE Transactions on Visualization and Computer Graphics 23(1) (2017), pp. 801–810.

VISUALIZATION AND EXTRACTION OF CARVINGS FOR HERITAGE CONSERVATION

IN the 13th century, artists have carved illustrations of the episodes from the church history and depictions of important persons into the plaster of cathedrals and churches. One prominent example are the carvings on the walls of the cloister of the Cathedral of Saints Catherine and Maurice in Magdeburg. Due to the open structure of the cloister, these carvings were subjected to weathering. Nowadays, the illustrations in parts of the plaster are hard to recognize and many structures are only visible with extreme side light. Therefore, restorers and historians are in need of a software tool that allows to analyze digital copies of the plasters, and to explore and edit the carvings virtually. In addition, virtual images visualizing the carvings and illustrations of the original shape of the carvings are attractive for visitors of the historic site.

In this chapter, techniques addressing the technical challenges are proposed that have to be solved for creating such a tool and an evaluation of the associated software is presented by the relevant target group.

This approach is in contrast from methods for general sunken-relief extraction. The carvings are special since the structures are canal-like and flat. Aging and weathering have additionally flattened the structures so that they can barely be perceived with the naked eye. Furthermore, the surface of the plaster exhibits irregular structures and the data contains noise. The amplitude of the noise and the height of the structures almost match the depth of the carvings. Both factors complicate the extraction of the carving structures.

To enable the processing of sunken-relief extraction, the plasters and carvings were first digitalized using a passive stereo vision system, resulting in a triangle mesh approximating the surface of the plaster. For analyzing the digitized plaster, a detection filter for carvings in surfaces was developed that searches for parts with a distribution of principal curvatures that exhibits characteristic for carvings, i.e., locally cylinder-like structures. This approach is based on the Frangi filter [FNVV98] from medical image processing, which is used for detecting vessel-like structures in images. This method was adapted to detect carvings in 3D surfaces. The novel method is specifically designed for detecting canal-like structures and is robust against noise. For filtering the set of detected carving structures and selecting the most salient structures, a scheme is proposed that combines connected component analysis and an importance measure

for the saliency of the carvings. Finally, visualization techniques are employed to emphasize the detected carving structures, see Figure 58. The software implementing these methods has been tested by domain experts in a user study that highlights the effectiveness and utility of the new approach. The experiments demonstrate that state-of-the-art methods for relief extraction, as well as feature line extraction and exaggerated shading, fail on the type of data that were considered. In summary, the following contributions were made:

- An algorithm for detecting carving structures in surfaces based on an adaptation of the Frangi filter and a scheme for ordering connected components of the detected carving structures by saliency
- A framework for interactive exploration of carving structures in surfaces that allows for selecting and deselecting structures, adding new carvings and visually highlighting them

7.1 RELATED WORK

One relevant topic in this field is the depiction of strong features via lines.

FEATURE LINES Methods for feature line extraction can be divided into view-dependent and view-independent approaches. Although, for carving extraction it is important to illustrate all features independent of the camera position, it is worth mentioning view-dependent approaches as well. Ridge and valley lines have the property that every point is an extremum of the principal curvature on the corresponding curvature line [Por94; IFP97]. The computation of these lines is delicate since it involves estimating third-order derivatives on the surfaces. Prominent approaches use surface fitting [OBS04; YBS05] or a combination of discrete curvature approximation and smoothing techniques [HPW05]. A technique for view-dependent feature line detection was introduced by DeCarlo et al. [DFRS03]. They extended the contour definition to determine salient features. Another view-dependent feature line approach was presented by Xie et al. [XHTS+07]. Their method depends on the underlying surface shading and allows for compensating noise by adding light sources. This approach was later improved [ZHS10]. A method that combines ridges and valleys with a view-dependent approach was proposed by Judd et al. [JDA07]. They defined view-dependent curvature quantities to apply their technique. Kolomenkin et al. [KST08] proposed a view-independent approach to determine the transitions from flat regions to ridges. Their approach can be used to enhance small features on archeological objects. Zhang et al. [ZHXX+11] introduced Laplacian lines, which are used to illustrate prominent features. The



Figure 58: On the left the original model with standard headlight. On the right the same dataset with the new visualization technique.

Laplacian of the surface normal field is pre-computed and the result is used for illustrating the features. Weinkauff and Günther [WG09] used the Morse–Smale complex of a curvature-based indicator function to construct networks of feature-lines. Robustness against noise is improved by simplifying the Morse–Smale complex using the concept of persistence [EHZ01]. Rieck et al. [RML12] developed persistence-based filters and topological signatures for the multivariate analysis of data and applied the techniques to the analysis of cultural heritage artifacts. While the method performs very well on artificial data, the processing of 3D-scans requires manual segmentation of the digital object and adjustment of parameters (see Sections 6.2.1 and 6.2.2 in [RML12]). A technique to illustrate surface features and enhance the spatial impression with lines was presented in our previous work [LMP13].

FEATURE-BASED SHADING A technique for enhancing surface features by modifying the surface normals was presented by Cignoni et al. [CST05]. Their idea is to smooth the normals and to compare them with the original ones: the larger the differences, the more salient the underlying geometric feature. Another approach for enhancing geometric features by modifying the surface normals was introduced by Rusinkiewicz et al. [RBD06]. They iteratively smooth the surface normals and store the results. Afterwards, the stored sequences of normals are combined to determine the shading. Lee et al. [LHV06] used an approach that automatically determines lighting positions to enhance features. Here, the surface is first segmented using a curvature-based watershed method. These patches are used for the placements of lights. Vergne et al. [VPBG+09] introduced a local shape descriptor that enhances salient features. They warped the incoming lighting in order to increase the view-dependent surface curvature information. This leads to a perceptually enhanced surface shape. For analyzing salient features on archaeological surfaces, Kolomenkin et al. [KST11] presented an approach to automatically process scanned artifacts. They defined a *prominent field*, which is a smooth field with certain properties at surface features to enhance geometrical features. Vergne et al. [VPBG+11] presented a technique called *Radiance Scaling* that improves the classical shading by enhancing features on the surface. They scaled light intensities based on curvature and material characteristics. An exaggeration technique to enhance relief structures was proposed by Miao et al. [MFWP12]. This approach is based on three features maps comprising local height, normal difference and mean curvature, which are then combined. Furthermore, the normals are adjusted to gain an enhanced rendering of the surface. Ammann et al. [ABGG+12] introduced an approach that analyzes several surface features per vertex. These features are used to combine different shading techniques to enhance salient structures on the surface. As this approach enhances all features, it may be prone to noise and even highlight these undesired regions. For the illustration of complex archaeological artifacts, Gilboa et al. [GTSK13] reported on the development of computerized automatic systems. Through automatic feature detection, they highlighted salient areas as line drawings, colored images or enhanced 3D models. Sharma and Kumar [SK14] presented a way to investigate artifacts with machine-learning techniques. Their system employs the user's preferred visualization settings to generate a reasonable result. Several techniques that depend on various vertex features, e.g., curvature, ambient occlusion factor, are applied to gain an appropriate shading result.

RELIEF EXTRACTION AND GENERATION There is a large body of literature on relief extraction and generation and the reader is referred to the survey by Kerber et al. [KWCZ+12] for an overview. A

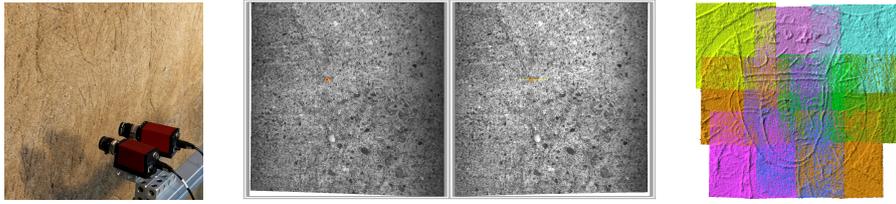


Figure 59: Stereo camera setup (left) and corresponding images (middle) for the computation of a 3D point cloud (patch). On the right the registration result for nine false-colored patches is shown.

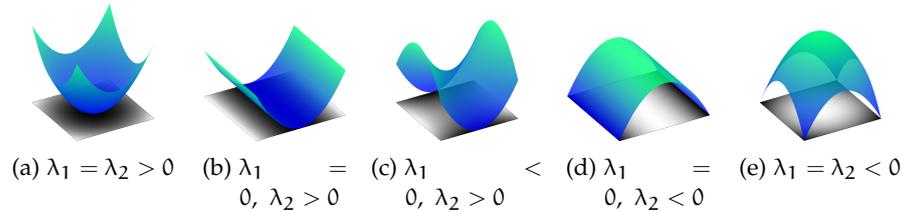


Figure 60: Different greyscale images are illustrated with their corresponding height surfaces. The height surfaces' principal curvatures are listed below. For $\lambda_1 = \lambda_2 = 0$ the image is constant and the height map is a plane.

general introduction to 3D shape analysis for archaeological surfaces was provided by Tal [Tal14]. Closest related to this work is the relief extraction scheme by Zatzarinni et al. [ZTS09]. They represented the surface containing the relief as a height field over a base surface. The relief is found through an analysis of the height field. In contrast to this scheme, their method aims at reliefs with enough height (or depth). In the case of the carvings considered here, the ratio between depth of the reliefs, amplitude of noise and natural structures in the surface and the (very low, but present) curvature of the surface make it difficult for their method to extract the carvings.

7.2 DATA ACQUISITION

The incisions cover an area of 37 m by 5 m outside on a wall to the courtyard of the Magdeburg cathedral. Because of current restoration work, an existing scaffold could be used for the data capturing. Varying lighting conditions and unavoidable vibrations of the scaffold require a rapid measurement process. A passive stereo vision system was used because it requires only a pair of synchronously captured images.

The measurement area was chosen based on the required resolution of less than 0.1 mm. Using two grayscale cameras with a resolution of 2048×2048 pixels the system was designed to capture an area of about 170 mm by 170 mm, see Figure 59. The calculation of

calibration parameters was accomplished with the OpenCV library. A precisely measured calibration plane with circular marks was used. The centers of the images of the circular marks can be localized with high accuracy. This leads to a reliable estimation of the intrinsic and extrinsic camera parameters. Testing the spatial intersection of the calibrated stereo vision system using a stochastic pattern on a plane surface (calibration standard) results in residual errors up to a maximum of 0.02 mm, which is acceptable for this application.

To compute pairs of corresponding points in both images of one capture, the acquisition was started with an evenly distributed field of points on one image. Then the task is to find the best correlation between a squared mask around that point on image one and the projection of this mask along the epipolar line on image two. To save computation time and to increase robustness, this line is trimmed by taking into account the depth of field (sharpness). Because of the slightly different viewing angles and varying imaging characteristics, the correlation algorithm has to be robust against differences in dynamics and the average intensity. Along with other approaches, the normalized cross-correlation introduced by Moravec [Mor97] gives reliable results. It is invariant against additive and multiplicative intensity changes in the search window. The spatial intersection for all rays of corresponding points for one capture leads to a resulting 3D point cloud of up to 4 million points.

To reconstruct a larger area, the results of a number of captures need to be combined. This is accomplished by manually picking three corresponding 3D points within two patches and a subsequent refinement using the Iterative Closest Point algorithm [BM92]. Since the patches overlap by about 40%, there are enough unique features in the point set to ensure a robust registration. A drawback of this method is that the overall shape of the wall becomes slightly deformed after registration of a large number of patches. This is not affecting the visual perception of the result, but a comparison to point data captured by 3D laser scanning shows a significant deviation of several centimeters on a wall with 2.5 m in length.

7.3 REQUIREMENT ANALYSIS

Based on a literature search and discussions with domain experts, the following list of requirements were devised:

- A method for detecting carving structures
- Interaction techniques that allow the user to modify the set of carving structures, e.g., erasing wrongly detected carvings or adding carvings
- A visualization technique that enhances the carvings

In the real-world plaster, the carving can barely be perceived with the naked eye. Only under special lighting conditions, the restorers can get an impression of the carving structures. To circumvent these difficulties, the whole plaster with high resolution was digitalized and the domain experts were allowed to process the data such that the carvings can be illustrated. Since with standard visualization of the scan data, it is still difficult to see carvings, the experts asked for visualization methods that can emphasize the carvings. The experts are interested in interactively modifying and controlling the visualization, e.g., by selecting additional structures to be emphasized or manually removing falsely detected carving structures.

7.4 EXTRACTION AND GENERATION OF CARVINGS

For the analysis and processing of carvings in plasters, schemes for carve extraction and ranking were developed as well as methods for generating new carvings.

7.4.1 *Carve Extraction and Ranking*

Though the plaster that should be analyzed is attached to one wall of the cloister, it is not planar. Therefore, the carvings cannot be described by a height field over the plane. The digitalized plaster is treated as a curved surface, in which the carvings should be detected.

CURVATURE COMPUTATION For detecting carvings, the distribution of the curvatures over the surface is analyzed. Since the plaster is described as triangle meshes, the principal curvatures and curvature directions of triangle meshes need to be computed. Following Rusinkiewicz [Rus04], first a discrete shape operator in every triangle is computed, which measures the deviation of the normals across the edges of the triangle. Then, a discrete shape operator at the vertices is constructed by transforming and averaging the shape operators of the triangles. The eigenvalues of the shape operator are the principle curvatures λ_1, λ_2 , which are ordered such that $|\lambda_1| \leq |\lambda_2|$. When computing the discrete shape operator (in particular for noisy data), it is recommended to smooth the surface beforehand. A bi-Laplacian smoothing is used, which results in a robust curvature computation in the experiments. An alternative to smoothing the surface is to use larger neighborhoods for the curvature computation, see [CSM03; HP11].

FRANGI FILTER FOR SURFACES The carving structures are “small canals” in the shape. At the bottom of each canal, the geometry equals that of a round tube or a cylinder. This means one principal curvature vanishes and the other one is approximately one over the radius of

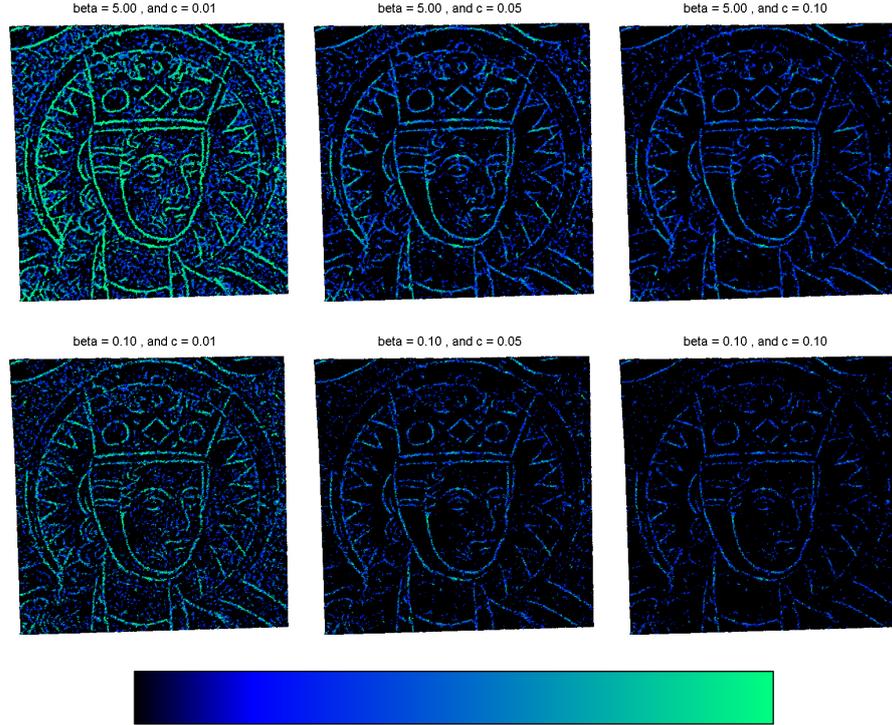


Figure 61: Results of the vesselness measure with different parameter settings. The colormap is shown at the bottom, where values range linearly from 0 to 1, which corresponds to black to cyan, respectively.

the cross section of the carving, i.e., $\lambda_1 = 0$ and $\lambda_2 \neq 0$. Figure 60 shows examples of surfaces and their principal curvatures. Since the carvings are engraved, the non-vanishing principal curvatures of all carvings have the same sign. Whether the sign is positive or negative depends on the orientation of the surface. The orientation is chosen such that the surface normals point outwards (out of the wall) and the curvature values are positive. For detecting the carving (or vessel) structures, analogous to the Frangi filter [FNVV98], the quantities $\mathcal{S} := \sqrt{\lambda_1^2 + \lambda_2^2}$ and $\mathcal{R}_{\mathcal{B}} := \lambda_1/\lambda_2$ are used and the *vesselness measure* is defined as

$$\mathcal{V}_0(s) = \begin{cases} 0 & \text{if } \lambda_2 < 0, \\ \exp\left(-\frac{\mathcal{R}_{\mathcal{B}}^2}{2\beta^2}\right) \cdot \left(1 - \exp\left(-\frac{\mathcal{S}^2}{2c^2}\right)\right). \end{cases} \quad (49)$$

In the discrete setting, the principal curvatures are computed at the vertices of the mesh and Eq. 49 is used to get a value of the vesselness measure at every vertex. Linear interpolation in the triangles yields a scalar field representing the vesselness measure on the triangle mesh. A blue to cyan colormap is used to display the influence of the parameters β and c . To depict the threshold, all values that are smaller than the threshold in black are additionally displayed.

PARAMETER ADJUSTMENT The vesselness measure, Eq. 49, involves two parameters, β and c . To help the users in choosing values for the two parameters, preview galleries were created. The users can choose preferred parameter settings by selecting an image. An example of an image gallery is shown in Figure 61. After a preview image has been selected, the corresponding parameter setting can be modified to fine-tune the result. The user-interface for this step is inspired by the work of Mönch et al. [MLKW+13]. When pressing the right mouse button, the user can alter the two parameters by mouse movements. Horizontal movements affect only one of the parameters and vertical movement the other parameter. The benefit of this interaction scheme is that the user can maintain focus on the image while changing the values of parameters.

RANKING THE CARVINGS Once the parameters β and c are set, the vesselness measure can be computed. For extracting the carvings, a connected component analysis is used, which clusters the vertices into components. Starting with an arbitrary vertex with $\mathcal{V} > 0$, all neighbors in the one-ring with $\mathcal{V} > 0$ were iteratively added to the same component until every vertex with positive vesselness measure is assigned to a component. For any component C_i , its saliency is defined as

$$\mathcal{V}_{C_i} := \sum_{\mathbf{p} \in C_i} A_V(\mathbf{p}) \cdot \mathcal{V}(\mathbf{p}), \quad (50)$$

where $A_V(\mathbf{p})$ denotes the Voronoi area of vertex \mathbf{p} , see [MDSBo2]. Finally, the components were ranked by their saliency.

ACTIVATING THE CARVINGS After connected components in the carvings structures are determined and the components are ranked by saliency, the user can use this information for selecting the components that will be enhanced in the rendering. The selection proceeds in two phases, first the user specifies a number n of components to be selected. The tool then selects the n components with highest saliency value. In the second phase, the user has full control over the selection and can manually add components to and remove components from the selection. To help the user in modifying the selection, The tool assigns colors to all components. Selected components are displayed in cyan and non-selected components receive a color that reflects their saliency value (cf. Eq. 50). The colormap for the non-selected components ranges from dark red to gray over orange, see Figure 62. Parts of the surface that do not belong to a carving structure receive a dark color. To further assist the user during the manual modification of the selection set, already activated components will be highlighted in red by a mouse over to indicate that clicking will delete the selection and not activated components will be colored in green to indicate that clicking will activate them.



Figure 62: Selection of carving structures and visualization with emphasized carvings. A standard rendering of the captured plaster is shown in (a). To help the user in selecting carving structures, the tool displays connected components of the detected carvings. Each component receives a color reflecting its saliency (b). The colormap is shown at the bottom, where values range from gray (minimum value) to red (maximum value). The minimum, maximum value correspond to the minimum, maximum saliency (cf. Eq. 50) in the set of detected components. The values in between are linearly interpolated. The user can activate the n components with highest saliency value. The result for $n = 4$ is shown in (c). Activated components are shown in cyan. To fine-tune the result, the user can further modify the selection set. In the final rendering, the selected carvings are emphasized (d).

VISUALIZING THE CARVINGS To emphasize the selected carving structures in the rendering, the shading of the selected components

is modified. The idea is to use the vesselness scalar field to weight the lighting in order to enhance the carving cues. Let \mathbf{l} denote the light vector. Then the shading is computed using $s = \langle \mathbf{l}, \mathbf{n} \rangle^{1+\alpha \cdot \mathcal{V}_0}$, where α expresses how strongly the carving cues are emphasized. For all results $\alpha = 2$ was used. Using a significantly higher value of α would result in stronger, but unnatural looking carvings. In addition to emphasizing the selected carvings, the tool allows users to deemphasize the non-selected structures. This is achieved by isotropically smoothing the normals of the surface (using 10 iterations of Laplace smoothing) and using the smoothed normal vectors to compute the shading everywhere except for the selected carving structures.

7.4.2 Carve Generation

In addition to visualizing the carvings, a tool for generating carving structures is provided. For example, to help with the recreation of lost parts of an illustration by adding new virtual carvings. Since editing the surface itself is a challenging problem which requires sophisticated modeling tools [FKSM+04]. To preserve the original surface, virtual carvings were added by overlaying the original plaster. This is realized by modifying the surface normals in the areas where new carvings are added. To specify the location and width of the virtual carvings, the user is provided with a brush tool. The size of the brush specifies the width of the carving. The result of the brushing is stored in a scalar field on the surface. Initially, this scalar field is 0 for every vertex. Whenever the brush is used, every vertex inside the brushed region in screen space is assigned a new scalar value. Denoting the distance of a vertex to the center of the brush by $dist$, the value $r_n = \frac{r-dist}{r}$ is assigned, which ensures that the midpoint assumes the maximum value. Whenever the user draws over brushed vertices, the already stored values are compared with the potential new values. The greater value is then assigned to the vertex. A scalar function $f(r_n)$ is used to determine the modified normals of the surface. The function serves as a desired height (or offset in normal direction) over the original surface. The height value at any point depends only on the value r_n assigned to the point during brushing. The function f depends on the carvings observed in the data. The goal is to make the newly created carvings look similar in style to the ones that are already in the data.

HEIGHT FIELD CALCULATION To obtain the generative model for new carvings, the existing carvings were analyzed, which are extracted using the scheme described in Sec. 7.4.1. The generative model represents the carvings as a height field over a smoothed version of the surface. The profile of the height field is analyzed and used to generate new carvings with a similar profile. The first step is to con-

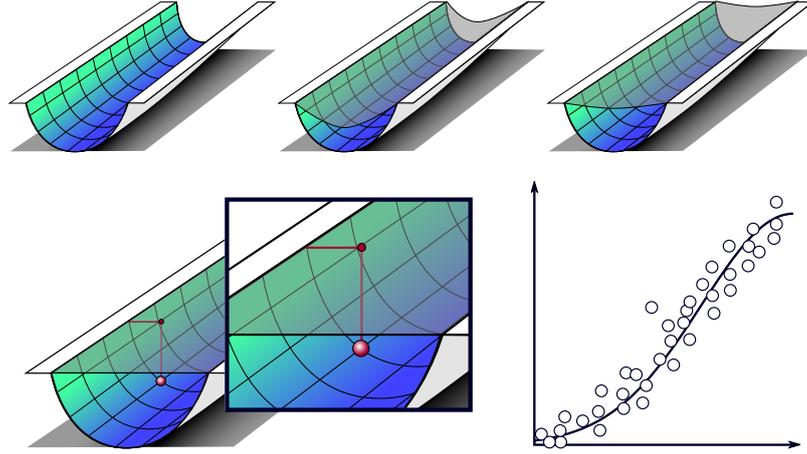


Figure 63: First, the carving is smoothed and the boundary is fixed (top row). Afterwards, every point on the smoothed carving is assigned the height distance to the original carving. The values are ordered according to their height and plotted. Finally an exponential function is fitted to the ordered height data.

struct a base surface over which the carvings are represented using a height field. For this, the surface is copied and the carvings is heavily smoothed by applying implicit Laplacian smoothing, i.e., mean curvature flow. The implicit smoothing is defined by the iteration

$$\mathbf{p}_i = (\mathbf{I} - \mu\mathbf{L})\mathbf{p}_{i+1}, \quad \mathbf{p} \in \mathcal{V}_e, \quad (51)$$

where \mathbf{I} is the identity matrix, \mathbf{L} is the Laplacian matrix, and μ is the weighting factor, which is set to 0.05. The iteration starts with the original vertex positions \mathbf{p}_0 . After each smoothing step, the Euclidean distance between the new and the old positions of the vertices is computed

$$D^{i+1} = \sum_{\mathbf{p} \in \mathcal{V}_e} \|\mathbf{p}_{i+1} - \mathbf{p}_i\|.$$

The smoothing flow converges when the quotient D^{i+1}/D^i tends to 1. The smoothing is stopped if $D^{i+1}/D^i > 0.99$. In practice, the smoothing process is terminated after 30-50 iterations. During the smoothing, the carvings evolve into their surface normal directions and towards the minimal surface spanned by the constrained boundaries.

The distance of the carving to the minimal surface is measured, see Figure 63. Explicitly, the intersection of the line generated by the point on the smoothed carving and the new normal with the triangles on the original carving is determined. The intersection point and the point on the smoothed carving yield the height distance. In case two or more intersections occur, the shortest distance is assigned to the point. Every point of the smoothed carving is assigned to a distance, i.e., the height scalar field. For every carving component, the distance to the x-axis is scaled in a range of $[0, 0.5]$.

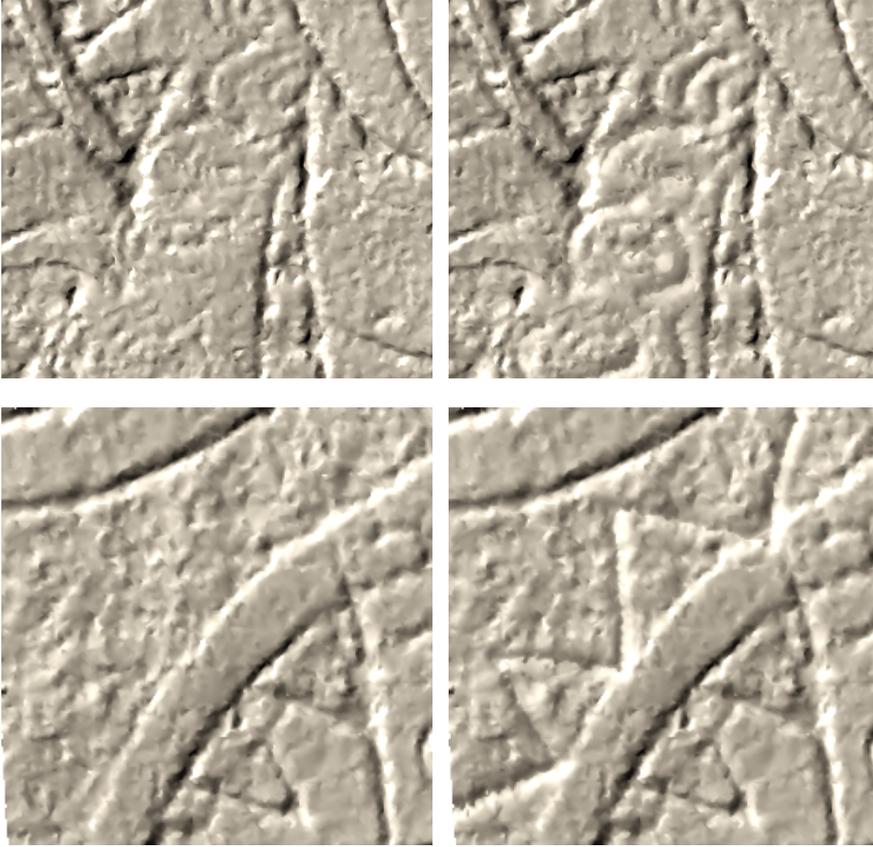


Figure 64: Two examples where carvings were added (right).

Then a Gaussian distribution $f(x) = a \cdot \exp(-((x - 0.5)/c)^2) + d$ in the range $x \in [0, 0.5]$ is fit to the data, see Figure 63 bottom right. The resulting function can be used as the input for the drawings on the surface: $f(r_n)$. This yields the scalar field \mathcal{C} . Based on this scalar field \mathcal{C} , the gradient per triangle is determined. Given a triangle $\Delta = \{i, j, k\}$, the gradient of every triangle is determined by:

$$\nabla \mathbf{g} = (\mathcal{C}_j - \mathcal{C}_i) \frac{(\mathbf{p}_i - \mathbf{p}_k)^\perp}{2A_\Delta} + (\mathcal{C}_k - \mathcal{C}_i) \frac{(\mathbf{p}_j - \mathbf{p}_i)^\perp}{2A_\Delta}, \quad (52)$$

where A_Δ is the area of the triangle Δ and \perp denotes a counter-clockwise rotation by 90° in the plane generated by the triangle, see [BKPA+10]. The gradient for every vertex is obtained by transforming and averaging the gradients of the triangles. This yields the vertex gradient of the scalar field \mathcal{C} .

The vertex gradients of the scalar field can be used to calculate the new normals of the carving structure. Given is a scalar field \mathcal{S} defined on a plane by $\mathcal{S} = f(x, y)$, $x, y \in \mathbb{R}$. It generates a height surface

by setting $z - f(x, y) = 0$. Then, the unit normal vector $\mathbf{n}(x_0, y_0)$ at (x_0, y_0) is

$$\mathbf{n}(x_0, y_0) = \frac{1}{\sqrt{1 + \left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}} \cdot \begin{pmatrix} -\frac{\partial f}{\partial x} \\ -\frac{\partial f}{\partial y} \\ 1 \end{pmatrix} \quad (53)$$

This can be simplified by determining the gradient of f : $\nabla \mathbf{g} = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)^T$ then, the normal is

$$\mathbf{n}(x_0, y_0) = \frac{1}{\sqrt{1 + \langle \nabla \mathbf{g}, \nabla \mathbf{g} \rangle}} \cdot \begin{pmatrix} -\nabla \mathbf{g}_x \\ -\nabla \mathbf{g}_y \\ 1 \end{pmatrix}. \quad (54)$$

As the height field is considered on the plane, the normal of this plane is $\mathbf{n}_p = (0, 0, 1)^T$. Therefore, the angle γ can be calculated between the plane normal \mathbf{n}_p and the normal determined by the height surface $\mathbf{n}(x_0, y_0)$:

$$\begin{aligned} \gamma &= \text{acos}(\langle \mathbf{n}_p, \mathbf{n}(x_0, y_0) \rangle) \\ &= \text{acos}\left(\frac{1}{\sqrt{1 + \langle \nabla \mathbf{g}, \nabla \mathbf{g} \rangle}}\right). \end{aligned} \quad (55)$$

Thus, this γ depends on the length of the gradient only. Hence, the normals of the generated carvings can be calculated with the scalar field. After the vertex gradients were determined, the smoothed normal is used to obtain the normals at the carvings. Eq. 55 is applied, which yields the angle. This angle is used to determine how much the smoothed normal needs to be rotated towards the vertex gradient in order to get the normal \mathbf{n}_e at the carving. Instead of using \mathbf{n}_e , it is mixed with the original normals. Otherwise, unnatural looking carvings may be obtained. A linear transformation of \mathbf{n}_e and \mathbf{n} is used to obtain the new normal $\mathbf{n}' = \mathcal{C} \cdot \mathbf{n}_e + (1 - \mathcal{C}) \cdot \mathbf{n}$, where \mathcal{C} denotes the scalar field obtained by the drawing on the surface $f(r_n)$. Thus, more natural looking carvings are obtained, see Fig. 64 (right).

7.5 EVALUATION WITH DOMAIN EXPERTS

To assess the utility of the proposed methods, an informal evaluation with three domain experts was conducted. Two of them are monument conservators and one is an expert in the digitization of cultural heritage sites and artifacts. The evaluation was conducted in three steps. First, the participants answered general questions concerning their experience in this field, and provided some personal data. In the second step, the possibilities that the proposed framework offers were explained to the experts. Then, they could use the software for analyzing the data set, selecting and editing carvings, and generating



Figure 65: A cutout of a model with different resolution, which corresponds to the number of vertices mentioned in Table 3.

new carvings. During this process, the participants were encouraged to think out loud and their comments were recorded. After testing the tool, the participants filled out a questionnaire concerning their assessment of the usefulness of the tool for their work and their impression of the usability and handling of the software. In the last step, the participants were asked to mention positive as well as negative aspects about the tool.

GENERAL ASPECTS The general aspects comprise questions about age, gender, and experience. The participants (P₁, P₂, P₃) are 65, 38, and 41 years old, respectively. P₁ and P₃ are males and P₂ is female. (P₁, P₂) work as monument conservators and have experience for 46 and 20 years, respectively. (P₃) is working for an institute that regularly acquires 3D data from cultural heritage objects. He has 20 years of experience. He is also responsible for preparing the data before the monument conservator can use it. The users were asked about the usefulness of having a tool that identifies the carvings. The participants' background in computer-assisted analysis of carvings and about their experience in analyzing 3D digitizations of carvings was inquired. For most of the questions, a 5-point Likert scale with answers (—, —, o, +, ++) was prepared. The question about the importance of the analysis of the carvings using digital 3D models was answered with '++' by all participants. The last question in this category was about the importance of the analysis of the carvings with this tool, which was rated with ++ by P₂ and P₃ and with + by P₁.

FRAMEWORK ASPECT After giving the participants the opportunity to use the tool, they were asked several questions concerning the usefulness and handling of this tool. In particular, they were asked about the possibilities for highlighting carvings and about the editing of the carvings. All participants stated that enhancing the visibility of the carvings is useful. P₁ and P₂ rated it with ++ and P₃ with +. Next, it was inquired whether the functionality to delete parts of the detected carvings is helpful. All participants rated this with ++. Another feature is that the illumination of the digital plaster can be modified by the positions and directions of the light sources to improve the perception of the carving structures. This feature was rated as useful and all participants judged it with +. All participants also confirmed the usefulness of the tool for creating new carvings. P₁ and P₃ rated it with ++ and P₂ with +. The question whether

Table 3: Timings of the preprocessing steps.

# Vertices	85k	40k	20k	10k	5k
Time (s)	198	95	48	24	12

the participants find the framework useful for analyzing the carvings was confirmed with ++ by P₁ and P₂ and with + by P₃. Finally, they were asked if the presented method can reliably detect the carvings. P₁ and P₂ rated it with +, whereas P₃ refrained from answering as he does not have enough knowledge about the datasets.

COMMENTS In the last stage of the evaluation, it was asked whether the experts would use the framework for analyzing the carvings in a professional setting. All participants replied with 'yes'. The question whether they could use the tool for exploration of digitized plasters was also answered with 'yes' by all three participants. Finally, experts were asked for additional feedback and for ideas and suggestions for additional features. They emphasized that some *annotation tools* would be helpful for sharing results with other domain experts and for working collectively on a dataset. Moreover, the experts would like to have a *history listing* who made which changes and the possibility to restrict the user's access, such that an inexperienced user could only analyze but not edit the data.

7.6 RESULTS

The considered plaster is on a wall with width of 2.22 m and height of 0.98 m. After digitization of the plaster, the data comprises around 10 million vertices and around 21 million triangles. To handle the data set, the digitalized plaster is split in three parts to analyze them separately. For the sake of efficiency, the algorithm was not applied on the fully resolved data set in most experiments. To create lower resolutions, the mesh was simplified using Edge Contraction with Quadric Error Metrics [GH97].

RESOLUTION EXPERIMENTS The experiments were performed on a mid-class notebook with an Intel Core i5 with 2.6 GHz, 8 GB RAM, and an nVidia GeForce 870M with 4 GB. In Table 3, the run times for different resolutions of the dataset is shown. The preprocessing includes the curvature estimation, the computation of the vesselness measure and the determination of the ordered connected components of the carvings. As expected, the run times are linear to the number of vertices. On different datasets with approximately the same number of vertices, similar run times were obtained. Figure 65 shows a comparison of results obtained using different resolutions of the mesh.

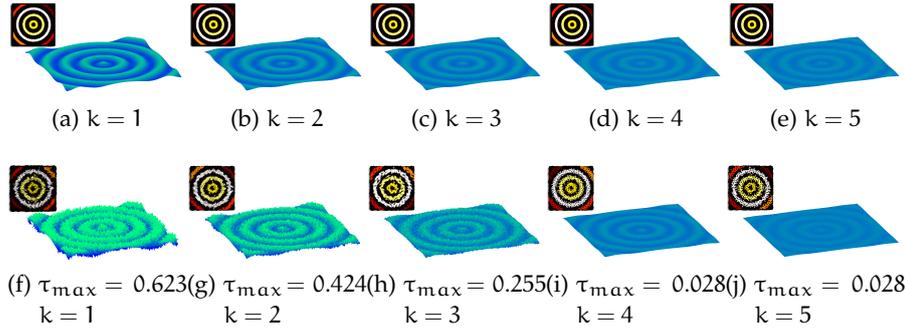


Figure 66: A parametric surface with different heights without noise (top row) and with noise (bottom row). The novel algorithm can detect the carvings well, see the small figures in the upper-left corner of the individual examples. The detected carvings are indicated by colors.

ROBUSTNESS EXPERIMENTS To test the robustness, the method was applied to a parametric surface: $z = 2^{-k} \cos(2 \cdot \sqrt{x^2 + y^2})$, with $x, y \in [-10, 10]$. Depending on k , the height varies ranging from $-1/2^k$ to $1/2^k$. For different k , it was tested if the algorithm can detect the carved structures, see Fig. 66 top row. The surface is depicted as well as the detected carvings top left. Thus, for different heights the carvings can be detected well. Afterwards, noise was added in normal direction, i.e., the points are moved along the normals: $\mathbf{p}_{\text{new}} = \mathbf{p}_{\text{new}} + \tau \mathbf{n}$, where τ is a random value in $[-\tau_{\max}, \tau_{\max}]$. Even in the presence of noise, the algorithm can detect the carvings, see Fig. 66 bottom row.

COMPARISONS WITH PAINTINGS In the last century, artists and photographers tried to depict the carvings in their own work. These images mostly inherit additional knowledge about the carvings. If a carving seems to be discontinuous, the artist can retouch the missing part. Unfortunately, this leads to interpretation and sometimes to wrongly repainted results. In Figure 67, different cutouts are shown: a crown, a hand, and a necklace. As the novel methods are free from personal interpretations, it can depict the correct illustrations. Note that additional carvings were not generated in the results. The other paintings may lead to misinterpretations. These paintings were used to show what the carvings on the wall looked like. Mostly, experts relied on these wrongly painted illustrations. The domain experts confirmed that these illustrations give rise to false interpretations and that the novel approach is quite useful to detect the carvings well. They also stated that they want to use the presented results to publish new insights in the future.



Figure 67: Different cutouts of the carvings. From left to right: photo, steel engraving (1852), painting (1891), painting (1998), and the new approach. First row: the four little circular carvings around the rhomb in the crown can only be perceived with the new approach. Second row: based on the domain experts, the fingers of the painting from 1891 and the presented approach are correct. Third row: the new approach and the painting from 1998 are correct.

DIFFERENT DATASETS For further evaluation of the approach, it was applied to different datasets. Thanks to the Laboratory of Computer Graphics & Multimedia at the Technion and to the Zinman Institute of Archaeology, University of Haifa for providing the datasets. In Figure 68, the new method is applied to different datasets. Thus, not only carvings can be detected and visualized, but also the letters and details on the ancient stone plates can be enhanced.

7.7 COMPARISON TO ALTERNATIVE METHODS

The novel approach was compared with the exaggerated shading technique by Rusinkiewicz et al. [RBD06] and with the D.o.R System [Dor], which is based on the work of Gilboa et al. [GTSK13] and Kolomenkin et al. [KST09]. The exaggerated shading emphasizes the features in the rendered surfaces well, and thereby highlights the carvings. However, since the technique highlights any surface feature it also emphasizes other parts in the digital plaster like noise or irregularities in the plaster. In contrast, the novel method separates the noise and natural structures in the plaster from the carvings and emphasizes only the carvings. Figure 69 shows a comparison of the plaster visualized with both approaches. Closest to the presented work is the method presented by Zatzarinni et al. [ZTS09]. Their tech-



Figure 68: Various models with normal shading (left) and the new method (right). The last two models show the identified carvings in different colors



Figure 69: Exaggerated shading [RBD06] applied to the dataset (left) and the DoR System [Dor] (middle). Details can be depicted well, but noise is also enhanced. With the novel method (right) the carvings (shown in the inset on the figure) are detected and can be used to decently enhance the features.

nique works well for coarser carvings, but in case of strong noise and slightly bended surfaces their method cannot fully extract the carvings. For extracting the carvings, they proposed an automatic approach to set the threshold for distinguishing between foreground and background structures. Since this approach may be error-prone

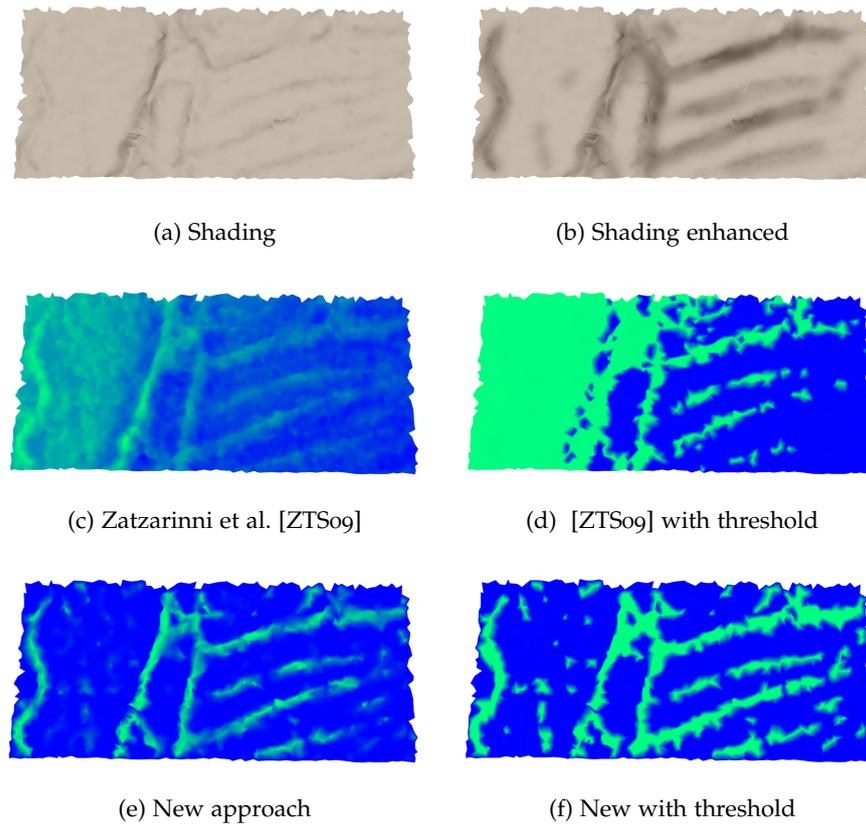


Figure 70: In (a) a dataset with Phong shading is shown, in (b) the carvings were manually enhanced. In (c) and (d) the technique by Zatarinni et al. is applied. Here, a threshold is set to detect the carvings. In (e) and (f) the new technique with correctly detected carvings is shown.

with noisy data, the threshold were manually determined to obtain more reliable results. Both the novel and Zatarinni et al.'s approach are representing the surface containing the carvings using a height field over a base surface. To better explain why the novel approach gives the expected result, the height fields and the components are visualized, see Figure 70 and 71. Whereas the new approach detects the carvings correctly, the approach by Zatarinni et al. results in transition of the height from left to right due to the noisy data and the slightly bended surface. Even if different thresholds were set, not all carvings can be detected.

In addition to the comparison to alternative exaggerated shading techniques and the method by Zatarinni et al., feature line techniques were tried to test whether these can be used for extracting and visualizing the carving structures. The experiments indicate that such methods are too sensitive to noise to compete with the presented method for the extraction of carving structures. Also, enhancing the carvings with feature line methods did not yield results comparable

to the proposed scheme. Examples of this kind are shown in Figure 72. The images illustrate that some features can be depicted well, but in all cases noise is also enhanced.

7.8 DISCUSSION & FUTURE WORK

A techniques was presented to support monument restorers working with heritage-protected carving structures in plasters. The methods enable the analysis, processing and visualization of carving structures in digital surfaces. A method for extracting carving structures in surfaces based on the Frangi filter from image processing was provided. The idea is to search in the surface for curvature configurations with a vanishing and a positive principal curvature. This concept is implemented by formulating a vesselness measure based on the principal curvatures. Furthermore, a techniques was proposed hat supports the user in finding appropriate parameters for the detection process. To robustly extract carvings, a connected component analysis is used in combination with a saliency measure for the components. The component analysis groups regions with the desired curvature pattern in connected carving structures. The saliency measure allows for ranking the groups. Setting a threshold on the saliency allows to robustly extract carvings. The component analysis ensures that coherent structures are obtained and noise is filtered out, since it receives a low saliency value.

The carving structures are used as input for the visualization which highlights the carving structure and thereby increases their visibility. The resulting tool allows users to better perceive the carvings, in particular in cases where the structures have flattened due to aging and weathering. Moreover, the user can interact with the carving selection process by adjusting a simple and effective set of thresholds. Furthermore, users may wish to add new carvings. An example-based approach for creating carvings in the “style” of the historic original carvings was proposed.

The robustness of the approach was confirmed by applying the techniques to artificially generated surfaces with different heights. With the method it was possible to show how the carvings look like without risking misinterpretation from artists. As future work, the experts asked for several extensions. In particular, they asked for annotation tools. This would enable collaborative editing and exploration. Moreover, the experts would like to have a log listing who made which changes. The experts also suggested to extend the framework such that datasets with textures are supported. Then, the analysis could also take the texture into account, in addition to the current geometric analysis.

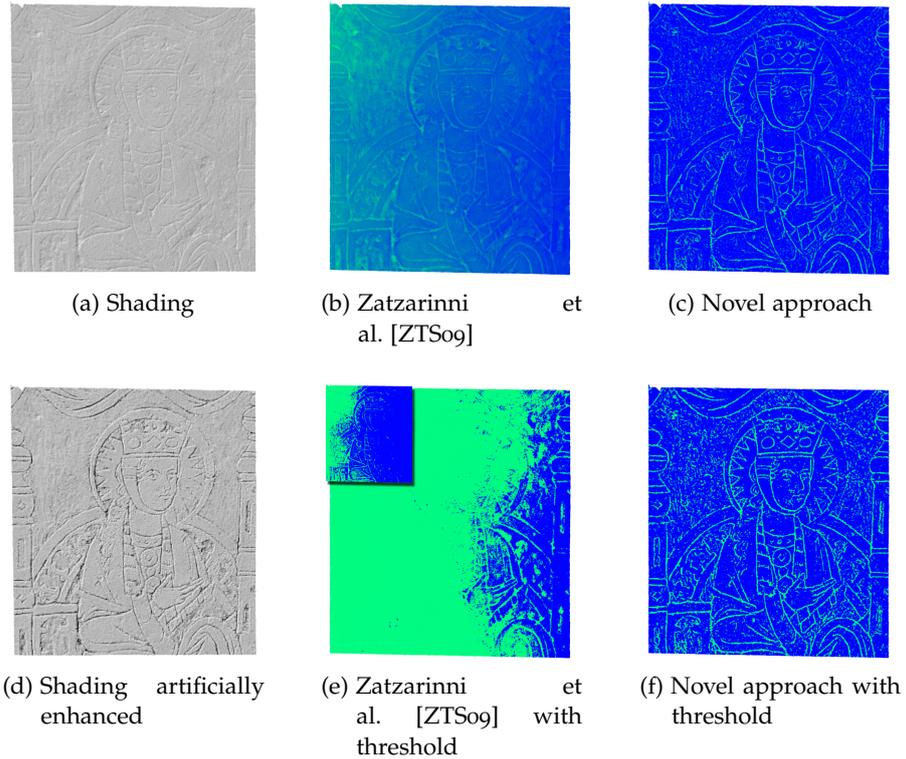


Figure 71: In (a) a dataset with Phong shading is shown. The carvings were manually enhanced to perceive more easily the features (b). In (b) and (e) the technique by Zatarinni et al. is applied. Setting different thresholds will either detect the carvings left or right. In (c) and (f) the novel technique with a threshold is shown and it detects the carvings correctly.

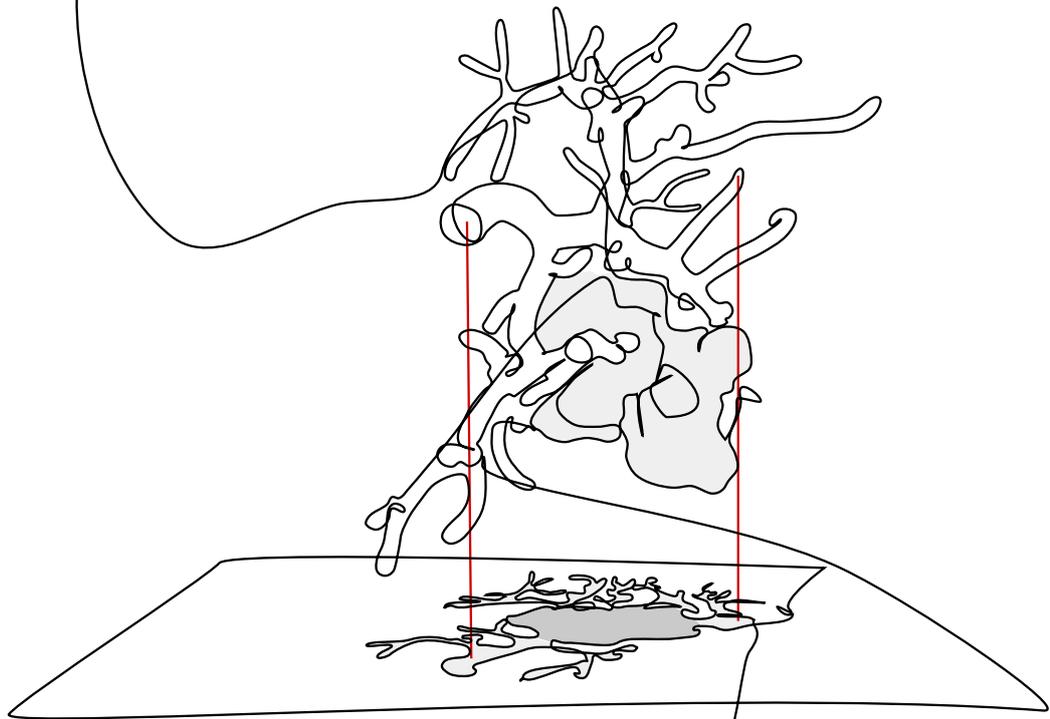


Figure 72: Different feature line methods are applied to the model. From left to right: shading, ridges and valleys, suggestive contours, apparent ridges, and photic extremum lines.

Altogether, based on the evaluation and the experiments, the work is considered as a successful step in the development of visualization techniques and tools for supporting monument restorers.

8

SUPPORTING LINES



This chapter is partly based on:

Kai Lawonn, Maria Luz, Bernhard Preim, and Christian Hansen.

“Illustrative Visualization of Vascular Models for Static 2D Representations.”

In: Proc. of Medical Image Computing and Computer- Assisted Intervention (MICCAI).

Springer, 2015, pp. 399–406.

IMPROVED DEPTH-PERCEPTION WITH SUPPORTING LINES

DURING surgery it is often important to understand the morphology of vascular structures and their spatial relation to surrounding risk structures such as tumors. Therefore, high quality vascular 3D models, reconstructed from CT, are applied using direct or indirect volume rendering techniques. The vascular models often portray complex geometries and therefore demand cognitive effort and user interaction. Physicians request a visualization technique that can encode essential information as the location, the spatialness, and the distances to important anatomical structures. The ideal solution might be to provide one or more static images that encode all relevant information at a glance. These static images might be printed out or displayed on a monitor in order to provide decision-support during surgery.

The work of Ritter et al. [RHDKV+07] is closest to this work and is guided by a similar set of requirements and application scenarios. They evaluated non-photorealistic rendering techniques for vascular structures and showed that these techniques facilitate shape and depth assessment. Wang et al. [WBDF+14] presented a method for depth perception on interventional X-ray images. They improved the rendering scheme and used depth cues that preserve the original X-ray intensities. Furthermore, methods for depth-based color-coding were proposed, e.g. pseudo-chromadepth by Ropinski et al. [RSHo6]. Kersten-Oertel et al. [KOCC13] investigated the effect of different perceptual cues for vascular volume visualization and showed that pseudo-chromadepth cues were stronger cues than stereopsis. Because color-coding of depth is not always applicable (e.g. printouts or AR projections on the organ surface) a color-sparse representation of spatial information would be beneficial. A visualization technique is proposed that uses as few colors as possible, see Fig. 73. To assess the benefit of this visualization technique, it was compared with Phong shading [Pho75] and pseudo-chromadepth rendering [RSHo6] as it was rated best of different depth cues in previous studies, c.f. [RSHo6; KOCC13].

8.1 RELATED WORK

Incorrect spatial interpretation is the most common perceptual problem in AR applications [KSIF10]. The need to assess spatial information of superimposed 3D models accurately during an intervention

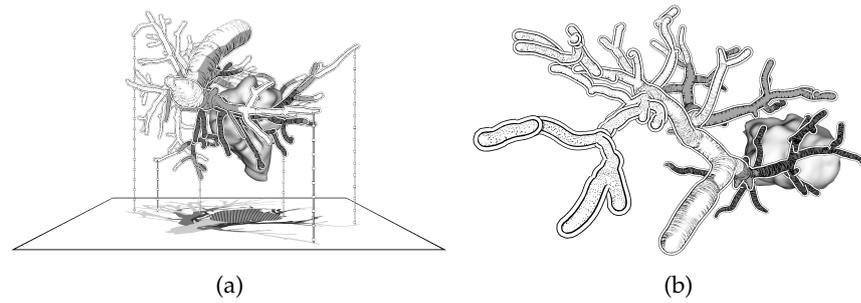


Figure 73: The visualization technique applied to hepatic vasculature with a tumor obtained by using [HPSP01]. In (a), illustrative shadows, supporting lines, and contours as depth cues were used. In (b), different hatching styles and a decent gray tone emphasize vessels at risk around the tumor.

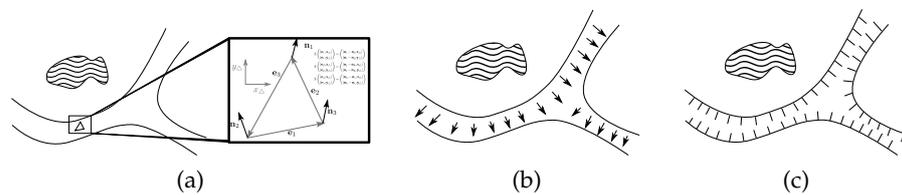


Figure 74: In the first step, the curvature tensor per triangle is determined, see (a). This is used to calculate the PCDs per vertex, see (b). Afterwards, the streamlines are determined, see (c).

has led to the development of several techniques to improve spatial perception in AR.

8.1.1 Improving Spatial Perception for Medical Visualization

Stereoscopic vision improves the spatial assessment of computer-generated images [BFHW+02; WVKJ+06; CKHST+12; BHSC+14]. Therefore, autostereoscopic monitors or head-mounted displays can be of great value in the operating room (OR). For many surgeons, additional hardware is cumbersome or difficult to use within the constraints of the OR [KCC14]. In addition, the benefit of stereoscopic vision is task-dependent, and the results of Kersten-Oertel et al. [KCC14] indicate that it does not provide effective spatial cues for the depth judgment of vascular 3D models.

Computer-generated monoscopic depth cues have been proposed to improve the spatial perception. Kinetic depth, especially motion parallax, provides depth cues in many AR applications [BSHN07; FMTP+09]. A virtual window was introduced, which simulates a view on the patient's innards [SBHN06]. The window plane is slightly textured to improve the spatial perception through partial occlusion when using a tracked head-mounted display. Projecting guidance

graphics, such as entry points and surgical plans, onto an organ surface while tracking the operator's head using projector-based AR have been proposed [GW03; WNCC13]. A disadvantage of kinetic depth is that the head of the operators needs to be tracked, which leads to additional hardware in the OR, and that operators are forced to continuously move their head to improve the spatial impression.

Color-coding can also be employed to encode depth and inter-object distance in AR. Ropinski et al. [RSH06] introduced pseudo-chromadepth rendering as an enhancement of the chromadepth rendering, where the number of hues were reduced. According to the user study of Kersten-Oertel et al. [KCC14], pseudo-chromadepth rendering allows for the best relative depth assessments for vascular structures. However, our own results indicate that illustrative rendering techniques could improve depth assessment results over pseudo-chromadepth rendering [LLPH15]. In addition, multiple viewports are proposed as a method to improve the spatial perception for AR. To visualize inter-object distances in AR environments, seamless switching between an AR view and a VR view was proposed [CCMH+15]. Bichlmeier et al. [BHFNO9] presented *virtual mirrors*, a render-to-texture approach, which provides additional views on the model. Particularly the interpretation of partial self-occlusions inside complex models is improved.

The correct display of *occlusion* between 3D virtual models and real objects (patient, operator, instruments, etc.) on a 2D surface is important to convey spatial relations. An information filter based on intra-operative distance measures was introduced to reduce occlusion of anatomical structures by virtual overlays [KSBCO+15]. Fischer et al. [FHS07] propose to provide a distance map of the real scene using a TOF camera. By performing a registration between distance map and virtual 3D models, the distance map is used to decide which objects are rendered at the front, and to control the rendering style. A common way is to change the rendering style for occluded objects, e.g., using a wireframe presentation [TYKO05] or transparency [EAT07].

FOCUS-AND-CONTEXT VISUALIZATION Illustrative rendering techniques can also improve the spatial perception for focus-and-context visualization. An approach that uses surface topology, viewing attributes, and the location of surgical instruments was presented to generate a transparency map to control the pixel transparency of video images [BSHN07]. A context-preserving focus region facilitates the spatial perception. Fischer et al. [FBS05] developed an illustrative rendering technique that generates a stylized augmented video stream. Based on an edge-detection algorithm, silhouettes are extracted and applied to both the camera image and the virtual objects. Visual realism of the virtual objects and the real camera image is reduced.

Both modalities become less distinguishable from each other and thus an improved immersion can be achieved. A similar approach is described by Kalkofen et al. [KMS07; KMS09]. However, they apply silhouettes to visualize hidden objects in the area of augmentation (X-ray visualization). Complexity of the superimposed silhouettes is controlled by an importance measure and an interactive user-controlled filter. Avery et al. [AST09] made additions to this concept for real scenes, whereas Lerotic et al. [LCMY07] analyzed surface features from stereo laparoscopic camera images and derive the orientation of the organ surface. This information is used to generate a translucent contour layer. Surfaces of the real scene parallel to the viewing plane are rendered transparently, whereas sloped structures are accentuated using strokes. An associated user study proved that the proposed visualization led to an improved depth assessment.

8.2 MEDICAL BACKGROUND

The medical visualization techniques described in this chapter are presented using the example of liver surgery. The liver is an essential organ with a complex vascular anatomy. Unfortunately, it is also an organ frequently affected by tumors. If tumors in the liver are removed by a surgeon, the vascular anatomy must be preserved as much as possible while ensuring an adequate safety margin around tumors. To this end, physicians employ x-ray computed tomography (CT) or magnetic resonance (MR) imaging to gain information about the patient's liver anatomy and pathology. Based on these images, surgical planning software provides 3D models of a patient. This additional information assists radiologists and surgeons in planning surgery and choosing a resection strategy [HZPS+14].

When a patient is diagnosed with a malignant liver tumor or a liver metastasis, a wide range of essential decisions has to be made by the physicians involved. One important part of the surgical decision pipeline is the preoperative risk analysis. A risk analysis includes the determination of proper safety margins around tumors. Based on the widths of safety margins, affected vascular branches and vascular territories are calculated. In tumor surgery, different safety margins are usually explored and visualized before surgery, e.g. 2mm, 5mm and 10mm for metastasis in liver surgery [PSSO+00].

Once the planning is performed, surgeons would like to utilize the 3D planning models during an intervention. However, the intraoperative presentation of planning models has thus far been limited. One reason is that these models were not developed to provide information in complex workspaces like the operating room. Intraoperative represented models often portray complex geometries and therefore demand cognitive effort and user interaction, especially if the surgeon wants to use them as a navigation aid. However, during surgery

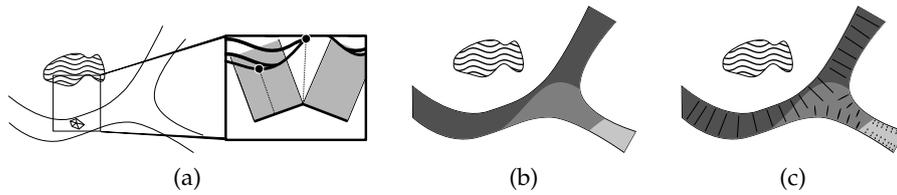


Figure 75: In the first step, the distance of every vertex of the vessel to the tumor is determined (a). If the vertex of the tumor lies inside the prism spanned by a triangle, the distance to the plane is used otherwise the distance to the vertex is used. The distances are used to classify the risk zones (b). The ConFIS method is applied. The zone determines the shape of the streamlines (c).

it is important that the surgeon focuses on the operating field instead of operating a computer. Therefore, it is necessary that the information required for surgery is encoded in a static image. This becomes obvious in case 3D planning models are projected onto the liver surface during open liver surgery using a light projector as described in [KHKD+08; GPOSW11], or superimposed with the images from a laparoscopic camera [MRAM+04; FMHN08]. In both cases, an augmentation of the surgical view by planning models results in limited depth perception, misinterpretation of spatial relations, colors, and contrast as well as unacceptable occlusions of the operation field.

For medical documentation purposes it is often desired to store a 3D model with a specific camera view point within a static 2D document, e.g., in a clinical Picture Archiving and Communication System (PACS), or for printouts in medical publications in education and science. An interaction with the model is not possible; a color-sparse representation that encodes spatial information and distances without the need to rotate the model would be beneficial in this case.

8.3 METHOD

The underlying vessel visualization technique is model-based with circular cross sections as a major model assumption. These vascular surface models are very smooth, not subject to noisy artifacts but abstract from details of the vascular shape. For applications, such as computer-assisted surgery planning, this is appropriate [OP05] but for vascular surgery, a more faithful visualization would be required. In this section, the various visualization techniques are described that are included in the novel approach. First, our hatching method [LMP13] is recapped. Second, the visualization of the distance of the vascular structures to the tumor is explained. Finally, various illustrative visualization techniques are described to enhance the depth perception.

8.3.1 Hatching - ConFIS

The ConFIS method (recall [LMP13]) for hatching is employed, see Figure 74 for an overview. The *ConFIS* method needs the principle curvature directions (PCDs) which are computed with the algorithm by Rusinkiewicz [Ruso4], see Section 2.1.1.3. This yields the PCDs, which are used to generate the streamlines on the surface. The authors suggested to explicitly calculate the streamlines by linearizing the vector field per triangle. Instead of explicit calculations, an explicit Euler scheme is applied to calculate the streamlines. Here, the same adaptive step size of the radius of the incircle per triangle is used as we suggested in our previous work [LMP13]. Furthermore, the feature regions were not determined like the authors suggested, as this is not essential to perceive the vascular structure. After calculating the streamlines, they were stored in an extra file. This has the advantage that for further examinations the calculation time is significantly decreased.

8.3.2 Distance Illustrations

First, the distances of the vascular structures to the tumor model are determined. Here, the tumor model is given as a triangulated surface mesh. For the calculation of the distance, the triangles of the tumor are used to generate a prism along the normal. Second, for every vertex of the vascular structure, the algorithm iterates over the triangles of the tumor and tests if the vertex lies inside the associated prism. If this is the case, the distance of the triangle with the vertex is determined. Here, the triangle is used to determine a plane and to calculate the distance to the vertex. Third, the distance of the vessel vertices with the tumor vertices is calculated. Finally, all distances are compared and the shortest distance is associated to the vertex of the vessel, see Figure 75a. This yields the distances of the vascular structure with the tumor, which is independent of the tumors' tessellation.

As described in Section 8.2, it is important to illustrate different safety margins around the tumor. To provide a color sparse representation, a dark grey, a bright grey, and white is used. To enhance vessel parts at risk zones, the *ConFIS* visualization was additionally altered at potentially affected vessels. Instead of generating streamlines, quads were generated in the *Geometry Shader*: The surface normal and the tangent vector of the streamline was exploited to generate a quad on the surface represented as triangle strips. For these strips a width is needed. Then, the tangent vector \mathbf{t} is determined by subtracting the adjacent point $\mathbf{p}_i, \mathbf{p}_{i+1}$ and normalizing the resulting vector: $\mathbf{t} = \frac{\mathbf{p}_{i+1} - \mathbf{p}_i}{\|\mathbf{p}_{i+1} - \mathbf{p}_i\|}$. The quad's extent is determined by computing

the cross-product of the normal vector \mathbf{n} with \mathbf{t} . For each adjacent point pair $\mathbf{p}_i, \mathbf{p}_{i+1}$, the following points are generated:

$$\begin{aligned} \mathbf{p}_{\text{out}1} &= \mathbf{p}_i + \text{width} \cdot \mathbf{t} \times \mathbf{n}, \\ \mathbf{p}_{\text{out}2} &= \mathbf{p}_i - \text{width} \cdot \mathbf{t} \times \mathbf{n}, \\ \mathbf{p}_{\text{out}3} &= \mathbf{p}_{i+1} + \text{width} \cdot \mathbf{t} \times \mathbf{n}, \\ \mathbf{p}_{\text{out}4} &= \mathbf{p}_{i+1} - \text{width} \cdot \mathbf{t} \times \mathbf{n}. \end{aligned} \tag{56}$$

Then, a coordinate system (x, y) is generated for every streamline in the range of $y \in [-\text{width}, \text{width}]$, where width is length of the quad's extent and $x \in [0, L]$, where L is the length of the current streamline. Afterwards, the coordinates are used to draw glyphs on the quads. Therefore, the coordinates are altered in:

$$\begin{aligned} x' &= \text{fract}(d \cdot x) - 0.5 \\ y' &= d \cdot y, \end{aligned} \tag{57}$$

where fract computes the fractional part of the argument and d is a value that specifies how many glyphs are drawn in the interval $[0, 1]$. Then, implicit ellipses are used which are drawn on the quad:

$$\alpha \cdot x'^2 + y'^2 \leq r^2.$$

Whenever this inequality is fulfilled the fragments are drawn, where r specifies the radius and α is a distance-dependent value that describes the width of the ellipse. The values of α are set as:

$$\alpha(\text{dist}) = \begin{cases} 0 & \text{if } z \leq 2 \\ \frac{z}{\text{dist}_{\text{max}}} & \text{otherwise,} \end{cases}$$

where dist describes the distance from the tumor and dist_{max} is the maximum distance. A threshold of 2mm is used as this describes the risk zone. α is set such that the streamlines are fully drawn at the risk zone. In addition, the glyphs continuously change from ellipses to circles depending on the distance to the tumor. Therefore, at the safe regions circles are used and with increasing distance to the tumor, the circles become more elliptical in shape, see Figure 76. Furthermore, the streamlines were drawn at the contour margin and opacity decreases depending on the length to ensure a consistent fade off during the interaction. This idea is altered at the most dangerous zone. Here, all streamlines are drawn on the surface to emphasize this important region. Figure 75 represents an overview of the method.

8.3.3 Depth Cues

Different illustrative visualization methods were used to provide depth cues, see Figure 77. Inspired by Ritter et al. [RSHSo3], illustrative

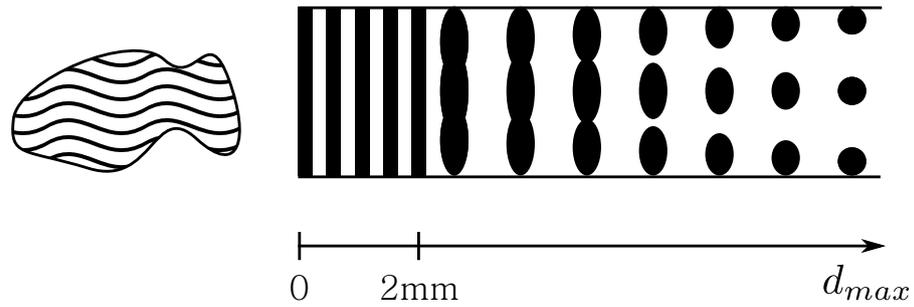


Figure 76: For the risk zones, streamlines are used for illustration. The more distant from the tumor the more the α increases. Therefore, at safe regions only circles are drawn.

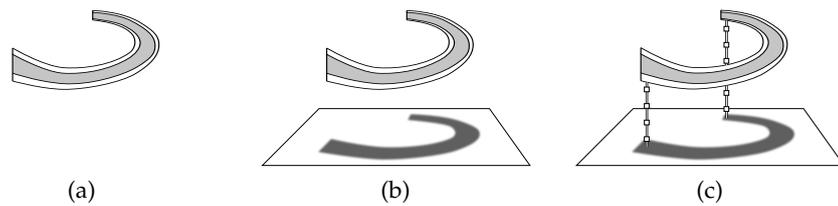


Figure 77: Different depth cues for vascular structures. In (a), view-aligned quads are used as a contour illustration. Furthermore, illustrative shadows are used to enhance the depth perception, see (b). The user can add supporting lines to improve the correspondence between a branch and the shadow.

shadows were employed. For this, a plane is placed under the planning model. The position of the plane is determined by the center point of the planning model and it is shifted along the up-vector of the camera. The up-vector of the camera serves also as the normal of the plane. Therefore, the plane stays at its position during the interaction with the model. Afterwards, a camera is placed on the opposite direction of the plane to create shadows by an orthographic projection of the vascular structure as well as the tumor. Frame buffer objects were used to save the scene of the camera and to texture the underlying plane. For the vessel shadows, the same grey values are used as for distance encoding between vessel and tumor in the 3D model in order to illustrate the plane. For the tumor, hatched strokes are applied to emphasize the projected tumor. For this, the coordinates of the plane and Equation 57 for the x -component are used. Whenever the x -component is around zero and the shadow corresponds to the tumor, the lines are drawn.

SUPPORTING LINES As an additional depth cue and reference for the shadow visualization, supporting lines are introduced. Here, the idea by Glueck et al. [GCAR+09] is used to create lines from the objects of interest to the plane. First, the medical expert can select a point on the surface by picking and a supporting line is generated.

Second, a view-aligned quad is provided. Third, the first point of the surface along the viewing ray and the up-vector of the camera is used to determine the intersection point with the plane. On this quad, texture coordinates are generated to draw a ruler-like illustration. Finally, (x, y) coordinates are generated, where x represents the width of the quad and y lies in the range $[0, L_{line}]$ with L_{line} represents the length of the quad and it starts from the plane with $y = 0$. The ruler-like illustration can also be used as an indicator for the distance to the plane. Here, a white line is drawn with a black contour such that it can be perceived independent of the background color. Therefore, if x lies in a margin around zero, the fragments are drawn white and for a larger margin, the black color is used. In addition, for a fixed distance, small white squares are illustrated which are also surrounded by a black contour. Conceptionally, the same coordinate system is utilized on the quad described in Equation 57, but instead of drawing circles, squares are drawn. The ruler-like supporting lines can also be used for precise depth comparisons. In case two supporting lines are nearly on the same horizontal line of the plane, it is hard to perform a depth comparison. However, ruler can be utilized to see at which point the supporting line intersects with the back line of the plane. The perceived height can serve as an indicator for depth. The more squares are under the back line of the plane the closer the corresponding point.

DEPTH-DEPENDENT HALO CONTOUR To further support depth perception, a depth-dependent halo contour is utilized. This idea was inspired by Bruckner and Gröller [BG07] who apply halos to increase the depth perception for volume data sets. In the first step, the shading is determined by multiplying the normals of the vascular structure with the view-vector. These values are used as a scalar field: $\varphi = \langle \mathbf{n}, \mathbf{v} \rangle$, where \mathbf{n} is the normal and \mathbf{v} is the normalized view-vector. Afterwards, the zero-crossings are determined on the surface. Here, the algorithm tests for every triangle if the sign of φ is different for the three vertices. In case the sign changes, e.g., $\varphi_i > 0$, $\varphi_j > 0$ and $\varphi_k < 0$, a zero-crossing inside the triangle occurs. The zero-crossing is on edge \mathbf{e}_1 , which connects vertex i and k and \mathbf{e}_2 , which connects vertex j and k . $t = \frac{\varphi_k}{\varphi_k - \varphi_j}$ is determined and multiplied with the edge \mathbf{e}_2 . This yields the position of the zero crossing on the edge and the position on the other edge is determined similarly. Finally, these positions are used to generate a view-aligned quad as described in Equation 56, but the normal vector is used to extend the quad. Because surgical planning models are usually based on world space coordinates, the perceived thickness of the contour is a measure for the distance between vessel position and viewport position.

8.4 EVALUATION

To assess the advantages of the new visualization concept, a qualitative evaluation was conducted. The evaluation was performed with 10 researchers (8 men, 2 women, age of 24 - 37 with an average of 27.6) who are familiar with medical visualization and 5 physicians with a medical focus on surgery (1 men, 4 women, age of 27-38 with an average of 30.8). The evaluation consists of three steps:

1. The participants were shown the new visualization technique with different 3D planning models. They could interact and explore the model until they were ready for the next task.
2. The participants were explained the new visualization techniques and they were shown how to place the supporting lines. In addition, the shadow and the contour representation were mentioned as depth cues.
3. The participants were shown 12 different printed scenes with 3D planning models rendered using the novel approach. The visualized models were generated out of 6 clinical CT datasets including 12 vascular trees – portal vein and hepatic vein for each dataset. The datasets contained between 1 to 10 tumors or metastasis. In addition, four red points were placed on different branches of each tree and labeled with A, B, C, and D. The scenes were presented to the participants and they were asked to order the labeled branches according to the distance from the viewing plane.

The experiments were recorded using a video camera. Afterwards, the participants were asked how they assessed the distances. Furthermore, they were asked which technique was the most useful. At the end, the participants could also state additional remarks.

The first task was about to convey a feeling for the new visualization technique. The participants should become familiar with the visualization. Then they were asked about the illustration techniques. Here, the goal was to figure out which of the techniques is self-explaining and where the participants needed an extra explanation. For this, the second task was about an explanation of the presented methods, in case the participants did not figure it out. Finally, for the last task, the participants were shown non-colored different scenes with different models. Here, the goal was to assess the quality of the new visualization concept if the participants can use this method for depth perception.

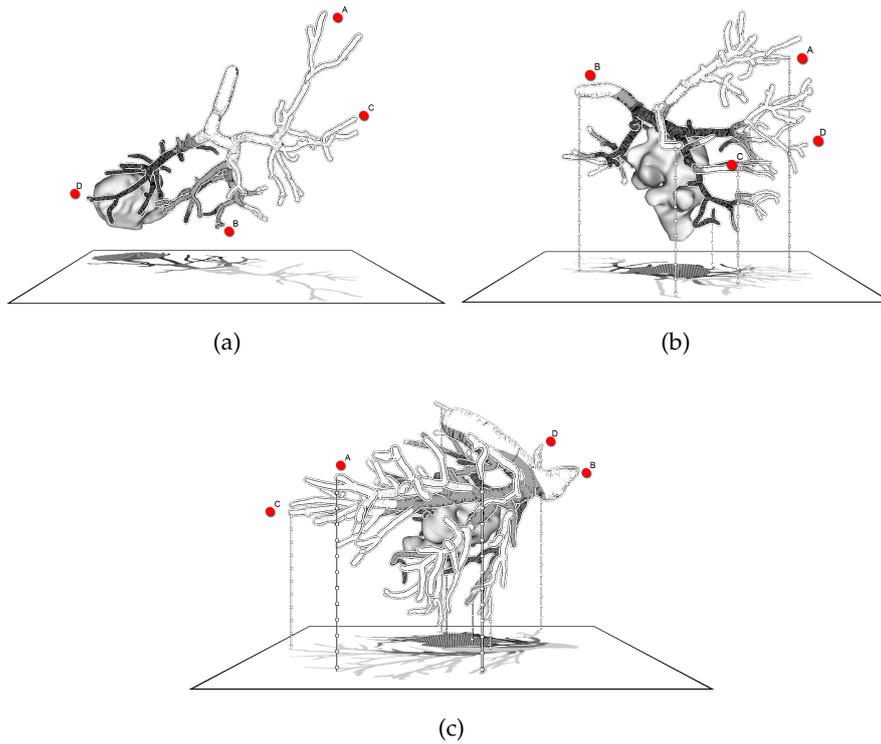


Figure 78: Only three participants ordered the points wrongly on the dataset 2 without supporting lines, (a). The most participants could not identify the depth of point D on dataset 4, see (b). Only three participants mixed the ordering of two points on dataset 8, see (c).

8.5 DISCUSSION AND RESULTS

Regarding the informal feedback, the right choices of points were determined concerning the distances to the view plane. To assess the given answers and to measure its quality, they were scored as follows: One point was given if the depth sorting was correct. In case the participants made one mistake, i.e., the ordering of two points was wrong, the users received only half a point. In case two or more orderings were wrong, they received zero points. The results of the evaluation were split in researchers with background in medical visualization and physicians.

The participant could receive a maximum of 12 points for the whole evaluation. The researchers received a minimum of 9 points and a maximum of 12 points. On average they had 10.4 points. The physicians had a minimum of 10 points and a maximum of 11 points as well. Their average was 10.7 points. Overall the participants had an average point of 10.5.

After the evaluation, the participants were asked which of the presented depth cues was the most supportive one. In general, the participants stated that a mix of shadows and supporting lines was the

	1	2	3	4	5	6	7	8	9	10	11	12	Average	
Researchers	1	1	1	0.5	1	1	1	1	1	1	1	1	11.5	
	1	1	1	1	1	1	1	1	1	1	1	1	12	
	0.5	0	1	1	1	1	1	1	1	1	0.5	1	10	
	1	1	1	0.5	0.5	1	1	1	1	1	1	0.5	1	10
	1	1	1	0.5	1	1	0.5	1	1	1	1	1	1	11
	0.5	1	0.5	0.5	1	1	1	1	1	1	1	1	1	10.5
	1	1	0.5	0	0.5	1	1	0.5	1	1	1	0.5	1	9
	1	1	1	0	0.5	1	1	0.5	1	1	1	0.5	1	9.5
	0.5	1	1	0	0.5	1	1	1	1	1	1	1	1	10
	1	0.5	1	0.5	1	1	1	0.5	1	1	1	1	1	10.5
	Physicians	0	1	1	1	1	1	1	1	1	1	1	1	11
0.5		0.5	1	0	1	1	1	1	1	1	1	1	10	
1		1	1	0.5	1	1	1	1	1	1	0.5	1	11	
0.5		1	1	0.5	1	1	1	1	1	1	1	1	11	
1		0.5	1	0.5	1	1	1	1	1	1	0.5	1	10.5	

(a)

Figure 79: The result of the evaluation. The columns correspond to the different datasets. The light color represents one mistake and the dark grey represents more than one mistake. On the average column a value less than 75% of 12 is colored in dark grey, points between 75% and 85% are colored in light grey and better than 85% is represented as the rest.

most useful combination. The contours were rarely used as the differences in width were hard to perceive. Two participants remarked that a non-linear function depending on the distances may be useful for the width of the contour. Three participants stated that the contour supports a spatial impression and supports the depth perception, but it depends on the distance. For vessel branches far away, the contour is very useful even the spatial impression, but for vessels very close to each other it is hard to perceive a difference.

For the evaluation with the datasets, different images with and without supporting lines were shown. The question if the participant could assess the depth with shadows only, depends on the complexity of the vessel. For one vascular tree with only a few branches, the participants were able to find the vessel corresponding to the shadow vessel and, therefore, could assess the correct depth, see Figure 93a. On another dataset, the points were not correctly sorted by four participants. Three of them mixed two points and one had two mistakes. If the points were placed on the regions with supporting lines, all participants were able to correctly sort the points. In cases where points were placed on regions with supporting lines and on regions without supporting lines, it strongly depends on the region and on the complexity of the vessel. In one case three participants correctly sorted the points, whereas eight had one mistake and four participants did two or more mistakes, see Figure 93b. On another dataset with points on supporting lines and without, three participants mixed one point and the other participants correctly ordered the points, see Figure 93c.

Summary: In general, if the points were placed on supporting lines the participants could order the depth correctly. In cases where the supporting lines were nearly on the same horizontal line, some participants focused on the squares on the ruler and counted how many squares on the line are under the back line of the plane to give the right ordering. Four participants asked for a grid on the floor to easily

see the intersection point of the supporting line with the plane. They stated that this would enhance and fasten the sorting process of the supporting lines. For lines nearly on the same depth, a grid may be an indicator for the correct depth.

If no supporting lines were placed, the result strongly depends on the complexity of the vascular tree. Therefore, the more complex the vascular structure the more supporting lines are needed to assess the correct depth. A reason for that might be that for more complex structures it is hard to associate the vessel branch with the corresponding shadow part. Thus, for the framework the medical expert can place the supporting lines manually. All participants stated that the contour is a visually pleasing element and that it supports the perception of the shape. Unfortunately, for depth cues it is not very useful as the width difference is not significant for close vessels.

After the evaluation, the participants were asked about the quality of the visualization. The researchers were satisfied with the visualization techniques. They found the result visually pleasing and stated that it is easy to identify areas of risk. Furthermore, they stated that the correspondence of the grey tones on the vessel with the shadows were helpful to assess the depth. They could use this as a hint to find the corresponding vessel in the shadow image. The transition from points to lines was also helpful to enhance dangerous regions.

The physicians with emphasis on operations were also satisfied with the visualization technique. One surgeon also stated that this can be used for intervention planning and that it is helpful to illustrate the scene without colors. Furthermore, they liked the idea of placing the supporting lines. They also liked the interaction and the illustration during the exploration, they could easily perceive the areas of risk. All physicians would use the visualization technique to discuss different cases.

EVALUATION: To assess the novel visualization concept, a qualitative evaluation was started. The evaluation was performed with ten researchers who are familiar with medical visualization and five physicians with a focus on surgery. First, the subjects were shown the novel visualization technique with different vascular models. Then, they could explore the model, until they were ready for the next task. Second, the participants were explained the new visualization technique and showed them how to place the supporting lines. Finally, twelve printed scenes were presented with 3D vascular models rendered using the novel approach. The models were generated out of six liver CT datasets including twelve vascular trees – portal vein and hepatic vein for each dataset. In addition, red points were placed on different branches of each tree. The scenes were presented to the subjects and asked them to order the labeled branches according to the

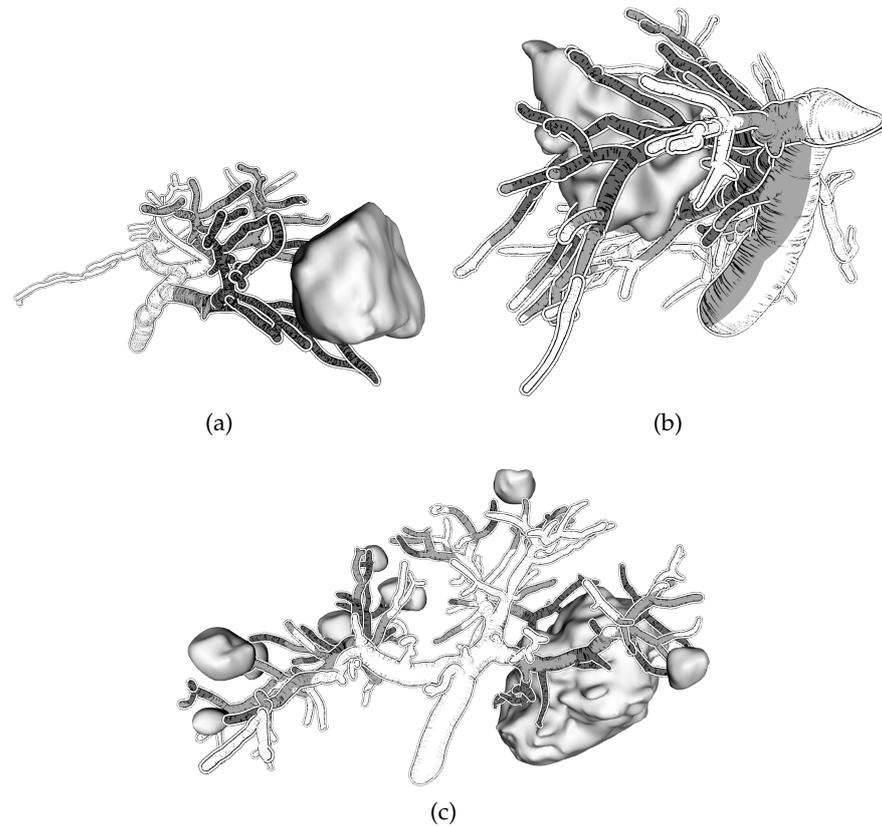


Figure 80: Different results of the visualization of vascular structures with tumors.

distance from the viewing plane, and asked them how they assessed the distances. The experiments were video-recorded.

Based on the results of the qualitative evaluation, a quantitative study was conducted using a web-based framework with 50 subjects (19 female, 31 male, age $M=30.9y$, range 17-48y). 24 subjects are computer scientists, others are physicians (8), engineers (6), natural scientists (5), and others (7). Most of them (31) have no experience with vascular visualizations. Because in the work by Kersten-Oertel et al. [KOCC13], pseudo-chromadepth rendering was rated as the best depth cue for vascular visualization, the novel illustrative visualization technique was compared with pseudo-chromadepth rendering [RSHo6], and Phong shading [Pho75] (no explicit depth cue). To probe the subjects' assessment of relative depth distances, 24 tasks were designed that require a precise judgement of depth. To assess the effect of the new illustration technique, eight visualizations of vascular models using the new techniques had to be compared with eight pseudo-chromadepth renderings and eight Phong renderings of the same model identical in every major aspect, except for the display algorithm. The models were presented to the subjects in pseudo-randomized order to minimize learning effects. For each task, subjects

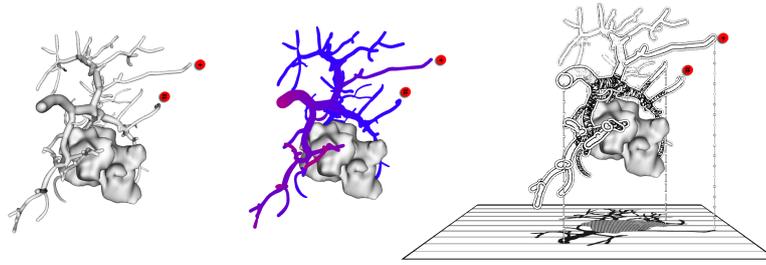


Figure 81: A scene with a tumor and liver vessel is illustrated using different rendering styles. From left to right: Phong shading, pseudochromadepth, and the novel technique.

had to sort the correct depth order of two labeled branches, see Fig. 93. Four conditions were defined: FF, FN, NF, and NN to specify the positions of the labels, where F means *far* and N means *near*. The first capital describes the pixels' distance of the labeled branches in the image space. The maximum distance D is considered connecting branches that are furthest away from each other. If the distance d of the labeled branches fulfills $d < \frac{D}{2}$ then it is defined as N, otherwise as F. For the second capital the z distance of the near and far plane is used. If the z distance of the labeled branches is greater than half of the distance of the far and near plane it is denoted as F otherwise as N.

The quantitative investigation included a training phase, in which all three visualization techniques were explained and subjects were asked to perform two test tasks (not included in the evaluation) for each modality. They were also informed that time is measured after the training phase for each task. After completing a task, subjects were asked to rate their confidence on a 5-point Likert confidence scale where 5 means very confident and 1 very unconfident, and to press a button to start the next task.

8.6 RESULTS

In the qualitative study, subjects stated that a mix of illustrative shadows and supporting lines was the most useful combination. Contours were rarely used as the differences in width were hard to perceive. Three subjects stated that the contour supports a spatial impression and supports the depth assessment, but it depends on the distance of labeled branches. If the points were placed on supporting lines the subjects could order the depth correctly. Four subjects asked for a grid on the floor to easily see the intersection point of the supporting line with the plane. They stated that this would enhance and fasten the sorting process of the supporting lines. This feature was added for the quantitative study. If no supporting lines were placed, the result strongly depends on the complexity of the vascular tree. Therefore,

the more complex the vascular structure the more supporting lines are needed to assess the correct depth.

For the quantitative study, the data was analyzed by a combination of the three visualization techniques (illustrative, pseudo-chromadepth, and Phong) each with 2 distances each for xy and z distance (F far, N near) using ANOVA for repeated measurements (see Table 5 for the statistical results).

ASSESSMENT: Subjects were able to assess the relative distances more precisely with the illustrative visualization, followed by pseudo-chromadepth (Fig. 94). The given answers for Phong seem to be randomly; only $\approx 48\%$ of the distance estimation tasks were performed correctly. With Phong, subjects achieved better results when the distance in the xy plane was small. For the illustrative visualization, the same but weaker effect was assumed, while for pseudo-chromadepth rendering the xy distance seems to have no effect. The opposite effect was observed for z distance: while subjects could perform the pseudo-chromadepth tasks significantly better in case the z distance was large, this effect was not observed for Phong and the illustrative visualization.

CONFIDENCE: Subjects were most confident with the illustrative visualization, followed by pseudo-chromadepth, and most unconfident with Phong (Fig. 83). They felt more confident with conditions FF and NF. The significant triple interaction shows, that subjects are most confident with condition FF using illustrative visualization or pseudo-chromadepth, and condition NN using illustrative visualization. Subjects felt most unconfident with condition NN using pseudo-chromadepth, and with Phong in general.

TIME: As shown in Fig. 84, subjects decided significantly faster using pseudo-chromadepth ($M = 10.95$ s), followed by Phong ($M = 13.21$ s) and illustrative ($M = 14.04$ s). Particularly, the FN condition seems to be difficult and delays decisions ($M = 13.95$ s), while for all other conditions subjects needed approximately the same amount of time. The significant triple interaction seems to be caused by conditions FN and NF which are the best for pseudo-chromadepth and worst for illustrative, while NN and FF are approximately equal for all visualization techniques.

8.7 DISCUSSION

Illustrative visualization has a high potential to improve the spatial perception of complex vascular models. The proposed visualization method provides strong monocular depth cues without the require-

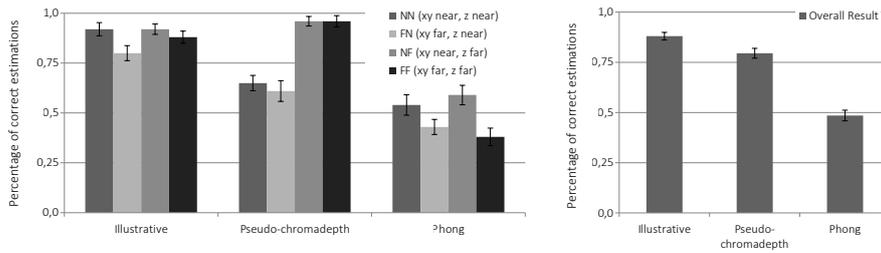


Figure 82: Results for the distance estimation on a scale from 0 (false) to 1 (right). The errors bars represent the standard error.

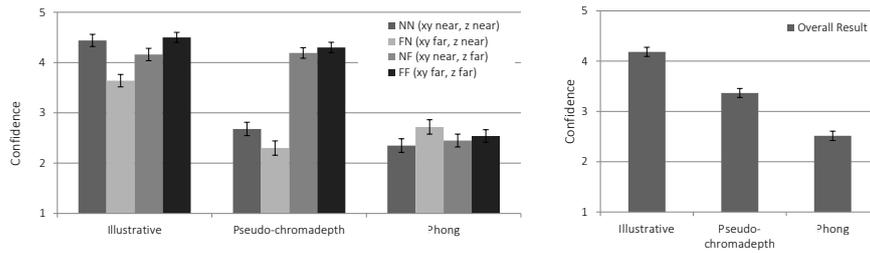


Figure 83: Confidence rated on a 5 point Likert scale (5 = very confident, 1 = very unconfident). The errors bars represent the standard error.

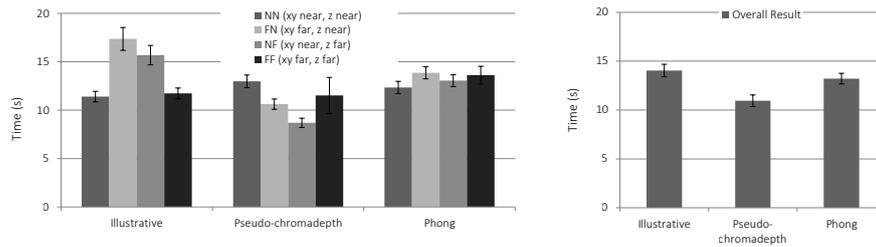


Figure 84: Time for task completion. The errors bars represent the standard error

ment of a medium able to display color and intensity gradients, however, color coding can be used to encode additional information. Especially if relative distances between vascular structures are small, supporting lines and illustrative shadows are a great help to perceive the depth.

Regarding the qualitative study, minor alterations were made and a web-based quantitative evaluation with 50 subjects was conducted. The results indicate that the new techniques enhances depth assessment of vascular models for static 2D representations. The new approach leads to more correct depth assessments compared to pseudo-chromadepth rendering which was rated best in previous studies [RSH06; KOCC13]. Moreover, the subjects were even more confident in their decision. As a drawback, the subjects needed more time to assess the depth of the illustrative technique. This is because the supporting lines and the shadows need more cognitive effort to map from

the 3D scene to the plane. Therefore, it is a tradeoff between a fast depth assessment and a correct estimation.

According to the results of the qualitative study, the depth-dependent contour was rarely used as differences in contour width were hard to perceive. The contour might not be the main depth cue when illustrative shadows and supporting lines are available. One idea was to use the contour to encode additional information, e.g. instead of using a white contour, color-coding can be applied to encode distances. However, contours can yield a false assessment of vessel diameters which might be critical in some cases.

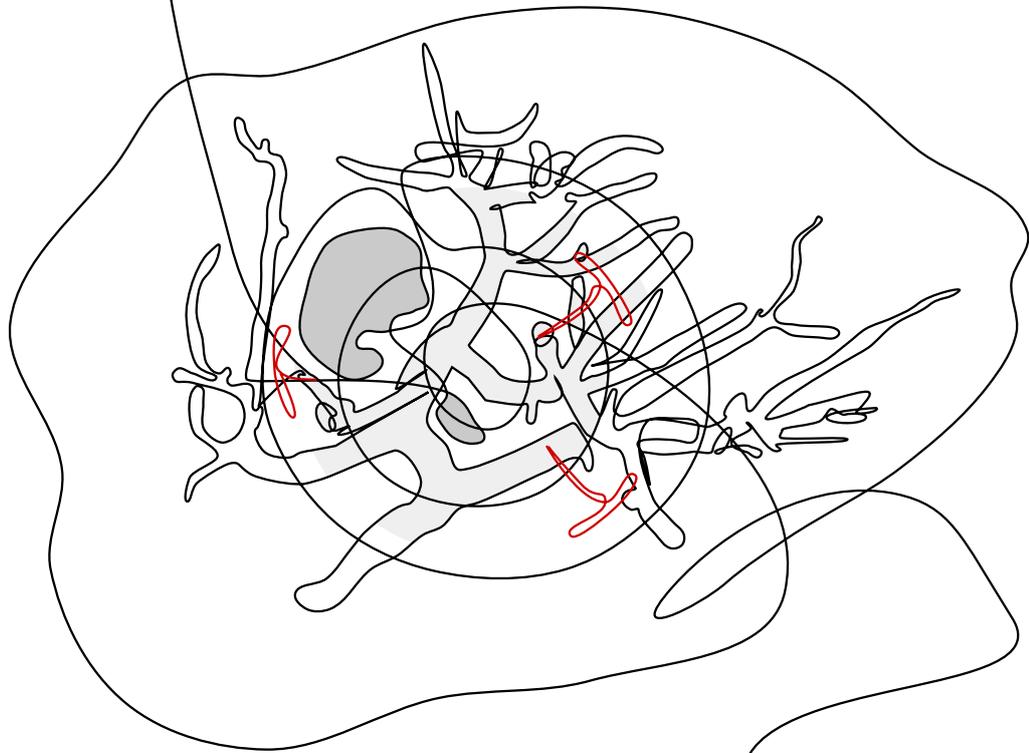
Table 4: ANOVA results of different effects.

Effect	df1		df2		Assessment		Confidence		Time	
					F	η^2	F	η^2	F	η^2
Visualization techniques	2		98		93.25**	.66	118.61**	.71	16.71**	.25
xy distances	1		49		22.86**	.32	.77	.02	2.62	.05
z distances	1		49		27.31**	.36	110.43**	.69	1.90	.04
Visualization techniques * xy	2		98		4.42*	.08	7.05**	.13	.37	.01
Visualization techniques * z	2		98		20.16**	.29	85.94**	.64	2.00	.04
xy * z	1		49		.03	.00	12.92**	.21	4.12*	.08
Visualization techniques * xy * z	2		98		2.12	.04	16.38**	.25	28.38**	.37

*p < 0.05 and **p < 0.01

9

SUPPORTING ANCHORS



This chapter is partly based on:

Kai Lawonn, Maria Luz, and Christian Hansen.

“Improving Spatial Perception of Vascular Models using Supporting Anchors and Illustrative Visualization.”

In: Computers & Graphics (2017), in print.

IMPROVED DEPTH-PERCEPTION WITH SUPPORTING ANCHORS

THIS chapter introduces novel ideas to tackle the problem presented in Chapter 8. In contrast to the previous approach, the new method can be applied for possible augmented reality applications as no orientation is assumed. The orientation was implicitly given by the underlying plane. Furthermore, supporting anchors help to assess the depth, contrarily to the previous approach, the anchors are optimally placed such that no user input is necessary.

9.1 METHOD

In this section, the various visualization techniques are described that are included in this approach. The visualization techniques are illustrated in a sketched overview, see Fig. 85. First, the determination of potential candidates for anchor placement is described. Second, it is explained how the optimal positions for the supporting anchors are obtained. Finally, a novel hatching technique is introduced that is used for the visualization of vascular structures.

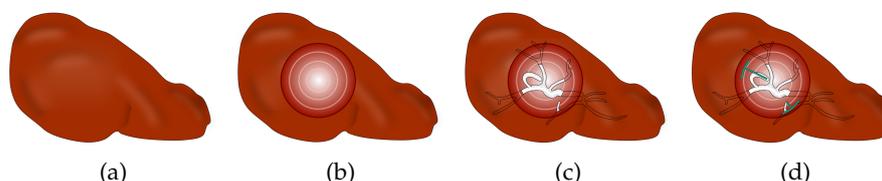


Figure 85: Sketched overview of the visualization approach. In the liver model, (a) a cylinder is generated (b). Inside the cylinder, the vessels are drawn with hatching strokes (c). Finally, supporting anchors are automatically placed for improved depth assessment (d). The supporting anchors are connected to the cylinder and rings on the cylinder indicate the depth of these anchors.

9.1.1 From Voxels to 3D Anatomical Models

The segmentation of vascular structures from voxel data (CT or MRI) is an essential step for the proposed visualization technique. To this end, a semi-automatic region-growing algorithm is used. Subsequently, a 3D representation based on convolution surfaces [OP05] is generated. This approach is model-based and features circular cross sections as a major model assumption. Due to the properties of this im-

PLICIT surface reconstruction technique, the resulting vascular surface models are very smooth and not subject to noise artifacts. Tumors are segmented using a threshold-based segmentation approach, combined with model-based morphological processing [MBKD+09].

9.1.2 Visualization Approach

To improve the understanding of spatial relations, illustrative visualization methods are employed that are tailored for complex 3D vascular models. For this, the vascular 3D model is indicated using contours, and tumors are shown using toon shading. To illustrate a resection volume, a cylinder is used that is placed inside the vascular model and serves as a cutaway surface. This cylinder is used to define the borders of a wedge resection, a type of resection often performed in laparoscopic liver surgery to resect a single tumor or metastasis. Additionally, equidistant rings are generated on the cylinder that encode the depth. The larger the ring, the closer it is to the viewer. Inspired by the visualization concept of interactive lenses [BSPB+93], vessels inside the cylinder are highlighted with a novel and fast hatching scheme, which is referred as the focus region. The hatching strokes encode both, the distance to the tumor, as well as the shape of vessels, cf. Fig. 91. In addition, supporting anchors are displayed. They connect important vessel endpoints and the closest point on the cylinder surface. These supporting anchors provide an intuitive depth cue, such that projected lines on the cylinder indicate which vessel is closer to the viewer. For this, the rings on the cylinder support the depth assessment. In case of two supporting anchors, e.g., Fig. 85d, the anchors and the proximity to the rings give the crucial hint. In Fig. 85d, the bottom anchor is closer to the first ring than the upper anchor. Therefore, the bottom anchor is closer to the viewer. The size of the cylinder, and thus the wedge resection, depends on the size of the lesion and the safety margin, which is usually between 4-8 cm. This radius is used as a preset, but the user can interactively change the radius.

9.1.3 Candidate Selection

Before the optimal supporting anchor placement is explained, a set of reasonable candidates, i.e., positions on the anatomical 3D model is needed. Vessel endpoints are assumed as suitable points to place an anchor. The reconstruction of the vessels, see Sec. 9.1.1, yields smooth cap-shaped ends in the vessel geometry. Because a cylinder is used to cut into the liver, the observer either sees the vessel ends or a part of a vessel segment that enters the cylinder and leaves again, see Fig. 85d. Mostly, vessels do not bend strongly such that the course of the vessel can be approximately assessed and, therefore, the vessel

ends serve as ideal candidates. Because the surface meshes are used, the ends need to be extracted. For this, the algorithm is based on the curvature measures. See Sec. 2.1.1.3, for an explanation how the curvature values κ_1, κ_2 are determined. Afterwards, the shape index is calculated

$$s = \frac{1}{2} - \frac{1}{\pi} \tan^{-1} \left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2} \right), \quad (58)$$

with $\kappa_1 \geq \kappa_2$. This representation of the shape index is chosen such that it ranges from 0 to 1, where 0 stands for a cup-shaped and 1 for a cap-shaped form. The ends of the vessel are cap-shaped, and thus correspond to a shape index of 1. A suitable threshold t is needed such that the points p_i can be taken into account for which $s(p_i) \geq t$ holds. Therefore, the shape index for every point is calculated, and afterwards the threshold is determined by Otsu's method [Ots79]. For this, the algorithm tries various thresholds such that two classes are formed. Then, the algorithm searches for the threshold that minimizes the intra-class variance or, equivalently, tries to maximize the inter-class variance. The threshold with highest inter-class variance is used to define the regions of interest. This results in various regions with connected points, fulfilling $s(p_i) \geq t$. The next step is to identify one candidate for each region. For this, the clustered regions are identified via a connected component analysis and look for the point with the highest shape index. This point is then used as a potential candidate for the supporting anchor.

9.1.4 *Optimal view-dependent Candidate Picking*

In this section, the choice of the optimal points from the previously determined set of candidates is explained. In contrast to the previous method 8, where the anchor selection was done manually by a user, a way to select the optimal candidates in a view-dependent manner is proposed. For this, all candidates are identified that are within the range of the cylinder, which is used as the cutaway surface (see Sec. 9.1.2). As selecting all the candidates would lead to visual clutter, a reduced set of optimal candidates needs to be picked. To this end, optimal candidates fulfill the following requirements:

Req. 1 DEPTH DISTRIBUTION: Candidates are equally distributed in depth.

Req. 2 RING DISTRIBUTION: The projection of the candidates to the cylinder is equally distributed on the ring.

The number of optimal points can be set by the user. 15 points turned out to be an optimal number for the data, which were encountered. More points would increase the visual clutter and less would miss

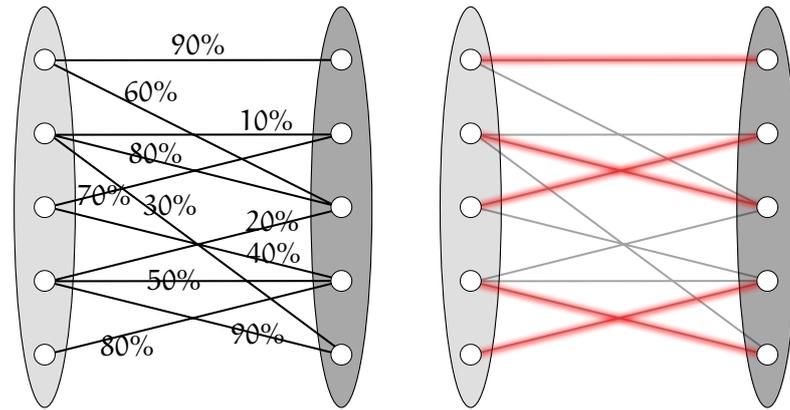


Figure 86: A weighted bipartite graph (left) and the optimal matching (right). For this problem, the left set of points describes the depth distribution, and the right set of points the ring distribution, respectively.

important parts. The number may vary depending on the complexity of the vascular structures, the underlying image quality, the quality of the segmentation, and the kind of organ that is visualized.

The challenge is to select candidates that fulfill both Req. 1 and Req. 2. In case only requirement Req. 1 is fulfilled, it may happen that the lines are projected to a narrow region on the cylinder or that lines overlap each other. In case only Req. 2 is fulfilled, candidates may only be placed at a specific depth, hindering perception of the full depth extent of the vessel tree. When both requirements are fulfilled, an optimal anchor candidate set are obtained that is evenly distributed, both in terms of depths and angles. The problem of satisfying both requirements is solved with algorithms that can be found in the mathematical field of graph theory.

GRAPH THEORY AND BIPARTITE GRAPHS In graph theory, a bipartite graph is written as $G = (U, V, E)$ and it is a graph consisting of nodes and edges whose nodes can be divided into disjoint sets U, V such that there exist no edge connecting two points in U or V . Thus, every edge connects a node from U and V . One well-known problem in the context of bipartite graphs is Hall's marriage theorem. Assume given are two groups of men and women. Every woman states with whom she may be happily married. The problem is now to wed the women such that every woman is married to a man she agrees to, and that only one woman is married to only one man. Hall's theorem states that in case certain properties are fulfilled, a perfect matching exists, such that every woman is married at the end with a man she agrees to. The complexity of the problem can be raised by assigning a percentage to every potential marriage, which indicates the chance of not getting a divorce. Assume that a woman W_1 likes man M_1 (and vice versa), and the edge is weighted with 100%, then they will live

happily ever after. In case the assigned percentage is 10%, both parties should set up a prenup before marriage. This problem, also called the assignment problem, is to find a maximal weighted matching, i.e., a set of marriages such that the probability of a lasting marriage is high (see Fig. 86). For a given bipartite graph with weighted edges, the maximum matching can be found by applying the algorithm of Galil [Gal86].

DEPTH-BASED SELECTION In the following, a weighted bipartite graph with the requirements Req. 1 and Req. 2 is created. Afterwards, the algorithm yields the optimal candidates such that both Req. 1 and Req. 2 are fulfilled in an optimal way. Given the view direction and the position of the camera, first the minimal d_{\min} and maximal d_{\max} distance of the candidates to the camera's view plane are determined. With the given candidates c_1, \dots, c_n a function d is obtained assigning every candidate c_i its depth ranging in $[d_{\min}, d_{\max}]$. The depth-wise optimal position of the $n_c = 15$ optimal candidates would be

$$p_d(j) = d_{\min} + \frac{j}{n_c - 1}(d_{\max} - d_{\min}) \quad (59)$$

with $j \in \{0, 1, \dots, n_c - 1\}$. The distance of two perfect consecutive points in depth would be $1/(n_c - 1)(d_{\max} - d_{\min})$. For every candidate c_i the closest depth index $\text{ind}_d(i)$ is determined with

$$\text{ind}_d(i) = \text{argmin}_k |d_i - p_d(k)|. \quad (60)$$

Then, c_i is assigned a depth score value

$$s_d(i) = \frac{1}{2 \cdot (n_c - 1)}(d_{\max} - d_{\min}) - |d_i - p_d(\text{ind}_d(i))|. \quad (61)$$

Thus, if c_i lies on a perfect depth value it is assigned $1/(2 \cdot (n_c - 1))(d_{\max} - d_{\min})$ (half of the distance of two consecutive perfect depth values) and if the candidate lies exactly in between, then the score 0 is assigned. This value $s(i)$ is then normalized between 0 and 1.

ANGLE-BASED SELECTION Afterwards, a similar value is computed based on the angle when the candidates are projected on the cylinder. For this, the view plane of the camera is used and the candidates are projected onto the plane. Next, a coordinate system is created on the view plane. The center point of the cylinder is also projected on the plane and yields the origin of the coordinate system. The (normalized) up vector of the camera yields the y-axis and the (normalized) cross-product of the up-vector and the view direction yield the x-axis. With this coordinate system, the projected candidates are normalized. Then, for every candidate the clockwise angle starting from

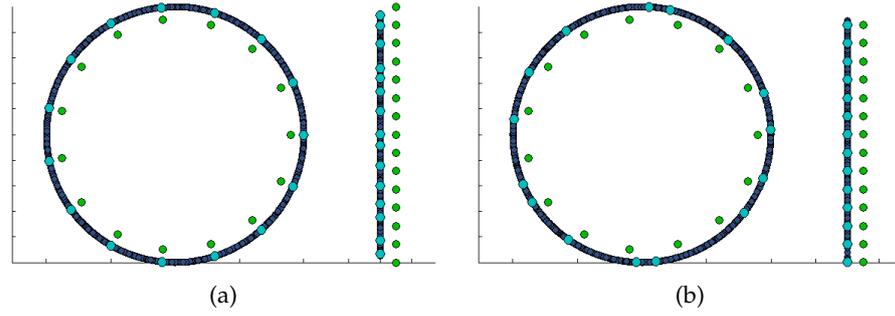


Figure 87: In this example, random points are generated to demonstrate the influence of φ . The ring shows the angle distribution, whereas the bar shows the depth distribution. The green points show the optimal positions. If $\varphi = 0$ is set, (a) is obtained. Here, the points are selected such that they are distributed equally around the ring ((a) left), fulfilling Req. 2. If $\varphi = 1$ is set, then points are selected such that they are evenly distributed in depth ((b) right) and Req. 1 is fulfilled.

the y-axis is determined. Here again, a perfect distribution regarding requirement Req. 2 of the angles would be

$$p_\alpha(j) = 2\pi \cdot \frac{j}{n_c - 1} \quad (62)$$

with $j \in \{0, 1, \dots, n_c - 1\}$. Then, for every candidate c_i its angle α_i is assigned and the closest angle index $\text{ind}_\alpha(i)$ is determined with

$$\text{ind}_\alpha(i) = \text{argmin}_k |\alpha_i - p_\alpha(k)|. \quad (63)$$

Afterwards, every candidate c_i has an angle score value

$$s_\alpha(i) = \frac{2\pi}{2 \cdot (n_c - 1)} - |\alpha_i - p_\alpha(\text{ind}_\alpha(i))|, \quad (64)$$

which is normalized to $[0, 1]$.

BIPARTITE WEIGHTED GRAPH Now, a bipartite weighted graph can be constructed. For this, two disjoint vertex sets of n_c points are generated. Then, for every candidate c_i , an edge connecting the point $\text{ind}_d(i)$ in the first set with the point $\text{ind}_\alpha(i)$ in the second set is constructed. The corresponding edge connecting these two points is then assigned the weight

$$\varphi \cdot s_d(i) + (1 - \varphi) \cdot s_\alpha(i). \quad (65)$$

Here, φ allows to choose points that are more distributed in depth, or more in the angle, see Fig. 87. For equal consideration $\varphi = 0.5$ is set. In case an edge is constructed that already exists, it is replaced if it has a higher score than the existing edge. This assignment yields

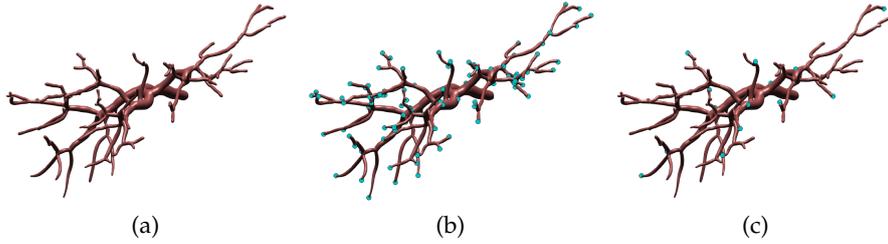


Figure 88: The vessel surface (a) is used to determine the optimal candidates. First, the vessel endpoints are detected (b) and based on the camera properties (camera position and view vector) the optimal candidates are calculated (c).

a weighted bipartite graph. Applying the algorithm of Galil [Gal86] yields a maximum weighted matching. As the optimal set of edges are determined, the corresponding point candidates can be identified, and this gives the optimal selection of candidates that fulfill requirement Req. 1 and Req. 2. In the case that no candidates lie inside the cylinder, no optimal candidates can be determined and thus no anchor is generated. This also means that no vessel endpoint lies inside the cylinder, which in the cases never occurred. This approach is view-dependent and, thus, whenever camera interaction stops, new optimal candidates are determined and the supporting anchors are updated. In Fig. 88, the results of the different steps are described. Furthermore, the liver surface is added to depict the visualization technique in context, see Fig. 89.

9.1.5 Hatching

To apply hatching methods, the PCDs $\mathbf{v}_1, \mathbf{v}_2$ are required, see Section 2.1.1.3. For the results, the PCDs were not smoothed.

Instead of generating streamlines c by propagating them along the PCD \mathbf{v}_1 , such that $c' = \mathbf{v}_1$ holds, a scalar field is determined. The scalar field f should fulfill the property $\nabla_c f = 0$, meaning that c is an isoline of f . In order to obtain f , the integrated divergence is needed associated with the vertex i , which can be determined by:

$$\nabla \cdot \mathbf{W} = \frac{1}{2} \sum_j \cot \alpha_1 (\mathbf{e}_1 \cdot \mathbf{w}_j) + \cot \alpha_2 (\mathbf{e}_2 \cdot \mathbf{w}_j), \quad (66)$$

where the sum is taken over incident triangles j with \mathbf{e}_1 and \mathbf{e}_2 are the two edge vectors starting from vertex i , α_1 and α_2 are the opposing angles of vertex i , and \mathbf{w}_j as the vector in the tangent space associated to the triangle j , see Fig. 90 left. Therefore, the divergence of the PCD \mathbf{v}_2 is determined. Note that the PCD associated with the smaller curvature is used. Especially if tubular surfaces, e.g., vessels, are given, the streamlines need to be placed along the curvature direction \mathbf{v}_1 associated with the largest curvature. Thus, streamlines are

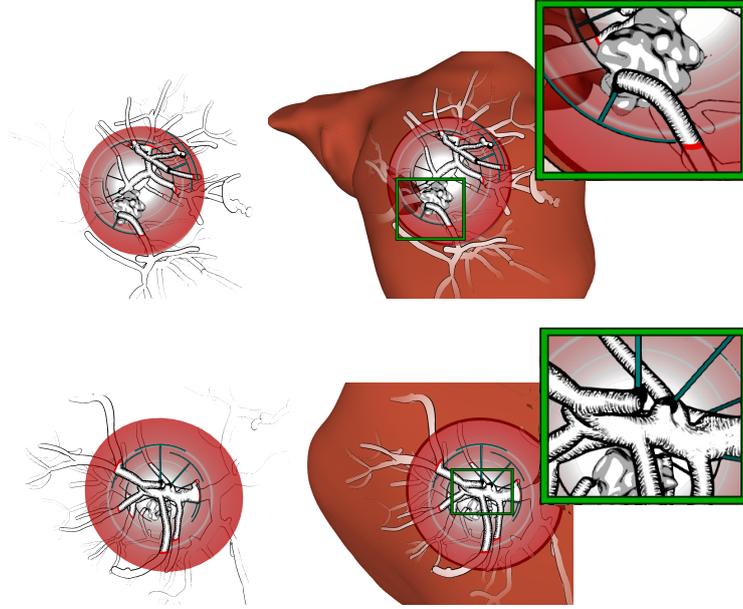


Figure 89: The result of the novel illustrative visualization approach. First, the vessel with the tumor and anchors is shown. The close-up shows that the hatching is consistent and the anchor points are set at reasonable positions, in terms of depth and circular distribution.

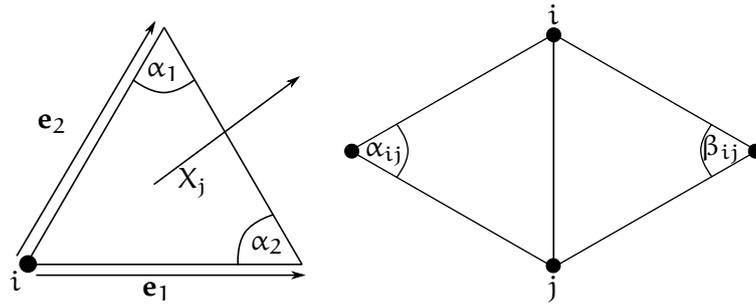


Figure 90: Triangles of the mesh with the notation for determining the divergence (left) and the Laplacian (right).

obtained that wrap around the vessel. Since the gradient of a scalar field is perpendicular to the isolines, it should be parallel to \mathbf{v}_2 . Let $\mathbf{b} \in \mathbb{R}^{|V|}$ be the vector of divergences of the PCD \mathbf{v}_2 . The scalar field f is computed by solving the Poisson problem

$$Lf = \mathbf{b}. \quad (67)$$

Here L is the discretized Laplacian [Mac49]

$$(Lf)_i = \sum_j (\cot \alpha_{ij} + \cot \beta_{ij})(f_j - f_i), \quad (68)$$

where the sum is taken of all neighboring vertices j , and α_{ij} , β_{ij} are the opposing angles of the corresponding edge, see Fig. 90 right. A

least squares method is applied to solve the Poisson problem. The resulting scalar field has isolines along the PCD \mathbf{v}_1 . The problem is that the difference of the scalar field at two different positions does not correspond to the geodesic distance between both positions. This means, if two isolines of certain values are drawn, e.g., u_1, u_2 , the distance between them may strongly vary. It would be desired that the distance should be proportional to $|u_1 - u_2|$. Thus, it should be ensured that neighbored isolines have approximately the same distance as neighboring isolines, at a different region on the vessel.

9.1.6 Real-time Hatching with Isolines

To generate the hatching lines, the approach by Hummel et al. is applied [HGHH+10]. The scalar field is analyzed in the fragment shader on the GPU to determine if the derivative of the scalar field varies strongly. Therefore, the isolines can be adapted depending on the change. This yields regular isolines on the surface. Similar to our previous work [LMP13], the isolines should illustrate the vessel only at the contour margin. Here, the method by Kindlmann et al. [KWTM03] is used to ensure a margin at the contour that has nearly the same thickness everywhere. For this, the previously computed PCDs and the curvatures are needed. First, the view direction is projected onto the tangent plane given by the normal for every vertex. This direction can be used to determine the curvature value along this direction. For this, the distribution of the projected vector with the PCDs is determined and thus the (view) curvature term κ_d is obtained. Then the isolines are drawn whenever

$$|\langle \mathbf{d}, \mathbf{n} \rangle| \leq \sqrt{T \cdot \kappa_d (2 - T \cdot \kappa_d)} \quad (69)$$

is fulfilled. T is the thickness, which is set to 2 in the results and \mathbf{d} is the view direction. Additionally, the isolines are faded out if they get close to the contour margin. In case the vessel parts are within a region of 5mm around the tumor (tumor safety margin), the isolines are darken completely, see Fig. 91. This yields a spatial understanding that provides insights on where critical vessel parts are located.

9.2 IMPLEMENTATION

In this section, selected details from the implementation of the visualization approach are described, see Fig. 92 for a visual overview using a real data-set. The method consists of two major steps, i.e., the pre-processing step and the run-time, which are again divided in different parts:

1. Pre-processing step:
 - Determine the possible candidates for the supporting anchors, Sec. 9.1.3

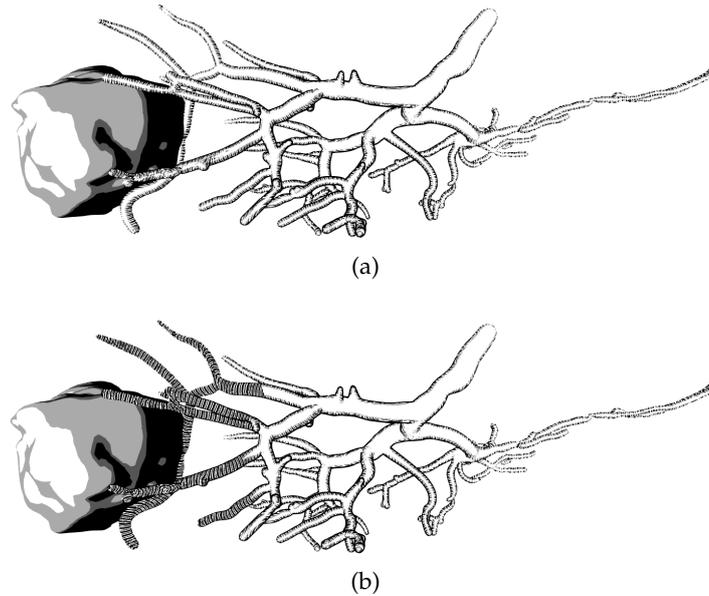


Figure 91: In this example, the difference between hatching strokes that do not show the distance to the tumor (a) and where additional distance cues are activated (b) are shown.

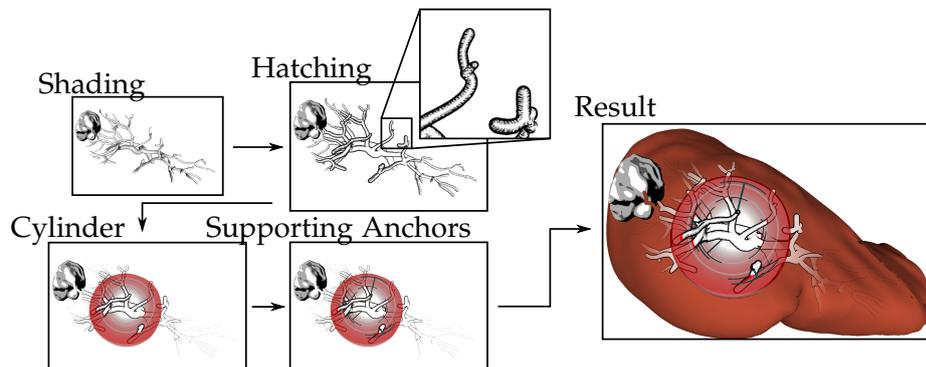


Figure 92: The pipeline of the anchor-based visualization technique. First, it starts with the shaded mesh. Afterwards, the hatching technique is applied to represent the model without colors. Next, the cutting cylinder is visualized. The optimization algorithm determines the supporting anchors for improved spatial perception. Finally, the liver model can be visualized.

- Compute the scalar field for the hatching illustration, Sec. 9.1.5

2. Run-time:

- Visualize the tumor, liver, and the cylinder, Sec. 9.1.2
- Pick the optimal view-dependent candidates, Sec. 9.1.4
- Illustrate the surface with hatching, Sec. 9.1.6

As described, the pre-processing steps consists of the determination of reasonable candidates and the computation of the scalar field for

the hatching illustration. The candidates are stored as a vertex ID list. The scalar field is stored as a vertex buffer object such that every vertex is assigned a value. Both lists are saved in a separate file. The pre-processing step is done on the CPU. During run-time, the tumor, liver, and the cylinder are visualized on the GPU. With the precomputed scalar field, the vessels are illustrated as well and depicted with hatching strokes. The strength of the strokes determines the distance to the tumor. The hatching illustration is also visualized with the GPU. The computation of the view-dependent candidates is optional: either the candidates are calculated for every frame, or the candidates are newly determined whenever no interaction is performed. The calculation of the candidates is done on the CPU with no noticeable drop in frame rates. The latter approach is preferred, such that the candidates are only determined when there is no interaction. This means that the candidates are maintained during the interaction, but the supporting lines from the candidates to the cylinder change such that the distance from the candidate to the cylinder surface is minimal. The calculation of the optimal candidates after the interaction stops causes less visual clutter and is thus preferred. As stated previously, the number of candidates are set to 15, but the user can change this number during run-time.

9.3 QUANTITATIVE EVALUATION

To assess the advantages of the visualization concept, a quantitative study using a web-based questionnaire was conducted. A web-based study was used, because this allows to include a wider range of participants in the study. Because in Kersten-Oertel et al. [KOCC13], pseudo-chromadepth rendering was rated as the best depth cue for vascular visualization, the illustrative visualization technique was compared with pseudo-chromadepth rendering [RSHo6], and Phong shading [Pho75] without an explicit depth cue. The subject pool consists of 81 subjects (36 female, 45 male, age $M=30.20$, range 18-61). 38 subjects are computer scientists, and the others are physicians (15), natural scientists (8), economists (5), engineers (3), and from other backgrounds (12). Most of them (56) have no experience with vascular visualizations.

To probe the subjects' assessment of relative depth distances, different tasks were designed that require a precise judgment of depth. For each task, the subjects had to sort the correct depth order of two labeled branches, see Fig. 93. Tasks were defined where the z distance between two branches is low (<20 mm), because such close spatial relations are difficult to assess for vascular visualizations. These tasks featured various vessel models that were rendered with three visualization techniques comprising the new visualization technique, pseudo-chormadepth rendering, and Phong shading. The models were

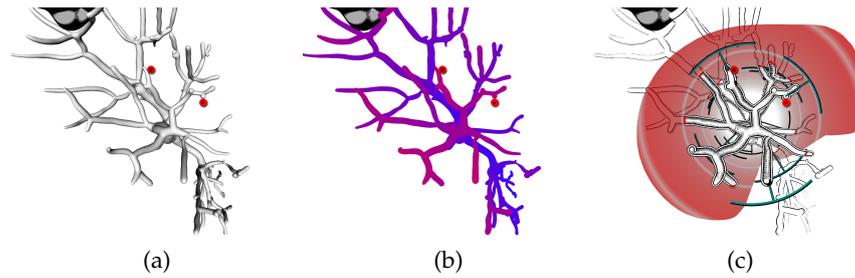


Figure 93: Three scenes with a tumor and a liver vessel are illustrated using different renderings styles (representing 3 of the 24 tasks). From left to right: Phong shading, pseudo-chromadepth, and the novel technique. The labels for the distance assessment tasks are colored in red.

presented to the subjects in pseudo-randomized order to minimize learning effects. For every visualization technique, eight tasks were performed resulting in 24 tasks in overall.

The data collection procedure included a training phase of participants, in which all three visualization techniques were explained, and subjects were asked to perform two test tasks (not included in the evaluation) for each visualization technique. The participants were instructed to complete the tasks as fast and as correct as possible. They were also informed that time is measured. To avoid guessing, which can bias the results, participants were instructed not to guess and could refrain from answering. After completing a task, subjects were asked to rate their confidence on a 5-point Likert confidence scale, where 5 means very confident and 1 very unconfident. If the participants stated *no idea* they could proceed with the next task without rating the confidence. For evaluation all *no idea* statements were given value 1 *very unconfident*. Dependent variables for the study are the distance assessment as a percentage of correct estimations, time needed to complete the task, and the confidence of subjects.

9.4 RESULTS

Data from the evaluation study was analyzed by within-subject one-way ANOVA with the only factor representing the three experimental conditions (visualization techniques: illustrative, pseudo-chromadepth, and Phong) (see Table 5 for the statistical results). ANOVA is an extension of two-sample t-test and is used to compare means and data variation of several groups/conditions. Because the ANOVA was significant, three t-tests were performed for paired samples (within-design) to locate exactly between which of the three visualizations the differences exist. The results of the statistical analysis are described in the following.

Table 5: ANOVA and t-test results for dependent variables distance assessment, confidence and time needed to perform assessment. Alpha error for ANOVA is .05 and for pairwise comparisons .05/3 (Bonferroni corrected). Significant results are identified by *.

	df	F	t	p	η^2	d	Effect
Distance assessment							
ANOVA	2, 160	215.53		<.01 *	.73		large
Pairwise comparisons							
Illustrative vs. Pseudo-chromadepth	80		10.48	<.01 *		1.16	large
Illustrative vs. Phong shading	80		19.73	<.01 *		2.19	large
Pseudo-chromadepth vs. Phong shading	80		10.93	<.01 *		1.21	large
Confidence							
ANOVA	2, 160	131.04		<.01 *	.62		large
Pairwise comparisons							
Illustrative vs. Pseudo-chromadepth	80		14.50	<.01 *		1.61	large
Illustrative vs. Phong shading	80		10.56	<.01 *		1.17	large
Pseudo-chromadepth vs. Phong shading	80		5.10	<.01 *		.57	medium
Time							
ANOVA	2, 160	7.88		<.01 *	.09		medium
Pairwise comparisons							
Illustrative vs. Pseudo-chromadepth	80		3.56	<.01 *		.40	small
Illustrative vs. Phong shading	80		1.50	.14		.17	no effect
Pseudo-chromadepth vs. Phong shading	80		2.69	<.01 *		.30	small

Distance Assessment. Study participants were able to assess the relative distances more precisely with the illustrative visualization ($M = 0.84$, $SD = 0.19$), followed by pseudo-chromadepth ($M = 0.54$, $SD = 0.19$) (see Fig. 94, top). They obtained the worst task performance with Phong shading ($M = 0.26$, $SD = 0.16$). These differences were reflected in significance in the ANOVA and t-tests with large effects.

Exploratively, wrong assessments were looked at and the stimuli where participants refrained from answering. For illustrative visualization, participants assessed only 10% of stimuli incorrectly, and in only in 5% of stimuli they refrained from answering. For pseudo-

chromadepth, participants assessed the relative distances incorrectly in 23% of cases and refrained from answering in 23% of cases. For Phong, the most responses (61%) were wrong and only in 13% of cases participants refrained from answering. These results show that participants followed the instructions and did not guess.

Confidence. Participants were most confident with the illustrative visualization ($M = 3.85$, $SD = 0.73$), followed by Phong ($M = 2.78$, $SD = 0.68$), and least confident with pseudo-chromadepth ($M = 2.40$, $SD = 0.51$) (Fig. 94, middle). The ANOVA revealed significant results. The difference between pseudo-chromadepth and Phong shading was highly significant, but the effect was moderate. The difference between illustrative visualization and both other techniques revealed large effects.

Time. As shown in Fig. 94 (bottom), subjects decided significantly faster using pseudo-chromadepth ($M = 9.38$, $SD = 4.08$), followed by Phong ($M = 10.65$, $SD = 3.92$) and illustrative visualization ($M = 11.39$, $SD = 4.21$). The ANOVA, as well as all pairwise comparisons between illustrative visualization and Pseudo-chromadepth and between Pseudo-chromadepth and Phong shading, revealed significant results. However, there was no significant difference in timing between illustrative visualization and Phong shading.

Neither previous experience with vascular visualization nor being a physician had an influence on the performance. However, there was an influence from gender differences: male participants feel more confident estimating distances with illustrative visualization than female participants. Moreover, there was a significant correlation between age and time needed to complete the task for pseudo-chromadepth ($r = 0.35$, $p < 0.01$), as well as a marginal significant correlation for illustrative visualization ($r = 0.20$, $p = 0.08$) and for Phong shading ($r = 0.21$, $p = 0.06$) emerged. The relevance of age was also evident in distance assessment of pseudo-chromadepth (marginal significant correlation: $r = 0.21$, $p = 0.06$).

In summary, the results show that participants were able to assess relative distances more precisely and were most confident using our illustrative visualization approach. However, they reached their answers significantly slower compared to pseudo-chromadepth rendering.

9.5 DISCUSSION AND FUTURE WORK

In this work, an illustrative visualization technique was introduced that aims to improve spatial perception of complex vascular 3D models. As demonstrated by the presented user study, the novel approach has the potential to improve spatial perception for complex anatomical planning models. The results are consistent with the findings of Lawonn et al. [LLPH15], and show that subjects have serious prob-

lems using color-coded depth visualization and Phong shading when the z distance between relevant objects is small.

Supporting anchors were employed that serve as depth cues. Instead of manually setting these anchors, an automatic approach was proposed by reformulating the problem of finding optimal anchor position candidates as finding the maximum weighted matching in a bipartite graph.

For the vascular illustration, a novel hatching technique was developed that gives comparable results to state-of-the-art methods, which reduces the computation time and the memory usage. Instead of propagating streamlines over the surface to illustrate the hatches, a scalar field with the method of *least squares* was calculated. This scalar field has the property that the hatching lines are the isolines of this scalar field.

The suitability of a visualization method for 3D anatomical models depends strongly on the complexity of the underlying planning model. The complexity of a vascular tree model is characterized by many factors, e.g., branching depth, number of overlapping branches, the size and shape of branches. In addition, tumors and metastases could show different complexity, e.g., a model with one tumor is obviously of lower complexity than a model with 30 liver metastases. In the future, anatomical models might be classified into different complexity categories in order to propose appropriate illustrative visualization techniques for each category. Furthermore, the goal was to integrate meaningful interaction concepts. For instance, the possibility to mark a region on the vessel combined with a visualization that reveals which branches would be affected if this part is cut off should be added. Here, an illustration that not only highlights the distances, but also the affected branches at risk zones might be a solution.

Different distinguishable rendering styles depending on the level of surgical risk could be investigated in the future. Therefore, work in the field of surgical risk assessment [DBW11; HZRL+13], which introduces different glyphs and shadow-like distance indicators could be combined with the novel illustrative AR visualization approach. Besides visualization methods, auditory feedback could enhance illustrative visualizations and facilitate spatial assessment during image-guided interventions. On the one hand, a redundant encoding of distance information could be implemented. On the other hand, sound models could provide the user with additional information that is not encoded (or hard to perceive) in the visualization.

The new visualization approach was demonstrated and evaluated using the example of liver surgery. However, the feedback from the participants in the study lead us to conclude that the techniques can be transferred to other medical fields where planning models are applied, e.g., lung (see Fig. 95), pancreas and kidney surgery. With this work, the foundation for future integration and evaluation of the

novel visualization techniques in a surgical augmented reality simulator was laid, see Figure 96. This could be utilized to reduce incorrect spatial interpretation in medical AR, to enhance established AR visualization methods, and to aid physicians to reduce risks during image-guided interventions.

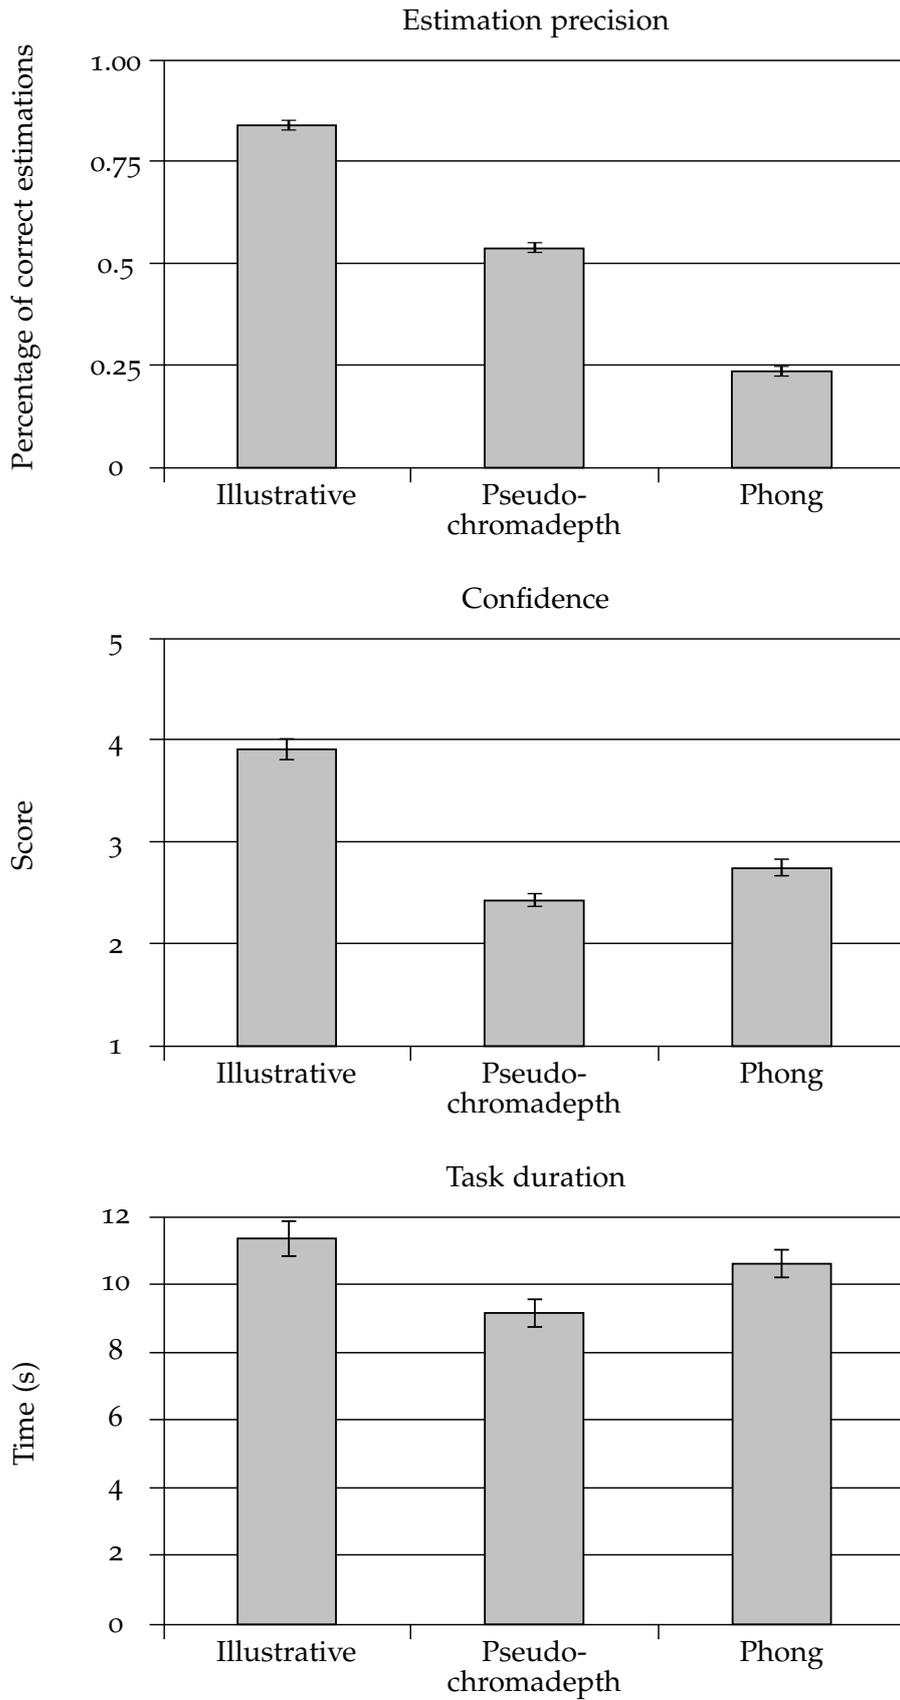


Figure 94: Results for the distance estimation on a scale from 0 (false) to 1 (right) (top). The confidence rated on a 5-point Likert scale (5 = very confident, 1 = very unconfident) (middle) and the time for task completion (bottom). The errors bars represent the standard error.

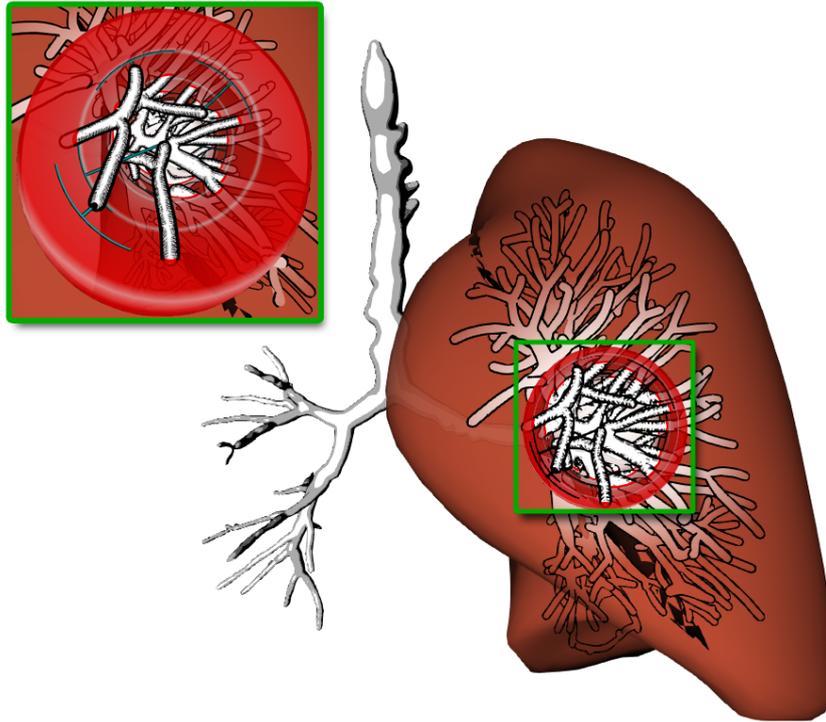


Figure 95: The supporting anchors with our illustration method applied to a lung data set (left lobe) with a vessel tree.

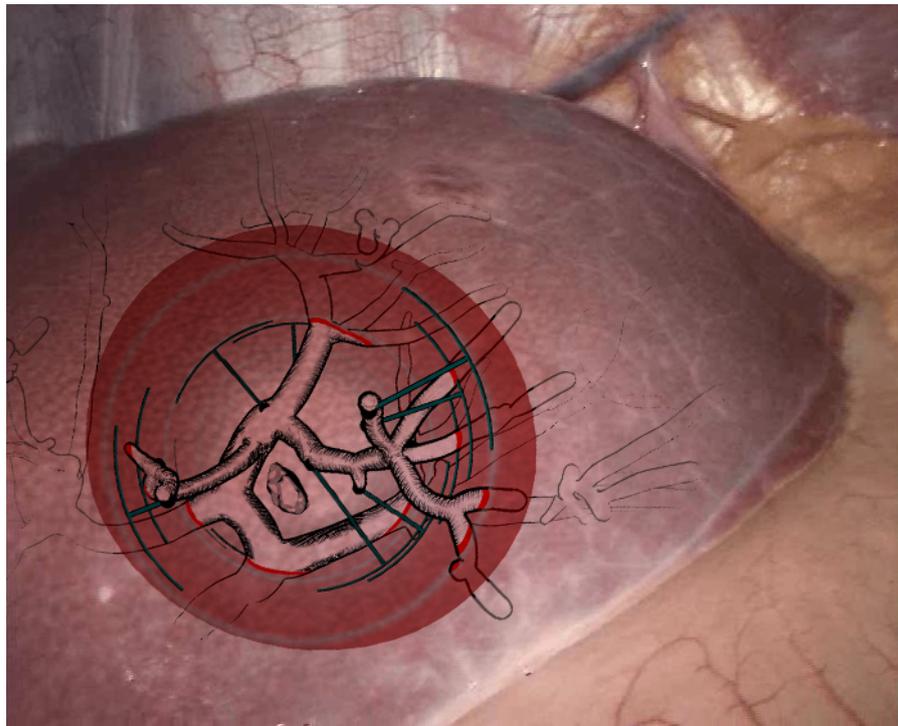


Figure 96: The illustrative visualization technique applied to a laparoscopic view on a liver data set with a tumor. Supporting anchors for depth cues were employed and a hatching scheme to illustrate a spatial impression was designed.

Part V

CONCLUSION

10 CONCLUSION

EOIL

CONCLUSION

THE objective of enhancing certain features, the focus structures, may be simply formulated, but the solution may be anything but easy. The difficulty strongly depends on the formulation of the problem and the underlying data. This habilitation focused on the question of how to enhance focus structures in selected applications ranging from medical applications to monument restoration. For this, the thesis was roughly divided into five parts. Part i gives motivation for the problem of enhancing certain structures and outlines the thesis. A comprehensive background in differential geometry and discrete differential geometry is described in Part ii. This part provides the reader with the necessary foundational knowledge. First, a mathematical background is presented; afterwards, this field is applied to discrete surface meshes. The concepts of continuous measurements, e.g., curvatures and derivatives, are applied to discrete surfaces. Determination of these values on piecewise continuous surfaces is demonstrated. Afterwards, basic visualization approaches are introduced in Part iii. In this chapter, two overviews are presented. The first overview provides a state-of-the-art report on illustrative visualization techniques. It encompasses line drawing techniques on surface meshes, as well as shading techniques. Various evaluation methods that demonstrate how the merit of these techniques can be assessed are additionally presented. The second overview presents visualization techniques for multimodal datasets. It describes smart visibility methods that can be applied to visualize focus structures in the context of PET/CT datasets. Furthermore, various visualization techniques are listed and an overview is presented that comprises the paper in the context of multimodal data visualization, with focus on different smart visibility techniques. Part iv presents the main part of this habilitation. Five applications with different tasks of focus enhancement are introduced. First, an application in the field of PET/CT rendering is shown. Here, the PET structures need to be enhanced with the simultaneous depiction of the surrounding CT anatomical context. Second, an application for the visualization of blood vessels, including wall thickness and animated blood flow, is shown. The blood flow needs to be depicted whereas the user needs additional information about the vessel morphology and the wall thickness. The third application presents an algorithm that analyzes carvings in plasters. The carvings can be extracted and visualized. The last application deals with visualization techniques that improve the depth perception of vessels.

Although much work remains to be done, this habilitation showed different visualization techniques that enhances focus structures in various applications. Novel methods were developed that fill a gap in the literature and all the advantages were confirmed in evaluations. Furthermore, I believe that the presented algorithm is the basis for further improvements and developments.

BIBLIOGRAPHY

- [Dor] <http://webee.technion.ac.il/labs/cgm/Computer-Graphics-Multimedia/Software/EdgeDetection/EdgeDetection.html>.
- [ATGP13] Pascual Abellán, Dani Tost, Sergi Grau, and Anna Puig. "Regions-based illustrative visualization of multimodal datasets." In: *Computerized Medical Imaging and Graphics* 37.4 (2013), pp. 263–271.
- [ARSK15] R. Al-Rousan, M. S. Sunar, and H. Kolivand. "Stylized line drawings for shape depiction." In: *Proc. ICIDM*. Piscataway, NJ, USA: IEEE, 2015, 28:1–28:5.
- [ATHL15] Ahmed Al-Taie, Horst K Hahn, and Lars Linsen. "Uncertainty Estimation and Visualization for Multi-modal Image Segmentation." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine*. The Eurographics Association, 2015, pp. 21–30.
- [ACSDL+03] Pierre Alliez, David Cohen-Steiner, Olivier Devillers, Bruno Lévy, and Mathieu Desbrun. "Anisotropic polygonal remeshing." In: *ACM Transactions on Graphics* 22.3 (July 2003), pp. 485–493.
- [ABGG+12] Lucas Ammann, Pascal Barla, Xavier Granier, Gael Guennebaud, and Patrick Reuter. "Surface Relief Analysis for Illustrative Shading." In: *Computer Graphics Forum* 31.4 (June 2012), pp. 1481–1490.
- [APBEI+08] Luca Antiga, Marina Piccinelli, Lorenzo Botti, Bogdan Ene-Iordache, Andrea Remuzzi, and David A. Steinman. "An image-based modeling framework for patient-specific computational hemodynamics." In: *Medical & Biological Engineering & Computing* 46.11 (2008), pp. 1097–1112.
- [AMLP+08] S. Appanaboyina, F. Mut, R. Löhner, CM Putman, and J.R. Cebral. "Computational Fluid Dynamics of Stented Intracranial Aneurysms Using Adaptive Embedded Unstructured Grids." In: *Numerical Methods in Fluids* 57.5 (2008), pp. 475–493. ISSN: 1097-0363.
- [App67] Arthur Appel. "The Notion of Quantitative Invisibility and the Machine Rendering of Solids." In: *Proc. ACM National Conference*. New York: ACM, 1967, pp. 387–393.
- [AST09] Ben Avery, Christian Sandor, and Bruce H. Thomas. "Improving Spatial Perception for Augmented Reality X-Ray Vision." In: *Virtual Reality Conference*. IEEE Computer Society, 2009, pp. 79–82.
- [BGCP11] Alexandra Baer, Rocco Gasteiger, Douglas Cunningham, and Bernhard Preim. "Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow." In: *Computer Graphics Forum* 30.3 (2011), pp. 811–820.
- [BHSC+14] Alexandra Baer, Antje Hübler, Patrick Saalfeld, Douglas Cunningham, and Bernhard Preim. "A Comparative User Study of a 2D and an Autostereoscopic 3D Display for a Tympanoplastic Surgery." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM)*. Wien, 2014, pp. 181–190.
- [BLSP15] Alexandra Baer, Kai Lawonn, Patrick Saalfeld, and Bernhard Preim. "Statistical Analysis of a Qualitative Evaluation on Feature Lines." In: *Bildverarbeitung für die Medizin (BVM)*. Springer Verlag, 2015, pp. 71–76.
- [BTBP07] Alexandra Baer, Christian Tietjen, Ragnar Bade, and Bernhard Preim. "Hardware-Accelerated Stippling of Surfaces Derived from Medical Volume Data." In: *IEEE/Eurographics Symposium on Visualization (EuroVis)*. 2007, pp. 235–242.
- [BW13] Dale L Bailey and Kathy P Willowson. "An evidence-based review of quantitative SPECT imaging and potential clinical applications." In: *Journal of Nuclear Medicine* 54.1 (2013), pp. 83–89.
- [BSD09] Michael Balzer, Thomas Schlömer, and Oliver Deussen. "Capacity-Constrained Point Distributions: A Variant of Lloyd's Method." In: *ACM Transactions on Graphics* 28.3 (Aug. 2009), 86:1–86:8.
- [Bano8] Isaac Bankman. *Handbook of medical image processing and analysis*. Academic Press, 2008.
- [BTM06] Pascal Barla, Joëlle Thollot, and Lee Markosian. "X-toon: An Extended Toon Shader." In: *Proc. NPAR*. Annecy, France: ACM, 2006, pp. 127–132. ISBN: 1-59593-357-3.
- [BHIM+05] Dirk Bartz, Hans Hagen, Victoria Interrante, Kwan-Liu Ma, and Bernhard Preim. "Illustrative Rendering Techniques for Visualization – Future of Visualization or Just Another Technique?" In: *Proc. Visualization. Panels 4*. Los Alamitos: IEEE Computer Society, 2005, pp. 715–718.
- [BSW08] Mikhail Belkin, Jian Sun, and Yusu Wang. "Discrete Laplace operator on meshed surfaces." In: *Proc. SCG*. College Park, MD, USA: ACM, 2008, pp. 278–287. ISBN: 978-1-60558-071-5.
- [BA05] Alexander Belyaev and Elena Anoshkina. "Detection of Surface Creases in Range Data." In: *Mathematics of Surfaces XI: 11th IMA International Conference, Loughborough, UK, September 5-7, 2005. Proceedings*. Berlin/Heidelberg: Springer, 2005, pp. 50–61. ISBN: 978-3-540-31835-4.
- [BOBMR+07] M Benamor, L Ollivier, H Brisse, G Moulin-Romsee, V Servois, and S Neuenschwander. "PET/CT imaging: what radiologists need to know." In: *Cancer Imaging* 7.Special issue A (2007), S95.
- [BCGF10] Pierre Bénard, Forrester Cole, Aleksey Golovinskiy, and Adam Finkelstein. "Self-Similar Texture for Coherent Line Stylization." In: *Proc. NPAR*. New York: ACM, 2010, pp. 91–97.
- [BHK14] Pierre Bénard, Aaron Hertzmann, and Michael Kass. "Computing Smooth Surface Contours with Accurate Topology." In: *ACM Transactions on Graphics* 33.2 (Apr. 2014), 19:1–19:21.
- [BLCF+12] Pierre Bénard, Jingwan Lu, Forrester Cole, Adam Finkelstein, and Joëlle Thollot. "Active Strokes: Coherent Line Stylization for Animated 3D Models." In: *Proc. NPAR*. Aire-la-Ville, Switzerland: Eurographics Association, 2012, pp. 37–46.
- [BE99] F. Benichou and G. Elber. "Output Sensitive Extraction of Silhouettes from Polygonal Geometry." In: *Proc. Computer Graphics and Applications*. IEEE Computer Society, 1999, pp. 60–69. ISBN: 0-7695-0293-8.

- [BBKN14] C. Schulte zu Berge, M. Baust, A. Kapoor, and N. Navab. "Predicate-Based Focus-and-Context Visualization for 3D Ultrasound." In: *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014), pp. 2379–2387.
- [BCTS+12] Anne Berres, Mathias Goldau, Marc Tittgemeyer, Gerek Scheuermann, and Hans Hagen. "Tractography in Context: Multimodal Visualization of Probabilistic Tractograms in Anatomical Context." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2012, pp. 9–16.
- [BM92] Paul J. Besl and Neil D. McKay. "A Method for Registration of 3-D Shapes." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14.2 (1992), pp. 239–256.
- [BHP15] Johanna Beyer, Markus Hadwiger, and Hanspeter Pfister. "State-of-the-Art in GPU-Based Large-Scale Volume Visualization." In: *Computer Graphics Forum* 34.8 (2015), pp. 13–37.
- [BHWB07] Johanna Beyer, Markus Hadwiger, Stefan Wolfsberger, and Katja Bühler. "High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions." In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1696–1703.
- [BTBK+00] T. Beyer, D. W. Townsend, T. Brun, P. E. Kinahan, M. Charron, R. Roddy, J. Jerin, J. Young, L. Byars, and R. Nutt. "A combined PET/CT scanner for clinical oncology." eng. In: *Journal of Nuclear Medicine: Official Publication, Society of Nuclear Medicine* 41.8 (2000), pp. 1369–1379. ISSN: 0161-5505.
- [BHFN09] C. Bichlmeier, S. M. Heining, M. Feuerstein, and N. Navab. "The virtual mirror: a new interaction paradigm for augmented reality environments." In: *IEEE Trans Med Imaging* 28.9 (2009), pp. 1498–1510.
- [BSHN07] C. Bichlmeier, T. Sielhorst, S. M. Heining, and N. Navab. "Improving Depth Perception in Medical AR." In: *Proc. of Bildverarbeitung für die Medizin*. 2007, pp. 217–221.
- [BSPB+93] Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. "Toolglass and Magic Lenses: The See-through Interface." In: *ACM SIGGRAPH*. 1993, pp. 73–80.
- [BBBV12] Åsmund Birkeland, Stefan Bruckner, Andrea Brambilla, and Ivan Viola. "Illustrative membrane clipping." In: *Computer Graphics Forum* 31.3pt1 (2012), pp. 905–914.
- [BFHW+02] W. Birkfellner et al. "A head-mounted operating binocular for augmented reality visualization in medicine—design and initial evaluation." In: *IEEE Trans Med Imaging* 21.8 (2002), pp. 991–997.
- [BBMN+07] Jorik Blaas, Charl P Botha, Charles Majoie, Aart Nederveen, Frans M Vos, and Frits H Post. "Interactive visualization of fused fMRI and DTI for planning brain tumor resections." In: *Proc. of SPIE Medical Imaging*. International Society for Optics and Photonics (SPIE). 2007, 65091P–65091P.
- [BBP07] Jorik Blaas, Charl P Botha, and Frits H Post. "Interactive Visualization of Multi-Field Medical Data Using Linked Physical and Feature-Space Views." In: *Proc. of EuroVis*. 2007, pp. 123–130.
- [BLEL+13] Alexander Bock, Norbert Lang, Guillermo Evangelista, Ralph Lehrke, and Timo Ropinski. "Guiding Deep Brain Stimulation interventions by fusing multimodal uncertainty regions." In: *Proc. of IEEE Pacific Visualization Symposium (PacificVis)*. IEEE. 2013, pp. 97–104.
- [BS05] Udepta D Bordoloi and Han-Wei Shen. "View selection for volume rendering." In: *Proc. of IEEE Visualization*. 2005, pp. 487–494.
- [Bor86] Gunilla Borgefors. "Distance transformations in digital images." In: *Computer vision, graphics, and image processing* 34.3 (1986), pp. 344–371.
- [BGPB+11] Michelle A. Borkin, Krzysztof Z. Gajos, Amanda Peters, Dimitrios Mitsouras, Simone Melchionna, Frank J. Rybicki, Charles L. Feldman, and Hanspeter Pfister. "Evaluation of Artery Visualizations for Heart Disease Diagnosis." In: *IEEE Transactions on Visualization and Computer Graphics* 17 (Dec. 2011).
- [BJHS+09] Silvia Born, Werner Jaipek, Mario Hlawitschka, Gerek Scheuermann, Christos Trantakis, Jürgen Meixensberger, and Dirk Bartz. "Multimodal visualization of DTI and fMRI data using illustrative methods." In: *Bildverarbeitung für die Medizin*. Springer, 2009, pp. 6–10.
- [BMGS13] Silvia Born, Michael Markl, Matthias Gutberlet, and Gerek Scheuermann. "Illustrative Visualization of Cardiac and Aortic Blood Flow from 4D MRI Data." In: *IEEE Pacific Visualization*. 2013, pp. 129–136. ISBN: 9781467347990.
- [BPMS12] Silvia Born, Matthias Pfeifle, Michael Markl, and Gerek Scheuermann. "Visual 4D MRI Blood Flow Analysis with Line Predicates." In: *IEEE Pacific Visualization*. 2012, pp. 105–112.
- [BRLK+04] Tobias Boskamp, Daniel Rinck, Florian Link, Bernd Kümmerlen, Georg Stamm, and Peter Mildnerberger. "New Vessel Analysis Tool for Morphometric Quantification and Visualization of Vessels in CT and MR Imaging Data Sets 1." In: *Radiographics* 24.1 (2004), pp. 287–297.
- [BSBK+11] Andreas Boss, Lars Stegger, Sotirios Bisdas, Armin Kolb, Nina Schwenzer, Markus Pfister, Claus D Claussen, Bernd J Pichler, and Christina Pfannenberger. "Feasibility of simultaneous PET/MR imaging in the head and upper neck area." In: *European radiology* 21.7 (2011), pp. 1439–1446.
- [BKPA+10] Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Levy. *Polygon Mesh Processing*. AK Peters, 2010. ISBN: 978-1-56881-426-1.
- [BCPV+12] Andrea Brambilla, Robert Carnecky, Ronald Peikert, Ivan Viola, and Helwig Hauser. "Illustrative Flow Visualization: State of the Art, Trends and Challenges." In: *Eurographics State of the Art Reports*. Goslar, Germany: Eurographics Association, 2012, pp. 75–94.
- [BBBR+12] Roger Bramon, Imma Boada, Anton Bardera, Joaquim Rodriguez, Miquel Feixas, Josep Puig, and Mateu Sbert. "Multimodal data fusion based on mutual information." In: *Visualization and Computer Graphics, IEEE Transactions on* 18.9 (2012), pp. 1574–1587.
- [BRBB+13a] Roger Bramon, Marc Ruiz, Anton Bardera, Imma Boada, Miquel Feixas, and Mateu Sbert. "An Information-Theoretic Observation Channel for Volume Visualization." In: *Computer Graphics Forum* 32.3pt4 (2013), pp. 411–420.
- [BRBB+13b] Roger Bramon, Marc Ruiz, Anton Bardera, Imma Boada, Miquel Feixas, and Mateu Sbert. "Information Theory-Based Automatic Multimodal Transfer Function Design." In: *Biomedical and Health Informatics, IEEE Journal of* 17.4 (2013), pp. 870–880.

- [BBPHR08] Ralph Brecheisen, Anna Vilanova Bartroli, Bram Platel, and Bart M ter Haar Romeny. "Flexible GPU-Based Multi-Volume Ray-Casting." In: *Proc. of VMV*. 2008, pp. 303–312.
- [BSMB+07] Simon Breslav, Karol Szerszen, Lee Markosian, Pascal Barla, and Joëlle Thollot. "Dynamic 2D Patterns for Shading 3D Scenes." In: *ACM Transactions on Graphics* 26.3 (July 2007), 20:1–20:6.
- [BSS04] John Brosz, Faramarz Samavati, and Mario Costa Sousa. "Silhouette Rendering Based on Stability Measurement." In: *Proc. Conference on Computer Graphics*. Budmerice, Slovakia: ACM, 2004, pp. 157–167. ISBN: 1-58113-967-5.
- [BG06] S. Bruckner and M.E. Gröller. "Exploded Views for Volume Data." In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 1077–1084.
- [BG07] Stefan Bruckner and Eduard Gröller. "Enhancing Depth-Perception with Flexible Volumetric Halos." In: *IEEE Trans. Vis. Comput. Graph.* 13.6 (2007), pp. 1344–1351.
- [BGMP+10] Stefan Bruckner, Eduard Gröller, Klaus Mueller, Bernhard Preim, and Deborah Silver. "Illustrative Focus+Context Approaches in Interactive Volume Visualization." In: *Scientific Visualization: Advanced Concepts*. 2010. Chap. 10.
- [BNZB+08] Andreas K Buck, Stephan Nekolla, Sibylle Ziegler, Ambros Beer, Bernd J Krause, Ken Herrmann, Klemens Scheidhauer, Hans-Juergen Wester, Ernst J Rummeny, Markus Schwaiger, et al. "SPECT/CT." In: *Journal of Nuclear Medicine* 49.8 (2008), pp. 1305–1319.
- [BF08] Michael Burns and Adam Finkelstein. "Adaptive cutaways for comprehensible rendering of polygonal scenes." In: *ACM Transactions on Graphics* 27.5 (2008), 154:1–154:7.
- [BHWV+07] Michael Burns, Martin Haidacher, Wolfgang Wein, Ivan Viola, and Eduard Gröller. "Feature Emphasis and Contextual Cutaways for Multimodal Medical Visualization." In: *Proc. of EuroVis*. Vol. 7. 2007, pp. 275–282.
- [BKRF+05] Michael Burns, Janek Klawe, Szymon Rusinkiewicz, Adam Finkelstein, and Doug DeCarlo. "Line Drawings from Volume Data." In: *ACM Transactions on Graphics* 24.3 (July 2005), pp. 512–518. ISSN: 0730-0301.
- [BVW08] Stef Busking, Anna Vilanova, and Jarke J. van Wijk. "Particle-based non-photorealistic volume visualization." In: *The Visual Computer* 24.5 (2008), pp. 335–346.
- [CL93] Brian Cabral and Leith Casey Leedom. "Imaging Vector Fields Using Line Integral Convolution." In: *Proc. SIGGRAPH*. Anaheim, CA: ACM, 1993, pp. 263–270. ISBN: 0-89791-601-8.
- [CS99] Wenli Cai and Georgios Sakas. "Data Intermixing and Multi-volume Rendering." In: *Computer Graphics Forum* 18.3 (1999), pp. 359–368.
- [CD05] Y. Cai and F. Dong. "Surface hatching for medical volume data." In: *Proc. of International Conference on Computer Graphics, Imaging and Visualization*. 2005, pp. 232–238.
- [CMS99] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann, 1999. ISBN: 978-1-55860-533-6.
- [CS15] Luis Cardona and Suguru Saito. "Hybrid-space Localized Stylization Method for View-dependent Lines Extracted from 3D Models." In: *Proc. NPAR*. Istanbul, Turkey: Eurographics Association, 2015, pp. 79–89.
- [Car76] Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall, 1976.
- [Car92] Manfredo P. do Carmo. *Riemannian Geometry*. Boston, MA: Birkhäuser, 1992.
- [CRSH+13] Onofrio A Catalano, Bruce R Rosen, Dushyant V Sahani, Peter F Hahn, Alexander R Guimaraes, Mark G Vangel, Emanuele Nicolai, Andrea Soricelli, and Marco Salvatore. "Clinical impact of PET/MR imaging in patients with cancer undergoing same-day PET/CT: initial experience in 134 patients - a hypothesis-generating exploratory study." In: *Radiology* 269.3 (2013), pp. 857–869.
- [CP05] F. Cazals and M. Pouget. "Estimating differential quantities using polynomial fitting of osculating jets." In: *Computer Aided Geometric Design* 22.2 (Feb. 2005), pp. 121–146.
- [CCAP+05] Juan R Cebal, Marcelo Adrián Castro, Sunil Appanaboyina, Christopher M Putman, Daniel Millan, and Alejandro F Frangi. "Efficient pipeline for image-based patient-specific analysis of cerebral aneurysm hemodynamics: Technique and sensitivity." In: *IEEE Transactions on Medical Imaging* 24.4 (2005), pp. 457–467.
- [CMWP11a] Juan R. Cebal, F. Mut, J. Weir, and Christopher M. Putman. "Quantitative Characterization of the Hemodynamic Environment in Ruptured and Unruptured Brain Aneurysms." In: *American Journal of Neuroradiology* 32.1 (2011), pp. 145–151.
- [CMWP11b] Juan R. Cebal, Fernando Mut, J. Weir, and Christopher M. Putman. "Association of Hemodynamic Characteristics and Cerebral Aneurysm Rupture." In: *American Journal of Neuroradiology* 32.2 (2011), pp. 264–270.
- [CZLX15] D. Chen, Y. Zhang, H. Liu, and P. Xu. "Real-Time Artistic Silhouettes Rendering for 3D Models." In: *Proc. Computational Intelligence and Design*. Vol. 1. 2015, pp. 494–498.
- [CS92] Xin Chen and Francis Schmitt. "Intrinsic Surface Properties from Surface Triangulation." In: *Proc. the ECCV*. 1992, pp. 739–743. ISBN: 3-540-55426-2.
- [Che09] Simon R Cherry. "Multimodality imaging: Beyond PET/CT and SPECT/CT." In: *Seminars in nuclear medicine* 39.5 (2009), pp. 348–353.
- [CCMH+15] H. Choi, B. Cho, K. Masamune, M. Hashizume, and J. Hong. "An effective visualization technique for depth perception in augmented reality-based surgical navigation." In: *Int J Med Robot* 12(1) (2015), pp. 62–72.
- [CCGP+08] Alan Chu, Wing-Yin Chan, Jixiang Guo, Wai-Man Pang, and Pheng-Ann Heng. "Perception-aware Depth Cueing for Illustrative Vascular Visualization." In: *International Conference on Biomedical Engineering and Informatics*. IEEE Computer Society, 2008, pp. 341–346.
- [CST05] Paolo Cignoni, Roberto Scopigno, and Marco Tarini. "A Simple Normal Enhancement Technique for Interactive Non-photorealistic Renderings." In: *Computer & Graphics* 29.1 (2005), pp. 125–133. ISSN: 0097-8493.

- [CG07] Gregory Cipriano and Michael Gleicher. "Molecular Surface Abstraction." In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1608–1615.
- [CJG09] Gregory Cipriano, George N. Phillips Jr., and Michael Gleicher. "Multi-Scale Surface Descriptors." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1201–1208.
- [CKHST+12] Dane Coffey, Fedor Korsakov, Haleh Hagh-Shenas, Lauren Thorson, Arin Ellingson, David Nuckley, and Daniel F. Keefe. "Visualizing Motion Data in Virtual Reality: Understanding the Roles of Animation, Interaction, and Static Presentation." In: 31.3pt3 (2012), pp. 1215–1224.
- [CSM03] David Cohen-Steiner and Jean-Marie Morvan. "Restricted Delaunay triangulations and normal cycle." In: *Proc. SCG*. San Diego, California, USA: ACM, 2003, pp. 312–321. ISBN: 1-58113-663-3.
- [Coh06] Marcelo Cohen. "Focus and Context for Volume Visualization." PhD thesis. University of Leeds, 2006.
- [CGLB+08] Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz. "Where Do People Draw Lines?" In: *ACM Transactions on Graphics* 27.3 (Aug. 2008), 88:1–88:11. ISSN: 0730-0301.
- [CSDF+09] Forrester Cole, Kevin Sanik, Doug DeCarlo, Adam Finkelstein, Thomas Funkhouser, Szymon Rusinkiewicz, and Manish Singh. "How Well Do Line Drawings Depict Shape?" In: *ACM Transactions on Graphics* 28.3 (July 2009), 28:1–28:9. ISSN: 0730-0301.
- [CRL01] Derek Cornish, Andrea Rowan, and David Luebke. "View-dependent Particles for Interactive Non-photorealistic Rendering." In: *Proc. Graphics Interface*. Ottawa, Ontario, Canada: CIPS, 2001, pp. 151–158. ISBN: 0-9688808-0-0.
- [CM11] Carlos D Correa and Kwan-Liu Ma. "Visibility histograms and visibility-driven transfer functions." In: *Visualization and Computer Graphics, IEEE Transactions on* 17.2 (2011), pp. 192–204.
- [CSFWS03] Mario Costa Sousa, Kevin Foster, Brian Wyvill, and Faramarz Samavati. "Precise Ink Drawing of 3D Models." In: *Computer Graphics Forum* 22.3 (Sept. 2003), pp. 369–379.
- [DK99] X Dai and S Khorram. "Data fusion using artificial neural networks: a case study on multitemporal change analysis." In: *Computers, Environment and Urban Systems* 23.1 (1999), pp. 19–31.
- [DFR04] Doug DeCarlo, Adam Finkelstein, and Szymon Rusinkiewicz. "Interactive Rendering of Suggestive Contours with Temporal Coherence." In: *Proc. NPAR*. Annecy, France: ACM, 2004, pp. 15–145. ISBN: 1-58113-887-3.
- [DFRS03] Doug DeCarlo, Adam Finkelstein, Szymon Rusinkiewicz, and Anthony Santella. "Suggestive Contours for Conveying Shape." In: *ACM Transactions on Graphics* 22.3 (July 2003), pp. 848–855.
- [DR07] Doug DeCarlo and Szymon Rusinkiewicz. "Highlight Lines for Conveying Shape." In: *Proc. NPAR*. San Diego, California: ACM, 2007, pp. 63–70. ISBN: 978-1-59593-624-0.
- [DMSB99] Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H Barr. "Implicit fairing of irregular meshes using diffusion and curvature flow." In: *Proc. of ACM SIGGRAPH*. 1999, pp. 317–324.
- [DHRS+99] Oliver Deussen, Jörg Hamel, Andreas Raab, Stefan Schlechtweg, and Thomas Strothotte. "An Illustration Technique Using Intersections and Skeletons." In: *Proc. Graphics Interface*. San Francisco: Morgan Kaufmann, 1999, pp. 175–182.
- [DHOS00] Oliver Deussen, Stefan Hiller, Cornelius van Overveld, and Thomas Strothotte. "Floating Points: A Method for Computing Stipple Drawings." In: *Computer Graphics Forum* 19.3 (Sept. 2000), pp. 41–50.
- [DI13] Oliver Deussen and Tobias Isenberg. "Halftoning and Stippling." In: *Image and Video-Based Artistic Stylisation*. Ed. by Paul Rosin and John Collomosse. London/Heidelberg: Springer, 2013. Chap. 3, pp. 45–61. ISBN: 978-1-4471-4519-6.
- [DBW11] Christian Dick, Rainer Burgkart, and Rüdiger Westermann. "Distance Visualization for Interactive 3D Implant Planning." In: *IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE Visualization)* 17.12 (2011), pp. 2173–2182.
- [DHSK+13] Stefan Diepenbrock, Sven Hermann, Michael Schäfers, Michael Kuhlmann, and K Hinrichs. "Comparative visualization of tracer uptake in vivo small animal PET/CT imaging of the carotid arteries." In: *Computer Graphics Forum* 32.3pt2 (2013), pp. 241–250.
- [DPLB+11] Stefan Diepenbrock, Jörg-Stefan Prašni, Florian Lindemann, Hans-Werner Bothe, and Timo Ropinski. "2010 IEEE Visualization Contest Winner: Interactive Planning for Brain Tumor Resections." In: *IEEE Computer Graphics and Applications* 5 (2011), pp. 6–13.
- [DWE03] Joachim Diepstraten, Daniel Weiskopf, and Thomas Ertl. "Interactive Cutaway Illustrations." In: *Computer Graphics Forum* 22.3 (2003), pp. 523–532.
- [Dij59] Edsger W Dijkstra. "A note on two problems in connexion with graphs." In: *Numerische Mathematik* 1.1 (1959), pp. 269–271.
- [DMLB12] Andre F van Dixhoorn, Julien Milles, Baldur van Lew, and Charl P Botha. "BrainCove: A Tool for Voxel-wise fMRI Brain Connectivity Visualization." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2012, pp. 99–106.
- [DSCLO7] Belma Dogdas, David Stout, Arion F Chatzioannou, and Richard M Leahy. "Digimouse: a 3D whole body mouse atlas from CT and cryosection data." In: *Physics in medicine and biology* 52.3 (2007), p. 577.
- [Dö7] Jürgen Döllner. "Non-Photorealistic 3D Geovisualization." In: *Multimedia Cartography*. Ed. by William Cartwright, Michael P Peterson, and Georg Gartner. Berlin/Heidelberg: Springer, 2007, pp. 229–240.
- [DCLK03] Feng Dong, Gordon J. Clapworthy, Hai Lin, and Meleagros A. Krokos. "Nonphotorealistic Rendering of Medical Volume Data." In: *IEEE Comput. Graph. Appl.* 23.4 (2003), pp. 44–52. ISSN: 0272-1716.
- [DZHF09] Jiang Dong, Dafang Zhuang, Yaohuan Huang, and Jingying Fu. "Advances in multi-sensor data fusion: algorithms and applications." In: *Sensors* 9.10 (2009), pp. 7771–7784.
- [DC90] Debra Dooley and Michael F. Cohen. "Automatic Illustration of 3D Geometric Models: Lines." In: *ACM SIGGRAPH Computer Graphics* 24.2 (Mar. 1990), pp. 77–82.

- [DCH88] Robert A Drebin, Loren Carpenter, and Pat Hanrahan. "Volume rendering." In: *ACM Siggraph Computer Graphics*. Vol. 22. 4. ACM, 1988, pp. 65–74.
- [DSEB+12] Alexander Drzezga, Michael Souvatzoglou, Matthias Eiber, Ambros J Beer, Sebastian Fürst, Axel Martinez-Möller, Stephan G Nekolla, Sibylle Ziegler, Carl Ganter, Ernst J Rummeny, et al. "First clinical experience with integrated whole-body PET/MR: comparison to PET/CT in patients with oncological diagnoses." In: *Journal of Nuclear Medicine* 53.6 (2012), pp. 845–855.
- [EHZ01] Herbert Edelsbrunner, John Harer, and Afra Zomorodian. "Hierarchical Morse Complexes for Piecewise Linear 2-manifolds." In: *Proc. of Symposium on Computational Geometry*. Medford, Massachusetts, USA: ACM, 2001, pp. 70–79. ISBN: 1-58113-357-X.
- [EWE04] M. Eissele, D. Weiskopf, and T. Ertl. "Frame-to-frame coherent halftoning in image space." In: *Proc. Theory and Practice of Computer Graphics*. 2004, pp. 188–195.
- [ERBH+12] Mathias Eitz, Ronald Richter, Tamy Boubekeur, Kristian Hildebrand, and Marc Alexa. "Sketch-based Shape Retrieval." In: *ACM Transactions on Graphics* 31.4 (July 2012), 31:1–31:10. ISSN: 0730-0301.
- [Elb99] Gershon Elber. "Interactive Line Art Rendering of Freeform Surfaces." In: *Computer Graphics Forum* (1999). ISSN: 1467-8659.
- [EC06] Gershon Elber and Elaine Cohen. "Probabilistic silhouette based importance toward line-art non-photorealistic rendering." In: *The Visual Computer* 22.9 (2006), pp. 793–804. ISSN: 1432-2315.
- [EAT07] N. Elmqvist, U. Assarsson, and P. Tsigas. "Employing dynamic transparency for 3D occlusion management: design issues and evaluation." In: *Proc. of the 11th International Conference on Human-Computer Interaction*. Springer, 2007, pp. 532–545.
- [ESMH+05] F. Enders, N. Sauber, D. Merhof, P. Hastreiter, C. Nimsky, and M. Stamminger. "Visualization of white matter tracts with wrapped streamlines." In: *Proc. of IEEE Visualization*. 2005, pp. 51–58.
- [EBRI09] Maarten H. Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. "Depth-Dependent Halos: Illustrative Rendering of Dense Line Data." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (Nov. 2009), pp. 1299–1306. ISSN: 1077-2626.
- [EBRI15] Maarten H. Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. *Interactive Illustrative Line Styles and Line Style Transfer Functions for Flow Visualization*. Tech. rep. 1503.05787. arXiv.org, Mar. 2015.
- [EBRI11] Maarten Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. "Illustrative Line Styles for Flow Visualization." Anglais. In: *Proc. Pacific Graphics*. Los Alamitos: IEEE Computer Society, 2011, pp. 105–110.
- [FPJT+14] Adam J Fairfield, Jonathan Plasencia, Yun Jang, Nicholas Theodore, Neil R Crawford, David H Frakes, and Ross Maciejewski. "Volume curtaining: a focus+ context effect for multimodal volume visualization." In: *Proc. of SPIE Medical Imaging*. International Society for Optics and Photonics, 2014, pp. 903527–903527.
- [FMTP+09] V. Ferrari, G. Megali, E. Troia, A. Pietrabissa, and F. Mosca. "A 3-D mixed-reality system for stereoscopic visualization of medical dataset." In: *IEEE Trans Biomed Eng* 56.11 (2009), pp. 2627–2633.
- [FPT04] Maria Ferre, Anna Puig, and Dani Tost. "A framework for fusion methods and rendering techniques of multimodal volume data." In: *Computer Animation and Virtual Worlds* 15.2 (2004), pp. 63–77.
- [FMHN08] M. Feuerstein, T. Mussack, S. M. Heining, and N. Navab. "Intraoperative laparoscope augmentation for port placement and resection planning in minimally invasive liver resection." In: *IEEE Transactions on Medical Imaging* 27.3 (2008), pp. 355–369.
- [FBS05] J. Fischer, D. Bartz, and W. Straßer. "Stylized Augmented Reality for Improved Immersion." In: *Proc. of IEEE Virtual Reality* (2005), pp. 195–202. ISSN: 1087-8270.
- [FHS07] J. Fischer, B. Huhle, and A. Schilling. "Using Time-of-Flight Range Data for Occlusion Handling in Augmented Reality." In: *Proc. of Eurographics Symposium on Virtual Environments (EGVE)*. 2007, pp. 109–116.
- [FJLS+16] Jakub Fišer, Ondřej Jamříška, Michal Lukáč, Eli Shechtman, Paul Asente, Jingwan Lu, and Daniel Šykora. "StyLit: Illumination-Guided Example-Based Stylization of 3D Renderings." In: *ACM Transactions on Graphics* 35.4 (2016).
- [Flo03] Michael S. Floater. "Mean value coordinates." In: *CAGD* 20.1 (2003), pp. 19–27. ISSN: 0167-8396.
- [FH05] Michael S. Floater and Kai Hormann. "Surface Parameterization: A Tutorial and Survey." In: *Advances in Multiresolution for Geometric Modelling*. Ed. by Neil A. Dodgson, Michael S. Floater, and Malcolm A. Sabin. Berlin/Heidelberg: Springer, 2005, pp. 157–186. ISBN: 978-3-540-26808-6.
- [FVWK+11] O Fluck, Christoph Vetter, Wolfgang Wein, Ali Kamen, Bernhard Preim, and Rüdiger Westermann. "A survey of medical image registration on graphics hardware." In: *Computer methods and programs in biomedicine* 104.3 (2011), e45–e57.
- [FNVV98] AlejandroF Frangi, Wiroj Niessen, KoenL Vincken, and MaxA Viergever. "Multiscale vessel enhancement filtering." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 1496. 1998, pp. 130–137. ISBN: 3-540-65136-5.
- [FPPK+04] Juhana Frösen, Anna Piippo, Anders Paetau, Marko Kangasniemi, Mika Niemelä, Juha Hernesniemi, and Juha Jääskeläinen. "Remodeling of Saccular Cerebral Artery Aneurysm Wall Is Associated with Rupture: Histological Analysis of 24 Unruptured and 42 Ruptured Cases." In: *Stroke* 35.10 (2004), pp. 2287–2293.
- [FH09] Raphael Fuchs and Helwig Hauser. "Visualization of Multi-Variate Scientific Data." In: *Computer Graphics Forum* 28.6 (2009), pp. 1670–1690.
- [FKSM+04] Thomas Funkhouser, Michael Kazhdan, Philip Shilane, Patrick Min, William Kiefer, Ayellet Tal, Szymon Rusinkiewicz, and David Dobkin. "Modeling by Example." In: *Proc. of ACM SIGGRAPH* 23.3 (Aug. 2004), pp. 652–663.
- [GP13] Ashwini Galande and Rahul Patil. "The art of medical image fusion: a survey." In: *Proc. of Advances in Computing, Communications and Informatics (ICACCI), International Conference on*. IEEE, 2013, pp. 400–405.

- [Gal86] Zvi Galil. "Efficient Algorithms for Finding Maximum Matching in Graphs." In: *ACM Comput. Surv.* 18.1 (Mar. 1986), pp. 23–38. ISSN: 0360-0300.
- [GJ12] A.M. Gambaruto and A.J. João. "Flow structures in cerebral aneurysms." In: *Computers & Fluids* 65.0 (2012). Sixth International Conference on Computational Fluid Dynamics (ICCFD6), pp. 56–65. ISSN: 0045-7930.
- [GH97] Michael Garland and Paul S. Heckbert. "Surface Simplification Using Quadric Error Metrics." In: *Proc. of ACM SIGGRAPH*. 1997, pp. 209–216.
- [Gas14] Rocco Gasteiger. "Visual Exploration of Cardiovascular Hemodynamics." PhD thesis. Germany: Otto-von-Guericke University Magdeburg, 2014.
- [GNBP11] Rocco Gasteiger, Mathias Neugebauer, Oliver Beuing, and Bernhard Preim. "The FLOWLENS: A Focus-and-Context Visualization Approach for Exploration of Blood Flow in Cerebral Aneurysms." In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 2183–2192.
- [GNKP10] Rocco Gasteiger, Mathias Neugebauer, Christoph Kubisch, and Bernhard Preim. "Adapted Surface Visualization of Cerebral Aneurysms with Embedded Blood Flow Information." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2010, pp. 25–32.
- [GTBP08] Rocco Gasteiger, Christian Tietjen, Alexandra Baer, and Bernhard Preim. "Curvature- and Model-Based Surface Hatching of Anatomical Structures Derived from Clinical Volume Datasets." In: *Proc. SmartGraphics*. 2008, pp. 255–262.
- [GEB16] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. "Image Style Transfer Using Convolutional Neural Networks." In: *Proceedings of IEEE CVPR*. 2016.
- [GPOSW11] K. A. Gavaghan, M. Peterhans, T. Oliveira-Santos, and S. Weber. "A portable image overlay projection device for computer-aided open liver surgery." In: *IEEE Trans Biomed Eng* 58.6 (2011), pp. 1855–1864.
- [Gen10] Weidong Geng. *The Algorithms and Principles of Non-photorealistic Graphics: Artistic Rendering and Cartoon Animation*. Berlin/Heidelberg: Springer, 2010.
- [GHMT+17] Ajay Abisheck Paul George, Pascal Heimer, Raphael Menges, Veda Thota, Astrid Maass, Jan Hamaekers, Martin Hofmann-Apitius, Nils Lichtenberg, Kai Lawonn, and Diana Imhof. "Novel insights into μ -conotoxin conformations: A comparative study using molecular dynamics and computer visualization methods." In: *Scientific Reports* (2017), submitted.
- [GKKJ92] Guido Gerig, Olaf Kübler, Ron Kikinis, and Ferenc A. Jolesz. "Nonlinear Anisotropic Filtering of MRI Data." In: *IEEE Engineering in Medicine and Biology Society* 11.2 (Nov. 1992), pp. 221–32.
- [GL13] Moritz Gerl and Tobias Isenberg. "Interactive example-based hatching." In: *Computers & Graphics* 37.1-2 (2013), pp. 65–80. ISSN: 0097-8493.
- [GTSK13] Ayelet Gilboa, Ayellet Tal, Ilan Shimshoni, and Michael Kolomenkin. "Computer-based, automatic recording and illustration of complex archaeological artifacts." In: *Journal of Archaeological Science* 40.2 (2013), pp. 1329–1339. ISSN: 0305-4403.
- [GIHL00] Ahna Girshick, Victoria Interrante, Steven Haker, and Todd Lemoine. "Line Direction Matters: An Argument for the Use of Principal Directions in 3D Line Drawings." In: *Proc. NPAR*. Anney, France: ACM, 2000, pp. 43–52. ISBN: 1-58113-277-8.
- [GLHS+14] Sylvia Glaßer, Kai Lawonn, Thomas Hoffmann, Martin Skalej, and Bernhard Preim. "Combined visualization of wall thickness and wall shear stress for the evaluation of aneurysms." In: *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014), pp. 2506–2515.
- [GW03] Neil D. Glossop and Zhanhe Wang. "Laser projection augmented reality system for computer-assisted surgery." In: *Proc. of the 18th International Symposium on Computer Assisted Radiology and Surgery (CARS)*. 2003, pp. 65–71.
- [GCAR+09] Michael Glueck, Keenan Crane, Sean Anderson, Andres Rutnik, and Azam Khan. "Multiscale 3D Reference Visualization." In: *Proc. of Interactive 3D Graphics and Games*. Boston, Massachusetts, 2009, pp. 225–232. ISBN: 978-1-60558-429-4.
- [Glo4] Jack Goldfeather and Victoria Interrante. "A novel cubic-order algorithm for approximating principal direction vectors." In: *ACM Transactions on Graphics* 23.1 (Jan. 2004), pp. 45–63. ISSN: 0730-0301.
- [GMKB+09] Amy Berrington de González, Mahadevappa Mahesh, Kwang-Pyo Kim, Mythreyi Bhargavan, Rebecca Lewis, Fred Mettler, and Charles Land. "Projected cancer risks from computed tomographic scans performed in the United States in 2007." In: *Archives of Internal Medicine* 169.22 (2009), pp. 2071–2077.
- [GGSC98] Amy Gooch, Bruce Gooch, Peter Shirley, and Elaine Cohen. "A Non-photorealistic Lighting Model for Automatic Technical Illustration." In: *Proc. SIGGRAPH*. New York: ACM, 1998, pp. 447–452. ISBN: 0-89791-999-8.
- [GG01] Bruce Gooch and Amy A. Gooch. *Non-Photorealistic Rendering*. Natick: A K Peters, Ltd., 2001.
- [GSGS+99] Bruce Gooch, Peter-Pike J. Sloan, Amy Gooch, Peter Shirley, and Richard Riesenfeld. "Interactive Technical Illustration." In: *Proc. Interactive 3D Graphics*. Atlanta, Georgia, USA, 1999, pp. 31–38. ISBN: 1-58113-082-1.
- [GVH07] Todd Goodwin, Ian Vollick, and Aaron Hertzmann. "Isophote Distance: A Shading Approach to Artistic Stroke Thickness." In: *Proc. NPAR*. New York: ACM, 2007, pp. 53–62. ISBN: 978-1-59593-624-0.
- [GTBP+10] Leonid Goubergrits, Bente Thamsen, André Berthe, Jens Poethke, Ulrich Kertscher, Klaus Affeld, Christoph Petz, Hans-Christian Hege, Heinrich Hoch, and Andreas Spuler. "In Vitro Study of Near-Wall Flow in a Cerebral Aneurysm Model with and without Coils." In: *American Journal of Neuroradiology* 31:8 (2010), pp. 1521–1528.
- [GTDS04] Stéphane Grabli, Emmanuel Turquin, Frédo Durand, and François X. Sillion. "Programmable Style for NPR Line Drawing." In: *Proc. Conference on Rendering Techniques*. Eurographics Association, 2004, pp. 33–44. ISBN: 3-905673-12-6.

- [GTDS10] Stéphane Grabli, Emmanuel Turquin, Frédo Durand, and François X. Sillion. "Programmable Rendering of Line Drawing from 3D Scenes." In: *ACM Transactions on Graphics* 29.2 (Mar. 2010), 18:1–18:20.
- [Gri05] Landis K Griffeth. "Use of PET/CT scanning in cancer patients: technical and practical considerations." In: *Proceedings (Baylor University. Medical Center)* 18.4 (2005), p. 321.
- [GRBo8] S. Gupta, K.P. Ramesh, and E.P. Blasch. "Mutual Information Metric Evaluation for PET/MRI Image Fusion." In: *Proc. of Aerospace and Electronics Conference*. 2008, pp. 305–311.
- [HGP05] Horst K Hahn, Manfred Georg, and Heinz-Otto Peitgen. "Fractal aspects of three-dimensional vascular constructive optimization." In: *Fractals in biology and medicine*. Springer, 2005, pp. 55–66.
- [HPSPo1] Horst K. Hahn, Bernhard Preim, Dirk Selle, and Hans Otto Peitgen. "Visualization and Interaction Techniques for the Exploration of Vascular Structures." In: *IEEE Visualization*. 2001, pp. 395–402.
- [Hah05] Horst Karl Hahn. "Morphological Volumetry: Theory, Concepts, and Application to Quantitative Medical Imaging." PhD thesis. University of Bremen, 2005.
- [HBG11] Martin Haidacher, Stefan Bruckner, and M Eduard Gröller. "Volume analysis using multimodal surface similarity." In: *Visualization and Computer Graphics, IEEE Transactions on* 17.12 (2011), pp. 1969–1978.
- [HBKGo8] Martin Haidacher, Stefan Bruckner, Armin Kanitsar, and M Eduard Gröller. "Information-based Transfer Functions for Multimodal Visualization." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2008, pp. 101–108.
- [HSC12] Balázs Hajagos, László Szécsi, and Balázs Csébfalvi. "Fast Silhouette and Crease Edge Synthesis with Geometry Shaders." In: *Proc. Conference on Computer Graphics*. Budmerice, Slovakia: ACM, 2012, pp. 71–76. ISBN: 978-1-4503-1977-5.
- [HMD11] Ghassan Hamarneh, Chris McIntosh, and Mark S Drew. "Perception-based visualization of manifold-valued medical images using distance-preserving dimensionality reduction." In: *IEEE Transactions on Medical Imaging* 30.7 (2011), pp. 1314–1327.
- [HS03] E. Hameiri and I. Shimshoni. "Estimating the principal curvatures and the Darboux frame from real 3-D range data." In: *Trans. Sys. Man Cyber. Part B* 33.4 (Aug. 2003), pp. 626–637. ISSN: 1083-4419.
- [HWRR+10] C. Hansen, J. Wieferrich, F. Ritter, C. Rieder, and H.-O. Peitgen. "Illustrative Visualization of 3D Planning Models for Augmented Reality in Liver Surgery." In: *Int J Comput Assist Radiol Surg* 5.2 (2010), pp. 133–141. ISSN: 1861-6410.
- [HZHS+09] C. Hansen, S. Zidowitz, M. Hindennach, A. Schenk, H. Hahn, and H.-O. Peitgen. "Interactive Determination of Robust Safety Margins for Oncologic Liver Surgery." In: *Int J Comput Assist Radiol Surg* 4.5 (2009), pp. 469–474.
- [HZPS+14] C. Hansen, S. Zidowitz, P. Preim, G. Stravrou, K. J. Oldhafer, and H. K. Hahn. "Impact of model-based risk analysis for liver surgery planning." In: *Int J Comput Assist Radiol Surg* 9.2 (2014), pp. 473–480. ISSN: 1861-6410.
- [HZRL+13] C. Hansen, S. Zidowitz, F. Ritter, C. Lange, K. Oldhafer, and H. K. Hahn. "Risk maps for liver surgery." In: *Int J Comput Assist Radiol Surg* 8.3 (2013), pp. 419–428.
- [HJ05] Charles D. Hansen and Chris R. Johnson, eds. *The Visualization Handbook*. Amsterdam: Elsevier, 2005.
- [HCZW10] Wen Hao, Wujun Che, Xiaopeng Zhang, and Yinghui Wang. "3D Model Feature Line Stylization Using Mesh Sharpening." In: *Proc. Virtual-Reality Continuum and Its Applications in Industry*. Seoul, South Korea: ACM, 2010, pp. 249–256. ISBN: 978-1-4503-0459-7.
- [HE98] Peter Hastreiter and Thomas Ertl. "Integrated registration and visualization of medical image data." In: *Proc. of Computer Graphics International*. IEEE, 1998, pp. 78–85.
- [HMIBGo1] H. Hauser, L. Mroz, G. Italo Bisch, and E. Groller. "Two-level volume rendering." In: *Visualization and Computer Graphics, IEEE Transactions on* 7.3 (2001), pp. 242–252.
- [HS13] Lei He and Scott Schaefer. "Mesh Denoising via Lo Minimization." In: *ACM Transactions on Graphics* 32.4 (2013), 64:1–64:8. ISSN: 0730-0301.
- [Hec90] Paul S. Heckbert. "Adaptive Radiosity Textures for Bidirectional Ray Tracing." In: *ACM Transactions on Graphics* 24.4 (1990), pp. 145–154. ISSN: 0097-8930.
- [HSFM+08] Anja Hennemuth, Achim Seeger, Ola Friman, Stephan Miller, Bernhard Klumpp, Steffen Oeltze, and Heinz-Otto Peitgen. "A comprehensive approach to the analysis of contrast enhanced cardiac MR images." In: *Medical Imaging, IEEE Transactions on* 27.11 (2008), pp. 1592–1610.
- [Her99] Aaron Hertzmann. "Introduction to 3D Non-Photorealistic Rendering: Silhouettes and Outlines." In: *SIGGRAPH Course Notes*. New York: ACM, 1999.
- [HJOC+01] Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. "Image Analogies." In: *Proc. SIGGRAPH*. New York: ACM, 2001, pp. 327–340.
- [HZoo] Aaron Hertzmann and Denis Zorin. "Illustrating Smooth Surfaces." In: *Proc. SIGGRAPH*. New York: ACM, 2000, pp. 517–526. ISBN: 1-58113-208-5.
- [HH12] Rodney J Hicks and Michael S Hofman. "Is there still a role for SPECT-CT in oncology in the PET-CT era?" In: *Nature reviews Clinical oncology* 9.12 (2012), pp. 712–720.
- [HSTN+04] Fernando Vega Higuera, Natascha Sauber, Bernd Tomandl, Christopher Nimsy, Guenther Greiner, and Peter Hastreiter. "Automatic adjustment of bidimensional transfer functions for direct volume visualization of intracranial aneurysms." In: *Proc. SPIE Medical Imaging*. Bellingham, Washington: SPIE/IS&T, 2004, pp. 275–284.
- [HP11] Klaus Hildebrandt and Konrad Polthier. "Generalized shape operators on polyhedral surfaces." In: *Computer Aided Geometric Design* 28.5 (2011), pp. 321–343.
- [HPWo5] Klaus Hildebrandt, Konrad Polthier, and Max Wardetzky. "Smooth Feature Lines on Surface Meshes." In: *Proc. of Eurographics Symposium on Geometry Processing*. Vienna, Austria, 2005, pp. 85–90. ISBN: 3-905673-24-X.

- [HHD03] Stefan Hiller, Heino Hellwig, and Oliver Deussen. "Beyond Stippling – Methods for Distributing Objects on the Plane." In: *Computer Graphics Forum* 22.3 (Sept. 2003), pp. 515–522.
- [HKG00] Jiří Hladůvka, Andreas König, and Eduard Gröller. "Curvature-Based Transfer Functions for Direct Volume Rendering." In: *Proc. Spring Conference on Computer Graphics and its Applications*. 2000, pp. 58–65.
- [Hod03] Elaine RS Hodges. *The guild handbook of scientific illustration*. John Wiley & Sons, 2003.
- [HB86] Karl Heinz Höhne and Ralph Bernstein. "Shading 3D-images from CT using gray-level gradients." In: *Medical Imaging, IEEE Transactions on* 5.1 (1986), pp. 45–47.
- [HBTR88] Karl-Heinz Höhne, Michael Bomans, Ulf Tiede, and Martin Riemer. "Display of multiple 3D-objects using the generalized voxel-model." In: *Proc. of Medical Imaging II*. International Society for Optics and Photonics. 1988, pp. 850–854.
- [HBKS05] Helen Hong, Juhee Bae, Heewon Kye, and YeongGil Shin. "Efficient Multimodality Volume Fusion Using Graphics Hardware." English. In: *Computational Science*. Vol. 3516. Lecture Notes in Computer Science. 2005, pp. 842–845.
- [HBV10] Christian Hörr, Guido Brunnett, and Christian Vix. "Line Drawings vs. Curvature Shading: Scientific Illustration of Range Scanned Artefacts." In: *Proc. Computational Aesthetics in Graphics, Visualization and Imaging*. London, United Kingdom: Eurographics Association, 2010, pp. 41–49. ISBN: 978-3-905674-24-8.
- [HS15] Ingrid Hotz and Thomas Schultz. *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Springer, 2015.
- [HCW10] C.-H. Huang, K.-C. Chang, and C. Wen. "Non-iterative stippling of greyscale three-dimensional polygon meshed models." In: *IET Computer Vision* 4.2 (2010), pp. 138–148. ISSN: 1751-9632.
- [HBPE+07] Jingshu Huang, Brian Bue, Avin Pattath, David S. Ebert, and Krystal M. Thomas. "Interactive Illustrative Rendering on Mobile Devices." In: *IEEE Computer Graphics and Applications* 27.3 (2007), pp. 48–56. ISSN: 0272-1716 (print), 1558-1756 (electronic).
- [HGH+10] Mathias Hummel, Christoph Garth, Bernd Hamann, Hans Hagen, and Kenneth I. Joy. "IRIS: Illustrative Rendering for Integral Surfaces." In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1319–1328.
- [IFP97] Victoria Interrante, Henry Fuchs, and Stephen M. Pizer. "Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture." In: *IEEE Transactions on Visualization and Computer Graphics* 3.2 (1997), pp. 98–117.
- [IFP95] Victoria Interrante, Henry Fuchs, and Stephen Pizer. "Enhancing Transparent Skin Surfaces with Ridge and Valley Lines." In: *Proc. of IEEE Visualization*. 1995, pp. 52–59.
- [IFP96] Victoria Interrante, Henry Fuchs, and Stephen Pizer. "Illustrating Transparent Surfaces with Curvature-directed Strokes." In: *Proc. of IEEE Visualization*. 1996, pp. 211–218.
- [IFHS+03] T. Isenberg, B. Freudenberg, N. Halper, S. Schlechtweg, and T. Strothotte. "A developer's guide to silhouette algorithms for polygonal models." In: *IEEE Transactions on Visualization and Computer Graphics* 23.4 (2003), pp. 28–37. ISSN: 0272-1716.
- [Ise13] Tobias Isenberg. "Evaluating and Validating Non-photorealistic and Illustrative Rendering." In: *Image and Video-Based Artistic Stylisation*. Ed. by Paul Rosin and John Collomosse. Vol. 42. Computational Imaging and Vision. London/Heidelberg: Springer, 2013, pp. 311–331.
- [Ise15] Tobias Isenberg. "A Survey of Illustrative Visualization Techniques for Diffusion-Weighted MRI Tractography." In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by Ingrid Hotz and Thomas Schulz. Berlin/Heidelberg: Springer, 2015. Chap. 12, pp. 235–256.
- [Ise16] Tobias Isenberg. "Interactive NPAR: What Type of Tools Should We Create?" In: *Proc. NPAR*. Goslar, Germany: The Eurographics Association, 2016, pp. 89–96.
- [IB06] Tobias Isenberg and Angela Brennecke. "G-Strokes: A Concept for Simplifying Line Stylization." In: *Computers & Graphics* 30.5 (2006), pp. 754–766.
- [IHS02] Tobias Isenberg, Nick Halper, and Thomas Strothotte. "Stylizing Silhouettes at Interactive Rates: From Silhouette Edges to Silhouette Strokes." In: *Computer Graphics Forum* 21.3 (2002), pp. 249–258. ISSN: 1467-8659.
- [IICS+13] Tobias Isenberg, Petra Isenberg, Jian Chen, Michael Sedlmair, and Torsten Möller. "A Systematic Review on the Practice of Evaluating Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (Dec. 2013), pp. 2818–2827.
- [IMS00] Tobias Isenberg, Maic Masuch, and Thomas Strothotte. "3D Illustrative Effects for Animating Line Drawings." In: *Proc. IV*. Los Alamitos: IEEE Computer Society, 2000, pp. 413–418.
- [INCS+06] Tobias Isenberg, Petra Neumann, Sheelagh Carpendale, Mario Costa Sousa, and Joaquim A. Jorge. "Non-Photorealistic Rendering in Context: An Observational Study." In: *Proc. NPAR*. New York: ACM, 2006, pp. 115–126.
- [JBBS+08] W.M. Jaineck, S. Born, D. Bartz, W. Straßer, and J. Fischer. "Illustrative Hybrid Visualization and Exploration of Anatomical and Functional Brain Data." In: *Computer Graphics Forum* 27.3 (2008), pp. 855–862.
- [JD14] Alex Pappachen James and Belur V Dasarathy. "Medical image fusion: a survey of the state of the art." In: *Information Fusion* 19 (2014), pp. 4–19.
- [JFSG+00a] P Jannin, OJ Fleig, E Seigneuret, C Grova, X Morandi, and J Scarabin. "Multimodal and multi-informational neuro-navigation." In: *Proc. of CARS - Computer Assisted Radiology and Surgery* (2000), pp. 167–172.
- [JFSG+00b] Pierre Jannin, Oliver J Fleig, Eric Seigneuret, Christophe Grova, Xavier Morandi, and Jean-Marie Scarabin. "A data fusion environment for multimodal and multi-informational neuronavigation." In: *Computer Aided Surgery* 5.1 (2000), pp. 1–10.
- [JNMS+09] Firdaus Janoos, Boonthanome Nouanesengsy, Raghu Machiraju, Han Wei Shen, Steffen Sammet, Michael Knopp, and István Á Mórocz. "Visual analysis of brain activity from fMRI data." In: *Computer Graphics Forum* 28.3 (2009), pp. 903–910.

- [JNLM05] Kyuman Jeong, Alex Ni, Seungyong Lee, and Lee Markosian. "Detail control in line drawings of 3D meshes." In: *The Visual Computer* 21.8 (2005), pp. 698–706. ISSN: 1432-2315.
- [JL97] Bruno Jobard and Wilfrid Lefer. "Creating Evenly-Spaced Streamlines of Arbitrary Density." In: *Proc. EG Workshop on Visualization in Scientific Computing*. Vienna: Springer, 1997, pp. 43–55. ISBN: 978-3-7091-6876-9.
- [JDD03] Thouis R. Jones, Frédo Durand, and Mathieu Desbrun. "Non-iterative, Feature-preserving Mesh Smoothing." In: *ACM Trans. Graph.* 22.3 (July 2003), pp. 943–949. ISSN: 0730-0301.
- [JQDB+08] A. Joshi, Xiaoning Qian, D.P. Dione, K. Bulsara, C. Breuer, A.J. Sinusas, and X. Papademetris. "Effective visualization of complex vascular structures using a non-parametric vessel detection method." In: *Visualization and Computer Graphics, IEEE Transactions on* 14.6 (2008), pp. 1603–1610. ISSN: 1077-2626.
- [JSVS+08] Alark Joshi, Dustin Scheinost, Kenneth P Vives, Dennis D Spencer, Lawrence H Staib, and Xenophon Papademetris. "Novel interaction techniques for neurosurgical planning and stereotactic navigation." In: *Visualization and Computer Graphics, IEEE Transactions on* 14.6 (2008), pp. 1587–1594.
- [JDA07] Tilke Judd, Frédo Durand, and Edward Adelson. "Apparent Ridges for Line Drawing." In: *ACM Trans. Graph.* 26.3 (2007).
- [JWNC+08] Martin S Judenhofer, Hans F Wehrl, Danny F Newport, Ciprian Catana, Stefan B Siegel, Markus Becker, Axel Thielscher, Manfred Kneilling, Matthias P Lichy, Martin Eichner, et al. "Simultaneous PET-MRI: a new approach for functional and morphological imaging." In: *Nature medicine* 14.4 (2008), pp. 459–465.
- [JKEF+13] Younhyun Jung, Jinman Kim, Stefan Eberl, Micheal Fulham, and David Dagan Feng. "Visibility-driven PET-CT visualisation with region of interest (ROI) segmentation." In: *The Visual Computer* 29.6-8 (2013), pp. 805–815.
- [KDM13] Laith M Kadasi, Walter C Dent, and Adel M Malek. "Cerebral aneurysm wall thickness analysis using intraoperative microscopy: Effect of size and gender on thin translucent regions." In: *Journal of Neurointerventional Surgery* 5.3 (2013), pp. 201–206.
- [KGBH+09] Bernhard Kainz, Markus Grabner, Alexander Bornik, Stefan Hauswiesner, Judith Muehl, and Dieter Schmalstieg. "Ray casting of multiple volumetric datasets with polyhedral boundaries on manycore GPUs." In: *ACM Transactions on Graphics (TOG)* 28.5 (2009), p. 152.
- [Kaj09] Vladimir Kajalin. "Screen-Space Ambient Occlusion." In: *ShaderX⁷*. Ed. by Wolfgang Engel. Boston: Charles River Media, 2009, pp. 413–424.
- [KMS07] Denis Kalkofen, Erick Mendez, and Dieter Schmalstieg. "Interactive Focus and Context Visualization for Augmented Reality." In: IEEE Computer Society, 2007, pp. 1–10.
- [KMS09] Denis Kalkofen, Erick Mendez, and Dieter Schmalstieg. "Comprehensible Visualization for Augmented Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 15.2 (Mar. 2009), pp. 193–204. ISSN: 1077-2626.
- [KDMF03] Robert D. Kalnins, Philip L. Davidson, Lee Markosian, and Adam Finkelstein. "Coherent Stylized Silhouettes." In: *ACM Transactions on Graphics* 22.3 (July 2003), pp. 856–861. ISSN: 0730-0301.
- [KMMK+02] Robert D. Kalnins, Lee Markosian, Barbara J. Meier, Michael A. Kowalski, Joseph C. Lee, Philip L. Davidson, Matthew Webb, John F. Hughes, and Adam Finkelstein. "WYSIWYG NPR: Drawing Strokes Directly on 3D Models." In: *ACM Transactions on Graphics* 21.3 (July 2002), pp. 755–762.
- [KNBH12] Evangelos Kalogerakis, Derek Nowrouzezahrai, Simon Breslav, and Aaron Hertzmann. "Learning Hatching for Pen-and-ink Illustration of Surfaces." In: *ACM Transactions on Graphics* 31.1 (Feb. 2012), 1:1–1:17.
- [KNSM+09] Evangelos Kalogerakis, Derek Nowrouzezahrai, Patricio Simari, James Mccrae, Aaron Hertzmann, and Karan Singh. "Data-driven Curvature for Real-time Line Drawing of Dynamic Scenes." In: *ACM Transactions on Graphics* 28.1 (Feb. 2009), 11:1–11:13. ISSN: 0730-0301.
- [KH11] Kevin Karsch and John C. Hart. "Snaxels on a Plane." In: *Proc. NPAR*. New York: ACM, 2011, pp. 35–42. ISBN: 978-1-4503-0907-3.
- [KWT88] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. "Snakes: Active contour models." In: *International journal of computer vision* 1.4 (1988), pp. 321–331.
- [KSBCO+15] D. Katic, P. Spengler, S. Bodenstedt, G. Castrillon-Oberndorfer, R. Seeberger, J. Hoffmann, R. Dillmann, and S. Speidel. "A system for context-aware intraoperative augmented reality in dental implant surgery." In: *Int J Comput Assist Radiol Surg* 10.1 (2015), pp. 101–108.
- [KMBS+13] Vincent Keereman, Pieter Mollet, Yannick Berker, Volkmar Schulz, and Stefaan Vandenbergh. "Challenges and current methods for attenuation correction in PET/MR." In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 26.1 (2013), pp. 81–98.
- [KWCZ+12] Jens Kerber, Meili Wang, J. Chang, Jian J. Zhang, Alexander G. Belyaev, and Hans-Peter Seidel. "Computer Assisted Relief Generation - A Survey." In: *Computer Graphics Forum* 31.8 (2012), pp. 2363–2377.
- [KOCC13] Marta Kersten-Oertel, Sean Jy-Shyang Chen, and D. Louis Collins. "An Evaluation of Depth Enhancing Perceptual Cues for Vascular Volume Visualization in Neurosurgery." In: *IEEE Trans. Vis. Graph.* 20.3 (2013), pp. 391–403.
- [KCC14] Marta Kersten-Oertel, Sean Jy-Shyang Chen, and D. Louis Collins. "An Evaluation of Depth Enhancing Perceptual Cues for Vascular Volume Visualization in Neurosurgery." In: *IEEE Transactions on Visualization and Computer Graphics* 20.3 (2014), pp. 391–403.
- [KKSS+12] Rostislav Khlebnikov, Bernhard Kainz, Markus Steinberger, Marc Streit, and Dieter Schmalstieg. "Procedural Texture Synthesis for Zoom-Independent Visualization of Multivariate Data." In: *Computer Graphics Forum* 31.3pt4 (2012), pp. 1355–1364.
- [KKLS+16] Dong-Joon Kim, Bohyoung Kim, Jeongjin Lee, Juneseuk Shin, Kyoung Won Kim, and Yeong-Gil Shin. "High-quality slab-based intermixing method for fusion rendering of multiple medical objects." In: *Computer Methods and Programs in Biomedicine* 123 (2016), pp. 27–42.
- [KCL09] Hyun Soo Kim, Han Kyun Choi, and Kwan H. Lee. "Feature Detection of Triangular Meshes Based on Tensor Voting Theory." In: *Comput. Aided Des.* 41.1 (2009), pp. 47–58. ISSN: 0010-4485.

- [KEF07] Jinman Kim, Stefan Eberl, and David Dagan Feng. "Visualizing Dual-Modality Rendered Volumes Using a Dual-Lookup Table Transfer Function." In: *Computing in Science and Engineering* 9.1 (2007), pp. 20–25.
- [KMIA+09] SungYe Kim, Ross Maciejewski, Tobias Isenberg, William M. Andrews, Wei Chen, Mario Costa Sousa, and David S. Ebert. "Stippling By Example." In: *Proc. NPAR*. New York: ACM, 2009, pp. 41–50.
- [KYYL08] Yongjin Kim, Jingyi Yu, Xuan Yu, and Seungyong Lee. "Line-art Illustration of Dynamic and Specular Surfaces." In: *ACM Transactions on Graphics* 27.5 (Dec. 2008), 156:1–156:10. ISSN: 0730-0301.
- [KNST+12] Taichi Kin, Hirofumi Nakatomi, Masaaki Shojima, Minoru Tanaka, Kenji Ino, Harushi Mori, Akira Kunimatsu, Hiroshi Oyama, and Nobuhito Saito. "A new strategic neurosurgical planning tool for brainstem cavernous malformations using interactive computer graphics with multimodal fusion images: Clinical article." In: *Journal of Neurosurgery* 117.1 (2012), pp. 78–88.
- [KWTM03] Gordon Kindlmann, Ross Whitaker, Tolga Tasdizen, and Torsten Moller. "Curvature-Based Transfer Functions for Direct Volume Rendering: Methods and Applications." In: *IEEE Transactions on Visualization and Computer Graphics*. 2003, pp. 513–520. ISBN: 0-7695-2030-8.
- [KGSAY+14] Hortense A Kirişli, Vikas Gupta, Rahil Shahzad, Imad Al Younis, Anoeskha Dharampal, Robert-Jan van Geuns, Arthur J Scholte, Michiel A de Graaf, Raoul MS Joemai, Koen Nieman, et al. "Additional diagnostic value of integrated analysis of cardiac CTA and SPECT MPI using the SMARTVis system in patients with suspected coronary artery disease." In: *Journal of Nuclear Medicine* 55.1 (2014), pp. 50–57.
- [KGLR+15] Paul Klemm, Sylvia Glaßer, Kai Lawonn, Marko Rak, Henry Völzke, Katrin Hegenscheid, and Bernhard Preim. "Interactive Visual Analysis of Lumbar Back Pain." In: *Proc. of the 6th International Conference on Information Visualization Theory and Applications (IVAPP)*. 2015, pp. 85–92.
- [KLG+15] Paul Klemm, Kai Lawonn, Sylvia Glaßer, Uli Niemann, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. "3D Regression Heat Map Analysis of Population Study Data." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 22 (1).1 (2015), pp. 81–90.
- [KOJLH+14a] Paul Klemm, Steffen Oeltze-Jafra, Kai Lawonn, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. "Interactive Visual Analysis of Image-Centric Cohort Study Data." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2014), pp. 1673–1682.
- [KOJLH+14b] Paul Klemm, Steffen Oeltze-Jafra, Kai Lawonn, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. "Interactive visual analysis of image-centric cohort study data." In: *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014), pp. 1673–1682.
- [KMKV+00] Paul G Klutz, Carolyn Cidis Meltzer, Victor L Villemagne, Paul E Kinahan, Subhash Chander, Marsha A Martinelli, and David W Townsend. "Combined PET/CT Imaging in Oncology: Impact on Patient Management." In: *Clinical Positron Imaging* 3.6 (2000), pp. 223–230.
- [KHGR02] Joe Kniss, Charles Hansen, Michel Grenier, and Tom Robinson. "Volume Rendering Multivariate Data to Visualize Meteorological Simulations: A Case Study." In: *Proc. of the EG Visualization Symposium*. Eurographics Association, 2002, pp. 189–194.
- [KKH01] Joe Kniss, Gordon Kindlmann, and Charles Hansen. "Interactive Volume Rendering Using Multi-dimensional Transfer Functions and Direct Manipulation Widgets." In: *Proc. of IEEE Visualization*. 2001, pp. 255–262.
- [KSWW+04] Joe Kniss, Jürgen P. Schulze, Uwe Wössner, Peter Winkler, Ulrich Lang, and Charles D. Hansen. "Medical Applications of Multi-field Volume Rendering and VR Techniques." In: *Proc. of the EG Visualization Symposium*. 350. Eurographics Association, 2004, pp. 249–254.
- [KHG09] Sebastian Knödel, Martin Hachet, and Pascal Guittou. "Interactive Generation and Modification of Cutaway Illustrations for Polygonal Models." In: *Proc. Smart Graphics*. Heidelberg/Berlin: Springer, 2009, pp. 140–151. ISBN: 978-3-642-02114-5.
- [KBPP+15] Benjamin Köhler, Silvia Born, Roy FP van Pelt, Uta Preim, and Bernhard Preim. "A Survey of Cardiac 4D PC-MRI Data Processing." In: *The Eurographics Association*, 2015, pp. 139–148.
- [KBKG08] Peter Kohlmann, Stefan Brückner, Armin Kanitsar, and M Eduard Gröller. "LiveSync++: Enhancements of an interaction metaphor." In: *Proc. of Graphics Interface*. Canadian Information Processing Society. 2008, pp. 81–88.
- [KBKK07] Peter Kohlmann, Stefan Brückner, Armin Kanitsar, and A Kanitsar. "LiveSync: Deformed viewing spheres for knowledge-based navigation." In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1544–1551.
- [KST09] M. Kolomenkin, I. Shimshoni, and A. Tal. "On edge detection on surfaces." In: *Proc. CVPR*. Los Alamitos: IEEE Computer Society, 2009, pp. 2767–2774.
- [KLST13] Michael Kolomenkin, George Leifman, Ilan Shimshoni, and Ayellet Tal. "Reconstruction of Relief Objects from Archeological Line Drawings." In: *Computing and Cultural Heritage* 6.1 (2013), 3:1–3:19. ISSN: 1556-4673.
- [KST08] Michael Kolomenkin, Ilan Shimshoni, and Ayellet Tal. "Demarcating Curves for Shape Illustration." In: *ACM Transactions on Graphics* 27.5 (Dec. 2008), 157:1–157:9. ISSN: 0730-0301.
- [KST11] Michael Kolomenkin, Ilan Shimshoni, and Ayellet Tal. "Prominent Field for Shape Processing and Analysis of Archaeological Artifacts." English. In: *International Journal of Computer Vision* 94.1 (2011), pp. 89–100. ISSN: 0920-5691.
- [KCODL06] Johannes Kopf, Daniel Cohen-Or, Oliver Deussen, and Dani Lischinski. "Recursive Wang Tiles for Real-Time Blue Noise." In: *ACM Transactions on Graphics* 25.3 (July 2006), pp. 509–518.
- [KMH01] Robert Kosara, Silvia Miksch, and Helwig Hauser. "Semantic depth of field." In: *Proc. of the IEEE Symposium on Information Visualization (InfoVis)*. IEEE, 2001, pp. 97–104.
- [KSF11] Bernhard Krauss, Bernhard Schmidt, and Thomas G Flohr. "Dual source CT." In: *Dual energy CT in clinical practice*. Springer, 2011, pp. 11–20.
- [KK99] Kevin Kreeger and Arie Kaufman. "Mixing translucent polygons with volumes." In: *Proc. of IEEE Visualization*. IEEE Computer Society Press. 1999, pp. 191–198.

- [KHKD+08] R. Krempien, H. Hoppe, L. Kahrs, S. Daeuber, O. Schorr, G. Eggers, M. Bischof, M. Munter, J. Debus, and W. Harms. "Projector-Based Augmented Reality for Intuitive Intraoperative Guidance in Image-Guided 3D Interstitial Brachytherapy." In: *International Journal of Radiation Oncology*Biophysics* 70.3 (2008), pp. 944–952. ISSN: 0360-3016.
- [KSW06] Jens Krüger, Jens Schneider, and Rüdiger Westermann. "Clearview: An interactive context preserving hotspot visualization technique." In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 941–948.
- [KW07] Jens Krüger and Rüdiger Westermann. "Efficient Stipple Rendering." In: *Proc. IADIS Computer Graphics and Visualization*. 2007.
- [KSIF10] E. Kruijff, J. E. Swan II, and S. Feiner. "Perceptual Issues in Augmented Reality Revisited." In: *Proc. of IEEE International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, 2010, pp. 3–12.
- [KTP10] Christoph Kubisch, Christian Tietjen, and Bernhard Preim. "GPU-based Smart Visibility Techniques for Tumor Surgery Planning." In: *International Journal of Computer Assisted Radiology and Surgery* (2010), pp. 667–678.
- [Küh06] W. Kühnel. *Differential Geometry: Curves - Surfaces - Manifolds*. Student mathematical library. American Mathematical Society, 2006. ISBN: 9780821839881.
- [KGMD+10] Anja Kuß, Maria Gensel, Björn Meyer, Vincent J. Dercksen, and Steffen Prohaska. "Effective Techniques to Visualize Filament-Surface Relationships." In: *Computer Graphics Forum* 29.3 (2010), pp. 1003–1012.
- [KYM12] Yunmi Kwon, Heekyung Yang, and Kyungha Min. "Pencil rendering on 3D meshes using convolution." In: *Computers & Graphics* 36.8 (2012), pp. 930–944. ISSN: 0097-8493.
- [KCWI13] Jan Eric Kyprianidis, John Collomosse, Tinghui Wang, and Tobias Isenberg. "State of the "Art": A Taxonomy of Artistic Stylization Techniques for Images and Video." In: *IEEE Transactions on Visualization and Computer Graphics* 19.5 (May 2013), pp. 866–885.
- [LMHB00] Adam Lake, Carl Marshall, Mark Harris, and Marc Blackstein. "Stylized Rendering Techniques for Scalable Real-time 3D Animation." In: *Proc. NPAR*. Annecy, France: ACM, 2000, pp. 13–20. ISBN: 1-58113-277-8.
- [Law15] Kai Lawonn. "Illustrative visualization of medical data sets." PhD thesis. University of Magdeburg, 2015.
- [LBS14] Kai Lawonn, Alexandra Baer, Patrick Saalfeld, and Bernhard Preim. "Comparative Evaluation of Feature Line Techniques for Shape Depiction." In: *Proc. Vision, Modeling and Visualization*. 2014, pp. 31–38. DOI: 10.2312/vmv.20141273.
- [LGP14a] Kai Lawonn, Rocco Gasteiger, and Bernhard Preim. "Adaptive Surface Visualization of Vessels with Animated Blood Flow." In: *Computer Graphics Forum* 33(8) (2014), pp. 16–27.
- [LGP14b] Kai Lawonn, Sylvia Glaßer, Anna Vilanova, Bernhard Preim, and Tobias Isenberg. "Occlusion-free Blood Flow Animation with Wall Thickness Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 22 (1).1 (2015), pp. 728–737.
- [LGP14b] Kai Lawonn, Tobias Günther, and Bernhard Preim. "Coherent View-Dependent Streamlines for Understanding Blood Flow." In: *EuroVis - Short Papers*. 2014, 5 pages.
- [LKEP14] Kai Lawonn, Michael Krone, Thomas Ertl, and Bernhard Preim. "Line Integral Convolution for Real-Time Illustration of Molecular Surface Shape and Salient Regions." In: *Computer Graphics Forum* 33.3 (2014), pp. 181–190.
- [LLH17] Kai Lawonn, Maria Luz, and Christian Hansen. "Improving Spatial Perception of Vascular Models using Supporting Anchors and Illustrative Visualization." In: *Computers & Graphics* (2017), in print.
- [LLPH15] Kai Lawonn, Maria Luz, Bernhard Preim, and Christian Hansen. "Illustrative Visualization of Vascular Models for Static 2D Representations." In: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2015, pp. 399–406.
- [LMP13] Kai Lawonn, Tobias Mönch, and Bernhard Preim. "Streamlines for Illustrative Real-time Rendering." In: *Computer Graphics Forum* 32.3 (2013), pp. 321–330.
- [LP16] Kai Lawonn and Bernhard Preim. "Feature Lines for Illustrating Medical Surface Models: Mathematical Background and Survey." In: *Visualization in Medicine in Life Sciences III*. Ed. by Lars Linsen, Bernd Hamann, and Hans-Christian Hege. Berlin/Heidelberg: Springer, 2016, pp. 93–132.
- [LSP14] Kai Lawonn, Patrick Saalfeld, and Bernhard Preim. "Illustrative Visualization of Endoscopic Views." In: *Bildverarbeitung für die Medizin (BVM)*. 2014, pp. 276–281.
- [LSBP17] Kai Lawonn, Noeska Smit, Katja Bühler, and Bernhard Preim. "A Survey on Multimodal Medical Data Visualization." In: *Computer Graphics Forum* (2017), accepted with minor revision.
- [LSPV15] Kai Lawonn, Noeska Smit, Bernhard Preim, and Anna Vilanova. "Illustrative Multi-volume Rendering for PET/CT Scans." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2015, pp. 103–112.
- [LTPH17] Kai Lawonn, Erik Trostmann, Bernhard Preim, and Klaus Hildebrandt. "Visualization and Extraction of Carvings for Heritage Conservation." In: *IEEE Transactions on Visualization and Computer Graphics* 23(1) (2017), pp. 801–810.
- [LVP17] Kai Lawonn, Ivan Viola, Bernhard Preim, and Tobias Isenberg. "A Survey of Surface-Based Illustrative Rendering for Visualization." In: *Computer Graphics Forum* (2017), submitted.
- [LFH10] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. "Interviews and Focus Groups." In: *Research Methods in Human-Computer Interaction*. John Wiley & Sons Ltd, 2010. Chap. 8, pp. 177–216.
- [LBMPC+01] Denis Le Bihan, Jean-François Mangin, Cyril Poupon, Chris A Clark, Sabina Pappata, Nicolas Molko, and Hughes Chabriet. "Diffusion tensor imaging: concepts and applications." In: *Journal of magnetic resonance imaging* 13.4 (2001), pp. 534–546.
- [LHV06] Chang Ha Lee, Xuejun Hao, and Amitabh Varshney. "Geometry-Dependent Lighting." In: *IEEE Transactions on Visualization and Computer Graphics* 12.2 (2006), pp. 197–207.

- [LKL06] Hyunjun Lee, Sungtae Kwon, and Seungyong Lee. "Real-time Pencil Rendering." In: *Proc. NPAR*. Annecy, France: ACM, 2006, pp. 37–45. ISBN: 1-59593-357-3.
- [LMLH07] Yunjin Lee, Lee Markosian, Seungyong Lee, and John F. Hughes. "Line Drawings via Abstracted Shading." In: *ACM Transactions on Graphics* 26.3 (July 2007), 18:1–18:5. ISSN: 0730-0301.
- [LUS13] Zoltán Lengyel, Tamás Umenhoffer, and László Szécsi. "Screen space features for real-time hatching synthesis." In: *Képfeldolgozók és Alakfelismerők Társaságának*. 2013, pp. 82–94.
- [LUS14] Zoltán Lengyel, Tamás Umenhoffer, and László Szécsi. "Realtime, coherent screen space hatching." In: *Proc. Computer Graphics and Geometry*. 2014, pp. 131–137.
- [LCMY07] Mirna Lerotic, AdrianJ. Chung, George Mylonas, and Guang-Zhong Yang. "PQ-space Based Non-Photorealistic Rendering for Augmented Reality." In: *MICCAI*. Vol. 4792. Lecture Notes in Computer Science. 2007, pp. 102–109.
- [LHTG89] David N Levin, XP Hu, Kim Kiat Tan, and Simranjit Galhotra. "Surface of the brain: three-dimensional MR images created with volume rendering." In: *Radiology* 171.1 (1989), pp. 277–280.
- [Lev88] Marc Levoy. "Display of Surfaces from Volume Data." In: *IEEE Computer Graphics Applications* 8.3 (1988), pp. 29–37.
- [LPRM02] Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. "Least Squares Conformal Maps for Automatic Texture Atlas Generation." In: *ACM Transactions on Graphics* 21.3 (July 2002), pp. 362–371. ISSN: 0730-0301.
- [LZLX+12] Lei Li, Yongqiang Zhou, Chunxiao Liu, Yingying Xu, and Jie Fu. "State-of-the-Art Line Drawing Techniques." In: *Computer, Informatics, Cybernetics and Applications: Proceedings of the CICA 2011*. Ed. by Xingui He, Ertian Hua, Yun Lin, and Xiaozhu Liu. Springer Netherlands, 2012, pp. 1249–1257. ISBN: 978-94-007-1839-5.
- [LACS08] Wilmot Li, Maneesh Agrawala, Brian Curless, and David Salesin. "Automated generation of interactive 3D exploded view diagrams." In: *ACM Transactions on Graphics (TOG)*. Vol. 27. 3. ACM. 2008, p. 101.
- [LAS04] Wilmot Li, Maneesh Agrawala, and David Salesin. "Interactive Image-based Exploded View Diagrams." In: *Proc. of Graphics Interface*. 2004, pp. 203–212.
- [LRAC+07] Wilmot Li, Lincoln Ritter, Maneesh Agrawala, Brian Curless, and David Salesin. "Interactive Cutaway Illustrations of Complex 3D Models." In: *ACM Transactions on Graphics* 26.3 (2007), 31:1–31:11.
- [LCKM05] Ronghua Liang, Gordon Clapworthy, Meleagros A. Krokos, and R. Mayoral. "Real-Time Predefined Shape Cutaway with Parametric Boundaries." In: *Proc. Computer Graphics, Imaging and Vision*. 2005, pp. 227–231.
- [LCC+01] Ping-Sung Liao, Tse-Sheng Chen, Pau-Choo Chung, et al. "A fast algorithm for multilevel thresholding." In: *J. Inf. Sci. Eng.* 17.5 (2001), pp. 713–727.
- [LSHL16] Nils Lichtenberg, Noeska Smit, Christian Hansen, and Kai Lawonn. "Slime: Seamless Line Illustration for Interactive Biomedical Visualization." In: *Proc. Visual Computing for Biology and Medicine*. 2016, to appear.
- [LHV12] Endre M. Lidal, Helwig Hauser, and Ivan Viola. "Design Principles for Cutaway Visualization of Geological Models." In: *Proc. of the Spring Conference on Computer Graphics*. 2012, pp. 47–54.
- [LHGT+06] G Sharat Lin, Horace H Hines, Genine Grant, Kimberly Taylor, and Carl Ryals. "Automated quantification of myocardial ischemia and wall motion defects by use of cardiac SPECT polar mapping and 4-dimensional surface rendering." In: *Journal of nuclear medicine technology* 34.1 (2006), pp. 3–17.
- [LFSB+15] Stefan Lindholm, Martin Falk, Erik Sundén, Alexander Bock, Anders Ynnerman, and Timo Ropinski. "Hybrid data visualization based on depth complexity histogram analysis." In: *Computer Graphics Forum* 34.1 (2015), pp. 74–85.
- [LLHY09] Stefan Lindholm, Patric Ljung, Markus Hadwiger, and Anders Ynnerman. "Fused Multi-Volume DVR using Binary Space Partitioning." In: 28.3 (2009), pp. 847–854.
- [LVLR08] Lars Linsen, Tran Van Long, Paul Rosenthal, and Stephan Rosswog. "Surface extraction from multi-field particle volume data using multi-dimensional cluster visualization." In: *Visualization and Computer Graphics, IEEE Transactions on* 14.6 (2008), pp. 1483–1490.
- [LML16] Xianyong Liu, Lizhuang Ma, and Yanping Liu. "Global Tone: using tone to draw in Pen-and-Ink illustration." In: *Multimedia Tools and Applications* (2016), pp. 1–17. ISSN: 1573-7721.
- [LC87] William E Lorensen and Harvey E Cline. "Marching cubes: A high resolution 3D surface construction algorithm." In: *ACM siggraph computer graphics*. Vol. 21. 4. ACM. 1987, pp. 163–169.
- [LWM15] Liming Lou, Lu Wang, and Xiangxu Meng. "Stylized strokes for coherent line drawings." In: *Computational Visual Media* 1.1 (2015), pp. 79–89. ISSN: 2096-0662.
- [LMTE+03] Aidong Lu, C. J. Morris, J. Taylor, D. S. Ebert, C. Hansen, P. Rheingans, and M. Hartner. "Illustrative interactive stipple rendering." In: *IEEE Transactions on Visualization and Computer Graphics* 9.2 (2003), pp. 127–138. ISSN: 1077-2626.
- [LTHE+02] Aidong Lu, Joe Taylor, Mark Hartner, David S. Ebert, and Charles D. Hansen. "Hardware Accelerated Interactive Stipple Drawing of Polygonal Objects." In: *Proc. the Vision, Modeling, and Visualization*. 2002, pp. 61–68.
- [LMo2a] Eric B. Lum and Kwan-Liu Ma. "Hardware-accelerated Parallel Non-photorealistic Volume Rendering." In: *Proc. of Non-photorealistic Animation and Rendering*. 2002, pp. 67–74.
- [LMo2b] Eric B. Lum and Kwan-Liu Ma. "Interactivity is the Key to Expressive Visualization." In: *ACM SIGGRAPH Computer Graphics* 36.3 (Aug. 2002), pp. 5–9.
- [LLZ11] Tao Luo, Renju Li, and Hongbin Zha. "3D Line Drawing for Archaeological Illustration." In: *International Journal of Computer Vision* 94.1 (2011), pp. 23–35. ISSN: 1573-1405.
- [LF09] Christopher Lux and Bernd Fröhlich. "GPU-based ray casting of multiple multi-resolution volume datasets." In: *Proc. of Advances in Visual Computing*. Springer, 2009, pp. 104–116.

- [Mac49] R. MacNeal. *The Solution of Partial Differential Equations by Means of Electrical Networks*. PhD thesis. California Institute of Technology, 1949.
- [MWCE09] Ross Maciejewski, Insoo Woo, Wei Chen, and David S Ebert. "Structuring feature space: A non-parametric method for volumetric transfer function generation." In: *Visualization and Computer Graphics, IEEE Transactions on* 15.6 (2009), pp. 1473–1480.
- [MV98] JB Antoine Maintz and Max A Viergever. "A survey of medical image registration." In: *Medical image analysis* 2.1 (1998), pp. 1–36.
- [MFO+02] Isabel H Manssour, Sergio S Furuie, Silvia D Olabarriaga, et al. "Visualizing inner structures in multimodal volume data." In: *Proc. of Computer Graphics and Image Processing*. IEEE. 2002, pp. 51–58.
- [MRAM+04] J. Marescaux, F. Rubino, M. Arenas, D. Mutter, and L. Soler. "Augmented-Reality-Assisted Laparoscopic Adrenalectomy." In: *JAMA: The Journal of the American Medical Association* 292.18 (2004), pp. 2214–2215.
- [MBKK+10] Giuliano Mariani, Laura Bruselli, Torsten Kuwert, Edmund E Kim, Albert Flotats, Ora Israel, Maurizio Dondi, and Naoyuki Watanabe. "A review on the clinical uses of SPECT/CT." In: *European journal of nuclear medicine and molecular imaging* 37.10 (2010), pp. 1959–1985.
- [MFKH+12] Michael Markl, Alex Frydrychowicz, Sebastian Kozerke, Mike Hope, and Oliver Wieben. "4D flow MRI." In: *Journal of Magnetic Resonance Imaging* 36.5 (2012), pp. 1015–1036. ISSN: 1522-2586.
- [MKGT+97] Lee Markosian, Michael A. Kowalski, Daniel Goldstein, Samuel J. Trychin, John F. Hughes, and Lubomir D. Bourdev. "Real-time Nonphotorealistic Rendering." In: *Proc. SIGGRAPH*. New York: ACM, 1997, pp. 415–420. ISBN: 0-89791-896-7.
- [Mar76] D. Marr. "Early Processing of Visual Information." In: *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 275.942 (1976), pp. 483–519. ISSN: 0080-4622.
- [MH80] D. Marr and E. Hildreth. "Theory of Edge Detection." In: *Proceedings of the Royal Society of London B: Biological Sciences* 207.1167 (1980), pp. 187–217. ISSN: 0080-4649. eprint: <http://rspb.royalsocietypublishing.org/content/207/1167/187.full.pdf>.
- [MAL11] Domingo Martín, Germán Arroyo, M. Victoria Luzón, and Tobias Isenberg. "Scale-Dependent and Example-Based Stippling." In: *Computers & Graphics* 35.1 (2011), pp. 160–174.
- [MBNL+12] L Massoptier, W Birkfellner, U Nestle, A Laprie, U Sabatini, K Bühler, and A Freudenthal. "Multimodal morphologic and functional images for radiotherapy by the research training network called SUMMER." In: *Proc. of IEEE International Symposium on Biomedical Imaging (ISIB)* (2012).
- [MS98] M. Masuch and T. Strothotte. "Visualising ancient architecture using animated line drawings." In: *Proc. Information Visualization*. 1998, pp. 261–266.
- [MTHG03] Oliver Mattausch, Thomas Theußl, Helwig Hauser, and Eduard Gröller. "Strategies for Interactive Exploration of 3D Flow Using Evenly-Spaced Illuminated Streamlines." In: *Proc. SCCG*. New York: ACM, 2003, pp. 230–241. ISBN: 158113861X.
- [MH04] Morgan McGuire and John F. Hughes. "Hardware-determined Feature Edges." In: *Proc. NPAR*. Annecy, France: ACM, 2004, pp. 35–47. ISBN: 1-58113-887-3.
- [MSVF09] Jonatas Medeiros, Mario Sousa, Luiz Velho, and Carla Maria Dal Sasso Freitas. "Perspective Contouring in Illustrative Visualization." In: *Proc. Computer Graphics and Image Processing*. 2009, pp. 48–55.
- [MSEN+06] Dorit Merhof, Markus Sonntag, Frank Enders, Christopher Nimsy, Peter Hastreiter, and Guenther Greiner. "Hybrid visualization for white matter tracts using triangle strips and point sprites." In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 1181–1188.
- [MPFS03] Oscar Meruvia Pastor, Bert Freudenberg, and Thomas Strothotte. "Real-Time Animated Stippling." In: *IEEE Computer Graphics and Applications* 23.4 (July 2003), pp. 62–68. ISSN: 0272-1716.
- [MPS02] Oscar Meruvia Pastor and Thomas Strothotte. "Frame-Coherent Stippling." In: *Proc. Eurographics Short*. 2002.
- [MPS04] Oscar Meruvia Pastor and Thomas Strothotte. "Graph-based point relaxation for 3D stippling." In: *Proc. Mexican International Conference in Computer Science*. 2004, pp. 141–150.
- [MEBP+17] Monique Meuschke, Wito Engelke, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Automatic Viewpoint Selection for Exploration of Time-dependent Cerebral Aneurysm Data." In: *Bildverarbeitung für die Medizin (BVM)*. 2017, in print.
- [MKPP+16] Monique Meuschke, Benjamin Köhler, Uta Preim, Bernhard Preim, and Kai Lawonn. "Semi-automatic Vortex Flow Classification in 4D PC-MRI Data of the Aorta." In: *Computer Graphics Forum* 35.3 (2016), pp. 351–360.
- [MLKP+16] Monique Meuschke, Kai Lawonn, Benjamin Köhler, Uta Preim, and Bernhard Preim. "Clustering of Aortic Vortex Flow in Cardiac 4D PC-MRI Data." In: *Bildverarbeitung für die Medizin (BVM)*. 2016, pp. 182–187.
- [MVBP+17a] Monique Meuschke, Samuel Voss, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Combined Visualization of Vessel Deformation and Hemodynamics in Cerebral Aneurysms." In: *IEEE Transactions on Visualization and Computer Graphics (Proceedings of the Scientific Visualization 2016)* 23(1) (2017), pp. 761–770.
- [MVBP+17b] Monique Meuschke, Samuel Voss, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Glyph-based Comparative Stress Tensor Visualization in Cerebral Aneurysms." In: *Computer Graphics Forum* (2017), conditionally accepted.
- [MDSB02] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr. "Discrete Differential-Geometry Operators for Triangulated 2-Manifolds." In: *Proc. VisMath*. 2002, pp. 35–57.
- [MFWP12] Yong-Wei Miao, Jie-Qing Feng, Jin-Rong Wang, and Renato Pajarola. "A Multi-Channel Saliency Based Detail Exaggeration Technique for 3D Relief Surfaces." English. In: *Journal of Computer Science and Technology* 27.6 (2012), pp. 1100–1109. ISSN: 1000-9000.
- [Min13] Kyungha Min. "Drawing Features of 3D Meshes in Pencil-drawing Style." In: *Software Engineering and its Applications* 7.6 (2013), pp. 359–366.

- [Min15] Kyungha Min. "Feature-guided convolution for salient rendering of 3D meshes." In: *Engineering Systems Modelling and Simulation* 7.1 (2015), pp. 1–5.
- [MBC02] Jason L. Mitchell, Chris Brennan, and Drew Card. "Real-time Image-space Outlining for Non-photorealistic Rendering." In: *Proc. SIGGRAPH*. San Antonio, Texas: ACM, 2002, pp. 239–239. ISBN: 1-58113-525-4.
- [MBKD+09] J. H. Moltz et al. "Advanced Segmentation Techniques for Lung Nodules, Liver Metastases, and Enlarged Lymph Nodes in CT Scans." In: *IEEE Journal of Selected Topics in Signal Processing* 3.1 (2009), pp. 122–134.
- [MLKW+13] Tobias Mönch, Kai Lawonn, Christoph Kubisch, Rüdiger Westermann, and Bernhard Preim. "Interactive Mesh Smoothing for Medical Applications." In: *Computer Graphics Forum* 32(8) (Dec. 2013), pp. 110–121.
- [MNP11] Tobias Mönch, Mathias Neugebauer, and Bernhard Preim. "Optimization of Vascular Surface Models for Computational Fluid Dynamics and Rapid Prototyping." In: *Proc. International Workshop on Digital Engineering*. 2011, pp. 16–23.
- [Mor97] H. P. Moravec. "Toward automatic visual obstacle avoidance." In: *Proc. of Conf. on Artificial Intelligence*. 1997, pp. 584–94.
- [MMBU92] E. Mortensen, B. Morse, W. Barrett, and J. Udupa. "Adaptive boundary detection using 'live-wire' two-dimensional dynamic programming." In: *Proc. of Computers in Cardiology*. 1992, pp. 635–638.
- [MPF10] F. de Moura Pinto and C.M.D.S. Freitas. "Importance-Aware Composition for Illustrative Volume Rendering." In: *Proc. of Conference on Graphics, Patterns and Images (SIBGRAPI)*. 2010, pp. 134–141.
- [MNP07] Konrad Mühler, Mathias Neugebauer, Christian Tietjen, and Bernhard Preim. "Viewpoint Selection for Intervention Planning." In: *Proc. of EuroVis*. 2007, pp. 267–274.
- [MPSV14] Mathieu Le Muzic, Julius Parulek, Anne-Kristin Stavrum, and Ivan Viola. "Illustrative Visualization of Molecular Reactions using Omniscient Intelligence and Passive Agents." In: *Computer Graphics Forum* 33.3 (2014), pp. 141–150.
- [MWPV15] Mathieu Le Muzic, Manuela Waldner, Julius Parulek, and Ivan Viola. "Illustrative Timelapse: A Technique for Illustrative Visualization of Particle Simulations on the Mesoscale Level." In: *IEEE Pacific Visualization*. Zijiang Campus, Zhejiang University, Hangzhou, China, 2015, pp. 247–254.
- [NTSW+10] Nikhil V Navkar, Nikolaos V Tsekos, Jason R Stafford, Jeffrey S Weinberg, and Zhigang Deng. "Visualization and planning of neurosurgical interventions with straight access." In: *Proc. of Information Processing in Computer-Assisted Interventions*. Springer, 2010, pp. 1–11.
- [NMWW+04] André Neubauer, Lukas Mroz, Stefan Wolfsberger, Rainer Wegenkittl, Marie Thérèse Forster, and Katja Bühler. "STEPS—an application for simulation of transsphenoidal endonasal pituitary surgery." In: *Proc. of IEEE Visualization*. IEEE. 2004, pp. 513–520.
- [NGBD+09] Mathias Neugebauer, Rocco Gasteiger, Oliver Beuing, Volker Diehl, Martin Skalej, and Bernhard Preim. "Map Displays for the Analysis of Scalar Data on Cerebral Aneurysm Surfaces." In: *Computer Graphics Forum* 28.3 (2009), pp. 895–902.
- [NLB+13] Mathias Neugebauer, Kai Lawonn, Philipp Berg, Gabor Janiga, Oliver Beuing, and Bernhard Preim. "AmniVis – A System for Qualitative Exploration of Near-Wall Hemodynamics in Cerebral Aneurysms." In: *Computer Graphics Forum* 32.3 (2013), pp. 251–260.
- [NIC07] Petra Neumann, Tobias Isenberg, and Sheelagh Cpendale. "NPR Lenses: Interactive Tools for Non-Photorealistic Line Drawings." In: *Proc. Smart Graphics*. Vol. 4569. Lecture Notes in Computer Science. Berlin/Heidelberg: Springer, 2007, pp. 10–22.
- [NOEH+10] Tan Khoa Nguyen, Henrik Ohlsson, Anders Eklund, Frida Hernell, Patric Ljung, Camilla Forsell, Mats Andersson, Hans Knutsson, and Anders Ynnerman. "Concurrent volume visualization of real-time fMRI." In: *Proc. of IEEE/EG International Symposium on Volume Graphics*. Eurographics-European Association for Computer Graphics. 2010, pp. 53–60.
- [NJLM06] Alex Ni, Kyuman Jeong, Seungyong Lee, and Lee Markosian. "Multi-scale Line Drawings from 3D Meshes." In: *Proc. Interactive 3D Graphics and Games*. Redwood City, California: ACM, 2006, pp. 133–137. ISBN: 1-59593-295-X.
- [ND05] Marc Nienhaus and Jürgen Döllner. "Blueprint Rendering and Sketchy Drawings." In: *GPU Gems* 2. Ed. by Matt Pharr. Addison-Wesley, 2005, pp. 235–252.
- [NM00] J. D. Northrup and Lee Markosian. "Artistic Silhouettes: A Hybrid Approach." In: *Proc. NPAR*. Annecy, France: ACM, 2000, pp. 31–37. ISBN: 1-58113-277-8.
- [NRSK+14] Miguel Nunes, Benjamin Rowland, Matthias Schlachter, Soleakhena Ken, Kresimir Matkovic, Anne Laprie, and Katja Bühler. "An integrated visual analysis system for fusing MR spectroscopy and multi-modal radiology imaging." In: *Proc. of Visual Analytics Science and Technology (VAST), IEEE Conference on*. IEEE. 2014, pp. 53–62.
- [Osi] OSIRIX. (last accessed: 09/06/2015).
- [OJCP15] Steffen Oeltze-Jafra, Juan Cebral, Gabor Janiga, and Bernhard Preim. "Cluster Analysis of Vortical Flow in Simulations of Cerebral Aneurysm Hemodynamics." In: *IEEE Transactions on Visualization and Computer Graphics* 22 (1).1 (2015), pp. 757–766.
- [OJMNS+17] Steffen Oeltze-Jafra, Monique Meuschke, Mathias Neugebauer, Sylvia Saalfeld, Kai Lawonn, Gabor Janiga, Hans-Christian Hege, Stefan Zachow, and Bernhard Preim. "Generation and Visual Exploration of Medical Flow Data: Survey, Research Trends, and Future Challenges." In: *Computer Graphics Forum* (2017), submitted.
- [OKGH+06] Steffen Oeltze, Anja Kuß, Frank Grothues, Anja Hennemuth, and Bernhard Preim. "Integrated Visualization of Morphologic and Perfusion Data for the Analysis of Coronary Artery Disease." In: *Proc. of EuroVis*. 2006, pp. 131–138.
- [OP05] Steffen Oeltze and Bernhard Preim. "Visualization of vasculature with convolution surfaces: method, validation and evaluation." In: *IEEE Trans. Med. Imaging* 24.4 (2005), pp. 540–548.
- [OBS04] Yutaka Ohtake, Alexander Belyaev, and Hans-Peter Seidel. "Ridge-Valley Lines on Meshes via Implicit Surface Fitting." In: *ACM Transactions on Graphics* 23.3 (Aug. 2004), pp. 609–612.

- [Ots79] Nobuyuki Otsu. "An automatic threshold selection method based on discriminate and least squares criteria." In: *Denshi Tsushin Gakkai Ronbunshi* 63 (1979), pp. 349–356.
- [PKSP+01] D. L. Page, A. Koschan, Y. Sun, J. Paik, and M. A. Abidi. "Robust Crease Detection and Curvature Estimation of Piecewise Smooth Surfaces from Triangle Mesh Approximations Using Normal Voting." In: *CVPR*. 2001, pp. 162–167.
- [PCGC+11] B. Pan, X. Chen, X. Guo, W. Chen, and Q. Peng. "Interactive Expressive Illustration of 3D City Scenes." In: *Proc. Computer-Aided Design and Computer Graphics*. 2011, pp. 406–410.
- [Pati16] Giuseppe Patané. "Laplacian Spectral Kernels and Distances for Geometry Processing and Shape Analysis." In: *Computer Graphics Forum* 35.2 (2016), pp. 599–624. ISSN: 1467-8659.
- [PVWo8] R.F.P. van Pelt, A. Vilanova, and H.M.M. van de Wetering. "GPU-based Particle Systems for Illustrative Volume Rendering." In: *Proc. of IEEE/ EG Symposium on Volume and Point-Based Graphics*. The Eurographics Association, 2008.
- [PBBC+10] Roy van Pelt, Javier Olivan Bescos, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, Bart ter Haar Romenij, and Anna Vilanova. "Exploration of 4D MRI Blood Flow using Stylistic Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1339–1347. ISSN: 1077-2626.
- [PBBC+11] Roy van Pelt, Javier Oliván Bescós, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, Bart M. ter Haar Romenij, and Anna Vilanova. "Interactive Virtual Probing of 4D MRI Blood-Flow." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2153–2162.
- [Pho75] Bui Tuong Phong. "Illumination for Computer Generated Pictures." In: *Commun. ACM* 18.6 (1975), pp. 311–317. ISSN: 0001-0782.
- [PKNS10] Bernd J Pichler, Armin Kolb, Thomas Nägele, and Heinz-Peter Schlemmer. "PET/MRI: paving the way for the next generation of clinical multimodality imaging applications." In: *Journal of Nuclear Medicine* 51.3 (2010), pp. 333–336.
- [PPCP13] Cyprien Pindat, Emmanuel Pietriga, Oliver Chapuis, and Claude Puech. "Drilling into Complex 3D Models with Gimlenses." In: *Proc. VRST*. Singapore: ACM, 2013, pp. 223–230. ISBN: 978-1-4503-2379-6.
- [PMV03] Josien PW Pluim, JB Antoine Maintz, and Max A Viergever. "Mutual-information-based registration of medical images: a survey." In: *Medical Imaging, IEEE Transactions on* 22.8 (2003), pp. 986–1004.
- [PVG98] Cle Pohl and John L Van Genderen. "Review article multisensor image fusion in remote sensing: concepts, methods and applications." In: *International journal of remote sensing* 19.5 (1998), pp. 823–854.
- [PDBG+01] Mihai Pop, Christian Duncan, Gill Barequet, Michael Goodrich, Wenjing Huang, and Subodh Kumar. "Efficient Perspective-accurate Silhouette Computation and Applications." In: *Proc. Computational Geometry*. New York: ACM, 2001, pp. 60–68.
- [Por94] Ian R. Porteous. *Geometric Differentiation for the intelligence of curves and surfaces*. Cambridge University Press, 1994.
- [PFH00] Emil Praun, Adam Finkelstein, and Hugues Hoppe. "Lapped Textures." In: *Proc. SIGGRAPH*. New York: ACM, 2000, pp. 465–470. ISBN: 1-58113-208-5.
- [PHWF01] Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. "Real-time Hatching." In: *Proc. SIGGRAPH*. New York: ACM, 2001, pp. 581–586. ISBN: 1-58113-374-X.
- [PBCI+16] Bernhard Preim, Alexandra Baer, Douglas Cunningham, Tobias Isenberg, and Timo Ropinski. "A Survey of Perceptually Motivated 3D Visualization of Medical Image Data." In: *Computer Graphics Forum* 22.5 (June 2016), pp. 501–525. ISSN: 1467-8659.
- [PB13] Bernhard Preim and Charl P Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. Morgan Kaufmann Publishers Inc., 2013.
- [PO07] Bernhard Preim and Steffen Oeltze. "3D Visualization of Vasculature: An Overview." In: ed. by Lars Linsen and Hans Hagen. Springer Verlag, 2007. Chap. Visualization in Medicine and Life Sciences, pp. 19–39.
- [PSSO+00] Bernhard Preim, Dirk Selle, Wolf Spindler, KarlJ. Oldhafer, and Heinz-Otto Peitgen. "Interaction Techniques and Vessel Analysis for Preoperative Planning in Liver Surgery." English. In: *Medical Image Computing and Computer-Assisted Intervention 2000*. Ed. by ScottL. Delp, AnthonyM. DiGoia, and Branislav Jaramaz. Vol. 1935. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2000, pp. 608–617. ISBN: 978-3-540-41189-5.
- [RSBS+13] Jens Raab, Henry Schäfer, Alexander Brost, Marc Stamminger, and Marcus Pfister. "Non-photorealistic rendering for minimally invasive procedures." In: *Proc. SPIE Medical Imaging: Image-Guided Procedures, Robotic Interventions, and Modeling*. Vol. 8671. Lake Buena Vista, Florida, USA, 2013.
- [RCL16] Roberto Ranon, Marc Christie, and Christophe Lino. "Algorithms and Techniques for Virtual Camera Control." In: ed. by Augusto Sousa and Kadi Bouatouch. Lisbon, Portugal: Eurographics Association, 2016, t5. ISBN: -.
- [Ras01] Ramesh Raskar. "Hardware Support for Non-photorealistic Rendering." In: *Proc. Workshop on Graphics Hardware*. New York: ACM, 2001, pp. 41–47. ISBN: 1-58113-407-X.
- [RC99] Ramesh Raskar and Michael Cohen. "Image Precision Silhouette Edges." In: *Proc. Interactive 3D Graphics*. Atlanta, Georgia, USA, 1999, pp. 135–140. ISBN: 1-58113-082-1.
- [RBGV08] Peter Rautek, Stefan Bruckner, Eduard Gröller, and Ivan Viola. "Illustrative Visualization: New Technology or Useless Tautology?" In: *ACM SIGGRAPH Computer Graphics* 42.3 (Aug. 2008), 4:1–4:8.
- [RRWG+08] Jean Raymond, Daniel Roy, Alain Weill, Francois Guilbert, Thanh Nguyen, Andrew J Molyneux, Allan J Fox, S Claiborne Johnston, Jean-Paul Collet, and Isabelle Rouleau. "Unruptured intracranial aneurysms and the Trial on Endovascular Aneurysm Management (TEAM): The principles behind the protocol." In: *Journal of Vascular and Interventional Neurology* 1.1 (2008), pp. 22–26.
- [RSHSGoo] Christof Rezk-Salama, Peter Hastreiter, Jörg Scherer, and Günther Greiner. "Automatic Adjustment of Transfer Functions for 3D Volume Visualization." In: *Proc. of VMV*. 2000, pp. 357–364.

- [RML12] B. Rieck, H. Mara, and H. Leitte. "Multivariate Data Analysis Using Persistence-Based Filtering and Topological Signatures." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2382–2391.
- [RRRPO8] Christian Rieder, Felix Ritter, Matthias Raspe, and Heinz-Otto Peitgen. "Interactive visualization of multimodal volume data for neurosurgical tumor treatment." In: *Computer Graphics Forum* 27.3 (2008), pp. 1055–1062.
- [RSHPO8] Christian Rieder, Michael Schwier, Horst K Hahn, and Heinz-Otto Peitgen. "High-Quality Multimodal Volume Visualization of Intracerebral Pathological Tissue." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2008, pp. 167–176.
- [RHDKV+07] Felix Ritter, Christian Hansen, Volker Dicken, Olaf Konrad-Verse, Bernhard Preim, and Heinz-Otto Peitgen. "Real-Time Illustration of Vascular Structures." In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (Jan. 4, 2007), pp. 877–884.
- [RSHSo3] Felix Ritter, Henry Sonnet, Knut Hartmann, and Thomas Strothotte. "Illustrative Shadows: Integrating 3D and 2D Information Displays." In: *Proc. of Intelligent User Interfaces*. Miami, Florida, USA, 2003, pp. 166–173. ISBN: 1-58113-586-6.
- [RT06] Guodong Rong and Tiow-Seng Tan. "Jump Flooding in GPU with Applications to Voronoi Diagram and Distance Transform." In: *Proc. of the Symposium on Interactive 3D Graphics and Games*. 2006, pp. 109–116.
- [RHRS+09] Timo Ropinski, Sven Hermann, Rainer Reich, Michael Schäfers, and Klaus Hinrichs. "Multimodal vessel visualization of mouse aorta PET/CT scans." In: *Visualization and Computer Graphics, IEEE Transactions on* 15.6 (2009), pp. 1515–1522.
- [ROP11] Timo Ropinski, Steffen Oeltze, and Bernhard Preim. "Survey of glyph-based visualization techniques for spatial multivariate medical data." In: *Computers & Graphics* 35.2 (2011), pp. 392–401.
- [RSH06] Timo Ropinski, Frank Steinicke, and Klaus Hinrichs. "Visually Supporting Depth Perception in Angiography Imaging." English. In: *Proc. of Smart Graphics*. 2006, pp. 93–104. ISBN: 978-3-540-36293-7.
- [RVBH+09] Timo Ropinski, Ivan Viola, Martin Biermann, Helwig Hauser, and Klaus Hinrichs. "Multimodal visualization with interactive closeups." In: *Proc. of Theory and Practice of Computer Graphics Conference*. Eurographics Association. 2009, pp. 17–24.
- [RC13] Paul Rosin and John Collomosse, eds. *Image and Video based Artistic Stylisation*. Vol. 42. Computational Imaging and Vision. London/Heidelberg: Springer, 2013.
- [RSPRO6] Antoine Rossset, Luca Spadola, Lance Pysker, and Osman Ratib. "Navigating the Fifth Dimension: Innovative Interface for Multidimensional Multimodality Image Navigation 1." In: *Radiographics* 26.1 (2006), pp. 299–308.
- [RE92] Jarek R. Rossignac and Maarten van Emmerik. "Hidden Contours on a Frame-buffer." In: *Proceedings of the Seventh Eurographics Conference on Graphics Hardware*. Cambridge, UK: Eurographics Association, 1992, pp. 188–203.
- [RK00] C. Rössl and L. Kobbelt. "Line-art rendering of 3D-models." In: *Proc. Computer Graphics and Applications*. 2000, pp. 87–96.
- [RKS00] Christian Rössl, Leif Kobbelt, and Hans-Peter Seidel. "Line Art Rendering of Triangulated Surfaces using Discrete Lines of Curvature." In: *Proc. Computer Graphics, Visualization and Interactive Digital Media*. 2000, pp. 168–175. ISBN: 80-7082-612-6.
- [RBE08] Friedemann Rößler, Ralf P Botchen, and Thomas Ertl. "Dynamic shader generation for GPU-based multi-volume ray casting." In: *Computer Graphics and Applications* 28.5 (2008), pp. 66–77.
- [RTFE+06] Friedemann Rößler, Eduardo Tejada, Thomas Fangmeier, Thomas Ertl, and Markus Knauff. "GPU-based Multi-Volume Rendering for the Visualization of Functional Brain Images." In: *Proc. of SimVis*. 2006, pp. 305–18.
- [RVBB+08] Marc Ruiz, Ivan Viola, Imma Boada, Stefan Bruckner, Miquel Feixas, and Mateu Sbert. "Similarity-based Exploded Views." In: *Proc. of Smart Graphics*. 2008, pp. 154–165.
- [Rus04] Szymon Rusinkiewicz. "Estimating Curvatures and Their Derivatives on Triangle Meshes." In: *Symposium on 3D Data Processing, Visualization, and Transmission*. Sept. 2004.
- [RBD06] Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. "Exaggerated Shading for Depicting Shape and Detail." In: *ACM Transactions on Graphics* 25.3 (July 2006), pp. 1199–1205. ISSN: 0730-0301.
- [RCDFo8] Szymon Rusinkiewicz, Forrester Cole, Doug DeCarlo, and Adam Finkelstein. "Line Drawings from 3D Models." In: *ACM SIGGRAPH Classes*. Los Angeles, California: ACM, 2008, 39:1–39:356.
- [Rus89] Pramod Rustagi. "Silhouette Line Display From Shaded Models." In: *IRIS Universe* 1989.9 (1989), pp. 42–44.
- [SBPP+15] Patrick Saalfeld, Alexandra Baer, Uta Preim, Bernhard Preim, and Kai Lawonn. "Sketching 2D Vessels and Vascular Diseases with Integrated Blood Flow." In: *Proc. of the 10th International Conference on Computer Graphics Theory and Applications (GRAPP)*. Berlin, 2015, pp. 379–390.
- [SBPP+16] Patrick Saalfeld, Alexandra Baer, Uta Preim, Bernhard Preim, and Kai Lawonn. "A Sketch-Based Interface for 2D Illustration of Vascular Structures, Diseases, and Treatment Options with Real-Time Blood Flow." In: vol. 598. Springer International Publishing, 2016. Chap. Computer Vision, Imaging and Computer Graphics Theory and Applications: International Joint Conference, Revised Selected Papers, pp. 19–40. ISBN: 978-3-319-29971-6.
- [ST90] Takafumi Saito and Tokihiro Takahashi. "Comprehensible Rendering of 3-D Shapes." In: *ACM SIGGRAPH Computer Graphics* 24.4 (Sept. 1990), pp. 197–206. ISSN: 0097-8930.
- [STSo6] Natascha Sauber, Holger Theisel, and H-P Seidel. "Multifield-graphs: An approach to visualizing correlations in multifield scalar data." In: *Visualization and Computer Graphics, IEEE Transactions on* 12.5 (2006), pp. 917–924.
- [SFNB14] M. Schlachter, T. Fechter, U. Nestle, and K. Bühler. "Visualization of 4D-PET/CT, Target Volumes and Dose Distribution: Applications in Radiotherapy Planning." In: *Proc. of MICCAI Workshop on Image-Guided Adaptive Radiation Therapy*. 2014.

- [Sch97] J. Schöberl. "Netgen: An Advancing Front 2D/3D-Mesh Generator based on Abstract Rules." In: *Computing and Visualization in Science* 1.1 (1997), pp. 41–52.
- [SS11] Nicole Schubert and Ingrid Scholl. "Comparing GPU-based multi-volume ray casting techniques." In: *Computer Science-Research and Development* 26.1-2 (2011), pp. 39–50.
- [SS09] Gustav K von Schulthess and Heinz-Peter W Schlemmer. "A look ahead: PET/MR versus PET/CT." In: *European journal of nuclear medicine and molecular imaging* 36.1 (2009), pp. 3–9.
- [SNEAF+16] Christoph Schulz et al. "Generative Data Models for Validation and Evaluation of Visualization Techniques." In: *Beyond Time And Errors: Novel Evaluation Methods For Visualization*. 2016.
- [Sec02] Adrian Secord. "Weighted Voronoi Stippling." In: *Proc. NPAR*. New York: ACM, 2002, pp. 37–43.
- [SHSo2] Adrian Secord, Wolfgang Heidrich, and Lisa Streit. "Fast Primitive Distribution for Illustration." In: *Proc. EGWR*. Goslar, Germany: Eurographics Association, 2002, pp. 215–226.
- [SKLB+10] Heiko Seim, Dagmar Kainmueller, Hans Lamecker, Matthias Bindernagel, Jana Malinowski, and Stefan Zachow. "Model-based auto-segmentation of knee bones and cartilage in MRI data." In: *Proc. of Medical Image Analysis for the Clinic: A Grand Challenge*. 2010, pp. 215–223.
- [SPSPo2] Dirk Selle, Bernhard Preim, Andrea Schenk, and Heinz-otto Peitgen. "Analysis of vasculature for liver surgical planning." In: *Medical Imaging, IEEE Transactions on* 21.11 (2002), pp. 1344–1357.
- [STKD12] Amir Semmo, Matthias Trapp, Jan Eric Kyprianidis, and Jürgen Döllner. "Interactive Visualization of Generalized Virtual 3D City Models using Level-of-Abstraction Transitions." In: *Computer Graphics Forum* 31.3 (2012), pp. 885–894.
- [Ser82] Jean Serra. *Image analysis and mathematical morphology, v. 1*. Academic Press, 1982.
- [SKGH+98] Luis Serra, Ralf A Kockro, Chua Gim Guan, Ng Hern, Eugene CK Lee, Yen H Lee, Chumpon Chan, and Wieslaw L Nowinski. "Multimodal volume-based tumor neurosurgery planning in the virtual workbench." In: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 1998, pp. 1007–1015.
- [SS12] S Kamaledin Setarehdan and Sameer Singh. *Advanced algorithmic approaches to medical image segmentation: state-of-the-art applications in cardiology, neurology, mammography and pathology*. Springer Science & Business Media, 2012.
- [SK14] Ankit Sharma and Subodh Kumar. "User-guided modulation of rendering techniques for detail inspection." In: *Computer Graphics Theory and Applications*. 2014, pp. 1–8.
- [SBHN06] T. Sielhorst, C. Bichlmeier, S. M. Heining, and N. Navab. "Depth perception—a major issue in medical AR: evaluation study by twenty surgeons." In: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 9 (2006), pp. 364–372.
- [SFCP12] Stephan Sigg, Raphael Fuchs, Robert Carnecky, and Ronald Peikert. "Intelligent Cutaway Illustrations." In: *Proc. PacificVis*. Los Alamitos: IEEE Computer Society, 2012, pp. 185–192. ISBN: 978-1-4673-0865-6.
- [SS10] Mayank Singh and Scott Schaefer. "Suggestive Hatching." In: *Proc. Computational Aesthetics in Graphics, Visualization and Imaging*. London, United Kingdom: Eurographics Association, 2010, pp. 25–32. ISBN: 978-3-905674-24-8.
- [SMGG01] Peter-Pike Sloan, William Martin, Amy Gooch, and Bruce Gooch. "The Lit Sphere: A Model for Capturing NPR Shading from Art." In: *Proc. Graphics Interface*. Toronto: CIPS, 2001, pp. 143–150.
- [SLKD+17] Noeska Smit, Kai Lawonn, Annelot Kraima, Marco DeRuiter, Hessam Sokooti, Stefan Bruckner, Elmar Eisemann, and Anna Vilanova. "PelVis: Atlas-based Surgical Planning for Oncological Pelvic Surgery." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 741–750.
- [SZKCo6] Noah Snaveley, C. Lawrence Zitnick, Sing Bing Kang, and Michael Cohen. "Stylizing 2.5-D Video." In: *Proc. NPAR*. Annecy, France: ACM, 2006, pp. 63–69. ISBN: 1-59593-357-3.
- [SCBW14] Justin Solomon, Keenan Crane, Adrian Butscher, and Chris Wojtan. "A General Framework for Bilateral and Mean Shift Filtering." In: *ArXiv e-prints* 32 (2014). arXiv: 1405.4734 [cs.GR].
- [Sor05] Olga Sorkine. "Laplacian Mesh Processing." In: *Eurographics – State of the Art Reports*. Ed. by Yiorgos Chrysanthou and Marcus Magnor. Goslar, Germany: Eurographics Association, 2005, pp. 53–70.
- [SB16] Augusto Sousa and Kadi Bouatouch, eds. *EG 2016 - Tutorials*. Lisbon, Portugal: Eurographics Association, 2016. ISBN: -.
- [SESS05] Mario Costa Sousa, David S. Ebert, Don Stredney, and Nikolai A. Svakhine. "Illustrative Visualization for Medical Training." In: *Proc. Eurographics Workshop on Computational Aesthetics in Graphics*. 2005, pp. 201–208.
- [SP03] Mario Costa Sousa and Przemyslaw Prusinkiewicz. "A Few Good Lines: Suggestive Drawing of 3D Models." In: *Computer Graphics Forum* 22.3 (2003), pp. 381–390. ISSN: 1467-8659.
- [SH95] Detlev Stalling and Hans-Christian Hege. "Fast and Resolution Independent Line Integral Convolution." In: *Proc. SIGGRAPH*. New York: ACM, 1995, pp. 249–256. ISBN: 0-89791-701-4.
- [SZV01] R. Stokking, K.J. Zuiderveld, and M.A. Viergever. "Integrated volume visualization of functional image data and anatomical surfaces using normal fusion." In: *Human Brain Mapping* 12 (2001), pp. 203–218.
- [SZPR+97] Rik Stokking, Karel J Zuiderveld, Hilleke E Hulshoff Pol, Peter P van Rijk, and Max A Viergever. "Normal fusion for three-dimensional integrated visualization of SPECT and magnetic resonance brain images." In: *The Journal of Nuclear Medicine* 38.4 (1997), p. 624.
- [SCSL+02] David Stout, P Chow, R Silverman, Richard M Leahy, X Lewis, S Gambhir, and A Chatzioannou. "Creating a whole body digital mouse atlas with PET, CT and cryosection images." In: *Mol. Imaging Biol* 4.4 (2002), S27.
- [SMI99] T. Strothotte, M. Masuch, and T. Isenberg. "Visualizing knowledge about virtual reconstructions of ancient architecture." In: *Proc. Computer Graphics International*. 1999, pp. 36–43.

- [SPRS+94] T. Strothotte, B. Preim, A. Raab, J. Schumann, and D. R. Forsey. "How to Render Frames and Influence People." In: *Computer Graphics Forum*. Vol. 13. 3. Special issue on Eurographics '94. 1994, pp. 455–466.
- [SS02] Thomas Strothotte and Stefan Schlechtweg. *Non-Photorealistic Computer Graphics. Modeling, Animation, and Rendering*. San Francisco: Morgan Kaufmann Publishers, 2002.
- [SBB13] Jordane Suarez, Farès Belhadj, and Vincent Boyer. "GPU Real Time Hatching." In: *Computer Graphics, Visualization and Computer Vision* 21.2 (2013), pp. 97–105.
- [SBB16] Jordane Suarez, Farès Belhadj, and Vincent Boyer. "Real-time 3D rendering with hatching." In: *The Visual Computer* (2016), pp. 1–16. ISSN: 1432-2315.
- [SKR15] Erik Sunden, Sathish Kotravel, and Timo Ropinski. "Multimodal volume illumination." In: *Computers & Graphics* 50 (2015), pp. 47–60.
- [SEI10] Pjotr Svetachov, Maarten H. Everts, and Tobias Isenberg. "DTI in Context: Illustrating Brain Fiber Tracts In Situ." In: *Computer Graphics Forum* 29.3 (June 2010), pp. 1024–1032.
- [SW04] Graeme Sweet and Colin Ware. "View Direction, Surface Orientation and Texture Orientation for Perception of Surface Shape." In: *Proc. Graphics Interface*. Waterloo, ON, Canada: CHCCS, 2004, pp. 97–106. ISBN: 1-56881-227-2.
- [SSK16] László Szécsi, Marcell Szirányi, and Ágota Kacsó. "Tonal Art Maps with Image Space Strokes." In: *EG - Posters*. The Eurographics Association, 2016.
- [Tal14] Ayellet Tal. "3D Shape Analysis for Archaeology." English. In: *3D Research Challenges in Cultural Heritage*. Ed. by Marinou Ioannides and Ewald Quak. Vol. 8355. Lecture Notes in Computer Science. 2014, pp. 50–63. ISBN: 978-3-662-44629-4.
- [TPD06] Andreas Tappenbeck, Bernhard Preim, and Volker Dicken. "Distance-based transfer function design: Specification Methods and Applications." In: *Proc. of SimVis*. 2006, pp. 259–274.
- [Tau95a] G. Taubin. "Estimating the tensor of curvature of a surface from a polyhedral approximation." In: *Proc. ICCV*. IEEE Computer Society, 1995, pp. 902–. ISBN: 0-8186-7042-8.
- [Tau95b] Gabriel Taubin. "A signal processing approach to fair surface design." In: *Proc. SIGGRAPH*. New York: ACM, 1995, pp. 351–358. ISBN: 0-89791-701-4.
- [TGE97] Christian Teitzel, Roberto Grosso, and Thomas Ertl. "Efficient and Reliable Integration Methods for Particle Tracing in Unsteady Flows on Discrete Meshes." In: *Proc. Visualization in Scientific Computing*. Berlin/Heidelberg: Springer, 1997, pp. 31–41.
- [TBBV+07] Maurice Termeer, Javier Oliván Bescós, Marcel Breeuwer, Anna Vilanova, Frans Gerritsen, and M Eduard Gröller. "CoViCAD: comprehensive visualization of coronary artery disease." In: *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007), pp. 1632–1639.
- [TIP05] Christian Tietjen, Tobias Isenberg, and Bernhard Preim. "Combining Silhouettes, Surface, and Volume Rendering for Surgery Education and Planning." In: *IEEE/Eurographics Symposium on Visualization (EuroVis)*. 2005, pp. 303–310.
- [TMSP+] Christian Tietjen, Björn Meyer, Stefan Schlechtweg, Bernhard Preim, Ilka Hertel, and Gero Strauss. "Enhancing Slice-based Visualization of Medical Volume Data." In: *Proc. of EuroVis*.
- [TPBG+08] Christian Tietjen, Roland Pfisterer, Alexandra Baer, Rocco Gasteiger, and Bernhard Preim. "Hardware-Accelerated Illustrative Medical Surface Visualization with Extended Shading Maps." In: *Proc. Smart Graphics*. 2008, pp. 166–177. ISBN: 978-3-540-85412-8.
- [TBBE+99] Paul S Tofts, Gunnar Brix, David L Buckley, Jeffrey L Evelhoch, Elizabeth Henderson, Michael V Knopp, Henrik BW Larsson, Ting-Yim Lee, Nina A Mayr, Geoffrey JM Parker, et al. "Estimating kinetic parameters from dynamic contrast-enhanced T1-weighted MRI of a diffusible tracer: standardized quantities and symbols." In: *Journal of Magnetic Resonance Imaging* 10.3 (1999), pp. 223–232.
- [TGKD+14] Christian Tominski, Stefan Gladisch, Ulrike Kister, Raimund Dachselt, and Heidrun Schumann. "A Survey on Interactive Lenses in Visualization." In: *EuroVis State-of-the-Art Reports*. 2014.
- [TCYH04] D.W. Townsend, J.P. Carney, J.T. Yap, and N.C. Hall. "PET/CT today and tomorrow." In: *The Journal of Nuclear Medicine* 45.1 (2004), 4S–14S.
- [TC00] S. M. F. Treavett and M. Chen. "Pen-and-Ink Rendering in Volume Visualisation." In: *Proc. of IEEE Visualization*. 2000, pp. 203–210.
- [TCVP+13] Giorgio Treglia, Paola Castaldi, Maria Felicia Villani, Germano Perotti, Angelina Filice, Valentina Ambrosini, Nadia Cremonini, Annibale Versari, Stefano Fanti, Alessandro Giordano, et al. "Comparison of different positron emission tomography tracers in patients with recurrent medullary thyroid carcinoma: our experience and a review of the literature." In: *Theranostics, Gallium-68, and Other Radionuclides*. Springer, 2013, pp. 385–393.
- [TG80] Anne M Treisman and Garry Gelade. "A feature-integration theory of attention." In: *Cognitive psychology* 12.1 (1980), pp. 97–136.
- [TYKO05] T. Tsuda, H. Yamamoto, Y. Kameda, and Y. Ohta. "Visualization methods for outdoor see-through vision." In: *Proc. of the International Conference on Augmented Tele-Existence*. ACM, 2005, pp. 62–69. ISBN: 0-473-10657-4.
- [USB92] Jayaram K Udupa, Supun Samarasekera, and William A Barrett. "Boundary detection via dynamic programming." In: *Proc. of Visualization in Biomedical Computing*. International Society for Optics and Photonics. 1992, pp. 33–39.
- [VDHP10] Martijn P Van Den Heuvel and Hilleke E Hulshoff Pol. "Exploring the brain network: a review on resting-state fMRI functional connectivity." In: *European Neuropsychopharmacology* 20.8 (2010), pp. 519–534.
- [VBTS07] David Vanderhaeghe, Pascal Barla, Joelle Thollot, and Francois X. Sillion. "Dynamic Point Distribution for Stroke-based Rendering." In: *Rendering Techniques*. Goslar, Germany: Eurographics Association, 2007, pp. 139–146.
- [VVPS+16] Libor Váša, Petr Vaněček, Martin Prantl, Věra Skorkovská, Petr Martínek, and Ivana Kolingerová. "Mesh Statistics for Robust Curvature Estimation." In: *Computer Graphics Forum* 35.3 (2016), pp. 271–280. ISSN: 1467-8659.

- [VFSH01] Pere-Pau Vázquez, Miquel Feixas, Mateu Sbert, and Wolfgang Heidrich. "Viewpoint Selection using Viewpoint Entropy." In: *Proc. of VMV*. Vol. 1. 2001, pp. 273–280.
- [VPBG+11] R. Vergne, R. Pacanowski, P. Barla, X. Granier, and C. Schlick. "Improving Shape Depiction under Arbitrary Rendering." In: *IEEE Transactions on Visualization and Computer Graphics* 17.8 (2011), pp. 1071–1081. ISSN: 1077-2626.
- [VBCS08] Romain Vergne, Pascal Barla, Xavier Granier, and Christophe Schlick. "Apparent Relief: A Shape Descriptor for Stylized Shading." In: *Proceedings of Non-photorealistic Animation and Rendering*. 2008, pp. 23–29.
- [VPBG+09] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick. "Light Warping for Enhanced Surface Depiction." In: *ACM Transactions on Graphics* 28.3 (2009), 25:1–25:8.
- [VPPG+14] Anna Vilanova, Bernhard Preim, Roy van Pelt, Rocco Gasteiger, Mathias Neugebauer, and Thomas Wischgoll. "Visual Exploration of Simulated and Measured Blood Flow." In: *Scientific Visualization*. Berlin/Heidelberg: Springer, 2014, pp. 305–324.
- [VBSH+13] I. Viola, Å. Birkeland, V. Solteszova, L. Helljesen, H. Hauser, K. Kotopoulos, S. and Nylund, O. Ulvang, D. and Øye, T. Hausken, and O. Gilja. "High-Quality 3D Visualization of In-Situ Ultrasonography." In: *Proc. of Eurographics 2013 - Dirk Bartz Prize*. Ed. by Hans-Christian Hege and Anna Vilanova. The Eurographics Association, 2013.
- [VGHB+05] Ivan Viola, Eduard Gröller, Markus Hadwiger, Katja Bühler, Bernhard Preim, Mario Costa Sousa, David Ebert, and Don Stredney. "Illustrative Visualization." In: *IEEE Visualization Tutorials*. Los Alamitos: IEEE Computer Society, 2005. Chap. 4.
- [VG05] Ivan Viola and Meister E. Gröller. "Smart Visibility in Visualization." In: *Proc. of Computational Aesthetics in Graphics, Visualization and Imaging*. 2005, pp. 209–216.
- [VHE10] Ivan Viola, Helwig Hauser, and David Ebert. "Editorial Note for Special Section on Illustrative Visualization." In: *Computers & Graphics* 34.4 (Aug. 2010), pp. 335–336.
- [VKG04] Ivan Viola, Armin Kanitsar, and Meister Eduard Groller. "Importance-driven volume rendering." In: *Proc. of IEEE Visualization*. IEEE Computer Society, 2004, pp. 139–146.
- [VNØU+08] Ivan Viola, Kim Nylund, Ola Kristoffer Øye, Dag Magne Ulvang, Odd Helge Gilja, and Helwig Hauser. "Illustrated Ultrasound for Multimodal Data Interpretation of Liver Examinations." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2008, pp. 125–133.
- [WVKJ+06] F. K. Wacker, S. Vogt, A. Khamene, J. A. Jesberger, S. G. Nour, D. R. Elgort, F. Sauer, J. L. Duerk, and J. S. Lewin. "An augmented reality system for MR image-guided needle biopsy: initial results in a swine model." In: *Radiology* 238.2 (2006), pp. 497–504.
- [WMLK89] Jerold W Wallis, Tom R Miller, Charles A Lerner, and Eric C Kleerup. "Three-dimensional display in nuclear medicine." In: *Medical Imaging, IEEE Transactions on* 8.4 (1989), pp. 297–230.
- [WH10] Y. Wan and C. Hansen. "Fast Volumetric Data Exploration with Importance-based Accumulated Transparency Modulation." In: *Proc. of IEEE/EG International Conference on Volume Graphics*. 2010, pp. 61–68.
- [WBDF+14] Xiang Wang, Christian Schulte zu Berge, Stefanie Demirci, Pascal Fallavollita, and Nassir Navab. "Improved interventional X-ray appearance." In: *IEEE International Symposium on Mixed and Augmented Reality*. 2014, pp. 237–242.
- [WMKG07] Max Wardetzky, Saurabh Mathur, Felix Kaelberer, and Eitan Grinspun. "Discrete Laplace operators: No free lunch." In: *Geometry Processing*. 2007, pp. 33–37. ISBN: 978-3-905673-46-3.
- [WB01] Kouki Watanabe and Alexander G. Belyaev. "Detection of Salient Curvature Features on Polygonal Surfaces." In: *Computer Graphics Forum* 20.3 (2001), pp. 385–392.
- [WPFH02] Matthew Webb, Emil Praun, Adam Finkelstein, and Hugues Hoppe. "Fine Tone Control in Hardware Hatching." In: *Proc. NPAR*. Annecy, France: ACM, 2002, pp. 53–59. ISBN: 1-58113-494-0.
- [WJWP09] Hans F Wehrl, Martin S Judenhofer, Stefan Wiehr, and Bernd J Pichler. "Pre-clinical PET/MR: technological advances and new perspectives in biomedical research." In: *European journal of nuclear medicine and molecular imaging* 36.1 (2009), pp. 56–68.
- [WYPW+15] M. Wei, J. Yu, W. M. Pang, J. Wang, J. Qin, L. Liu, and P. A. Heng. "Bi-Normal Filtering for Mesh Denoising." In: *IEEE Transactions on Visualization and Computer Graphics* 21.1 (2015), pp. 43–55. ISSN: 1077-2626.
- [WZYW+15] Mingqiang Wei, Lei Zhu, Jinze Yu, Jun Wang, Wai-Man Pang, Jianhuang Wu, Jing Qin, and Pheng-Ann Heng. "Morphology-preserving Smoothing on Polygonized Isosurfaces of Inhomogeneous Binary Volumes." In: *Comput. Aided Des.* 58 (2015), pp. 92–98. ISSN: 0010-4485.
- [WRDW+11] Florian Weiler, Christian Rieder, C. A. David, C. Wald, and Horst K. Hahn. "AVM-Explorer: Multi-Volume Visualization of Vascular Structures for Planning of Cerebral AVM Surgery." In: (2011). Ed. by K. Buehler and A. Vilanova.
- [WKCK+07] Wolfgang Wein, Ali Khamene, Dirk-André Clevert, Oliver Kutter, and Nassir Navab. "Simulation and fully automatic multimodal registration of medical ultrasound." In: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2007, pp. 136–143.
- [Wei+07] Wolfgang Wein et al. "Multimodal Integration of Medical Ultrasound for Treatment Planning and Interventions." PhD thesis. Technische Universität München, 2007.
- [WG09] T. Weinkauff and D. Günther. "Separatrix Persistence: Extraction of Salient Edges on Surfaces Using Topological Methods." In: *Computer Graphics Forum* 28.5 (2009), pp. 1519–1528. ISSN: 1467-8659.
- [WBEF+05] D. Weiskopf et al. "Visualization in the Einstein Year 2005: A case study on explanatory and illustrative visualization of relativity and astrophysics." In: *IEEE Visualization*. 2005, pp. 583–590.
- [WBEF+06] Daniel Weiskopf et al. "Explanatory and Illustrative Visualization of Special and General Relativity." In: *IEEE Transactions on Visualization and Computer Graphics* 12.4 (July 2006), pp. 522–534.
- [WEE02] Daniel Weiskopf, Klaus Engel, and Thomas Ertl. "Volume clipping via per-fragment operations in texture-based volume visualization." In: *Proc. of IEEE Visualization*. IEEE Computer Society, 2002, pp. 93–100.

- [WNCC13] Rong Wen, Binh P. Nguyen, Chin-Boon Chng, and Chee-Kong Chui. "In situ spatial AR surgical planning using projector-Kinect system." In: *International Symposium on Information and Communication Technology*. 2013, pp. 164–171.
- [WLM02] Brett Wilson, Eric B. Lum, and Kwan-Liu Ma. "Interactive Multi-volume Visualization." English. In: *Computational Science*. Vol. 2330. 2002, pp. 102–110.
- [XHTS+07] Xuexiang Xie, Ying He, Feng Tian, Hock-Soon Seah, Xianfeng Gu, and Hong Qin. "An Effective Illustrative Visualization Framework Based on Photic Extremum Lines (PELs)." In: *IEEE Transactions on Visualization and Computer Graphics* 13 (2007), pp. 1328–1335.
- [XNYC04] Hui Xu, Minh X. Nguyen, Xiaoru Yuan, and Baoquan Chen. "Interactive Silhouette Rendering for Point-based Models." In: *Proceedings of the First Eurographics Conference on Point-Based Graphics*. Switzerland: Eurographics Association, 2004, pp. 13–18. ISBN: 3-905673-09-6.
- [YBYS07] Shin Yoshizawa, Alexander G. Belyaev, Hideo Yokota, and Hans-Peter Seidel. "Fast and Faithful Geometric Algorithm for Detecting Crest Lines on Meshes." In: *Proc. Pacific Graphics*. IEEE Computer Society, 2007, pp. 231–237.
- [YBS05] Shin Yoshizawa, Alexander Belyaev, and Hans-Peter Seidel. "Fast and Robust Detection of Crest Lines on Meshes." In: *Proc. Solid and Physical Modeling*. Cambridge, Massachusetts: ACM, 2005, pp. 227–232. ISBN: 1-59593-015-9.
- [YBYS08] Shin Yoshizawa, Alexander Belyaev, Hideo Yokota, and Hans-Peter Seidel. "Fast, Robust, and Faithful Methods for Detecting Crest Lines on Meshes." In: *Comput. Aided Geom. Des.* 25.8 (Nov. 2008), pp. 545–560. ISSN: 0167-8396.
- [YWQW+14] Jinze Yu, Mingqiang Wei, Jing Qin, Jianhuang Wu, and Pheng-Ann Heng. "Feature-preserving mesh denoising via normal guided quadric error metrics." In: *Optics and Lasers in Engineering* 62.0 (2014), pp. 57–68. ISSN: 0143-8166.
- [YCo4] Xiaoru Yuan and Baoquan Chen. "Illustrating Surfaces in Volume." In: *Proc. of EG/IEEE VisSym*. 2004, pp. 9–16.
- [YNZC05] Xiaoru Yuan, Minh X. Nguyen, Nan Zhang, and Baoquan Chen. "Stippling and Silhouettes Rendering in Geometry-image Space." In: *Proc. Eurographics Conference on Rendering Techniques*. Konstanz, Germany, 2005, pp. 193–200. ISBN: 3-905673-23-1.
- [ZMO+07] Habib Zaidi, Osama Mawlawi, Colin G Orton, et al. "Simultaneous PET/MR will replace PET/CT as the molecular multimodality imaging platform of choice." In: *Medical Physics* 34.5 (2007), pp. 1525–1528.
- [ZSo4] Nordin Zakaria and Hans-Peter Seidel. "Interactive Stylized Silhouette for Point-sampled Geometry." In: *Proc. Computer Graphics and Interactive Techniques in Australasia and South East Asia*. Singapore: ACM, 2004, pp. 242–249. ISBN: 1-58113-883-0.
- [ZISS04] Johannes Zander, Tobias Isenberg, Stefan Schlechtweg, and Thomas Strothotte. "High Quality Hatching." In: *Computer Graphics Forum* 23.3 (Sept. 2004), pp. 421–430.
- [ZTS09] Rony Zatzarinni, Ayellet Tal, and Ariel Shamir. "Relief Analysis and Extraction." In: *Proc. of ACM SIGGRAPH Asia*. Yokohama, Japan, 2009, 136:1–136:9. ISBN: 978-1-60558-858-2.
- [ZSLE+15] Changgong Zhang, Thomas Schultz, Kai Lawonn, Elmar Eisemann, and Anna Vilanova. "Glyph-based Comparative Visualization for Diffusion Tensor Fields." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 22 (1).1 (2015), pp. 797–806.
- [ZHS10] Long Zhang, Ying He, and Hock Soon Seah. "Real-time computation of photic extremum lines (PELs)." In: *The Visual Computer* 26.6-8 (June 2010), pp. 399–407.
- [ZHXX+11] Long Zhang, Ying He, Jiazhi Xia, Xuexiang Xie, and Wei Chen. "Real-Time Shape Illustration Using Laplacian Lines." In: *IEEE Transactions on Visualization and Computer Graphics* 17.7 (July 2011), pp. 993–1006.
- [ZXYH+12] Long Zhang, Jiazhi Xia, Xiang Ying, Ying He, Wolfgang Mueller-Wittig, and Hock-Soon Seah. "Efficient and Robust 3D Line Drawings Using difference-of-Gaussian." In: *Graph. Models* 74.4 (July 2012), pp. 87–98. ISSN: 1524-0703.
- [ZDZB+15] Wangyu Zhang, Bailin Deng, Juyong Zhang, Sofien Bouaziz, and Ligang Liu. "Guided Mesh Normal Filtering." In: *Computer Graphics Forum* 34.7 (Oct. 2015), pp. 23–34. ISSN: 0167-7055.
- [ZZLZ13] Yu-Wei Zhang, Yi-Qi Zhou, Xue-Lin Li, and Li-Li Zhang. "Line-based sunken relief generation from a 3D mesh." In: *Graphical Models* 75.6 (2013), pp. 297–304. ISSN: 1524-0703.
- [ZHT05] Lei Zhu, Steven Haker, and Allen Tannenbaum. "Flattening maps for the visualization of multi-branched vessels." In: *IEEE Trans. Med. Imaging* 24.2 (2005), pp. 191–198.
- [ZLBI11] Matthew van der Zwan, Wouter Lueks, Henk Bekker, and Tobias Isenberg. "Illustrative Molecular Visualization with Continuous Abstraction." In: *Computer Graphics Forum* 30.3 (June 2011), pp. 683–690. ISSN: 1467-8659.

LIST OF PUBLICATIONS

- [BLSP15] Alexandra Baer, Kai Lawonn, Patrick Saalfeld, and Bernhard Preim. "Statistical Analysis of a Qualitative Evaluation on Feature Lines." In: *Bildverarbeitung für die Medizin (BVM)*. Springer Verlag, 2015, pp. 71–76.
- [GHMT+17] Ajay Abisheck Paul George, Pascal Heimer, Raphael Menges, Veda Thota, Astrid Maass, Jan Hamaekers, Martin Hofmann-Apitius, Nils Lichtenberg, Kai Lawonn, and Diana Imhof. "Novel insights into μ -conotoxin conformations: A comparative study using molecular dynamics and computer visualization methods." In: *Scientific Reports* (2017), submitted.
- [KGLR+15] Paul Klemm, Sylvia Glaßer, Kai Lawonn, Marko Rak, Henry Völzke, Katrin Hegenscheid, and Bernhard Preim. "Interactive Visual Analysis of Lumbar Back Pain." In: *Proc. of the 6th International Conference on Information Visualization Theory and Applications (IVAPP)*. 2015, pp. 85–92.
- [KLG+15] Paul Klemm, Kai Lawonn, Sylvia Glaßer, Uli Niemann, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. "3D Regression Heat Map Analysis of Population Study Data." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 22 (1).1 (2015), pp. 81–90.
- [KOJLH+14a] Paul Klemm, Steffen Oeltze-Jafra, Kai Lawonn, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. "Interactive Visual Analysis of Image-Centric Cohort Study Data." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2014), pp. 1673–1682.
- [LBSP14] Kai Lawonn, Alexandra Baer, Patrick Saalfeld, and Bernhard Preim. "Comparative Evaluation of Feature Line Techniques for Shape Depiction." In: *Proc. Vision, Modeling and Visualization*. 2014, pp. 31–38. doi: 10.2312/vmv.20141273.
- [LGVP+15] Kai Lawonn, Sylvia Glaßer, Anna Vilanova, Bernhard Preim, and Tobias Isenberg. "Occlusion-free Blood Flow Animation with Wall Thickness Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 22 (1).1 (2015), pp. 728–737.
- [LGP14b] Kai Lawonn, Tobias Günther, and Bernhard Preim. "Coherent View-Dependent Streamlines for Understanding Blood Flow." In: *EuroVis - Short Papers*. 2014, 5 pages.
- [LLH17] Kai Lawonn, Maria Luz, and Christian Hansen. "Improving Spatial Perception of Vascular Models using Supporting Anchors and Illustrative Visualization." In: *Computers & Graphics* (2017), in print.
- [LLPH15] Kai Lawonn, Maria Luz, Bernhard Preim, and Christian Hansen. "Illustrative Visualization of Vascular Models for Static 2D Representations." In: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2015, pp. 399–406.
- [LSBP17] Kai Lawonn, Noeska Smit, Katja Bühler, and Bernhard Preim. "A Survey on Multimodal Medical Data Visualization." In: *Computer Graphics Forum* (2017), accepted with minor revision.
- [LSPV15] Kai Lawonn, Noeska Smit, Bernhard Preim, and Anna Vilanova. "Illustrative Multi-volume Rendering for PET/CT Scans." In: *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. The Eurographics Association, 2015, pp. 103–112.
- [LTPH17] Kai Lawonn, Erik Trostmann, Bernhard Preim, and Klaus Hildebrandt. "Visualization and Extraction of Carvings for Heritage Conservation." In: *IEEE Transactions on Visualization and Computer Graphics* 23(1) (2017), pp. 801–810.
- [LVPI17] Kai Lawonn, Ivan Viola, Bernhard Preim, and Tobias Isenberg. "A Survey of Surface-Based Illustrative Rendering for Visualization." In: *Computer Graphics Forum* (2017), submitted.
- [LSHL16] Nils Lichtenberg, Noeska Smit, Christian Hansen, and Kai Lawonn. "Slime: Seamless Line Illustration for Interactive Biomedical Visualization." In: *Proc. Visual Computing for Biology and Medicine*. 2016, to appear.
- [MEBP+17] Monique Meuschke, Wito Engelke, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Automatic Viewpoint Selection for Exploration of Time-dependent Cerebral Aneurysm Data." In: *Bildverarbeitung für die Medizin (BVM)*. 2017, in print.
- [MKPP+16] Monique Meuschke, Benjamin Köhler, Uta Preim, Bernhard Preim, and Kai Lawonn. "Semi-automatic Vortex Flow Classification in 4D PC-MRI Data of the Aorta." In: *Computer Graphics Forum* 35.3 (2016), pp. 351–360.
- [MLKP+16] Monique Meuschke, Kai Lawonn, Benjamin Köhler, Uta Preim, and Bernhard Preim. "Clustering of Aortic Vortex Flow in Cardiac 4D PC-MRI Data." In: *Bildverarbeitung für die Medizin (BVM)*. 2016, pp. 182–187.
- [MVBP+17a] Monique Meuschke, Samuel Voss, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Combined Visualization of Vessel Deformation and Hemodynamics in Cerebral Aneurysms." In: *IEEE Transactions on Visualization and Computer Graphics (Proceedings of the Scientific Visualization 2016)* 23(1) (2017), pp. 761–770.
- [MVBP+17b] Monique Meuschke, Samuel Voss, Oliver Beuing, Bernhard Preim, and Kai Lawonn. "Glyph-based Comparative Stress Tensor Visualization in Cerebral Aneurysms." In: *Computer Graphics Forum* (2017), conditionally accepted.
- [OJMNS+17] Steffen Oeltze-Jafra, Monique Meuschke, Mathias Neugebauer, Sylvia Saalfeld, Kai Lawonn, Gabor Janiga, Hans-Christian Hege, Stefan Zachow, and Bernhard Preim. "Generation and Visual Exploration of Medical Flow Data: Survey, Research Trends, and Future Challenges." In: *Computer Graphics Forum* (2017), submitted.
- [SBPP+15] Patrick Saalfeld, Alexandra Baer, Uta Preim, Bernhard Preim, and Kai Lawonn. "Sketching 2D Vessels and Vascular Diseases with Integrated Blood Flow." In: *Proc. of the 10th International Conference on Computer Graphics Theory and Applications (GRAPP)*. Berlin, 2015, pp. 379–390.

- [SBPP+16] Patrick Saalfeld, Alexandra Baer, Uta Preim, Bernhard Preim, and Kai Lawonn. "A Sketch-Based Interface for 2D Illustration of Vascular Structures, Diseases, and Treatment Options with Real-Time Blood Flow." In: vol. 598. Springer International Publishing, 2016. Chap. Computer Vision, Imaging and Computer Graphics Theory and Applications: International Joint Conference, Revised Selected Papers, pp. 19–40. ISBN: 978-3-319-29971-6.
- [SLKD+17] Noeska Smit, Kai Lawonn, Annelot Kraima, Marco DeRuiter, Hessam Sokooti, Stefan Bruckner, Elmar Eisemann, and Anna Vilanova. "PelVis: Atlas-based Surgical Planning for Oncological Pelvic Surgery." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 741–750.
- [ZSLE+15] Changgong Zhang, Thomas Schultz, Kai Lawonn, Elmar Eisemann, and Anna Vilanova. "Glyph-based Comparative Visualization for Diffusion Tensor Fields." In: *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 22 (1).1 (2015), pp. 797–806.

EHRENERKLÄRUNG

Ich versichere hiermit, dass ich die vorliegende Arbeit ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe; verwendete fremde und eigene Quellen sind als solche kenntlich gemacht. Dritte haben von mir weder unmittelbar noch mittelbar geldwerte Leistungen für Arbeiten erhalten, die im Zusammenhang mit dem Inhalt der vorgelegten Habilitation stehen.

Ich habe insbesondere nicht wissentlich:

- Ergebnisse erfunden oder widersprüchliche Ergebnisse verschwiegen,
- statistische Verfahren absichtlich missbraucht, um Daten in ungerechtfertigter Weise zu interpretieren,
- fremde Ergebnisse oder Veröffentlichungen plagiiert,
- fremde Forschungsergebnisse verzerrt wiedergegeben.

Die Arbeit wurde bisher weder im Inland noch im Ausland in gleicher oder ähnlicher Form als Habilitation eingereicht und ist als Ganzes auch noch nicht veröffentlicht."

Magdeburg, 02. November 2017

Kai René Hartmut Lawonn