UNIVERSITÄT BAYREUTH

On Three-Dimensional Reconstruction

Habilitationsschrift in Informatik

der

Fakultät für Mathematik, Physik und Informatik

vorgelegt von

Dr. rer. nat.

Oleg Lobachev

aus Bayreuth

Bayreuth, 2018

On Three-Dimensional Reconstruction

Oleg Lobachev

University Bayreuth

8. Januar 2018

Vorgelegt der Fakultät für Mathematik, Physik und Informatik der Universität Bayreuth als Habilitationsschrift in Informatik.

Fachmentorrat: Prof. Dr. Michael Guthe
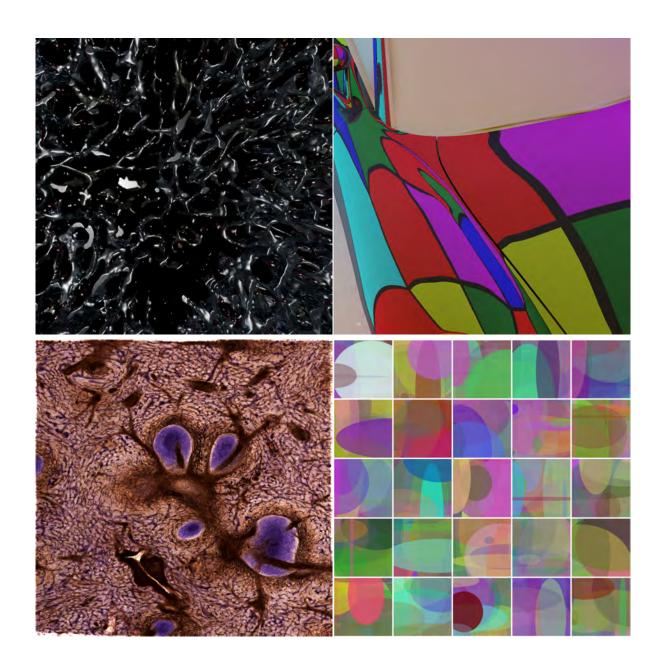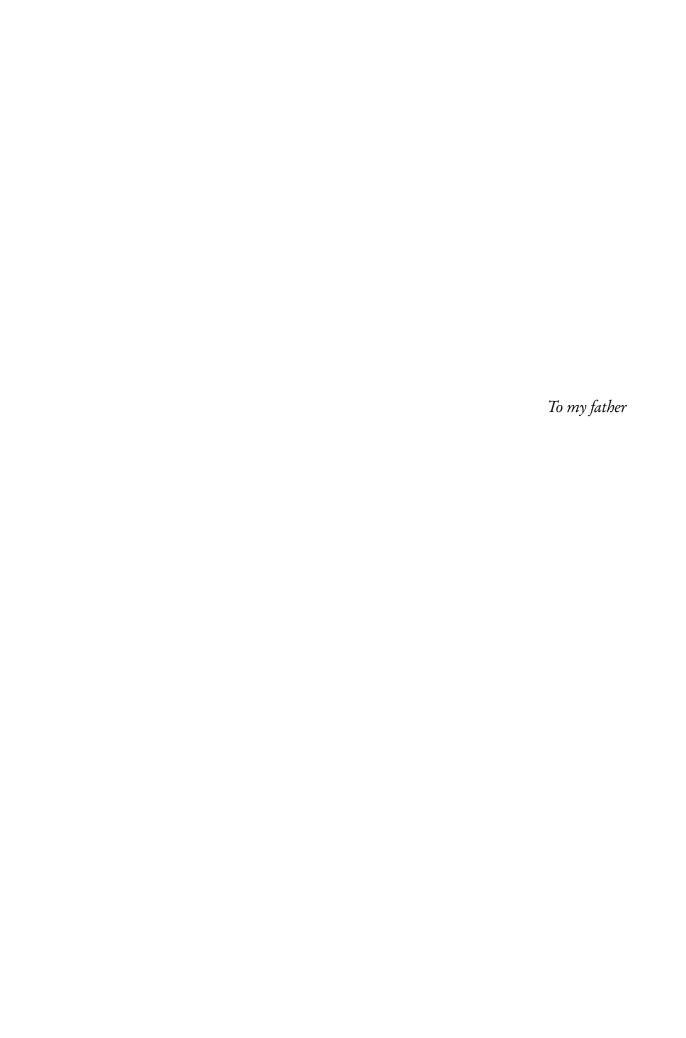Prof. Dr. Dominik Henrich
Prof. Dr. Thomas Rauber

C'est véritablement utile puisque c'est joli.

It is truly useful since it is beautiful.

A. de Saint Exupéry,
The Little Prince

To my father

CONTENTS

LIST OF FIGURES

'And what is the use of a book,'
thought Alice, 'without pictures or
conversations?'

L. Carroll, *Alice's Adventures in
Wonderland*

LIST OF TABLES

INTRODUCTION

'Begin at the beginning,' the King
said gravely, 'and go on till you
come to the end: then stop.'

L. Carroll, *Alice's Adventures in
Wonderland*

THIS WORK consolidates multiple years of experience the author gathered in 3D reconstruction from 2D data. Typical understanding of 3D reconstruction in the field of computer vision implies generating a point cloud or even a surface model from the photographs of the target object. These photographs are typically taken from multiple points of view. In a contrast, most of this thesis pursues a different approach. This work focuses on 3D reconstructions from *invasive* biological probes. Special techniques in medical research allow for generating of images at microscopic resolutions. The price for the high selectivity, sensibility, and resolution of our approach, the name is 'immunohistological staining', is that only relatively thin sections can be processed. With advances both in biological processing ('serial sections', immunohistology, details are in Appendix A) and in computer science (computer vision, medical image processing, volume and mesh filtering, Sections 0.1.1–0.1.3, Part I, Figure 0.1), the ability for 3D reconstructions progressed greatly. We can visualise microscopic blood vessel networks in human lymphatic organs, such as spleen and bone marrow at the level impossible before. Medical researchers could study form, structure, and interactions of different types of microvasculature with a greater clarity, for very long distance, and, in specific applications, for the first time ever. (As a preview, Figs. 0.2 and 0.3 (d), show 3D reconstructions of blood vessels in human bone marrow and spleen, respectively.)

In general, anatomy is a very 3D-driven study and research area. Real understanding requires a 3D perception of the subject, both in macro- and in microanatomy. Whereas there are other methods in macroanatomy, microanatomy up until recently lacked means of 3D perception. This
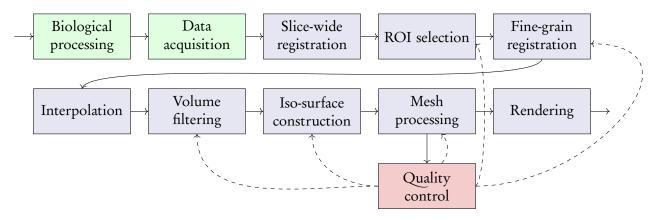
Figure 0.1: Our processing pipeline for 3D reconstruction from histological serial sections. Figs. 0.5 and 0.6 on page xxiii relate this diagram to Parts I and II.
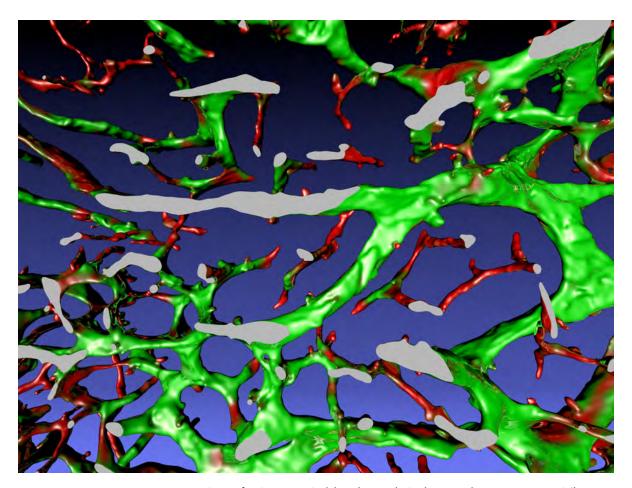
Figure 0.2: Our 3D reconstruction of microscopic blood vessels in human bone marrow. The blood vessels are colour-coded red to green according to their diameter. Grey shows the frontal surface plane, the dataset ends there. This colour annotation allows experts to discern between sinus (larger microscopic blood vessels) and capillaries. We found mostly sinus in blood production (haematopoietic) areas, and mostly capillaries in the surrounding fatty (adipose) tissue. This is R1 from Chapter 6.

thesis establishes methods for full-fledged 3D reconstructions. Such reconstructions are essential for microanatomy. They allow for the study of spatial positions and interactions of different types of cells and tissues. While we were mostly concerned with microscopic blood vessel networks, our methods are not limited to a specific organ or species. Nevertheless in this work we will be dealing with human spleen and human bone marrow tissues. We utilise virtual reality (VR). Domain experts can 'travel' through the mesh space, literally walk inside reconstructed vascular networks and thoroughly inspect them with the help of VR.

The digital processing methodology is the focus of Part I of this thesis. Part II presents the medical applications. Part III discusses further issues, Part IV concludes. In the following text, we present an overview of our digital processing pipeline (Section 0.1). Then we briefly showcase the results from Part II (Section 0.2) and discuss the contents of Part III (Section 0.3). Section 0.4 lists the publications of the author, Section 0.5 wraps up the overview of this thesis.

0.1   A 3D reconstruction pipeline for serial sections

3D reconstruction from serial sections is a quite long pipeline of methods. Initially, separate images do not correspond spatially, so first of all the sections need to form a volume. This is achieved with a registration. Then we need to interpolate and process the volume. Then we construct the mesh and process it. Highlighting important parts of the visualisation for medical researchers is also important. We invest a lot into quality control. Figure 0.1 visualises our pipeline, the below list gives details.

1. Biological processing: tissue acquisition, fixation, preparation of the probe, embedding, sectioning, staining, coverslipping.
2. Digital data acquisition: serial scanning in an optical scanning microscope.
3. Normalisation [Reinhard et al., 2001, Khan et al., 2014]. The intensity of sections varies initially due to slightly uneven thickness. In this step this is equalised.
4. Coarse registration: fitting the sections to each other (Ulrich et al. [2014b], Chapter 1).
5. Regions of interest (ROIs) selection.
6. Fine-grain registration (Lobachev et al. [2017a], Chapter 2). It is a decisive step for maintaining the connectivity of capillaries.
7. Colour processing, e.g., channel selection or colour deconvolution.
8. Interpolation to reduce anisotropy (Lobachev et al. [2017b], Chapter 3).
9. Volume filtering, e.g., a closing filter and a blur. Details are domain-specific, as Part II shows.
10. Mesh construction (Ulrich et al. [2014a], Chapter 4)
11. Mesh processing, e.g., decimation and smoothing. (Chapters 6–7) It often includes mesh colouring, e.g., separate components colouring or visualisation of shape diameter function. Details are again specific for each application.
12. Quality control (QC). If, for example, reconstruction interrupts microvessels or includes non-informative components, identify the cause and repeat from there. Chapter 5 details on quality control methodology based on virtual reality. We constantly verify our reconstructions in Part II.
13. Rendering.

The subject of Chapters 1 and 2 is registration of serial sections. Chapter 2 also provides an extensive quality assessment of our registration method. Our registration methods are innovative in the way of establishing correspondences. Our approaches are based on computer vision. We were the first to apply these methods for a 'sparse', i.e., feature-based, non-rigid registration of medical images. A scanned batch of serial sections after registration is in fact a 3D volume. These volume data exhibit strong anisotropy: typical resolution of a scanning microscope when using ×20 lens is 0.3–0.5 µm/pixel, however a typical section is 4–8 µm thick. Chapter 3 presents our interpolation approach to counter this issue. Chapter 4 improves mesh construction. Figure 0.3 shows images from the separate steps of the pipeline. For the QC in this thesis we compare the processed meshes (optionally coloured) with registered ROI images in an immersive setting using a VR head-mounted display.

Serial sections can be easily stored, can be repeatedly inspected, and allow for good penetration of the immunohistological staining solutions. While limited information contained in a single section would prevent mesoscopic overview, with modern advances in optical scanning microscopes, series of high-resolution digital images consisting of up to several hundreds of single sections can be acquired. As already mentioned, these images need to be registered, aligned to each other. Non-linear distortions of the sections pose a special problem. These distortions also need to be countered in the registration

(a)   (b)   (c)

(d)
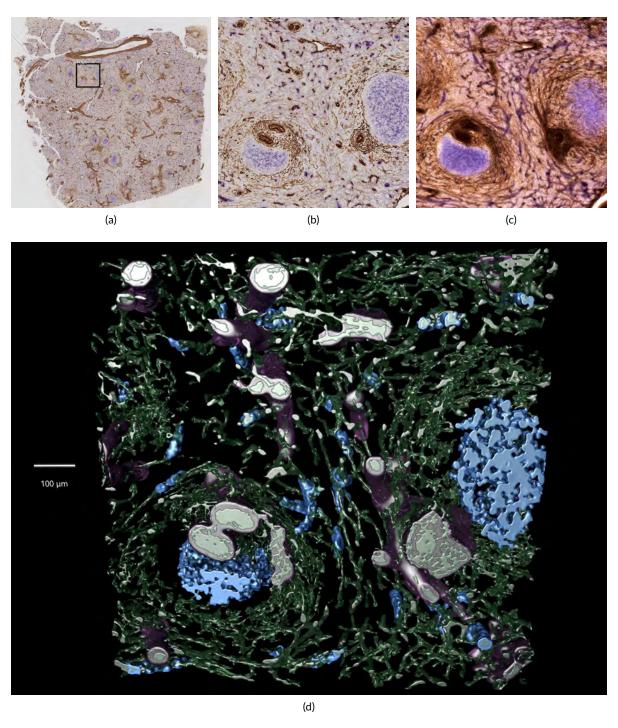
Figure 0.3: Our processing pipeline in action. A human spleen section is stained for CD34 plus SMA (both brown), and for CD271 (violet-blue). (a): An overview of a single section with a region of interest (ROI). (b): The ROI at larger magnification. (c): A volume rendering of the ROI after fine-grain registration of 24 serial sections. (d): Final meshes. The colours highlight the different functions: green shows blood vessel walls and supportive structures for capillaries. *(Continued on the next page.)*

Figure 0.3: (Continued from the previous page.) Blue represents CD271-positive cells, including special dendritic cells in follicles (large blue blobs) and capillary sheaths, the actual objects of interest. These appear as relatively small blue structures around green capillaries. Transparent raspberry rose hull shows a heuristics for arterioles, i.e., larger blood vessels.

procedure, as detailed below. Further problems such as anisotropy and methodology for quality control also ensue. Below we give an overview of the proposed solutions to these problems.

0.1.1 Registration

In contrast to most other approaches in medical registration, we do not use pixel-based methods, but utilise the classical computer vision approach: feature detection and matching. We detect features (typically using SURF) both in whole scans of serial sections—these are up to $30k \times 30k$ pixels in size—and in regions of interest. We match the features and compute a *rigid* registration using a weighted RANSAC [Ulrich et al., 2014b, Lobachev et al., 2017a]. This is, however, not enough. Sections are non-linearly distorted during the cutting procedures. To remedy this distortion, we introduce a non-linear undistortion mechanism. We express the undistortion as a movement of control points grid of a B-spline. The exact movement corresponds to a least-squares solution based on movement exhibited by radius matches of the close-by features. This is an unorthodox approach towards medical image registration. The optimisation procedure is performed iteratively and on all sections. Chapters 1 and 2 present our methods in full detail. In Chapter 2 we also iteratively match smaller features for better registration of smaller blood vessels. We were the first to use an growing hierarchy of feature sizes for the refinement of non-rigid registration.

Our presentation also includes a thorough evaluation of our registration method using multiple quality measures. Figure 0.4 visualises the distribution of SSIM [Wang et al., 2004] quality values in different registrations of bone marrow specimen. We use violin plots [Hintze and Nelson, 1998] to do so. Our method proves to be most beneficial w.r.t. this quality measure.

0.1.2 Volume and mesh processing

In our typical applications (Part II), after the registration we process volume data to obtain meshes. First of all, we *interpolate* the data to handle the severe anisotropy: typical resolutions are 0.5 μm/pixel vs. 7 μm/section. We compute dense optical flow [Farnebäck, 2003] between two consecutive slices to generate intermediate images [Lobachev et al., 2017b, Chapter 3]. With interpolation we reduce the anisotropy to $0.5 : 0.5 : 1$. We decided against fully isotropic representation as our typical regions of interest (ROIs) are squares with $2k$–$4k$ pixel sides.

Then we typically filter the volume data with a closing filter and minor Gaussian blur. With marching cubes, a corresponding mesh representation is devised. In production we use either our GPU implementation with tandem simplification [Ulrich et al., 2014a, Chapter 4] or a module from *3D Slicer* [Fedorov et al., 2012, Kikinis et al., 2014]. The meshes are further processed. In most cases we at least heal and remesh the representation using *PolyMender* [Ju, 2004], then use Taubin smoothing to obtain a nicer surface. After some further processing (e.g., removal of very small unconnected components or colouring) the mesh is ready for quality control, rendering, and evaluation by experts.

Bone marrow: SSIM
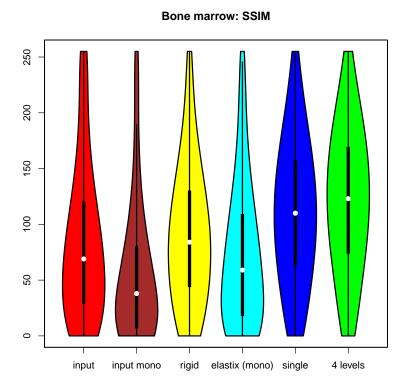
Figure 0.4: Violin plots [Hintze and Nelson, 1998] of the SSIM [Wang et al., 2004] quality measure for a registrations of bone marrow specimen. Here 'mono' means a single-channel data *elastix* [Klein et al., 2010] operates on, 'single' is our single-resolution method [Ulrich et al., 2014b, Chapter 1], '4 levels' stands for four iterations of our multi-resolution method [Lobachev et al., 2017a, Chapter 2]. The higher values are the better ones.

0.1.3   Quality control

Our 3D reconstruction pipeline features multiple parameters. The most prominent is, of course, the iso-value, but other parameters also play a role. Such parameters include registration settings, volume processing options (e.g., the settings of morphological filters), mesh filtering settings, etc. How can we be sure that our processing is correct?

Next issue: domain experts are highly trained to perceive histological stainings. A view of a stack of serial sections, i.e., a volume rendering, is a nice benefit. But very often the following question arises: 'I see something interesting in the reconstruction at this position. How does the original slice looks like at this point?'

These two problems motivated the development of our out-of-core volume rendering tool (used in Chapter 6) and also of our VR-based visualisation and quality control utility (Chapter 5). We used it for the remainder of Part II. Domain experts often had trouble with

- Pre-rendered visualisations,

- Usability of mesh viewers,
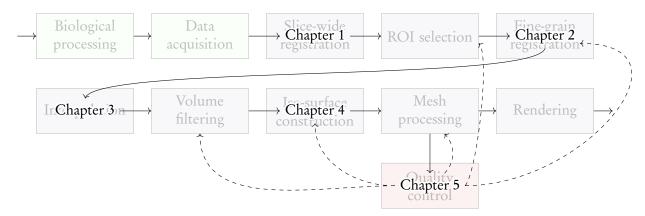
Figure 0.5: An overview of our pipeline (Fig. 0.1) with chapter numbers from Part I superimposed.
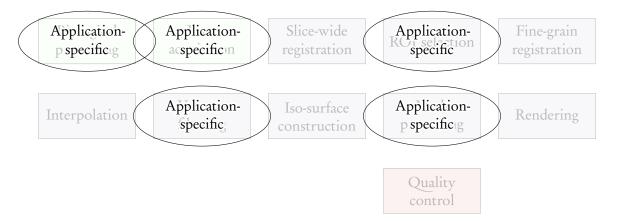
Figure 0.6: An overview of our pipeline with application-specific issues that are discussed in Part II.

- Communicating their findings to us,

- Verification and diagnosis of the reconstructions.

The work during the preparation stage of what would become Chapter 6 of this thesis emphasised the importance of a proper utility for quality control. Usability and navigation were major issues at that time. This work motivated the development of a new tool. We introduce a VR-based solution for navigation, annotation, and quality control of histological 3D reconstructions (Chapter 5). VR headset solves immediately the navigation issue, allowing the domain expert to move freely, both around and inside the reconstructions. The very intuitive navigation interface of the immersive VR solves the usability problem. The annotations are also intuitive ('use the controller to place a ball') and straightforward to implement. Our visualisation also allows the user to view both the final data and the registered serial sections. Thus a domain expert can verify the quality of our reconstructions and ensure correct diagnosis from the dataset.

0.2 Applications

Part II discusses medical research implications of our visualisations. Many of the conclusions presented in Chapters 6–7 would have been impossible without our computer science methods.

We discern larger and smaller microvessels in human bone marrow (Chapter 6). Our conclusion (drawn from 3D reconstructions) is that different types of microvessels predominantly occupy different regions. Connections between these two networks are rare. Comparing the reconstructions to the sections we found out that networks of larger blood vessels are drawn to the haematopoietic areas in the bone marrow. The smaller blood vessels are mostly associated with areas of pure adipose tissue. Such insight could have been drawn only from 3D reconstructions that were produced with our method.

Studying the the follicle surfaces in the human spleen (i.e., mostly 'white pulp', Chapter 7) we found out three distinct but interconnected capillary networks: a red pulp network, perifollicular capillaries, and an internal network in the white pulp. The perifollicular capillary network appears to give off branches to red pulp capillaries. This could not be followed further in single sections and required our reconstructions.

Summarising, our 3D reconstruction methods allowed for impossible before discoveries in medical research.

0.3 Further topics

The scope of this thesis is not limited to methodology of 3D reconstruction from serial sections and its applications in medical research. Chapters 8 and 9 present further topics.

Chapter 8 details on a very different 3D reconstruction approach, *deflectometry*. In an approach to rapidly capture surfaces of large specular objects—passenger cars—we developed a special device [Lobachev et al., 2013]. Basic idea is that a pattern is projected on the car's surface and multiple cameras capture the reflected image of the pattern. Knowing the pattern, we can discern the shape of the specular surface. The overall goal of the project was to detect damage to a car's hull e.g., from hail. Chapter 8 focuses on how to optimise camera placement in such a device. We maximise the area seen by cameras, while allowing for the capture of varying car models without reconfiguration of the device.

Chapter 9 presents a interesting application of image processing [Lobachev, 2017]. Encryption plays an ever-increasing role in the modern world. One possibility to compromise an encrypted communication is to swap a genuine key against a forged one before the communication channel is established. Chapter 9 focuses on the *visualisation* of public keys that makes them discernible by a human in a more effortless manner. Our approach enhances security because it visualises the keys directly. We are not completely dependent on the presence of hashes, but all other key visualisations are. Our visualisation also offer an aesthetic side benefit, Figure 0.7 gives a preview.

0.4 List of publications

At the time of writing the author of this thesis has coauthored 23 publications. We separate them here into three groups and mention to which chapters they correspond.

Postdoctoral, peer-reviewed publications, not part of another thesis:

1. B. S. Steiniger, C. Ulrich, M. Berthold, M. Guthe, and O. Lobachev. Capillary networks and follicular marginal zones in human spleens. three-dimensional models based on immunostained serial sections. *PLOS ONE*, 2018. doi: 10.1371/journal.pone.0191019. Accepted. (Chapter 7)
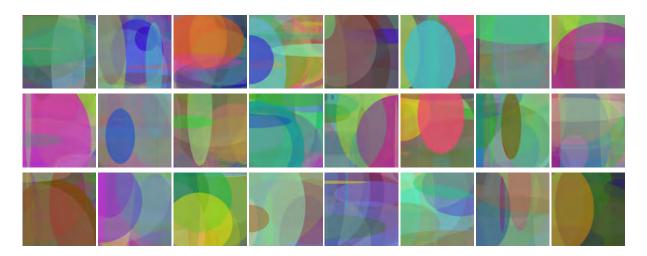
Figure 0.7: Visualisations of public keys from author's *PGP* keyring.

2. O. Lobachev. Direct visualization of cryptographic keys for enhanced security. *The Visual Computer*, 2017. ISSN 1432-2315. doi: 10.1007/s00371-017-1466-6. URL http://rdcu.be/Dtem. In press. (Chapter 9)

3. O. Lobachev, B. S. Steiniger, and M. Guthe. Compensating anisotropy in histological serial sections with optical flow-based interpolation. In I. Viola and J. Bittner, editors, *33rd Spring Conference on Computer Graphics*, SCCG '17, 2017b. In press. (Chapter 3)

4. A. Reger, J. Dumler, O. Lobachev, J. Neuberger, and R. Steinhilper. Machine allocation via pattern recognition in harmonic waves of manufacturing plants. In R. Teti, editor, *11th CIRP Conference on Intelligent Computation in Manufacturing Engineering*, ICME '17. CIRP, Elsevier, 2017.

5. B. S. Steiniger, V. Stachniss, V. Wilhelmi, A. Seiler, K. Lampp, A. Neff, M. Guthe, and O. Lobachev. Three-dimensional arrangement of human bone marrow microvessels revealed by immunohistology in undecalcified sections. *PLOS ONE*, 11(12):1–25, 12 2016. doi: 10.1371/journal.pone.0168173. (Chapter 6)

6. O. Lobachev, C. Ulrich, B. S. Steiniger, V. Wilhelmi, V. Stachniss, and M. Guthe. Feature-based multi-resolution registration of immunostained serial sections. *Medical Image Analysis*, 35: 288–302, 2017a. doi: 10.1016/j.media.2016.07.010. (Chapter 2)

7. N. Friedrich, O. Lobachev, and M. Guthe. Faking it: Simulating background blur in portrait photography using a coarse depth map estimation from a single image. In V. Skala, editor, *WSCG Short papers proceedings*, volume CSRN 2602 of *WSCG '16*, pages 15–24, 2016. ISBN 978-80-86943-58-9.

8. M. Schmidt, O. Lobachev, and M. Guthe. Coherent Metropolis light transport using speculative mutations. *Journal of WSCG*, 24(1):1–8, 2016. ISSN 1213-6964.

9. C. Ulrich, O. Lobachev, B. S. Steiniger, and M. Guthe. Imaging the vascular network of the human spleen from immunostained serial sections. In I. Viola, K. Buehler, and T. Ropinski,

editors, *Visual Computing for Biology and Medicine*. The Eurographics Association, 2014b. ISBN 978-3-905674-62-0. doi: 10.2312/vcbm.20141185. (Chapter 1)

10. M. Endler, O. Lobachev, and M. Guthe. Human-computer interaction using robust gesture recognition. In V. Skala, editor, *WSCG Posters proceedings*, WSCG '14, pages 45–48, 2014. ISBN 978-80-86943-72-5.

11. C. Ulrich, N. Grund, E. Derzapf, O. Lobachev, and M. Guthe. Parallel iso-surface extraction and simplification. In V. Skala, editor, *WSCG Communications proceedings*, WSCG '14, pages 361–368, 2014a. ISBN 978-80-86943-71-8. (Chapter 4)

12. O. Lobachev, M. Schmidt, and M. Guthe. Optimizing multiple camera positions for the deflectometric measurement of multiple varying targets. *Journal of WSCG*, 21(2):145–152, 2013. ISSN 1213-6964. (Chapter 8)

Submitted articles:

1. Martin Braun, Oleg Lobachev, and Phil Trinder. Arrows for parallel computation. Submitted 2018 to *Journal of Functional Programming*.

2. Oleg Lobachev, Moritz Berthold, Birte S. Steiniger, and Michael Guthe. Quality control of histological 3D reconstructions in virtual reality. Submitted 2017 to *IEEE Computer Graphics and Applications*. (Chapter 5)

Predoctoral publications and papers related to the PhD thesis:

1. O. Lobachev, M. Guthe, and R. Loogen. Estimating parallel performance. *Journal of Parallel and Distributed Computing*, 73(6):876–887, Elsevier, doi: 10.1016/j.jpdc.2013.01.011, 2013.

2. O. Lobachev. Parallel computation skeletons with premature termination property. In T. Schrijvers and P. Thiemann, editors, *Functional and Logic Programming*, FLOPS '12, LNCS 7294, pages 197–212, Springer-Verlag, doi: 10.1007/978-3-642-29822-6_17, 2012.

3. O. Lobachev. *Implementation and evaluation of algorithmic skeletons: Parallelisation of computer algebra algorithms*. PhD thesis. Philipps-University Marburg, 2011b.

4. O. Lobachev. *Multimodulare Arithmetik*. ISBN 978-3-8325-2881-2. Logos, Berlin, Germany, 2011a.

5. O. Lobachev and R. Loogen. Estimating parallel performance, a skeleton-based approach. In F. Loulergue, editor, *High-Level Parallel Programming and Applications*, HLPP '10, pages 25–34. ACM Press, 2010b. doi: 10.1145/1863482.1863489. ISBN 978-1-4503-0254-8.

6. O. Lobachev and R. Loogen. Implementing data parallel rational multiple-residue arithmetic in Eden. In V. P. Gerdt, W. Koepf, E. W. Mayr, and E. H. Vorozhtsov, editors, *Computer Algebra in Scientific Computing*, CASC' 10, LNCS 6244, pages 178–193. Springer-Verlag, 2010a. Extended and revised version as Technical Report bi2010-3, Philipps-Universität Marburg.

7. J. Berthold, M. Dieterle, O. Lobachev, and R. Loogen. Parallel FFT with Eden skeletons. In V. Malyshkin, editor, *Parallel Computing Technologies*, PaCT '09, LNCS 5698, pages 73–83. Springer-Verlag, doi: 10.1007/978-3-642-03275-2_8, 2009b. Extended version as Technical Report bi2009-2, Philipps-Universität Marburg.

8. J. Berthold, M. Dieterle, O. Lobachev, and R. Loogen. Distributed memory programming on many-cores – a case study using Eden divide-&-conquer skeletons. In K.-E. Großpitsch, A. Henkersdorf, S. Uhrig, T. Ungerer, and J. Hähner, editors, *Workshop on Many-Cores* at ARCS '09, pages 47–55. VDE-Verlag, 2009a.

9. O. Lobachev and R. Loogen. Towards an implementation of a computer algebra system in a functional language. In S. Autexier, J. Campbell, J. Rubio, V. Sorge, M. Suzuki, and F. Wiedijk, editors, *Intelligent Computer Mathematics*, AISC '08, LNAI 5144, pages 141–154. Springer-Verlag, doi: 10.1007/978-3-540-85110-3_12, 2008.

0.5 Byword

To gain a real understanding of microscopic anatomy, 3D reconstruction is a must. With the exception of single cells, structures in biological tissues span over multiple sections. Thus, it is impossible to reason about them from a 2D view only. This issue is even more dramatic in large, mesoscopic structures, i.e., at a scale of about 1 mm. In research questions on form and functionality of microscopic blood vessels, neither single slices nor serial sections can facilitate understanding without a full-fledged 3D reconstruction. The question of how the reconstructions are presented to domain experts matters, too. Our VR-based approach allows for better and more in-depth inspection of specimens. We developed two registration applications, an inter-section interpolation utility, a full-section annotation tool (not described here), a mesh generation and simplification framework, an out-of-core volume–mesh visualisation tool, a VR-based visualisation and annotation tool. We have also utilised some off-the-shelf volume and mesh processing applications, such as *3D Slicer* [Fedorov et al., 2012], *Fiji* [Schindelin et al., 2012], *MeshLab* [Cignoni et al., 2008], *PolyMender* [Ju, 2004]. The final renderings were created with *Cinema 4D*. Our tools allow users to prepare, perform, visualise, quality control, and annotate 3D reconstructions from histological serial sections. Our toolchain puts special emphasis on good communication with domain experts and annotations are especially helpful in that regard.

We are currently working on further and even larger-scale 3D reconstructions from serial sections, both in ROI resolution and in number of sections. Already obtained results are presented in Part II. Registration of consecutive sections with different stainings against each other would enable more versatile highlights of special kinds of cells. We aim to extend and further develop the VR-based visualisation tool from Chapter 5. Viable extensions include multi-user VR, better 3D selection, 3D sculpting, and advanced annotation options such as voice recording. Next, the author acknowledges help and support. Further on, Part I details on registration, interpolation, mesh construction, and quality control methods.

ACKNOWLEDGEMENTS

Part I

COMPUTER SCIENCE METHODS IN MEDICAL 3D RECONSTRUCTION

A mathematician, like a painter or a poet, is a maker of patterns. If his patterns are more permanent than theirs, it is because they are made with ideas.

G. H. Hardy, *A Mathematician's Apology*

FULL-SECTION REGISTRATION FOR THE IMAGING THE VASCULAR NETWORK OF THE HUMAN SPLEEN FROM IMMUNOSTAINED SERIAL SECTIONS

In which we present the coarse, section-wide registration of serial sections. (A fine-grain registration of spleen and also bone-marrow that extends this work is presented in Chapter 2. Insights are drawn from our reconstructions in Chapters 6–7.)

The spleen is one of the organs, where the micro-structure and the function on that level are not completely understood. It was for example only recently found that is has an open circulation, which distinguishes it from all other organs. Imaging the complete vascular network from the arteries to open-ended capillaries would greatly facilitate research in this area.

The structure of such tissue is best uncovered using immunohistological staining. This can however only be applied to thin tissue sections and larger structures span several slices. Due to the deformation induced when cutting the specimen, standard registration algorithms cannot be used to merge the images into a volume.

We propose a specialised matching algorithm to robustly determine corresponding regions in the images. After a rigid alignment of the scans, we use a cubic B-spline to deform and align the images. During this process we minimise the total deformation to produce as accurate results as possible.

This chapter was previously published as [Ulrich et al., 2014b], developed in a collaboration with Christine Ulrich, Birte S. Steiniger, and Michael Guthe.

Author Contributions: *Conceptualisation: OL, MG. Data curation: OL, MG, CU. Formal analysis: OL, MG. Funding acquisition: MG. Investigation: MG, OL, CU, BSS. Methodology: MG, OL. Project administration: MG. Resources: OL, MG, BSS. Software: OL, MG. Supervision: MG. Validation: OL, MG. Visualisation: OL. Writing—original draft: MG, OL. Writing—review & editing: OL, MG, CU, BSS. Corresponding author: OL. Presentation: OL.*

(a)   (b)   (c)   (d)
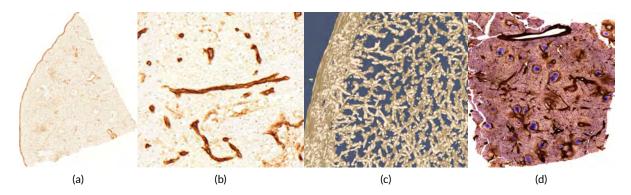
Figure 1.1: Typical input data and results of our registration method. *(Continued on the next page.)*

IMAGING blood vessels, nerves or other mesoscopic structures has several applications in biology and anatomy. Such structures may span regions of one or several millimetres and specimen of such size need to be digitised as a whole. On the other hand, small structures of less than 1 µm are also important to understand the structure and function of the tissue. Therefore, only microscopy based approaches are applicable in this context.

Our methods aim at providing insights into the course of arterial blood vessels in the human spleen. The spleenic vascular network is open: the capillaries have open ends and no connections to the venous sinuses. Thus, the blood flows through the connective tissue without any separating wall and—interestingly—without clotting. This behaviour is believed to be unique to the spleen, signifying the importance of the study of its blood vessel system. In addition, the human spleen has specialised capillary vessels (called sheathed capillaries, Chapter 7), which do not occur in any other organ. The location, shape and function of these vessels need to be clarified. Our goal is to enable answers to these questions by depicting the shape of the vessels in an aligned stack of serial sections. Registered serial sections are then used to obtain volume or mesh data of the capillary network. It does not make sense to regard a single section, as it is too thin to contain blood vessels connected at a mesoscopic dimension.

Injections or genetic modifications, which are commonly used in biological models, are not possible in humans. As a consequence, staining of specific cells to discriminate them, has to be performed after removal of the specimen. The only applicable technique that is versatile enough, is immunohistological staining. Unfortunately, staining substances do not penetrate a specimen further than a few dozen micrometers. Despite recent advances in confocal microscopy, it is therefore not possible to examine specimens with a thickness of more than approximately 50 µm using immunostaining. The only technique left is preparing 5 µm to 10 µm thick serial sections. This allows immunohistological staining, but disassembles the overall structure in mesoscopic dimension and produces deformations. As a result aligning the sections and reconstructing the volume is needed. Series with up to 400 sections are possible, amounting to a total thickness of up to 4 mm. Smaller series with up to 50 sections are simpler to produce and we thus currently work with this range. Our techniques are however applicable to these larger series as well.

So far only small details of serial sections have been used to examine blood vessels or nerves. With the recent introduction of automated slide scanning microscopes, large serial sections can be digitised with high resolution. However, the special properties of serial sections pose several challenges for such a reconstruction, which we will discuss in the following.

Immunohistological staining detects antigens in thin sections of tissue. We use the 'ABC'-system (avidin-biotinylated peroxidase complex) in our experiments. After sectioning, staining, and sealing with a cover slip, the section can be scanned for further digital processing. Figure 1.1, (a) shows a detail

Figure 1.1: (Figure on previous page.) Typical input data and results of our registration method. From left to right, (a): One section of an immunostained specimen of the human spleen using the antibody QBEnd10 (DAKO, Hamburg, no. M7165) which targets glycoprotein CD34 occurring in capillary endothelial cells, blood stem cells, and certain fibroblasts in the spleen the scan resolution is 25856 × 32000. (b): Detail of the section showing the stained capillaries. (c): Detail of the iso-surface reconstruction from the aligned serial sections. (d): Volume rendering of a further spleen immunostained for CD34, smooth muscle alpha actin and CD271 (first two targets coloured brown, the latter one coloured blue). Both (c) and (d) were aligned with our technique. Chapter 7 details on medical implications from 3D reconstruction of data shown in (a)–(c).

of a scanned immunostained specimen using diaminobenzidine (brown) only. This staining designates capillary endothelial cells, blood stem cells, and certain fibroblasts in the spleen. We use it to visualise the vascular network. Some specimen (like Figure 1.1, (d)) are also immunostained for further targets, shown in brown and blue colour. In our experiments, we utilise CD34 immunostaining as a standard.

We used a Leica SCN 400 slide scanner with an optical magnification factor of 20. With this magnification, the scanner produces images with a spatial resolution of about 0.5×0.5 µm per pixel. The section thickness is approximately 7 µm.

Aligning serial sections

After scanning the immunostained serial section, the images have to be aligned to assemble them back into a volume. This alignment process however poses several challenges.

- The sections are significantly deformed by the cutting knife of the microtome. While small deformations are tolerable for a single section, they need to be corrected for a 3D reconstruction. This means that some non-linear local deformations have to be applied to the scanned sections.

- Removing the embedding material has the consequence that no external markers can be used for the alignment. In contrast to many other applications, it is therefore not possible to use registration markers since they would be lost during staining.

- The specimen is human or animal tissue scanned at a high resolution. Therefore, the images contain highly repetitive structures—i.e., cells—and pose high demands on the processing algorithms. Among other things, this implies that gradient domain, block- and feature matching cannot be used directly.

The main contribution of this chapter is a novel, feature based registration and alignment algorithm that was specifically designed for immunostained serial sections of tissue. It solves several problems that render classical image registration techniques unsuitable in the context of serial section alignment. It is robust to highly repetitive smaller structures that otherwise lead to erroneous alignments. Using an iterative re-matching of the features, it is able to find correct correspondences for a large number of features and automatically removes outliers. It minimises the overall deformation of the specimen since we do not align to a single reference slice. Thus it reduces the influence of the cutting-induced distortions in the individual slices.

1.1   Related work

Shams et al. [2010] provide an extensive survey of parallelized algorithms for medical image registration. The methods they discuss are however all based on local difference measures and thus not directly suitable for our problem. Similarly, Hill et al. [2001] compared various approaches to register MRI and PET images with CT data.

There are also several books on medical image registration algorithms. Goshtasby [2005] gives a very broad overview of the methodology and applications of different techniques. More detailed descriptions of various approaches and a deeper mathematical background are provided by [Suru et al., 2005, Scherzer, 2006, Hajnal and Hill, 2010]. All of these books explicitly discuss non-rigid registration methods in the context of clinical imaging, i.e., CT, MRI, and PET data. An overview of bio-imaging software tools is given by Eliceiri et al. [2012].

1.1.1   Registration of serial sections

Only few methods to align serial sections of tissue exist. Among these are manual alignment techniques [van Krieken et al., 1985, Steiniger et al., 2003]. This approach was later improved by Steiniger et al. [2011b] using a joystick as input device.

One of the first practical approaches was the ICP algorithm [Besl and McKay, 1992]. Several semi-automatic approaches for the registration of immunostained serial sections of human capillaries in tumours targeting CD34 [Gijtenbeek et al., 2005, Gilhuis et al., 2006] were proposed. After an approximate manual pre-alignment using a rigid transformation, an affine transformation is computed, that minimises the pixel-wise normalised cross-correlation. These are however only suitable for small regions due to the pure linear correction and for small series because of the manual pre-alignment.

A fully automatic approach was proposed by Ourselin et al. [2001] that is based on block matching. Nikou et al. [2003] minimise a global energy function for auto-radiography sections, however, they operate on a very small MRT data sets. Ma et al. [2008a] developed a method to align sections of a mouse lymph node. It is based on first globally aligning a binary image generated by a foreground/background segmentation. Then a multi-resolution algorithm is used to refine the rigid alignment. A similar approach was proposed by Tanacs and Kato [2011] for MRT images. Ju et al. [2006] utilise the continuity of the biological matter to reduce 'jitter' in the registration of a section series with warp filtering.

As Steiniger et al. already noted: 'Slight distortions of the single sections during the cutting process (and other reasons) led to an irregular outline of the vessels' [Steiniger et al., 2003]. Rigid-only registration was studied in context of MRI and similar kinds of data [Roche et al., 2001, Jenkinson and Smith, 2001, Malandain et al., 2004].

1.1.2   Non-rigid registration

Cifor et al. [2011] register 2D histological scans using the smoothness assumption: the surface-to-reconstruct is a biological artefact that should not have non-continuous jumps. The elastic method is a quite popular approach [Guest et al., 2001, Chappelow et al., 2011]. Gefen et al. [2003] use a 3D wavelet to elastically transform histological images. Bajcsy and Kovačič [1989] developed a multi-resolution elastic algorithm to register CT data sets. Saalfeld et al. [2012] use a 2D elastic triangulation of Hookean springs for block matching. Wirtz et al. [2005] use higher-order image derivatives. Thirion [1998] used a diffusion model based on thermodynamic concepts to register two images. Bagci et al. [2012] review the modern state of the art approach based on multi-resolution methods. All multi-resolution methods however require data containing large structures which are not present in our specimen.

Rueckert et al. [1999] developed a non-rigid registration of 3D models, especially breast MRTs. It is similar to our work as it also uses a B-spline based deformation. The matching is however based on the mutual information of the images [Wells et al., 1996, Peng et al., 2005] and thus requires clear edges and larger uniform areas in the images. This approach was further improved to utilise multi-level B-splines [Schnabel et al., 2001]. This allows simulating a non-uniform control point grid, but the similarity measure is still not suitable for our context. Compared to [Xie and Farin, 2004], our work features orders of magnitude larger data sets with comparable execution time.

[Chui and Rangarajan, 2003] developed an algorithm for non-rigid registration of medical data sets, especially CT/MRT data. They also include a survey on different registration methods in their work. Their method circumvents the matching problem by reformulating the registration as point cloud matching of the features. This idea however does not work in our case, because the point cloud

formed by the features is a dense random point set without much structural information. In the context of CT data the results of the EMPIRE10 challenge are of interest [Murphy et al., 2011].

Wan et al. [2013] also use feature detection based on SURF and thin-plate splines for the non-rigid registration of a set of synthetic brain MRI images. This method, however, requires an exhaustive training step for feature selection and thus cannot be used for 'one-shot' specimen in the context of medical research. Kim et al. [1997] utilise thin-plate splines for alignment of auto-radiography images. They used a low-resolution undistorted video feed as guide images. Again, large structures are required for this method. In contrast to Auer et al. [2005] we operate on much larger images— complete histological scans at ×20 magnification. Auer et al. combine a multi-resolution rigid registration with thin-plate spline non-rigid step. We use one-step rigid transform and a B-spline-based non-rigid step. Chapter 2 presents the extension to multiple non-rigid step in a multi-resolution manner. Song et al. [2013] register a set of immunostained sections with alternated different staining.

Most of existing methods rely on a reference slice to which all others are deformed. As shown by Bagci and Bai [2010], selecting the 'best' reference is crucial for the final result. In a contrast, we perform an optimisation step to minimise the energy of the non-linear distortions over all sections in the series.

1.1.3 General image registration

There is an extensive research on image registration and its various applications. Early works [Lucas and Kanade, 1981] date back more than 30 years ago. Friston et al. [1995] describe the image registration process as a least-squares problem where the differences between images are minimised. Several surveys on image registration methods exist [Brown, 1992, Zitová and Flusser, 2003]. Image registration, as a process of overlaying different images of the same or similar objects typically consists of four phases: feature detection, feature matching, estimation of the transformation, application of the transformation. In our case we first estimate the rigid transformation and then the non-rigid one.

1.1.4 Feature detection and matching

Numerous image feature detection algorithms exist. Recent algorithms detect scale and rotation invariant features. Among them are the SIFT features [Lowe, 2004], SURF is their improvement in term of run-time efficiency [Bay et al., 2008]. BRISK [Leutenegger et al., 2011] aims to further reduce run-time compared to SURF without sacrificing the detection quality. Due to the low run-time and memory consumption, we chose to use this detector in our work. Another variant is the FREAK detector [Alahi et al., 2012] that is inspired by human vision and also targets at low processing times. Further feature detectors can be found in the survey of Heinly et al. [2012]. The n-SIFT [Cheung and Hamarneh, 2009] is a generalisation of SIFT to volume and otherwise multidimensional data. n-SIFT can be used in the registration of multiple volume data sets against each other.

After the features are detected, they need to be matched in pairs between two images. Typically, each feature is matched with the most similar one in the other image. Despite pruning strategies, like only keeping pairs where the matching in the other direction is the same, some mismatches almost always remain. These are usually filtered out using the random sample consensus algorithm [Fischler and Bolles, 1981] or one of its variants.

1.2 Overview

The overall registration and alignment algorithm is split into three steps. First, we determine suitable features in all scanned sections. Then we use those features to perform an initial pairwise rigid alignment of all sections. This alignment is used to iteratively deform the images and to identify the feature correspondences. The non-rigid alignment is performed in such a way, that the overall deformation of the images is minimised. Minimising the deformation leaves an overall rigid transformation of the images as only degree of freedom. We fix this by preventing rotation and translation of the first slice. This means that the remaining ones are rigidly aligned to the first, but all of them are deformed. This does *not* mean that the first slide is used as a reference for the distortion of all other slides. Finally, the deformed images are generated, forming an aligned stack. Figure 1.2 shows the overall work flow of our method.
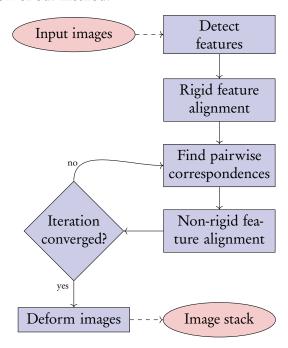
Figure 1.2: Overall work flow of our non-rigid registration and alignment algorithm.

1.3 Initial rigid alignment

The purpose of the initial rigid alignment step is to roughly register the slices and to facilitate the non-rigid alignment. First, we need to choose adequate feature candidates and then perform a pairwise matching. Based on these matches that still contain a high number of outliers, we then calculate a robust alignment.

1.3.1 Choosing candidate features

The first step we perform for each input image, is finding all feature points. The sections have been normalised prior to feature detection. For this purpose, we use the BRISK feature detector [Leutenegger et al., 2011]. Unfortunately, we cannot directly use *all* the detected features and compute pairs by searching for the most similar one in the reference image as shown in Figure 1.3. The reason

for this is that many small features cannot be reliably used because they originate from single cells or other common structures.
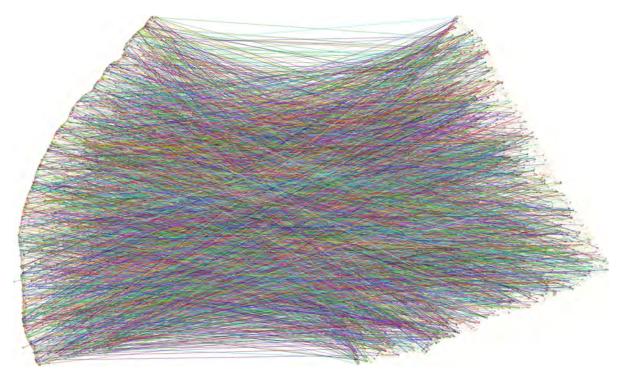
Figure 1.3: Using standard image registration results in almost random matches due to ambiguous features. Only each 10th match is shown to reduce visual clutter.

On the other hand, larger features—about an order of magnitude larger than a single cell—are usually much less ambiguous. Therefore, we select the N_{select} features in each input image i whose descriptors have the largest diameters. Then we compute pairs by selecting the most similar feature in the reference image $i-1$. The possible candidates can again be reduced by exploiting the fact that there is not much scaling between structures in adjacent sections. Thus we only consider features with a similar descriptor diameter.

$$\frac{1}{1+\varepsilon_l} \leq \frac{d_{tgt}}{d_{ref}} \leq 1+\varepsilon_l, \tag{1.1}$$

where d_{tgt} and d_{ref} are the target and reference diameters respectively, and the threshold ε_l the maximum local stretch. Figure 1.4 shows the matches found when only using larger features. Although the quality improved, the matches still contain a high amount of outliers and we have to perform a robust alignment procedure.

1.3.2 Pairwise alignment

After finding initial pairs, we perform a random sample consensus (RANSAC) [Fischler and Bolles, 1981] based registration. A rigid transformation in 2D is defined by a pair of points and its movement. The distance between points in a pair cannot change in a rigid transformation—in contrast to an affine one. In practice, this distance between two corresponding point pairs on neighbouring sections can
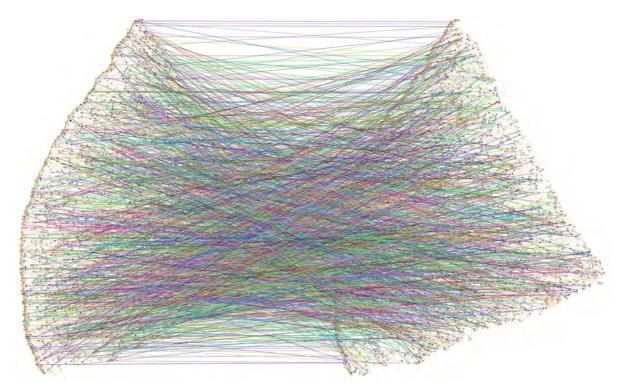
Figure 1.4: The relative number of false matches is reduced using larger features only. Nevertheless, still many outliers exist. As before, we show only each 10ᵗʰ match to reduce visual clutter.

only change due to deformation of the specimen. In this manner, we can restrict the candidate set. We choose two random matches a and b, and only use this pair, if

$$\frac{1}{1+\varepsilon_g} \leq \frac{\left\|P_{a,tgt}-P_{b,tgt}\right\|}{\left\|P_{a,ref}-P_{b,ref}\right\|} \leq 1+\varepsilon_g \tag{1.2}$$

$$\left\|P_{a,tgt}-P_{b,tgt}\right\| \geq d_{min}\cdot s_{image} \tag{1.3}$$

$$\left\|P_{a,ref}-P_{b,ref}\right\| \geq d_{min}\cdot s_{image}, \tag{1.4}$$

where $P_{a,tgt}$ is the position of a in the target image, etc., ε_g the maximum global stretch, with $\varepsilon_g \leq \varepsilon_l$, s_{image} the image size, and d_{min} the minimum relative distance between the points. Based on the selected pair, we compute a candidate transformation $\mathbf{R}_i$.

In addition to quickly rejecting futile candidates, we also augment computing the consensus score with the matching score of the pairs that are determined as inliers for the given transformation. Instead of simply counting all matches j with

$$\left\|\mathbf{R}_i P_{j,tgt}-\mathbf{R}_{i-1}P_{j,ref}\right\| \leq s_{image}\cdot\varepsilon_g,$$

we accumulate

$$\frac{1}{1+w_{match}\delta_j}, \tag{1.5}$$

where δ_j is the difference of the descriptors and w_{match} a weighting constant.

After finding a good transformation using the RANSAC algorithm, we again compute the inlier set. From this, we compute a rigid transformation with a weighted least squares error. The weighting is the same that we already used to compute the consensus score (see Equation 1.5). Figure 1.5 shows the found rigid transformation between two sections together with the matches in the consensus set.

Figure 1.5: Rigid registration between two sections with the consensus set of matching features. Notice that the upper area cannot be aligned due to local deformations.

1.4 Non-rigid alignment

The non-rigid alignment is an iterative process because each deformation of the images may change the inlier/outlier classification and thus the matching pairs. We first start by finding a corresponding feature in the reference images. Then we perform a least-squares deformation trying to align all found pairs. Finally, we compute the feature positions after deformation and check if the matches are still valid or if we can find a better correspondence that is now within the local search radius. If all pairs remain—or if we have performed a maximum number of iterations it_{max}—the registration terminates. We do not select a reference section and deform all sections to match it, but rely on global optimisation to minimise the energy of the deformation.

1.4.1 Finding pairwise matches

In each iteration, we first start finding new matches. For each feature j in image i we search a corresponding feature k in image $i-1$ with the most similar descriptor under the constraint that

$$\left\| T_i\left(P_{j,i}\right) - T_{i-1}\left(P_{k,i-1}\right) \right\| \leq s_{image} \cdot \varepsilon_g, \tag{1.6}$$

where T_i is the non-rigid transformation of image i. Note that in the first iteration, this is the rigid transformation $\mathbf{R}_i$. Figure 1.6 shows the matches found in the first iteration.
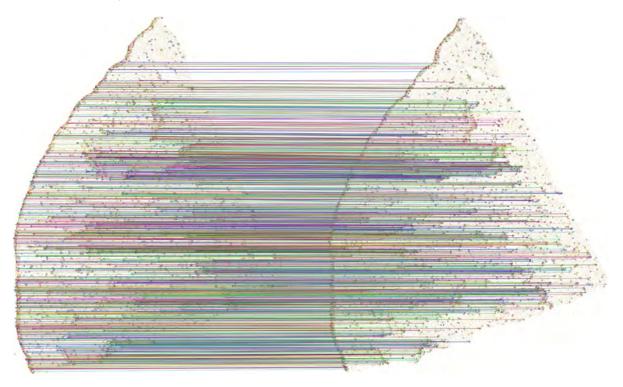
Figure 1.6: Pairwise matches found in the first iteration, where only high confidence matches— i.e., weight ≥ 0.2—are shown. Note that despite the increased density of matches, there are still very sparsely covered regions where large deformations are present, e.g., at the top of the image.

1.4.2 Computing image deformations

The non-linear deformation is based on a uniform bi-cubic B-spline over a 9×9 control point grid with three additional control points outside the image. The control point net initially covers the first image. The aligned images are constructed in the parameter domain of the B-spline that is scaled such that it has the same size as the first input image. This means that we can simply evaluate the B-spline to compute the source position in the deformed image for each pixel of the output.

We register each image to the previous one, starting with the registration of the second to the first. For each matching feature pair j, k in images i and $i - 1$ respectively, we have the following constraint of the control points:

$$\sum_{l=0}^{14}\sum_{m=0}^{14} B_l(u)B_m(v)\alpha^i_{l,m} = \mathbf{R}_i P_{j,i}, \tag{1.7}$$

$$u_i = u\left(T_i(P_{j,i})\right) \quad v_i = v\left(T_i(P_{j,i})\right), \tag{1.8}$$

where u and v are the mappings of a pixel in the aligned image to the parameter values, and $\alpha^i_{l,m}$ are the control points of the B-spline of image i. Again, we use the same confidence weighting described in Equation 1.5.

In addition, we want to prevent large deformations and smooth areas where we have no matches. We achieve this by adding a constraint for each edge of the control meshes:

$$\left(\alpha_{l+1,m}^{i} - \alpha_{l,m}^{i}\right) w_\alpha = \begin{pmatrix} d_x \\ 0 \end{pmatrix} \tag{1.9}$$

$$\left(\alpha_{l,m+1}^{i} - \alpha_{l,m}^{i}\right) w_\alpha = \begin{pmatrix} 0 \\ d_y \end{pmatrix}, \tag{1.10}$$

where d_x and d_y are the distances between two neighbouring control points in the output image in x and y direction and w_α is the smoothing weight. Note that this also induces a penalty for rotations but this is unproblematic because we already performed a global rigid alignment.

Combining all of the constraints leads to a sparse linear equation system. We solve this using the conjugate gradient method [Press et al., 2007] for sparse matrices. As initial value we use the control points from the previous iteration. For the first iteration, we simply set

$$\alpha_{l,m}^{i} = \left((l-3)d_x, (m-3)d_y\right)^T$$

which is a non-deformed control grid.

An overall deformation of all images has a degree of freedom. We use it to minimise the mean square deformation of all images after each registration pass. Moving each control point of all images by the same offset does not change the matching. We can thus minimise the mean square deformation by moving the control points, such that the average of each control point l, m over all images (including the first one) lies on a grid. In other words, we constrain the control points such that:

$$\sum_{i=1}^{N_{img}} \alpha_{l,m}^{i} = \begin{pmatrix} (l-3)d_x \\ (m-3)d_y \end{pmatrix}. \tag{1.11}$$

This uses the remaining degrees of freedom in the alignment to force the 'average deformation' to be zero. Previous approaches simply fixed a reference slice such that it was not deformed at all.

1.4.3   Deforming the images

To deform the images, we use the Open CV [Bradski and Kaehler, 2008] function remap. First, we compute the position in the source image for each pixel of the output image and store it in an array. Instead of evaluating the complete B-spline function for each pixel, we can exploit the fact that the image is a uniform grid. For each row of control points, we evaluate the 1D cubic B-spline for the x-coordinate of each pixel. Then we only need to compute 1D cubic B-splines instead of bi-cubic B-splines per pixel. As the number of image lines is significantly higher than the number of control point rows we only need to compute 4 basis functions instead of 8 per pixel and accumulate 4 instead of 16 control points. In total, this roughly saves 70% of the computation time. Then, the computed pixel location array is used to generate the aligned image. For the reconstruction, we use bi-linear filtering.

1.5   Results

Figure 1.7 shows an overview of the complete data set of the human spleen specimen consisting of 24 sections with a resolution of 25856×32000 pixels each. The thickness of each section is approximately

7 µm, so the digitised specimen has a total thickness of about 170 µm. The image is generated using a volume renderer that combines all sections into a single thick specimen. Such view would be impossible without the registration.

Figure 1.8 shows two interesting areas of the spleen specimen. On the left, it shows the shape of a larger artery, dividing up further until the capillaries, and the right side shows the blood supply of a follicle. Single sections do not contain longer portions of blood vessel, as Figure 1.1, (b) visually depicts. On a larger scale, the shape of the capillary network in the spleen can be reconstructed from our volume data using standard visualisation algorithms like marching cubes (c). We use this registration method as a preparation step throughout this thesis. A further refined registration (Chapter 2) was then used to obtain medical research results in Part II.

1.5.1   Alignment

The complete registration and alignment pipeline is implemented in C++ using the Open CV library [Bradski and Kaehler, 2008] version 2.4.8. The run-times are measured on an Intel Core i7 with 16 GB of RAM running Windows 7 64 bit and the executable is also compiled as 64 bit. We use the following constants for the alignment:

- Number of features selected per image $N_{select} = 20,000$

- Minimum RANSAC edge length $d_{min} = 0.1$

- Maximum local deformation $\varepsilon_l = 0.01$

- Maximum global deformation $\varepsilon_g = 0.005$

- Descriptor weight for matches $w_{match} = 10$

- Maximum number of non-rigid iterations $it_{max} = 20$

- Smoothing weight $w_\alpha = 0.2$

The final matches found for the alignment of two sections are shown in Figure 1.9. In contrast to Figure 1.6, matching feature pairs have been found for the whole area.

In our experiments we also encountered teared or otherwise damaged sections. Because of our global optimisation approach it was possible to register the batch with sufficient accuracy. The final aligned data set is shown in Figure 1.1, (d). Although the damage was not 'repaired' by our method, it is robust enough to align all parts.

Figure 1.10 compares our registration with the standard ICP algorithm [Besl and McKay, 1992]. We used the implementation of Pomerleau et al. [2013]. Starting from the same feature points from the original images as our approach, the ICP algorithm converges to a local minimum (top left). Using a manual pre-alignment, it obtains better results (top right). Due to the ambiguities, caused by the repetitive structures in human tissue, a consecutive non-rigid alignment, e.g., using the method of Chui and Rangarajan [2003] cannot correctly align the scans. In contrast to that, our method achieves much better results on the same data with the rigid stage only (bottom left). The final non-rigid stage (bottom right) drastically improves the alignment. Chapter 2 extends the current method and compares it to *elastix*.

Table 1.1 shows the timings for the different steps of the algorithm. The initial number of features detected per image is between $430k$ and $450k$, so we keep the 4–5% largest ones. Most of the computation time is spent in the iterative non-rigid alignment procedure (the CGNE solver calls)
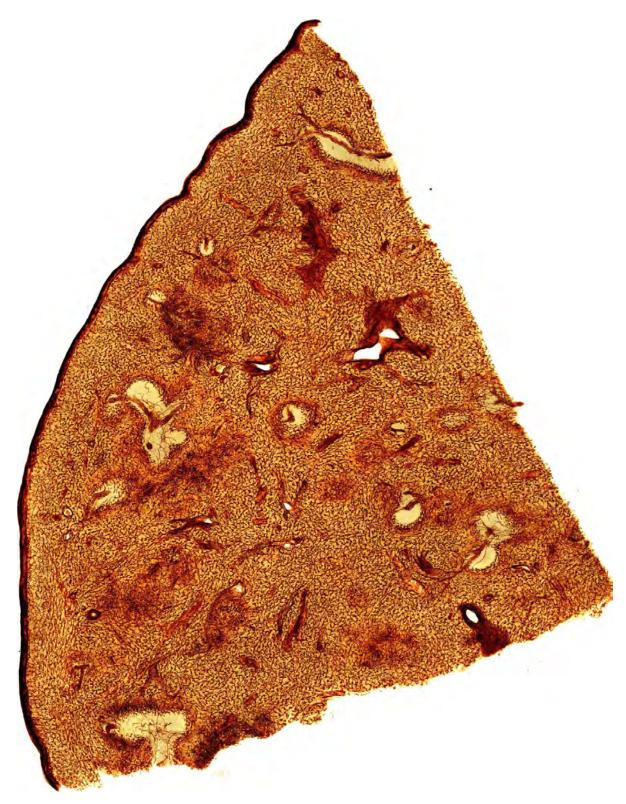
Figure 1.7: Overview of the complete aligned data set of the human spleen specimen. Capillary and larger vessel endothelia and certain fibroblasts are coloured brown, visualising the vascular network. Chapter 7 presents medical research results drawn from the specimen shown.

(a)

(b)

(c)

(d)
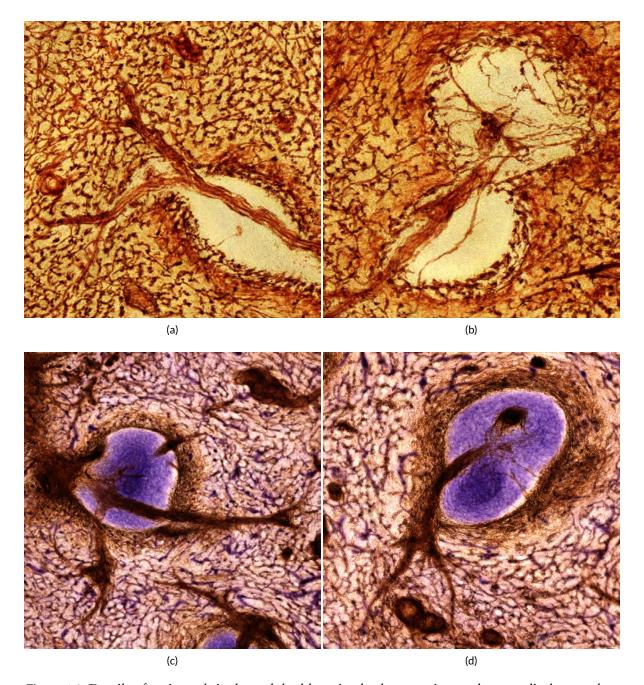
Figure 1.8: Details of registered single- and double-stained spleen specimens show medical research insights. (a): A larger artery develops into a network of smaller vessels. The shape of the blood vessel network in human spleens has been unexplored before, except for the fact that it has open ends. (b) shows the blood supply of a lymphatic follicle. (c)–(d) show similar structures in the two-colour specimen. Such in-depth views were impossible so far. In Part II we draw insights from 3D reconstructions of registered series.
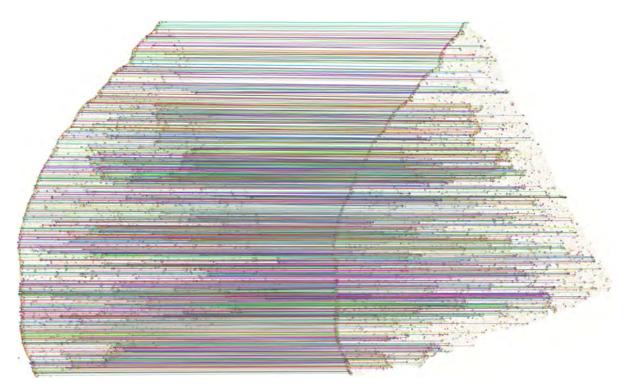
Figure 1.9: Final matches found during alignment of two sections, where only high confidence matches—i.e., weight ≥ 0.2—are shown. The whole images are densely covered with correctly matched feature pairs.

and in the actual deformation of the images. Note that using thin plate splines, this step would take significantly longer due to the high number of matched features. The run times for both data sets were identical since the images have the same size and both contain 24 sections. Due to the pairwise matching and solving a sparse linear system for the non-rigid alignment, our method scales linearly in the number of sections.

time (mm:ss)	feature detection	rigid transf.	non-rig. transf.	image recon.
image^a/pair^b	0:42^a	0:08^b	1:03^b	1:17^a
total	16:58	3:20	24:02	30:35

Table 1.1: Execution times for the spleen data sets with 24 sections. The first line shows the average time for various phases of our method that is required per image or pair. The second line shows total time for each phase.

1.6 Conclusion and limitations

After aligning stacks of immunostained spleen serial sections, we have produced volume renderings of two different specimen, as well as a 3D-model of the vessels (Figures 1.1 and 1.7). These enable the domain experts to analyse and discuss the shape and function of capillaries and arterial blood vessels

<div align="center">(a)</div>

<div align="center">(b)</div>

<div align="center">(c)</div>

<div align="center">(d)</div>
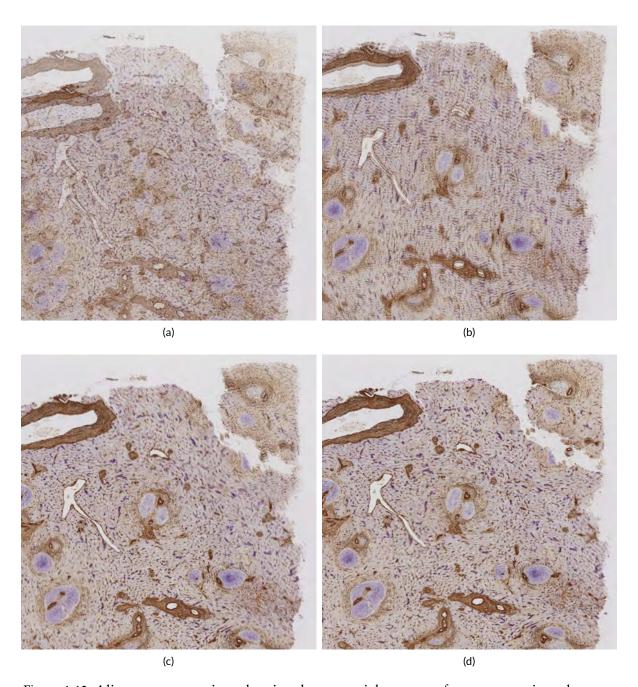
Figure 1.10: Alignment comparison showing the upper right corner of two consecutive spleen sections stained for CD34 and smooth muscle alpha actin in brown, and for CD271 in blue. (a): The result of the ICP algorithm applied to the original data. (b): ICP with initial manual positioning of the sections. (d): The result of our rigid registration. (d): The result of our non-rigid registration (final output of our method).

in the spleen at a level impossible without our method. Some details of interesting areas are shown in Figure 1.8. Chapters 6–7 discuss the medical results in-depth.

To achieve better understanding of histological data, we have presented a method that is able to

robustly register and align immunostained serial sections. This is an especially challenging task due to the highly repetitive cellular structures in the tissue. Our algorithm tolerates relatively large local and global deformations. We measured a global deformation of up to 2% of the image diagonal and a local stretch of up to 5%.

The most time consuming part of our method is the reconstruction of the deformed images. Therefore, we plan to implement this step on the GPU in the future. A second candidate for a GPU implementation is the conjugate gradient solver which needs most of the time in the non-rigid alignment phase.

Although we have developed and used our approach for human spleen specimens, it is much more widely applicable. Due to the use of feature detection to find similar structures in adjacent sections, our approach is only limited to tissue that contains at least some special histological structures. At a mesoscopic scale this is however no limitation, as blood vessels always provide sufficient landmarks in human organs. We have applied this method to human bone marrow specimen in Chapter 6.

While our method appears to be robust even to teared or damaged sections, explicitly repairing the damage before alignment would have several advantages. The alignment will be more precise close to the fracture as there will be consistent feature pairs on both sides. On the other hand, the reconstructed volume will also not be fractured any more. We therefore want to develop a method to find and 'repair' the fractures in the scanned sections.

Acknowledgements

FEATURE-BASED MULTI-RESOLUTION REGISTRATION OF IMMUNOSTAINED SERIAL SECTIONS

In which we discuss the fine-grain registration of serial sections and its quality. We are able to register the serial sections up to smallest capillaries, despite a lot of self-similarity. Supplementary material has been uploaded to Zenodo and is available under doi: 10.5281/zenodo.164404 and https://gdv-server.inf.uni-bayreuth.de/gdvcloud/index.php/s/NnSov0O65n9Gp01.

This is joint work [Lobachev et al., 2017a] with Christine Ulrich, Birte S. Steiniger, Verena Wilhelmi, Vitus Stachniss, and Michael Guthe. **Author Contributions:** *Conceptualisation: MG, OL. Data curation: OL, BSS, VW. Formal analysis: OL. Funding acquisition: MG. Investigation: OL, BSS, VW, VS, CU. Methodology: OL. Project administration: OL. Resources: OL, BSS, MG, VS, CU. Software: OL. Supervision: MG. Validation: OL, BSS, MG. Visualisation: OL. Writing—original draft: OL. Writing—review & editing: MG, BSS, VW, VS, OL. Corresponding author: OL.*

The form and exact function of the blood vessel network in some human organs, like spleen and bone marrow, are still open research questions in medicine. In this chapter we propose a method to register the immunohistological stainings of serial sections of spleen and bone marrow specimens to enable the visualisation and visual inspection of blood vessels. As these vary much in calibre, from mesoscopic (millimetre-range) to microscopic (few micrometers, comparable to a single erythrocyte), we need to utilise a multi-resolution approach.

Our method is fully automatic; it is based on feature detection and sparse matching. We utilise a rigid alignment and then a non-rigid deformation, iteratively dealing with increasingly smaller features. Our tool pipeline can already deal with series of complete scans at extremely high resolution, up to 620 megapixels (Chapter 1). The improvement presented increases the range of represented details up to smallest capillaries. This chapter provides details on the multi-resolution non-rigid registration approach we use. Our application is novel in the way the alignment and subsequent deformations are computed (using features, i.e., 'sparse'). The deformations are based on all images in the stack ('global').

We also present volume renderings and a 3D reconstruction of the vascular network in human spleen and bone marrow on a level not possible before. Our registration allows for easy tracking of even smallest blood vessels possible, thus granting experts a better comprehension.

A quantitative evaluation of our method and related state of the art approaches with seven different quality measures shows the efficiency of our method. We also provide z-profiles and enlarged volume renderings from three different registrations for visual inspection.

Figure 2.1: Teaser of this chapter: before and after our registration.

N THE mesoscopic scale, structures like blood vessels or nerves may span over millimetres in biological specimens and thus pose a significant problem for medical imaging. In this context large dimensions are as important as fine details. Such details can only be revealed by microscopy. We propose a 'sparse' registration process to solve a fundamental problem concerning 3D reconstruction in microscopic anatomy on a mesoscopic scale. This problem is relevant in immunohistological investigations necessitating serial sections.

Immunostaining

Immunohistological staining (immunostaining) is the technique of choice to discriminate cell types in human organ specimens. However, the staining does not penetrate specimens thicker than 50 µm. Such a thickness is also actually too large for transmitted light microscopy. With respect to human specimens, two choices are possible at this point: confocal microscopy or serial sections. Confocal microscopy has the limitations that it cannot be applied to thicker sections and that the sections cannot be stored for longer periods of time. The thickness of a single serial section is 5 µm to 10 µm. These sections can be easily stored, can be repeatedly inspected, and allow good penetration of the immunohistological staining solutions. The resolution of a common scanning microscope is 0.3–0.5 µm/pixel. The only drawback is that the limited information in a single section prevents a mesoscopic overview. However, with the modern advances in conventional scanning microscopes, series of high-resolution digital images consisting of up to several hundreds of single sections can be acquired. These images need to be registered. The special problem of this registration is the non-linear distortion of the sections. These distortions are inevitably associated with the sectioning process, because the embedding materials for the organ specimens are relatively soft. This happens in every immunohistological serial section and necessitates the matching of large and small landmarks. We use 'sparse', off-the-shelf feature detection combined with novel matching and warping in non-rigid transform instead of image-based methods and minimise the global distortion over the whole series of sections.

This work focuses on multi-resolution non-rigid registration of serial sections. 'Multi-resolution' applies to the varying *size* of the features used, not to an image pyramid. The non-rigid registration is only a part of the larger production pipeline consisting of image acquisition, registration (see also previous chapter), segmentation, interpolation (Chapter 3), 3D reconstruction (Chapter 4), visualisation (details in Part II), and quality control (Chapter 5).

Contributions

We present an approach for a 'sparse', feature detection-based multi-resolution registration of microscopic images. Feature-based rigid registration does not pose a significant problem, but it is not sufficient for our purposes, as already detailed earlier [Ulrich et al., 2014b, Chapter 1]. The same publication introduced a section-wide spline-based non-rigid registration that allowed to successfully register complete serial sections at their maximal available resolution.

However, further processing stages uncovered a central problem of 'simple' non-rigid registration. Previously, we used only large features. Now, we iterate over multiple feature sizes in the non-rigid registration part. We register the regions of interest (ROIs) after a coarse non-rigid registration using the method Chapter 1 used for human spleen. We stress that using the 'first stage' method alone, or even iterating it again on a ROI yields worse results than our new method. Section 2.4 details on this.

This chapter presents the iterated method consisting of feature detection, matching, filtering, and 'warping'. The latter is the process of adapting the B-spline control point grid for actual image
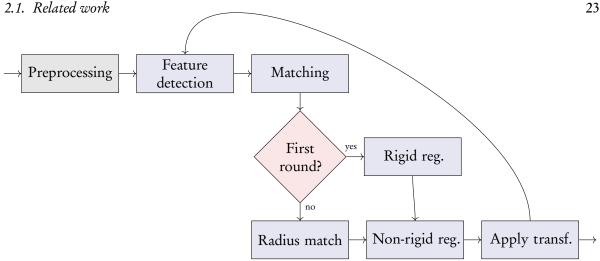
Figure 2.2: The general outline of our application. The loop is the focus of this paper: the non-rigid alignment. We do not discuss preprocessing.

(un-)distortion based on similarity of features with close spatial position. This is done at once for all images in the stack. The final step, applying the computed transformation to the actual images, is quite straightforward and is not regarded here in detail. We show the final result of our imaging pipeline in Fig. 2.10 and also in Chapters 6–7. We compare our approach to the state of the art [Ulrich et al., 2014b, Klein et al., 2010] using visual methods such as volume renderings and z-stacks (Figs. 2.8–2.12). Further, we quantitatively evaluate our method and related approaches using seven quality measures (Figs. 2.13–2.16).

Investigating micro-vessel density and distribution is relevant in a wide range of experimental therapies, especially in tumour treatment. Fully-automatic image registration on mesoscopic and microscopic scale is instrumental in such research. The main contribution of our paper is the registration based on a pyramid of decreasing sizes of features with a gradually refined control point grid. The 'search radius' for spacial location of the features also declines. The result is the non-rigid registration process that fits not only the large features, but also smaller ones. Our larger features are still too small to be the large landmarks, required for, e.g., the method of Bagci et al. [2012]. Smaller features represent smaller blood vessels and capillaries, or other microscopic-scale entities of the serial sections. Our registration and undistortion of serial sections allow to answer current research questions such as:

- 'What is the branching pattern of larger blood vessels?'

- 'Where is the capillary network localised and do regions without capillaries exist?'

- 'Do special cells exist in the vicinity of the capillaries?'

2.1   Related work

Early work on image registration and its applications dates back more than 30 years [Lucas and Kanade, 1981]. Friston et al. [1995] describe the image registration process as a least-squares problem where the differences between images are minimised. A survey on early methods was made by Brown [1992]. A more recent one by Zitová and Flusser [2003] includes more advanced image registration methods.
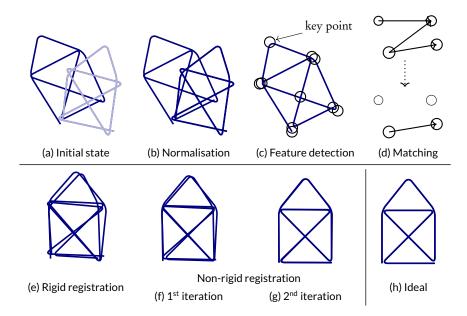
Figure 2.3: Illustrating the steps of our method. We show overlays of two simple synthetic sketches after each step of our pipeline for steps (a), (b), (e), (f), (g), and symbolise per image actions for steps (c) and (d). h) shows the ideal overlay. This chapter focuses on steps yielding (f) and (g).

Shams et al. [2010] survey methods for medical registration based on local differences. Unfortunately, these methods are not applicable for our data (cf. Figs. 2.5–2.7) that feature a lot of self-similarity. Eliceiri et al. [2012] provide an overview of bioimaging software. Sotiras et al. [2013] present a recent overview of deformable medical registration methods.

Medical registration is a large research topic. It can be applied to serial sections (as in our case, see Sections 2.1.2 and 2.1.3 for related work), MRI, CT, ultrasound, and other kinds of data (Section 2.1.4). However, one property of our method radically differs from most other methods. We use feature detection and then operate on the detected key points (i.e., in a 'sparse' manner) and not on original images. Only after the transformation is computed we apply it to actual images. Few methods [Can et al., 2002, Urschler et al., 2006, Ruiz et al., 2009, Han, 2010, Saalfeld et al., 2010, Ulrich et al., 2014b] use feature detection for medical registration (see Section 2.1.5). Some methods [Likar and Pernuš, 2001, Nikou et al., 2003, Saalfeld et al., 2010] compute the transformations not pair-wise, but for the whole data set ('globally') as we do, but in a different manner. Typically, a global affine transformation is computed first, then local (typically, pair-wise) as-rigid-as-possible or non-rigid improvements are introduced. Our method computes a pair-wise rigid registration, then global, non-rigid distortions are introduced to counter the sectioning-introduced distortions. These are, however, limited by their magnitude. To our knowledge, only Ulrich et al. [2014b] (i.e., Chapter 1) pursue a similar approach. We detail on differences between this publication and our method in Section 2.4.1. Note that 'global' here does not mean a global optimisation (finding a global maximum, not merely a local one), but processing the whole data set at once.

2.1.1   Feature detection

The general feature-based registration pipeline consists of feature detection and description, feature matching and pruning mismatches. Current feature detectors provide scale and rotation invariant features. SIFT [Lowe, 2004] is a well-known method, SURF [Bay et al., 2008] is similar to SIFT, it

uses faster approximations. Recent developments include FREAK [Alahi et al., 2012] and KAZE [Alcantarilla et al., 2012]. Heinly et al. [2012] survey the feature detectors. The *n*-SIFT [Cheung and Hamarneh, 2009] generalises SIFT to multidimensional data, this does not concern our application. We utilise SURF and remove mismatches using a random sample consensus algorithm [Fischler and Bolles, 1981].

2.1.2   Registration of serial sections

Serial sections have been aligned manually [van Krieken et al., 1985, Steiniger et al., 2003]. There is an approach to utilise computer input for better precision [Steiniger et al., 2011b], but the decision procedure is still carried out by an expert.

One of the first practical computer-based methods is the ICP algorithm [Besl and McKay, 1992]. The ICP method requires an approximate manual pre-alignment using a rigid transformation, then an affine transformation is computed, that minimises the pixel-wise normalised cross-correlation. Such methods are only suitable for small regions due to the purely linear correction and only for small series because of the manual pre-alignment. Semi-automatic approaches for the registration of immunostained serial sections of human capillaries in tumours [Gijtenbeek et al., 2005, Gilhuis et al., 2006] have been proposed. These methods feature low resolution of the registered images and have the problem of accumulating errors from the sequential transforms in further images.

Ourselin et al. [2001] propose a fully automatic method. They use block matching similar to MPEG [Le Gall, 1991] to estimate a rigid transformation between slices. A method to align sections of a mouse lymph node [Ma et al., 2008a] is based on first globally aligning a binary image generated by a foreground/background segmentation. Then a multi-resolution algorithm is used to refine the rigid alignment. In contrast, we perform a feature-based multi-resolution registration. Ju et al. [2006] utilise the continuity of the biological matter for registration using warp filtering. Xu et al. [2015] perform affine registration of serial sections from mice. They automatically extract landmarks for special cells based on staining and shape, but do not use a generic feature detection as we do. Further, we utilise a non-rigid registration.

In immunostained sections, the distortions introduced during the sectioning procedures make the rigid-only registration impractical [Steiniger et al., 2003].

2.1.3   Non-rigid registration of serial sections

Lester and Arridge [1999] provide an overview of non-rigid registration methods before 1999. Cifor et al. [2011] register 2D histological scans using the smoothness assumption: the surface to be reconstructed belongs to a biological specimen and should not change abruptly. The elastic method is a quite popular approach [Guest et al., 2001, Crum et al., 2004]. In Sections 2.4.2 and 2.4.3 we compare our approach to *elastix* [Klein et al., 2010]. Further methods include 2D elastic triangulation of Hookean springs [Saalfeld et al., 2012] and making use of higher-order image derivatives [Wirtz et al., 2005].

Some methods [Likar and Pernuš, 2001, Song et al., 2014, Peng et al., 2014] register multiple *different* stainings in non-rigid manner. Each series of our slices is stained utilising the same staining procedure.

Lorenz et al. [2012] extend the method of Rueckert et al. [1999] and apply it to microscopy imaging. Their method utilises a multi-resolution approach, but does not use feature detection. Further, they acquire images of a rat lung with intravital microscopy and observe motion artefacts from breathing, heartbeat, etc. We do, however, investigate *ex vivo* human organ specimens, the

distortion in our case is only introduced by the sectioning procedure. Two-photon microscopy and similar methods are not possible w.r.t. humans.

Many existing methods rely on a reference section to which all others are deformed, so that selecting the 'best' reference is crucial for the final result [Bagci and Bai, 2010]. No manual search for such a reference is needed in our approach.

2.1.4   Registration of other medical data

Methods used in registration of MRI, CT, or ultrasound data might also be useful for serial sections, but methodological differences need to be kept mind.

Rigid-only registration is well studied in the context of MRI and similar kinds of data [Roche et al., 2001, Jenkinson and Smith, 2001]. An automatic method by Nikou et al. [2003] minimises a global energy function for auto-radiography of sections, however, it operates on MRT and auto-radiography data sets with very low resolution. The elastic method was also applied to multi-modal registrations, e.g., Chappelow et al. [2011].

Bagci et al. [2012] review the multi-resolution methods. These, however, require data containing large structures which are not present in our specimens. We use the preceding rigid transformation based on correspondence of thousands of automatically detected features (i.e., sparse 'interesting image fragments') to obtain rough correspondence of two consecutive sections.

Thirion [1998] use a diffusion model based on thermodynamic concepts to register two images. Rueckert et al. [1999] developed a non-rigid registration of 3D models, especially breast MRTs. It is similar to our work as it also uses a B-spline based deformation. The matching of Rueckert et al. is based on the mutual information contained in adjacent images [Wells et al., 1996] and thus requires clear edges and larger uniform areas in the images, which is not given in our case. Their method has been further improved to utilise multi-level B-splines [Schnabel et al., 2001]. It was also applied to serial sections [Lorenz et al., 2012]. Xie and Farin [2004] also use B-splines, but their method is not feature detection-based and their input is human brain MRI data. Methods using thin-plate splines [Wan et al., 2013] also require large structures in the image pairs.

In contrast to the solutions mentioned above our problem is special, because the shape of the vascular network is only clear after the registration is complete. The slices contain *many* single sections of blood vessels, which are connected in a manner not evident from viewing a single large microscopical slide.

2.1.5   Feature-based registration of medical data

We are aware only of three methods [Ulrich et al., 2014b, Ruiz et al., 2009, Nagara et al., 2017] that use feature detection for the registration of serial sections in optical microscopy. Our method is sufficiently different from all of them, as discussed below. Sotiras et al. [2013] provide more information on applications of feature detection in the context of MRI and CT imaging. Some works, e.g., Can et al. [2002], utilise feature detection in registration of retina images.

The pipeline from the previous chapter [Ulrich et al., 2014b] is similar to current one. But a multi-resolution method is not employed there. Ulrich et al. [2014b] only investigated human spleen specimens only, we now register both human spleen and human bone marrow specimens. We show that smaller features are crucial for a successful registration of fine details. We detail on differences between Ulrich et al. (Chapter 1) and current method in Sections 2.4.1 and 2.4.3.

Ruiz et al. [2009] use feature detection, but only for the initial rigid registration. The non-rigid registration in their work is performed based on FFT-based similarity measures and increasing the

matched window. They register serial sections on the GPU basing on normalised cross-correlation. This is radically different from our approach, as we remain sparse: we also use detected image features ('key points') for non-rigid transformation and iterate on decreasing feature size (i.e., increasing number of features) for further steps.

Nagara et al. [2017] use a micro-CT image as a guide for the registration of serial sections. They also use feature detection for the initial registration, but resort to Markov random fields for the non-rigid phase. They scale the sections by the factor 100 in order to work with full sections. The largest section size tried with our current method was approx. 4 times larger than in Nagara et al. [2017]. We initially registered full serial sections at their original resolution in Chapter 1, there we worked with up to $30k$ pixels per side, whereas Nagara et al. [2017] were limited to a maximum of 2097 pixels per side. Nagara et al. [2017] appeared multiple months after our method [Lobachev et al., 2017a] was published.

Han [2010] [see also Murphy et al., 2011] combines 3D-SURF descriptors with a mutual information approach for lung CT registration. We operate only on SURF features and register serial sections. Urschler et al. [2006] estimate a thin-plate spline interpolation using SIFT features. Saalfeld et al. [2010] register a large electron microscopy data set as-rigidly-as-possible using features. We register optical microscopy images and also employ non-rigid registration. *Fiji* [Schindelin et al., 2012] includes an implementation of feature detection routines.

Some methods, e.g., Bağci and Bai [2008], Cifor et al. [2013], perform the registration (typically of MRI or CT data) in *feature space*, i.e., they separate multiple important properties of the images, like intensity or local phase. The registration then happens in the higher-dimensional space where all these properties are present. This is different from our method. We use *feature detection and matching*; characteristic properties of an image in a relatively small region are combined to a single key point. A feature detector generates hundreds of thousands of key points from our data. (We use words 'key point,' 'feature,' and 'landmark' as synonyms.) Key points are later matched based on their characteristics. We have implemented better rigid and non-rigid feature matching methods as a part of this work. But our main contribution is the 'sparse', 'global' multi-resolution feature detection-based medical registration and the way we implement the movement of distortion grids (Sections 2.2, 2.3.1).

2.2 Our method

2.2.1 Outline of the registration process

The digitised serial sections are pre-processed to enable better registration. We normalise the sections [Reinhard et al., 2001, Khan et al., 2014]. Depending on the specimen we may also use only a particular image channel for feature detection. The actual registration procedure can be divided into three phases (Fig. 2.2):

- Feature detection

- Rigid registration

- Non-rigid registration.

All steps except the second one are repeated allowing for an increasingly smaller size of feature points. We first perform feature detection and matching, automatically targeting important landmarks in the images. From corresponding matches a rough pair-wise rigid registration is computed. Because

of the sectioning-induced distortion, a rigid-only registration is not sufficient [Steiniger et al., 2003, Ulrich et al., 2014b]. In our method, smaller feature points are processed after the registration of the larger key points, thus providing a registration that is viable not only on a mesoscopic, but also on a microscopic scale. In other words, our multi-resolution method varies the minimal *size* of feature points used in the non-rigid registration step, not the size of input images or degree of Gaussian blur as other methods. We call each iteration of non-rigid registration a 'level' and typically use two or four levels. We showcase each step of the application flowchart on a synthetic example in Fig. 2.3.

Iterated matching of the feature key points in a region for every pair of consecutive sections is a prerequisite for our non-rigid transform. A local control point of the B-spline—a juncture of a rectilinear grid superimposed on the image—is then shifted to the direction where more key points can be matched. We express this procedure as an optimisation problem and solve it numerically. In the next outer loop the resolution of the grid is increased, and the match distance and the size of key points involved decrease. The inner loop runs over the set of serial sections, the optimisation of control point positions is an iterative method.

In the non-rigid phase the key points are shifted in the direction of a better match. This motion is carried out by the surrounding B-spline control points. The choice where to move the control point is expressed as a highly sparse least-squares problem solved with a Krylov subspace method, LSMR [Fong and Saunders, 2011] in our case.

2.2.2   Formal derivation

We present here a high-level overview of our method, then feature detection, description, and matching. RANSAC-based rigid transformation and non-rigid transformation follow.

Basic idea.   A typical two-image registration method searches for a function φ such that for two images, a fixed image $f := f_1$ and a moving image f_2, the difference between the moved image $m := \varphi(f_2)$ and f is minimal w.r.t. a distance function d:

$$d(f, m) := \|m - f\| + r,$$

with some regularisation term r and a norm $\|\cdot\|$. Basically, such φ is sought for that minimises $d(f_1, \varphi(f_2))$.

All our input images are distorted during the sectioning procedure, i.e., before the capture. Taking aside our rigid registration for now—details follow below—our non-rigid registration operates 'globally', at once, on a set of images $f_1, \ldots, f_k$. For these, a family of functions $\psi_1, \ldots, \psi_k$ is sought for, such that for a distance function d the terms

$$d\left(\psi_i(f_i), \psi_{i+1}(f_{i+1})\right) \quad \text{for all } i \in \{1, \ldots, k-1\} \tag{2.1}$$

are minimised. We search for all ψ_i at once. With distance function d we aim for minimal movement energy of functions ψ_i for all i.

To obtain this, we

1. Limit the energy and radius of the distortions,

2. Define the distance function in terms of the number of feature points matched between two regions of inputs,

3. Express the movements as a set of least-squares problems that are iterated over all n images multiple times to ensure a good *overall* solution is found.

The distance function d is implicitly defined through the weights of local matches. The direction of the (un-)distortion is guided by increased number and weight of matches. Fig. B.1 (on page 198) details on brute force matching vs. RANSAC vs. non-rigid transformation.

Another innovation of our method is that we use only feature detection and matching, i.e., a 'sparse' process to implement the above idea, not the 'dense' one as in most methods.

Features. In the first step of our method, we detect SURF features [Bay et al., 2008] in each pair of input images. To do so, we compute sum images $I_{\sum}(\cdot) = \sum_i \sum_j I(i,j)$ for an input image I. This enables the SURF detector to calculate the sum of intensities in an arbitrary rectangular region in an especially efficient manner. SURF would basically detect a feature in a region based on the values of the determinant of the Hessian matrix of an image. In the actual computation the determinant is approximated with three box filters. Filters of different size are applied to the same image to yield a pyramid of so-called octaves.

The feature description in SURF is based on the distribution of first-order Haar wavelets—an approximation to gradients used in SIFT descriptors. Let d_x be a wavelet response over a region in the direction of the x axis. SURF intensity structure for a region is then $(\sum d_x, \sum d_y, \sum |d_x|, \sum |d_y|)$.

A correspondence of two similar features is called a match. 'Wrong' not-fitting matches are eliminated using rigid registration (RANSAC) and radius matching phases (Fig. 2.2). We match the descriptions using a bi-directional brute-force matcher. Experiments have shown that such a simple matcher is sufficient for our goals.

RANSAC. For the RANSAC-computed rigid transformation we look for matches in a certain search radius around the current guess for the transform. The *weight w* of a match m is defined as

$$w(m) = \frac{1}{1 + \omega \Delta_m} \tag{2.2}$$

with a weighting constant ω and difference between the descriptors Δ_m. We use weights in our implementation of RANSAC [Fischler and Bolles, 1981] for rigid matches and in a radius match. Aside from this our rigid phase is straightforward.

Non-rigid transformation. We extend the formalism from Chapter 1 to multiple non-rigid transformations. Let $i-1$ and i be two consecutive images and j a feature in $i-1$ paired with some feature in i. Then, we can constrain a control point $\alpha^i_{l,m}$ of a B-spline of the image i with

$$\sum_{l=0}^{N} \sum_{m=0}^{N} B_l(u) B_m(v) \alpha^i_{l,m} = \psi_i^{(0)} \circ \cdots \circ \psi_i^{(n-1)}(P_{j,i}), \tag{2.3}$$

$$u_i = u\left(\psi_i^{(n)}(P_{j,i})\right) \quad v_i = v\left(\psi_i^{(n)}(P_{j,i})\right). \tag{2.4}$$

Above, P is the current pixel, $\psi_i^{(0)}$ is the i^{th} rigid transformation, further $\psi_i^{(x)}$ for $1 \leq x \leq n-1$ are the previous $n-1$ non-rigid transformations, $\psi_i^{(n)}$ is the current transformation. Further, N is the current number of control points in each dimension, and u and v map pixel to parameter values. We prevent large transformations using constraints on the edges of control meshes:

$$\begin{aligned}
\left(\alpha^i_{l+1,m} - \alpha^i_{l,m}\right) w_\alpha &= \begin{pmatrix} \delta_x & 0 \end{pmatrix}^T \\
\left(\alpha^i_{l,m+1} - \alpha^i_{l,m}\right) w_\alpha &= \begin{pmatrix} 0 & \delta_y \end{pmatrix}^T,
\end{aligned} \tag{2.5}$$

Table 2.1: Registration methods regarded in this chapter. In our typical pipeline we use A and then B. Fig. 2.4 shows that using A two times is inferior. Sections 2.4.1 and 2.4.3 also discuss the difference between the methods of Chapter 1 and this chapter. Sections 2.4.2 and 2.4.3 discuss using *elastix* (C) as a fine-grain method.

	Registration	Sparse	Resolutions	Method
A	coarse-grain	yes	single	Chapter 1, Ulrich et al. [2014b]
B	fine-grain	yes	multi	this chapter, Lobachev et al. [2017a]
C	fine-grain	no	multi	Klein et al. [2010]

where δ_x and δ_y are the distances between two neighbouring control points in the output image in both directions and w_α is the smoothing weight. This induces a penalty for rotations but we have performed a global rigid alignment before.

We iterate multiple times through the solutions of the individual optimisation problems for image pairs to ensure convergence over all sections. The starting value of the iteration is a non-distorted grid of control points.

The single step of the iteration for a single image pair is a large sparse linear system. The dimensions of the matrix are basically the number of key points and the total number of control points in the grid, viz. Eq. (2.3). In other words we solve

$$\mathbf{A}\hat{x} = b_x \qquad \mathbf{A}\hat{y} = b_y, \tag{2.6}$$

where $\hat{x}$ and $\hat{y}$ are the positions of the control points, b_x and b_y are the x and y coordinates of the key points. The matrix $\mathbf{A}$ is highly sparse, merely 0.1%–1.5% of the entries are non zero.

2.2.3 Pipeline

In our method we use a coarse-grain registration (Chapter 1) first on the whole series of sections, typically in their full resolution. Then the ROI is selected and registered again using a fine-grain registration. Table 2.1 gives an overview.
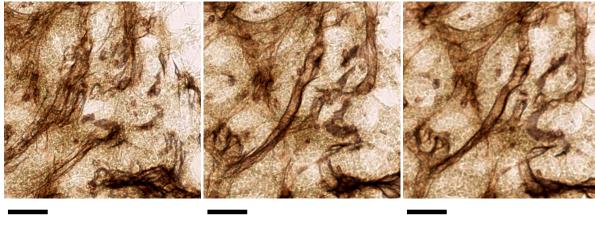
Experiments showed that for the fine-grain registration the method of Ulrich et al. [2014b] was not quite effective (Sections 2.4.1, 2.4.3). The other option for fine-grain registration, not based on feature detection at all, is *elastix* [Klein et al., 2010]. We compare our method to it in Sections 2.4.2–2.4.3.

Fig 2.4 shows a detail of the bone marrow data set (Section 2.3.2) registered with coarse-grain approach ('A' in terms of Table 2.1), using the method of Ulrich et al. for both coarse-grain and fine registration ('A–A'), and with our new approach ('A–B'). Small capillaries benefit from our method as seen in Fig. 2.4, (c), and Fig. B.7. These images show bone marrow specimen. We compare the methods 'A–A' (b) and 'A–B' (c) in detail in Sections 2.4.1 and 2.4.3. Figs. B.2 and B.3 discuss methods 'A,' 'A–A,' and 'A–B' w.r.t. a double stained spleen specimen. In Section 2.3.3 we also experiment with registration of the whole slice using our method (i.e., only 'B,' see Fig. 2.8).

2.3 Results

First we state the parameters of the registration we used to obtain the results below. Then we present the specimens that served as inputs to our method. In Section 2.3.3 we present volume renderings of those specimens.

(a) Coarse registration, 'A'   (b) Fine-grain approach, 'A–A'   (c) Fine-grain 'A–B', our method

Figure 2.4: Comparing the combination of different methods. Stack of 30 human bone marrow sections single-stained for CD34 and CD141 (both in brown, Section 2.3.2 provides details). Chapter 6 discusses the medical research findings from these data. Here we show a detail of a registered ROI, scale bar is 50 μm, see also Fig. B.7. A coarse registration (a) can be improved by a repeated application of the same method [Ulrich et al., 2014b] to ROI (b). We also align smaller key points with our method to yield a difference (c). The vessel contour is blurry in (b) and more precise in the result of our registration (c).

We show the whole specimens (Figs. 2.5–2.7) and registration results of regions of interest as volume renderings and as horizontal *z*-profiles, i.e., as a side views of the overlaid stack of images. Proper registration is absolutely mandatory to trace the vascular network in 30 serial sections.

2.3.1 Setup

We utilise the rather conservative SURF [Bay et al., 2008] feature detector, implemented as a part of Open CV[1]. A *direct* match of all-against-all key points is neither feasible nor produces good results. We rematch the key point pairs based on a weighted implementation of radius match with consecutive least-squares optimisation from Section 2.2.2.

For this chapter we implemented a modern LSMR solver [Fong and Saunders, 2011], using C++ and *Eigen*[2]. It drastically reduces the computation time spent in an individual optimisation step for identical inputs compared to a CG on normal equations solver. This improvement makes our non-rigid matching possible on an off-the-shelf personal computer despite significantly larger amounts of key points: $\approx 4\%$ in Ulrich et al. [2014b] vs. up to $\approx 50\%$ of all found key points in our case. We perform the non-rigid matching step multiple times, with increasing number of key points. We do not attempt to match *all* key points, we also remove too small ones.

We used the following fine-tuning parameters: the radius match distance was 0.5% of the largest image dimension; in our method it repeatedly scales down linearly with decreasing minimal key point size. The four-level process used key point sizes 100, 50, 20, and 10. We roughly halved the minimal size of a key point in each level. The minimal key point size of 100 was used in the coarse

---

[1]http://opencv.org/, we used the version 3.0.0-beta.
[2]http://eigen.tuxfamily.org/

(a) Overview   (b) A ROI   (c) A ROI
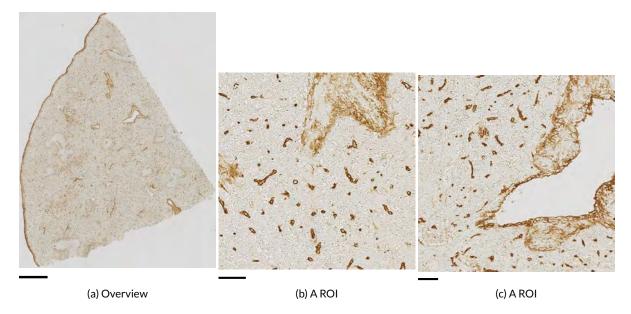
Figure 2.5: An overview of the specimens used: spleen, single stained. Left image shows the complete scan, the two additional images show typical ROI.
(a)–(c): Human spleen section single-stained for CD34 (brown colour). The CD34 molecule is present in membranes of endothelial cells lining capillaries and large vessels, in certain fibroblasts in the connective tissue surrounding large vessels, in so-termed splenic trabeculae and in the splenic capsule. (c) shows a trabecular vein.
Scale bar in (a) is 1 mm, scale bars in (b) and (c) are 50 µm.

single-level registration. Key point sizes less than 10 mostly represent single smaller cells that are not of interest for our registration. The sizes of key points used in a two-level process were 100 and 20. This represents a compromise between required large size for the first level and limited decrease of the key point size. The weighting constraint ω was 10. The control point grid had 9×9 entries (plus 3 knots on each edge); both dimensions were doubled in each iteration.

2.3.2   Specimens

We tested our approach on three series of serial sections comprising

- 'Spleen, single staining', Fig. 2.5 (a–c), 24 serial sections of a human spleen specimen, embedded in paraffin. The sections were immunostained [Steiniger et al., 2007] using an antibody which targets the membrane glycoprotein CD34 occurring in most blood vessel endothelial cells, blood stem cells, and certain fibroblasts in the spleen. This immunostaining produces an intense brown colour. The resolution is 22k × 20k pixel. Chapter 7 presents findings drawn from this section series.

- 'Spleen, double staining', Fig. 2.6 (a–c), 24 serial sections of a human spleen specimen, embedded in paraffin. The sections were immunostained for CD34-positive cells (brown), smooth muscle alpha actin (brown) and CD271-positive cells (violet-blue). Smooth muscle alpha actin is present in smooth muscle cells of arteries and arterioles, and in certain fibroblasts. CD271 is a membrane glycoprotein of fibroblast-related cells in certain regions of the spleen [Steiniger

(a) Overview   (b) A ROI   (c) A ROI
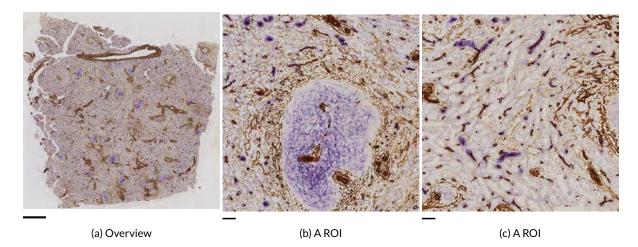
Figure 2.6: An overview of the specimens used: spleen, double stained. The left image shows the complete scan, the two additional images show typical ROI.

(a)–(c) Human spleen section double-stained for the cell membrane molecules CD34, smooth muscle alpha-actin (SMA, both in brown colour) and CD271 (blue). (b): A so-termed splenic follicle contains CD271$^+$ specialised fibroblasts (blue) and shows arteries and arterioles with thick SMA$^+$ walls and accompanying fibroblasts (brown). Branched SMA$^+$ fibroblasts are also present at the surface of the follicle. Capillary endothelia surrounding the follicle are also stained brown. (c): The tissue between splenic follicles contains capillaries (brown) with special sheaths (blue). Some fibroblasts between the capillaries are light blue.

Scale bar in (a) is 1 mm, scale bars in (b) and (c) are 50 μm.

(a) Overview   (b) A ROI   (c) A ROI
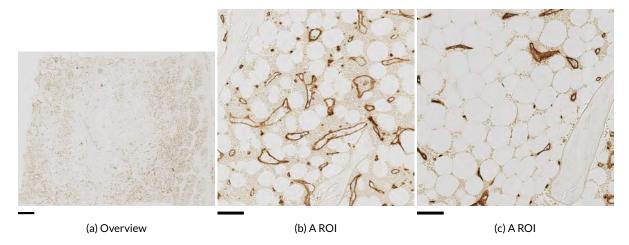
Figure 2.7: An overview of the specimens used: bone marrow. The left image shows the complete scan, the two additional images show typical ROI.

(a)–(c) Undecalcified bone sections from the human iliac crest single-stained for CD34 and CD141 (both stained brown). Endothelial cells in capillaries and large bone-marrow specific microvessels, so-termed sinuses, are visible. (b) Sinuses and capillaries are only present in bone marrow areas where blood formation occurs. (c) Areas where fatty tissue prevails, show capillaries, but no sinuses.

Scale bar in (a) is 1 mm, in (b) and (c) 100 μm.

et al., 2014a,b]. Staining for CD271 is also present in certain cells of follicles (see the large blue blob in Fig. 2.6 b)). These are not the actual focus of interest here, but due to their size the blue colour inside of follicles plays a large role when using larger features only, e.g., in a single-resolution registration. The resolution is 18.5k × 19.5k pixel.

- 'Bone marrow', Fig. 2.7 (a–c), 30 serial sections of undecalcified human bone marrow tissue from the iliac crest, embedded in methacrylate [Steiniger et al., 2013, 2016, Chapter 6]. They are immunostained for the glycoproteins CD34 and CD141 occurring in blood vessel endothelial cells. The immunostaining produces a dark brown colour in the blood vessel walls, but some other structures, most notably, single stem cells also gain some hue. The image resolution is 23k × 27k pixel.

Three basic types of a blood vessel section are possible: perpendicular sections cut the vessel in a close to orthogonal plane—the normal of the plane is close to the vessel axis. This yields small dots (collapsed capillaries) or circles (larger blood vessels) in a section. In longitudinal and tangential sections the plane is parallel to the long axis of the vessel. In longitudinal section also the interior of the vessel is in the cutting plane. This results in a single (capillary) or a double line (sinus, a larger blood vessel). A tangential section is similar to longitudinal, but the inner space of the blood vessel is not hit by the cutting plane, we see only the wall. This corresponds to a line of a possibly lighter intensity and varying width.

The objects resulting from blood vessel sections neighbour larger structures that can be more easily used for alignment. Larger blood vessels produce key points that are large enough to be processed by the method published previously (Chapter 1). Increasing the number of key points does not help, because there are too many *large* ones. Decreasing the maximal feature size disturbs the rigid registration before it has an effect. To tackle this problem, we rigidly register the sections using relatively few larger key points and then gradually increase the number of key points for the non-rigid registration.

The noise from the acquisition-induced contrast is especially well visible in Figs. 2.5, b), c), 2.6, c), and 2.7, b) and produces a large amount of smaller-sized features with both large self-similarity and very chaotic matches between the sections.

2.3.3   Volume renderings

Figure 2.8 shows the volume renderings of our registration of a spleen specimen after single staining (Fig. 2.5). We observe a very faithful matching of small capillaries. Here, we did not re-register the ROI after a coarse registration, but registered the whole image using our method and then just cropped out the ROI shown. The full registration is in Fig. B.4. All other registrations below were produced using the 'A–B' scheme from Section 2.2.3.

Fig. 2.9 shows a volume rendering of double staining of a spleen ROI based on our registration. Fig. 2.10 shows our final result: a rendering of a mesh consisting of segmented blood vessels from the same specimen. In both figures we see the network of smaller capillaries, in Fig. 2.9 these are partially covered by blue-stained special capillary sheath cells. Further, observe the follicles (large blue blobs) and their very special blood vessels: there are some large blood vessels going through the follicles, a much denser capillary network *around* the follicles (esp. well visible in Fig. 2.10), but no CD34-positive capillary endothelial cells *inside* the follicles. This corresponds to the results of others [Kusumi et al., 2015]. The histological reason for this fact needs to be investigated in more detail.
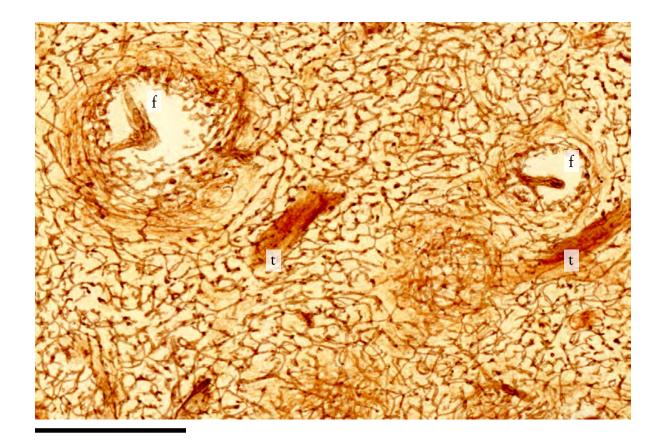
Figure 2.8: Volume rendering of spleen, single-stained for CD34. We overlay 24 images after 4 levels of our multi-resolution registration. The full image is in supplementary material (Figure B.4).
The detail is showing two follicles (f) and two trabeculae (t). It was not re-registered after cropping. The capillary network is especially dense at the surface of follicles. There are only very few capillaries within the follicles.
Scale bar is 500 μm.

Fig. 2.11 shows the vascular network in a bone marrow specimen (Fig. 2.7). Bone marrow is a completely different tissue with much less mesoscopic landmarks (no trabecular veins, no follicles). Our method allowed the registration of this tissue. This means that our method allows for effective multi-organ registrations of human tissues. Appendix B details on our registration results in Figs. B.5–B.7. The diameters of the bone marrow blood vessels change smoothly and both very thin blood vessels (capillaries) and larger ones are connected in a single network (Fig. 2.11). This image provided new histological data. It would not have been possible without our registration.

2.4   Discussion

We utilised feature detection and matching to non-rigidly register serial sections in a multi-resolution manner. Precise registration allows for a 3D reconstruction of the micro-vasculature. Our results show a substantial increase of visual and quantitative quality over a non-multi-resolution feature detection-based method [Ulrich et al., 2014b]. Further, our results are superior to the output of the
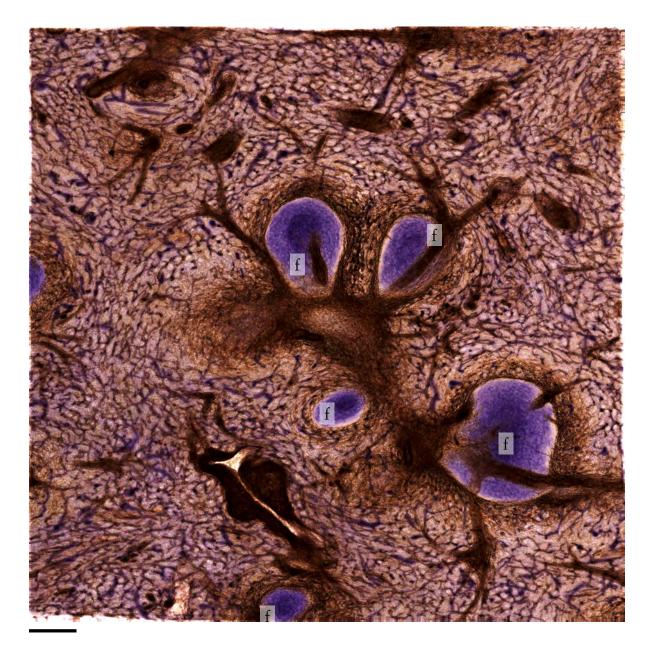
Figure 2.9: Stack of 24 serial sections of a human spleen specimen double-stained for CD34 plus SMA (brown) and CD271 (blue). This volume rendering provides a larger overview, using a different ROI. The branching of larger arterial vessels can be observed. In addition, the scarcity of capillaries in follicles (f) is evident in comparison to the surrounding tissue. The interior of follicles is positive for CD271, but a CD271-negative SMA-negative area is visualised at the surface of the follicles.

This is the volume rendering of our registration of another ROI of the specimen. We used four levels of our multi-resolution method.

Scale bar is 200 μm.
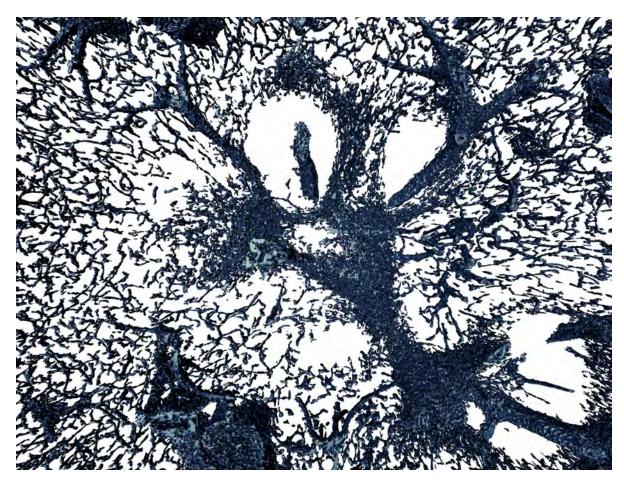
Figure 2.10: A rendering of the vascular system of a stack of human spleen sections corresponding to the centre of Fig. 2.9. The data originate from a stack of 24 serial sections of a human spleen specimen double-stained for CD34 plus SMA.

standard multi-resolution image-based method [Klein et al., 2010], and our method is faster. We have utilised seven quantitative measures to numerically express the quality of these registrations. Our method provides a foundation for 3D reconstruction of serial sections to enable the visualisation of vasculatures of normal and pathological human tissues at the microscopic to mesoscopic scale.

2.4.1 Visual and operational comparison with Ulrich et al.

Similar to Ulrich et al. [2014b], we perform an optimisation step over *all* slices at once to minimise the energy of the non-linear distortions based on detected features ('global' registration). In contrast to the previous work, we

1. Use a multi-resolution approach,

2. Use a different feature detector,

3. Use a different optimisation method, and

4. Target the microscopic-scale landmarks.
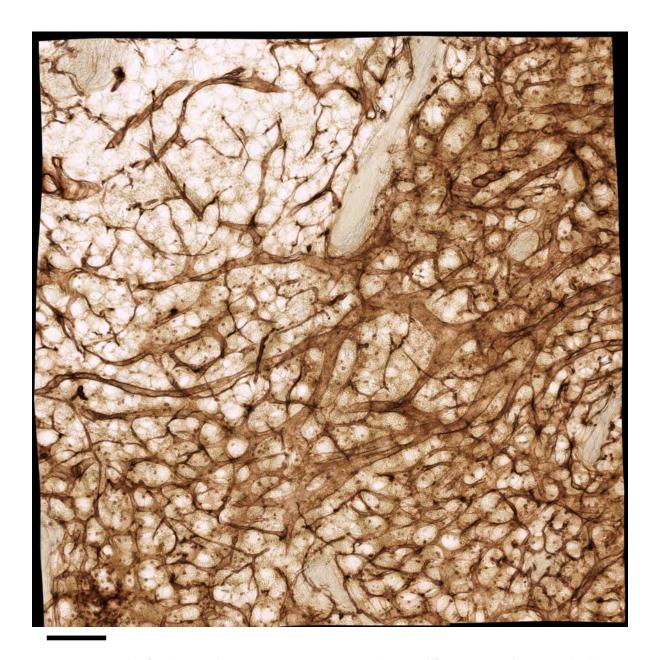
Figure 2.11: Stack of 30 human bone marrow sections single-stained for CD34 and CD141 (both in brown), showing a volume rendering of a ROI.
The volume rendering shows capillaries (microvessels with small diameter) and sinuses (microvessels with large diameter). Both microvessels form a network with smooth transitions from capillaries to sinuses. Sinuses may form round or flattened structures.
It is impossible to represent blood microvessels in such length without registering a series of sections. Chapter 6 details on the results of our findings using such registrations.
Figs. B.5–B.7 show more data. Scale bar is 100 μm.

Our new method processes not only the features of the scale of the initial rigid transform, but also iteratively much smaller features.

A non-multi-resolution approach would not perform the outer loop of Figure 2.2. Features of the same, relatively large size are used both for rigid mapping and for warping of control points for non-rigid transformation. Such registrations faithfully reproduce large objects in the immunostained scans, like follicles or larger blood vessels. However, the non-multi-resolution method does not care for smaller features, corresponding to, e.g., cross-sections of small capillaries. These were badly aligned with such a method and required a multi-resolution effort to be correctly registered.

The movement of the undistortion grid is similar both in the method of Ulrich et al. [2014b] and in this work. However, the equations (2.6) in Ulrich et al. [2014b] are orders of magnitude smaller, as they are solved using CG on normal equations [Press et al., 2007]. As the sparse matrix is not positive semidefinite, the method of Ulrich et al. solves

$$\mathbf{A}^T \mathbf{A} \hat{x} = \mathbf{A}^T b_x \qquad \mathbf{A}^T \mathbf{A} \hat{y} = \mathbf{A}^T b_y. \tag{2.7}$$

The equations (2.7) (i.e., the normalised version of Eq. (2.6)) can be basically solved with CG on normal equations or MINRES on normal equations—the latter is the essential idea of the LSMR method [Fong and Saunders, 2011] we use. Directly applying CG to Eq. (2.7) impairs the convergence time depending on the condition number of $\mathbf{A}$. Using LSMR we solve Eq. (2.6) directly; the solver we use is more robust to early termination and in practice converges much faster.

Our typical pipeline still uses single-resolution registration for a coarse, full-sized registration, as detailed in Section 2.2.3. We stress that using a coarse-grain registration as a fine-grain one (scheme 'A–A') is less efficient than using our multi-resolution approach instead ('A–B'). Figures 2.4 and 2.12 (and also Figs. B.2–B.3, B.7) showcase this.

2.4.2 Visual comparison with *elastix*

The large problem of a serial section registration with *elastix* [Klein et al., 2010] was the accumulating error from using a pair-wise registration. We showcase this in Figure 2.12. We show a detail of the data set from Figure 2.9, registered with *elastix*: Fig. 2.12, (a), (d), the method from Chapter 1 [Ulrich et al., 2014b]: (b), (e), and our method: (c), (f), as a volume rendering ((a)–(c)) and as a z-profile ((d)–(f)).

The z-profiles were consistently built though the middle of the slice stack. When comparing different registrations the lengths of the object cuts may not match, instead the differences in the *form* of the object sections shown in each image reveal the visual registration quality.

Notice that the capillaries in the *elastix* version are quite blurred. We attribute this to the accumulating error. In contrast, our method optimised the positions of control points of the undistortion grid for the whole set of slices at once (in a 'global' manner).

2.4.3 Quantitative quality measures

We discuss seven different quantitative measures of registration quality for each of the specimens (Figs. 2.5–2.7). These are:

- structural similarity (SSIM) [Wang et al., 2004].

- Jaccard measure (Appendix B.2.2).

- dense optical flow [Farnebäck, 2003].

(a) *elastix*   (b) Ulrich et al.   (c) Our method

(d) *elastix*   (e) Ulrich et al.   (f) Our method
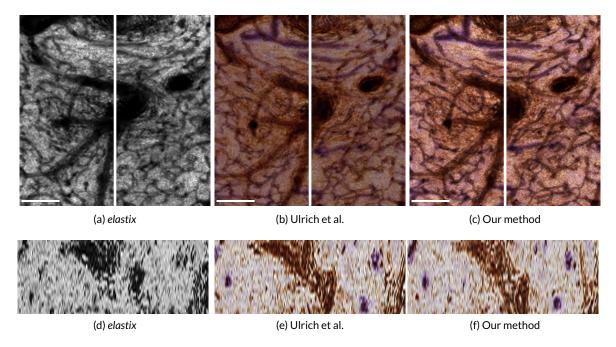
Figure 2.12: A comparison of the registrations for spleen, double staining. (a)–(c): crops of volume renderings, scale bar is 100 µm. (d)–(f): corresponding z-profiles at same scale. The vertical lines in the top images show where the z-profile was computed.

- pixel-wise sum of squared differences (SSD, in Appendix B).

- zero-mean sum of absolute differences (ZSAD).

- mutual information (MI) increase. We divide the mutual information in two registered consecutive images of a scan series by the mutual information in the original, input image pair.

- gradient cross-correlation (GCC). We compute the gradients in the direction of both axes separately and then blend the resulting images.

The exact definition of the quality measures used and more evaluation images are given in Appendix in Section B.2.1.

It is impossible to apply the classical '1–2–1' consistency error metric to our approach as we do not distort a *moving* image to fit a fixed one, but we move *both* images in the '1–2' pair in a manner that minimises movement energy. Our non-rigid method registers all images in the stack at once instead of pairwise registration. The movement energy is also iteratively minimised for *all* images in the stack. Thus, we compare the differences between the images in a '1–2' pair w.r.t. the seven quality measures. We regard pairs produced with Ulrich et al. [2014b] (Chapter 1), Klein et al. [2010], and current registration methods, also including the inputs to the registration.

To obtain the data for quality evaluation we registered the *whole* series with each method, picked a pair of consecutive images and compared them using the measures over the whole ROI. The result is presented in Fig. 2.13 (SSIM for all specimens), in Tables 2.2 (optical flow for all specimens, images in Figs. B.10–B.12) and Table 2.3 (Jaccard measure for all specimens), in Fig. 2.14 for spleen, single-stained, in Fig. 2.15 for spleen, double-stained, and in Fig. 2.16 for bone marrow. Supplementary
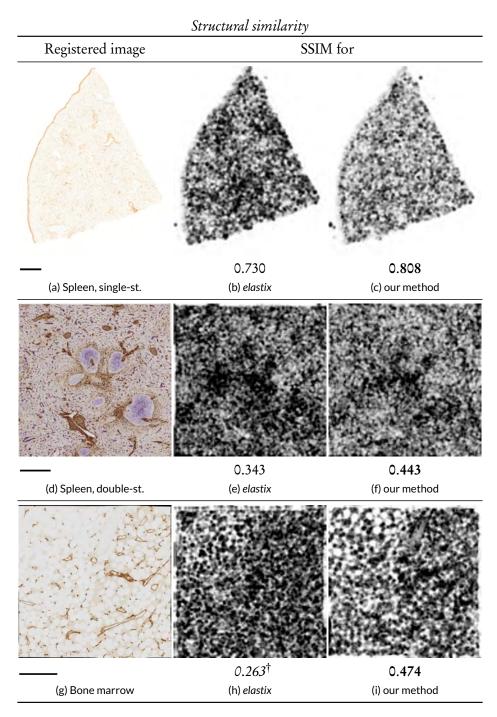
Figure 2.13: Structural similarity index (SSIM) for spleen, single-stained (top row, (a)–(c), scale bar = 1000 µm), spleen, double-stained (middle row, (d)–(f), scale bar = 500 µm), bone marrow (bottom row, (g)–(i), scale bar = 500 µm) specimens.

The left column, (a), (d), (g), shows a registered image, all other columns show SSIM images and SSIM values: the second column, (b), (e), (h), shows state of the art, *elastix*, last column, (c), f), (i), shows our method.

The higher the value and the lighter the image, the better is the quality. **Bold** shows the best value, *italics* with † denotes decrease from baseline. For full version see Figure B.9.

Spleen, single-stained

(a) Overview   (b) Full ZSAD   (c) 1st detail   (d) 2nd detail

Zero-mean sum of absolute differences

| 6.1 | 4.4 | 4.7† | **3.4** |
| (e) input | (f) rigid | (g) *elastix* | (h) our method |

| input | rigid | *elastix* | our method |
| 100% | 154.93% | 163.26%* | 158.23% |

(i) Mutual information increase

Gradient cross-correlation

0.79	0.68	0.69	**0.64**
0.80	0.68	0.70	**0.64**
(j) input	(k) rigid	(l) *elastix*	(m) our method
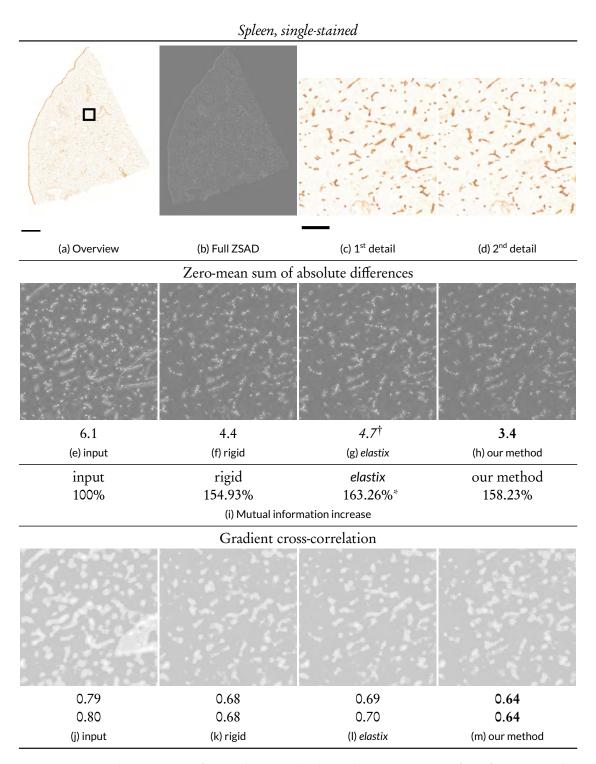
Figure 2.14: Quality measures for single-stained spleen. *(Caption continued on the next page.)*

Figure 2.14: (Continued from previous page.) Quality measures for single-stained spleen, entire section (same as Fig. 2.8). (a): Detail selection. Scale bar = 1 mm. (b): Zero-mean sum of absolute differences (ZSAD) measure for the whole image. Brightness and contrast for ZSAD images was corrected with +50.
(c, d): Detail of the first and second slide. Scale bar = 50 μm. (e)–(h): Zero-mean sum of absolute differences (ZSAD) for (e): input data, (f): rigid-only transform, (g): *elastix*, and (h): our method.
(i): Mutual information (MI) increase.
(j)–(m): Combined cross-correlation of gradient approximation (GCC). We show GCC for (j): the pair of original images, (k): rigid-only transformation, (l): *elastix*, and (m): our method. It is corrected with $\gamma = 6$.
The values below the images show the mean values of full image-wide quality measure. In case of GCC the first line states the mean of derivatives in the direction of the x axis and the second line shows the mean in the direction of y axis.
Aside from (i), the smaller the values and the darker the images the better. **Bold** is best, *italics* with † denotes decrease from the baseline, ∗ marks incomparable values.

material shows more cases in Figs. B.9–B.18. Images in Figs. 2.14–2.16 show crops to emphasise details of the registration. Below the images we show mean values for whole images, not for the crops shown. For the sake of presentation we adjusted the brightness and contrast to +50 in all presented ZSAD images. We also used gamma correction with $\gamma = 6$ in GCC images. The numerical values were not corrected.

The SSIM, Jaccard, and MI images and values are better when lighter and higher. SSD, ZSAD, GCC, and optical flow images and values are better when darker and lower.

SSIM for all specimens (Fig. 2.13)

Figure 2.13 shows the structural similarity index [Wang et al., 2004] for spleen, single-stained, in (b) and (c), for spleen, double-stained, in (e) and (f), and fir bone marrow in (h) and (i). Figs. 2.13, (a), (d), (g) show corresponding specimens registered using our method.

The SSIM images were most convincing for the single-stained spleen specimens. SSIM image for our method Fig. 2.13, (c) was the lightest in the series. The image for *elastix* (b) was darker. It was quite similar to rigid-only transformation (Fig. B.9), but still a bit *darker*, which corresponds to a minor decrease of the quality. Overall, our method was 45% better than input overlay, while *elastix* yielded 31% of improvement. The relative difference between SSIM values for our method and for *elastix* was over 10%.

The SSIM value of our method for double-stained spleen specimen (f) was improved 146% over baseline (Fig. B.9). Here *elastix* (e) was 90.5% better. Both non-rigid methods showed SSIM images much lighter than the rigid-only transformation (Fig. B.9).

For bone marrow specimens, our multi-resolution method with four levels (Fig. 2.13, (i)) improved SSIM by almost 51% over the input overlay and 35% over rigid-only registration, while *elastix* (h) produced a value 17% *worse* than single-channel rigid-only registration. Compared to our method, *elastix* was 44.43% worse. (Note, however, that *elastix* operated on the red channel only, while our method operated on coloured images, hence a direct comparison might be not appropriate.) We present an extended version of this evaluation and also a smaller ROI with only two levels of our

Table 2.2: The mean of motion magnitude in pixels for dense optical flow between two registered images. *elastix* was applied to a single channel only, these entries are marked with ⁂. **Bold** shows the best value, *italics* with † denote decrease from baseline. Lower values represent better quality.

Specimen	Size	Levels	Method				
			Input	Rigid	Chapter 1	*elastix*	Chapter 2
Spleen, single-stained	8000	4	4.09	1.49	*1.86*†	*1.60*⁂,†	**1.12**
Spleen, double-stained	5032	4	17.3	12.2	5.89	8.23⁂	**5.32**
Bone marrow	4032	4	9.74	7.47	**5.36**	6.16⁂	5.83

multi-resolution method in Appendix B (Fig. B.9). The smaller ROI shows the performance of the method from Chapter 1 in this case: the SSIM quality *decreased* by 8.6%.

Optical flow for all specimens (Table 2.2)

We measure the magnitude of dense optical flow (MDOF) between two consecutive serial sections based on Farnebäck [2003]. We use pairs of consecutive registered images as input and evaluate the motion vector field in terms of its magnitude and angle. Table 2.2 lists the means of MDOF between two consecutive registered images. The superiority of our method is shown for spleen specimens. Figs. B.10–B.12 present colour-coded vector fields for all specimens and each registration in Table 2.2.

For the spleen specimens our current method returned best MDOF results. Our method was up to 265% better than the input overlay. We improved the mean MDOF for the method of Ulrich et al. (Chapter 1) by over 33% for single-stained spleen specimen. As for *elastix*, the values were not quite comparable, as this program operates on single-channel images. Still, there was an improvement by almost 55% for double-stained spleen sections.

Unexpectedly, our old method from Chapter 1 was almost 9.7% better than our current method w.r.t. MDOF of bone marrow specimens. However, visual inspection (Section B) indicated than the result of the method Ulrich et al. was less usable than the result of our current registration. Our present method improves the registration of small capillaries which correspond to 7–21 pixels in perpendicularly sections of our bone marrow specimens. This size is due to the fact that capillaries tend to collapse during immersion fixation and thus they correspond to dots or ovals of up to 3–4 µm in such sections. The averaging window of the dense optical flow method used [Farnebäck, 2003] was however 15 pixels. Hence the finest-grain improvements through our method might be undetectable with this quality measure.

Jaccard measure (Table 2.3)

Jaccard measure assigns a value between zero and one for the similarity degree of two sets. We use it on our input images by summing up binary pixels. We threshold the data at a given value, yielding a coarse kind of segmentation. Then we compute the Jaccard measure for the resulting binary images. We summarise our findings in Table 2.3.

We chose a relatively high threshold value to obtain distinctive values. The Jaccard measure for our method applied to single-stained spleen specimens reached 1.0 at threshold 100. The stainings we used had a good contrast, thus we chose the threshold 220 for the same 8 bit specimens in Table 2.3. An extended version of Table 2.3 and examples of segmented images are in Section B.2.2 as Table B.1 and Figure B.13.

Table 2.3: Jaccard measure for all three specimens. The values were computed channel-wise where applicable. The value for the best channel is shown.
We denote the best overall value in **bold**, the values worse than baseline are marked with *italics* and †, values with possible compatibility problems (single-channel vs. best of three RGB channels) are marked with ∗. The values are scaled between 0 and 1.0, higher values represent better quality.

Specimen	Size	Levels	Thresh.	Input	Rigid	Chapter 1	*elastix*	Chapter 2
Spleen, single-stained	8000	4	220	0.9701	0.9785	*0.9762*†	*0.9774*∗,†	**0.9829**
Spleen, double-stained	5032	4	100	0.8622	0.8755	0.8917	*0.8158*∗,†	**0.9062**
Bone marrow	4032	4	150	0.9583	*0.9562*†	0.9612	*0.8083*∗,†	**0.9631**

Our method yielded best results for all specimens w.r.t. Jaccard measure. We improved over *elastix* by almost 10% for double-stained spleen and by over 16% for bone marrow specimens. The Jaccard measure for the *elastix* program was, however, computed for a single channel, while it was computed for every channel of a colour image for our method. Our method also outperformed Ulrich et al., but by a narrower margin.

Spleen, single-stained (Fig. 2.14)

The numerical values for SSD (in Appendix B, Fig. B.14) showed that our method produced a 14.5% better mean value than *elastix*.

ZSAD for our method in single-stained spleen was visually much clearer and of less intensity. The numerical values showed that ZSAD mean was the best in our new method, with 22.7% improvement.

MI showed a slight increase between rigid-only registration and our method. The value for *elastix* was formally better, but it is not really comparable as *elastix* used only the red channel of the images. Thus we evaluated the information increase for this colour channel only. (Section 2.4.3 discusses an experiment with grayscale-based *elastix* and grayscale-based MI on our registration to obtain comparable values.)

The GCC values—basically, the 'movement' in the image pair—was in all cases improved compared to the original overlay. Our method delivered the best average value overall: 0.64 vs. 0.80, i.e., a 25% improvement.

Our method showed best results for all quality measures with the exception of MI increase, where the data was not quite comparable.

Spleen, double-stained (Fig. 2.15)

Our method produced the best mean (Fig. B.15) and median (not shown) SSD values of all methods used.

The ZSAD image for our method looked less noisy and also visibly darker. Numerically this means that our method delivered the best results. The mean value of ZSAD for *elastix* was roughly 4% worse, while our method had a 21% better mean ZSAD value than rigid-only registration. The median for our method (not shown) was 24% better.

We obtained MI increase of 195% which was the best value across the methods tested. However, the value delivered by *elastix* (183%) was not really comparable to our results. The mutual information increase (i.e., a quotient of mutual information of the registered image pair and original image pair)

Spleen, double-stained

| (a) Overview | (b) Full ZSAD | (c) 1st detail | (d) 2nd detail |

Zero-mean sum of absolute differences

| 35.2 | 33.1 | $34.5^{\dagger}$ | **26.2** |
| (e) input | (f) rigid | (g) *elastix* | (h) our method |

| input | rigid | *elastix* | our method |
| 100% | 107.08% | 183.06%* | **195.37%** |

(i) Mutual information increase

Gradient cross-correlation

1.95	1.90	$1.96^{\dagger}$	**1.71**
1.99	1.94	1.98	**1.77**
(j) input	(k) rigid	(l) *elastix*	(m) our method
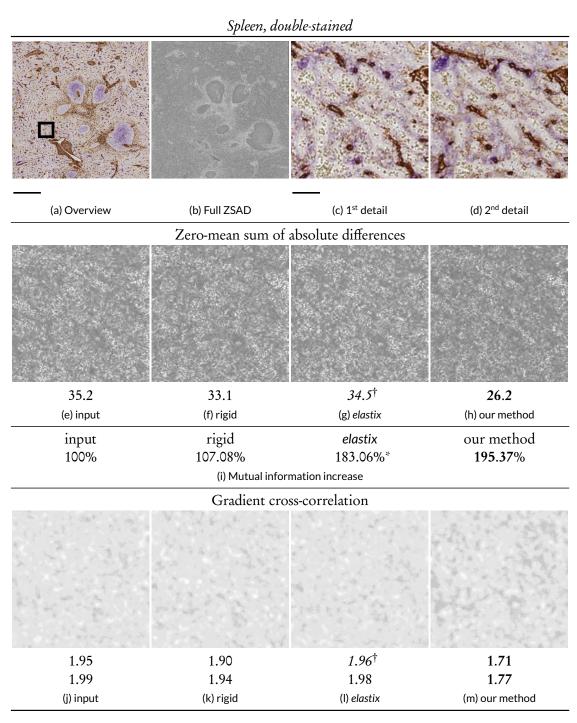
Figure 2.15: Quality measures for double-stained spleen (same ROI as in Figure 2.9.) (a): Detail selection. Scale bar = 500 μm. (b): Zero-mean sum of absolute differences (ZSAD) for full ROI. (See (e)–(h) for crops.) ZSAD images were enhanced with brightness and contrast +50. (c, d): Details of the first and second section, scale bar = 100 μm. (e)–(h): Zero-mean sum of absolute differences. *(Continued on the next page.)*

Figure 2.15: (Continued from the previous page.) (i): Mutual information (MI) increase with original image pair as a baseline. Note that *elastix* (marked with asterisk) shows data for blue channel only, but our method produces RGB images.
(j)–(m): Combined cross-correlation of gradient approximation (GCC). Images are +6 gamma-corrected.
We show mean values for each full image-wide quality measure. For GCC we state both the mean of the derivative in the direction of the x axis (first line) and y axis (second line).
Aside from (i), the smaller the values and the darker the images the better. **Bold** shows the best value, *italics* with † denote decrease from the baseline, ⁎ marks values not comparable.

was computed for RGB images in case of rigid-only transform and our method in Figure 2.15, but with blue channel only in both image pairs for *elastix*. (In Fig. 2.12 *elastix* operated in the same manner.) If we used only the blue channel for our method, we would obtain 176.48% increase. However, this would not be appropriate towards our method. If we compared the intensities only (and run *elastix* on intensity images), *elastix* would outperform our method by 3.6%.

The gradient cross-correlation (GCC) showed some darker areas in case of our method. These might guide us to perceive the whole image as darker, esp. when compared to baseline. As for mean values, the improvement of our method translated to 12.3%. The maximal values improved up to 61% (not shown).

Our method showed better results than *elastix* for *all* measures on the double-stained spleen specimen. Even if we disregard MI as the values are not quite comparable, our method outperformed *elastix* on the current data set w.r.t. six quality measures.

Bone marrow (Fig. 2.16)

Here we compare the method of Ulrich et al. [2014b], a non-multi-resolution feature-based registration method, to our new approach based on bone marrow images. We also show the initial coarse-grained registration ('input', 'A only' in terms of Table 2.1) and rigid-only registration as baselines. This discussion is based on four levels of our registration method. Fig. B.18 is an extended version of Fig. 2.16, Fig. B.17 shows a similar evaluation for only two levels of registration.

Our method showed the best overall mean SSD value. The SSD images, however, were not very informative (not shown here, but in Fig. B.18), they represented basically noise outside of bone marrow fat cells.

ZSAD showed an apparent matching problem w.r.t. the contours of the blood vessels quite clearly in a light colour, e.g., in Fig. 2.16, (e), (f). The ZSAD values seemed to decrease both for Ulrich et al. (g) and for our method (h), however our method showed less bright spots that correspond to a mismatch. This led to the best mean value (and also maximum and median, not shown). The difference between mean ZSAD for the method of Ulrich et al. and our method was almost 20%.

We compared MI for RGB images, thus all values were fully comparable. Here, the single-resolution registration was marginally ($< 0.7\%$) better than our method. Both showed improvements over rigid-only registration and input images.

The reason for the inferior performance of our method w.r.t. MI might be the capture of microscopic images. The contrast enhancement in original images visualises a lot small unstained cells that basically constitute background noise. Visual comparisons using the method of Ulrich et al. and our method (Sect. B, esp. Fig. B.7) underline the superiority of our method in the representation of small capillaries.

The GCC images (Fig. 2.16 (j)–(m)) in our method faded on a larger surface to darker grey, the same colour as the distinctively visible large unstained cells. We also noticed some black spots in Fig. 2.16, (m) not present in other images. This is surprising, as all images were adjusted with $\gamma = 6$, and thus dark grey in the images corresponds to a quite low GCC value. The numerical data on full ROIs not corrected for gamma showed the superiority of our method (Fig. 2.16, (m)). There, our method (m) improved up to 17.2% over input overlay (j). For comparison, the GCC for the method of Ulrich et al. (Fig. 2.16, (l)) improved up to 10.6%, and for *elastix* (not shown) it improved up to 9.2% over its single-channel baseline. Our method, Fig. 2.16, (m), performed best in both directions.

From the overwhelming majority of measures presented here (5 out of 7) and and from visual inspection (Appendix B) we are still able to conclude that our method outperforms Ulrich et al.

2.4.4   Execution times

We also evaluated the execution times of the three methods in question: the method of Ulrich et al. [2014b] (Chapter 1), *elastix* [Klein et al., 2010], and our current method. Table 2.4 shows the results. We performed our tests on a machine with Intel Core i7-3770K (3.5 GHz), 16 GB RAM, Windows 8.1 64 bit, using available build of *elastix* 4.7, a 64 bit build using Microsoft Visual Studio 2008 for Ulrich et al. and a 64 bit build using Microsoft Visual Studio 2013 for our method. Please note that *elastix* operated on single-channel images, hence, e.g., the warping time was reduced in this case. For a certain parameter set our method outperformed all other methods tested, as detailed below.

The registration of the double-stained spleen specimen with small minimal feature size took surprisingly long time with our method (marked with asterisks in Table 2.4). The possible reason is that the smaller size of the features (20, 10) fits well to numerous smaller structures in the scan. As the matrix **A** from Eq. (2.6) scales with the number of features involved, this produces huge matrices during the alignment procedure and thus slows the registration down. We have observed sizes of $\approx 140\text{k} \times 1\text{k}$. Generally it would be of benefit to stop at a larger feature size, but in this case we were interested in blood vessel sections of approximately the same feature size as noise. Better solvers, GPU computing, and a better filtering of candidate features could enable better performance in such cases. When the feature size was adapted to exclude such small structures—to be more exact: using two levels with feature size 100 and 50—our method performed well, in a bit more than 18 minutes for spleen, double-stained. This configuration of our method performed faster than both Ulrich et al. and *elastix* on all organ specimens tested (shown in **bold**).

2.4.5   Limitations

Our method is applicable to any kind of image data. We used it for the registration of micro-vessels in human spleen and bone marrow (Section 2.3.2), while the method of Ulrich et al. was only applied to human spleens. We have tested our method with further, non-medical test images and concluded that our method is also applicable to other image series (Section B.1).

Our method requires distinctive image regions that will be detected as features. These do not need to be large and many various features may be present across the images. In fact, our specimens already exhibit a lot of similar features across the images that do not contribute to a proper alignment.

Hence, our method is expected not to function properly on extremely homogeneous cell arrangements, but in most tissues blood vessels and nerves provide enough information for feature detection. Aside from this issue our method would function with factually any sufficiently large real-life images.
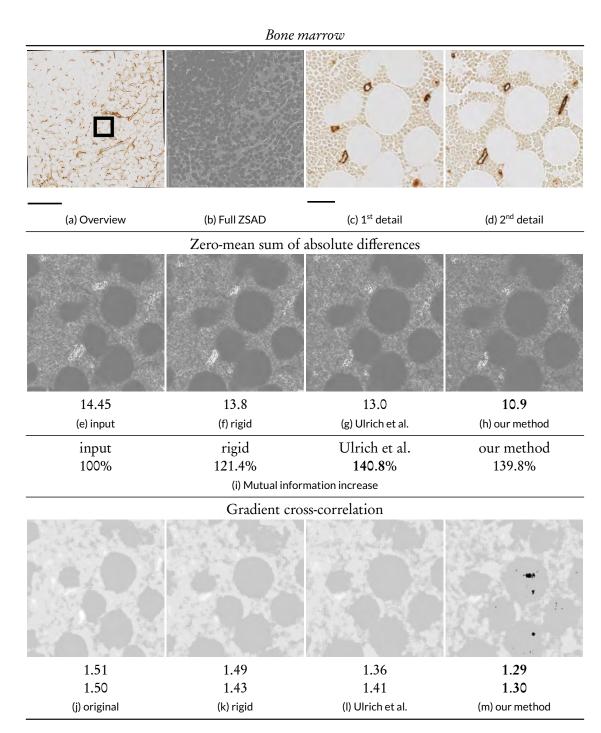
Figure 2.16: Quality measures for bone marrow specimen, same ROI as in Fig. 2.11. (a): Detail selection, scale bar = 500 μm. (b): Complete unedited ZSAD quality measure. (c), (d): The details of both sections, scale bar = 100 μm.
(e)–(h): Zero-mean sum of absolute differences (ZSAD). We enhanced the brightness and contrast. *(Continued on the next page.)*

Figure 2.16: (Continued from the previous page.) (i): Mutual information (MI) increase from appropriate baseline.

(j)–(m): Combined cross-correlation of gradient approximation (GCC), for (l): input images, (m): rigid-only transformation, (n): Ulrich et al., (o): our method. Brightness and contrast enhanced +50.

We state the mean values for each image-wide quality measures. For GCC the first (second) line states the mean of derivatives in the direction of the x (y) axis.

Smaller values and darker images indicate improvement, aside from (i). **Bold** shows the best value, *italics* with † denote decrease from baseline.

Table 2.4: Comparison of execution times. **Bold** font denotes best time, parameters in *italic* designate the configuration we used for evaluation in Sections 2.4.1–2.4.3.

'Scale' denotes the minimal key point size. An asterisk (*) marks execution times specially discussed in Section 2.4.4. In these cases phase contrast of unstained cells generates too many small features.

Specimen			Registration				Time
Name	Side	Slices	Method	Multi-resolution?	Levels	Scale	(hh:mm:ss)
Single-stained spleen (Figs. 2.5, 2.8, Chapter 7)	8000	24	Chapter 1	no	–	100	0:21:11
			Klein et al. [2010]	yes	4	–	01:01:05
					2	50	**0:12:06**
			our method	yes	*4*	*10*	0:49:20
Double-stained spleen (Figs. 2.6, 2.9)	5032	24	Chapter 1	no	–	100	0:18:31
			Klein et al. [2010]	yes	4	–	0:23:12
					2	50	**0:18:09**
			our method	yes	*2*	*20*	04:03:14*
					4	10	14:40:14*
Bone marrow (Figs. 2.7, 2.11, Chapter 6)	4032	30	Chapter 1	no	–	100	0:12:55
			Klein et al. [2010]	yes	4	–	0:17:43
					2	50	**0:07:07**
			our method	yes	*2*	*20*	0:13:00
					4	10	0:44:42

2.4.6   Answering medical research questions

Our volume renderings showed complex branching patterns of arterioles in human spleen (Fig. 2.9). These comprised dichotomous branching as well as penicillar branching. Further, we noticed the absence of capillaries in certain regions (follicles, coloured blue in Fig. 2.9) of the spleen. In contrast to previous assumptions, blood microvessels in the human bone marrow did not exhibit abrupt changes in calibre. Figure 2.11 (and also Figs. B.5–B.7) shows a large network beginning with smallest capillaries and ending in larger blood vessels, so-termed sinuses. Chapters 6–7 that are dedicated to medical research questions would have impossible without the presented registration method. Our research facilitated answering actual histological research questions about form and function of blood vessels in specific regions.

2.5 Conclusion

This work proposes a new multi-resolution approach to the registration of 2D images. It is used for the analysis of the vascular network obtained from serial sections stained with immunohistological methods that require thin sections. We were able to observe blood vessel networks of various organs in larger areas at a microscopic resolution. The sectioning procedure induces a distortion on the serial sections, which needs to be abolished through our non-rigid registration.

We demonstrate our approach to a 'sparse', 'global' feature detection-based multi-resolution registration of large histological sections, enabling more precise 3D reconstruction. Our method pays special attention to the smaller features that often correspond to objects that are of interest in medical research. We have also evaluated the visual and quantitative performance of our method. In all data series tested our approach outperforms the state of the art. Our registration method permits to visualise and compare organ-specific microvascular arrangements.

We register real data from histological research and show that our approach is competitive and cannot be circumvented by a repeated application of the method published previously.

2.5.1 Future work

Our implementation spends most of the time in LSMR solver calls that align the B-spline control points with the local feature points—in other words in the optimisation problem that arises in computing the (un-)distortion grid. The next two major time-consuming parts are the actual transformation of an image and the initial matching of the features. A better filtering approach for close matches would help. It seems feasible to utilise the GPU at least for the solver and for feature matches. A GPU-based SURF feature detector is also available. Using a different Krylov subspace solver may enable the sharing of precomputations in Eq. (2.6).

Another issue are the anisotropic dimensions of the sections. It could be of benefit to use the information from the feature matches for a better *z*-axis interpolation during the marching cubes stage. Chapter 3 solves the anisotropy problem of our series in a similar, but different manner.

Acknowledgements

COMPENSATING ANISOTROPY IN HISTOLOGICAL SERIAL SECTIONS WITH OPTICAL FLOW-BASED INTERPOLATION

In which we introduce a technique to reduce anisotropy is serial sections for better 3D reconstructions. We reported these results first as Lobachev et al. [2017b], a collaboration effort with Birte S. Steiniger and Michael Guthe.

Author Contributions: *Conceptualisation: MG, OL. Data curation: OL, BSS. Formal analysis: OL. Funding acquisition: MG. Investigation: OL, BSS. Methodology: OL. Project administration: OL. Resources: OL, BSS. Software: OL. Supervision: MG. Validation: OL, BSS. Visualisation: OL. Writing—original draft: OL. Writing—review & editing: OL, BSS, MG. Corresponding author: OL. Presentation: OL.*

Histological serial sections allow for 3D representation of anatomical structures in microscopic to mesoscopic range. However, due to the nature of the acquisition, they suffer from severe anisotropy: 14-to-1 in a single average microscopic paraffin section. We present an interpolation method based on optical flow and show that standard interpolation methods are less suited for serial sections.

With our non-linear interpolation approach we are able to represent the "movement" of image parts that are of interest. This allows for better 3D reconstructions and further insights in microanatomy.
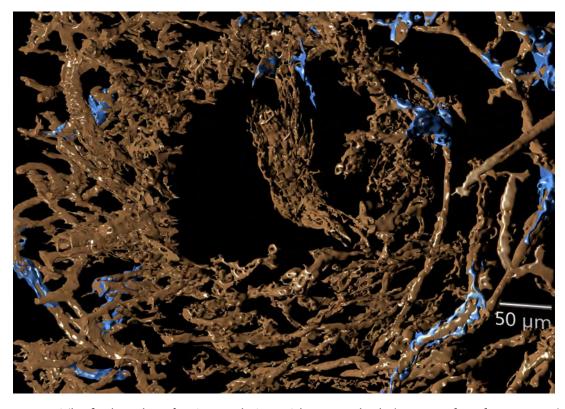
Figure 3.1: The final meshes after interpolation with our method. *(Continued on the next page.)*

EAL world data, especially acquired from some special modalities like MRT or histology, often suffer from anisotropy. This means that in some dimensions the data are not sampled as frequently as in others. In other words, they feature rectangular, not square pixels or non-cubic voxels. The most straightforward way to reduce or eliminate anisotropy is to interpolate the data forming a denser, ideally, uniform sampling grid. This paper discusses a *method* of interpolation that is more suitable for 3D reconstruction of histological data than standard approaches. Our approach maintains better connectivity of tissues across the sections despite small local spatial distortions.

Histology is one of the oldest medical imaging techniques, in contrast to many other approaches, it is invasive. However, in humans there is de facto no other way (Appendix A) to mark and investigate microscopic to mesoscopic cell structures of interest, such as microscopic blood vessels. With advances in histological staining (immunohistochemistry) and acquisition techniques (optical scanning microscopes), digitisation of large *series* of sections is possible. These need to be registered (Chapters 1 and 2). However, the resulting volumes have different resolution in the sectioning axis than in scanning plane.

The huge drawback of serial sections is an anisotropy. While the resolution of a scanning microscope can reach $0.25\,\mu m$/pixel in the xy plane, the z axis is limited to the thickness of a section. The minimal thickness of sections in hard plastic embedding obtained with a conventional microtome may be 0.5 to 3 µm/section. The thickness of paraffin serial sections lies in the range of 5 to 8 µm/section. The maximal section thickness with paraffin embeddings and a conventional microtome can reach 25 to 30 µm. We aim for better 3D reconstructions from "conventional" serial sections. Our presentation in this paper focuses on serial sections with typical anisotropy $0.5 : 0.5 : 7$ µm or 14-to-1. We reduce the anisotropy to $0.5 : 0.5 : 1$ µm or 2-to-1. *Any* kind of serial section data would profit from our approach, actually, any kind of anisotropic data does.

3.1   Related work

Standard interpolation methods like bilinear interpolation, bicubic interpolation [Keys, 1981], or Lanzos filter [Duchon, 1979] (see also [Burger and Burge, 2009]) are of course related to this paper. However, we focus on a more special and non-generic use case: we would like to interpolate medical data in a way that improves subsequent 3D reconstruction. Our method uses optical flow.

3.1.1   Optical flow

Optical flow as such was defined by Horn and Schunck [1981], while Lucas and Kanade [1981] pioneered a viable way for computing it that is still used today. Baker et al. [2011] present a nice overview of optical flow methods and their evaluation. A lot of effort was invested into variants and improvements of optical flow [Brox et al., 2004, Sun et al., 2010, Brox and Malik, 2011]. Recent improvements to the topic are presented by Revaud et al. [2015]. We use the *dense* optical flow computation [Farnebäck, 2003]. Keeling and Ring [2005] use optical flow as a similarity metric for

Figure 3.1: (Continued from previous page.) The final meshes for CD34$^+$ (brown) and CD271$^+$ cells (blue) in a rendering, data set 'spleen 2'. The threshold for blue mesh generation was set rather high: the inside of the follicle is empty. Input data for this mesh is in Figure 3.2. Chapter 7 reports on medical research implications of reconstructions from this data set.

(a) Frontal plane of the volume, "spleen 2"   (b)   (c)   (d)   (e)
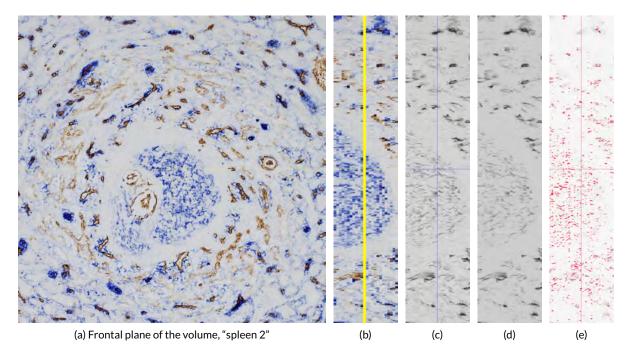
Figure 3.2: A quick summary. (a) shows the frontal plane of a volume obtained from histological serial sections, data set "spleen 2". (b), (c), and (d) show a z-profile of, correspondingly, registered data set, separated CD34$^+$ staining after interpolation, and final result of volume filtering. (c) and (d) are negated for presentation. (e) highlights differences between unfiltered and filtered data. The focus of this chapter is the interpolation (c) that facilitates better final meshes (Figures 3.1, 3.9).
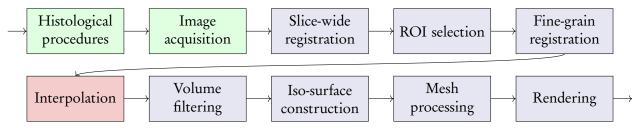
Figure 3.3: Our image processing pipeline, simplified. (Figure 0.1 shows full version.) All steps after acquisition are digital. The focus of this chapter is interpolation.

elastic image registration of medical images. It is in a sense dual to our approach: we also use optical flow as a measure for "movement" in two consecutive medical images. However, we use the detected motion to generate intermediates and not to decide what distortion needs to be applied to input images to make them more similar. Zitnick et al. [2005] combine image segmentation and optical flow computation.

3.1.2 Registration and 3D reconstruction of serial sections

An integrated suite for processing of serial sections is *TrackEM2* [Cardona et al., 2010], a part of *Fiji* [Schindelin et al., 2012]. 3D reconstruction from serial sections is an increasingly popular topic [Ourselin et al., 2001, Steiniger et al., 2003, Gijtenbeek et al., 2005, Ma et al., 2008b, Xu et al., 2015,

Saalfeld et al., 2012, Steiniger et al., 2016, Lobachev et al., 2017a]. It is especially of importance in the mesoscopic area which cannot be represented on with a single histological section. However, to align these sections, a *registration* method is needed. We used the methods of Lobachev et al. [2017a] (Chapter 2) for fine-grain registration and of Ulrich et al. [2014b] (Chapter 1) for coarse-grain, slice-wide registration. Other approaches include elastic registration [e.g., Guest et al., 2001], registration in feature space [e.g., Bağci and Bai, 2008], utilisation of mutual information [e.g., Rueckert et al., 1999, Chappelow et al., 2011, Schnabel et al., 2001], smoothness based registration [Cifor et al., 2011], thin-plate splines [e.g., Wan et al., 2013].

3.1.3   Interpolation of medical data

Optical flow was used for interpolation in live MRT data to reduce breathing artefacts [Ehrhardt et al., 2007]. However, there are many differences. Ehrhardt et al. [2007] work on 4D MRT data, we operate on a set of 2D microscopic images. They apply a demons-based registration method similar to Thirion [1998], we directly apply the distortion field to (pre-registered) input images (Section 3.2). Ehrhardt et al. [2007] operate on $512 \times 512 \times 200 \times 10$ data (i.e., 5.24×10^8 voxel), but they report more than 30 hours run time. Our method typically operates on 28.6–81.1% of their input volume, but produces results in a few minutes with a Python-based implementation. Decreased dimension helps in our case. Ehrhardt et al. [2007] use the method of Barron et al. [1994] for optical flow computation, we utilise the method of Farnebäck [2003]. A sparse interpolation for electron microscopy has been suggested before [Carata et al., 2013]. Unlike our method it uses a sparse optical flow. We compare our approach to standard interpolation methods in Section 3.3.3.

Thevenaz et al. [2000] discuss in-depth the quality and applications of some existing and modified interpolation methods, it is mainly focused on choosing better basis functions. Lehmann et al. [1999] provide an overview of existing interpolation methods in context of medical imaging and aspect ratio correction.

3.2   Method

3.2.1   Synopsis of the pipeline

We process immunohistochemically stained serial sections (though the method is also applicable to other data, as shown in Figs. 3.12, 3.13), hence we outline our complete imaging pipeline in this section to facilitate better understanding of the actual method. Figure 3.3 shows the overall pipeline.

A removed and fixed specimen was embedded in paraffin [Steiniger et al., 2014a,b]. Bone marrow, however, was embedded into methylmethacrylate [Steiniger et al., 2013, 2016]. The specimens were sectioned on a microtome. The embedding medium was removed, and the sections were stained using several immunohistological methods. Then, images were acquired. Scans of the entire sections were performed using a transmitted light scanning microscope (Leica SCN 400). From this point on the processing was completely digital.

Entire sections (typically very large images, ranging at $20k$–$30k$ pixels per side) are coarsely registered [Ulrich et al., 2014b]. This is needed for the subsequent selection of regions of interest (ROI). Potential irregularities in intensity caused by varying slice thickness are compensated with slice normalisation [Reinhard et al., 2001, Khan et al., 2014]. The ROIs are finely registered [Lobachev et al., 2017a] (Chapter 2), producing volume data. However, the anisotropy is severe. At this point the interpolation method presented here is used. After interpolation the volume can be further processed. We typically apply at least a closing filter and a minor Gaussian blur (Chapters 6–7). The

filtered volume is used for mesh construction with marching cubes [Lorensen and Cline, 1987]. The resulting mesh is post-processed. The actual processing depends on the objective, but it typically includes methods to close a few holes, to reduce the mesh size [Ju, 2004], and to correctly smooth the surface [Taubin, 1995]. The final mesh can be rendered with usual techniques.

We stress that our interpolation method allows for better meshes than standard approaches (Fig. 3.9 shows a comparison).

3.2.2 Dense optical flow

Farnebäck's dense optical flow [Farnebäck, 2003] is defined in the following way. We aim to compare two signals, f_1 and f_2. They are modelled with quadric polynomials, e.g., $f_1(\mathbf{x}) = \mathbf{x}^T \mathbf{A}_1 \mathbf{x} + \mathbf{b}_1^T \mathbf{x} + c_1$ for a symmetric matrix $\mathbf{A}_1$, a vector $\mathbf{b}_1$ and a scalar c_1. Assume f_2 is globally displaced by $\mathbf{d}$, it can be defined in a similar way in terms of $\mathbf{A}_2, \mathbf{b}_2$ and c_2. We can express $\mathbf{d}$ as $-\mathbf{A}_1^{-1}(\mathbf{b}_2 - \mathbf{b}_1)/2$. To make the computation more practical we use *local* polynomial approximations of images. In the following we omit the $\mathbf{x}$ argument, e.g., $\mathbf{A}$ always stands for a local approximation $\mathbf{A}(\mathbf{x})$. Let

$$\mathbf{A} := (\mathbf{A}_1 + \mathbf{A}_2)/2 \quad \text{and} \quad \Delta \mathbf{b} := -(\mathbf{b}_2 - \mathbf{b}_1)/2.$$

The distance is now a spatially varying displacement field $\mathbf{d}$ with $\mathbf{A}\mathbf{d} = \Delta \mathbf{b}$. With a weight function $\tilde{w}$—the edges of an image have a smaller weight—in a neighbourhood I holds:

$$\mathbf{d}(\mathbf{x}) = \sum_I w \mathbf{A}^T \Delta \mathbf{b} \bigg/ \sum_I w \mathbf{A}^T \mathbf{A}, \tag{3.1}$$

the sums run over all $\mathbf{x}$ in I, as $\mathbf{x}$ is an omitted for brevity argument to $\mathbf{A}$, $\Delta \mathbf{b}$, and w. In practice we compute $\mathbf{A}^T \mathbf{A}$, $\mathbf{A}^T \Delta \mathbf{b}$, and $\Delta \mathbf{b}^T \Delta \mathbf{b}$ (needed for the confidence value) pointwise.

The flow is computed on multiple scales to alleviate the problem with larger displacements. We compute the flow based on a classical 2^n image pyramid with 3 levels, averaging window size of 15 pixels, three iterations at each level, neighbourhood size for polynomial expansion of 5 pixels. The derivatives for polynomial expansion are smoothed with $\sigma = 1.2$.

Let us denote with $\mathscr{D}(\mathbf{A}, \mathbf{B})$ the complete computation of dense optical flow between $\mathbf{A}$ and $\mathbf{B}$. It is basically a vector field that consists of local approximations $\mathbf{d}(\mathbf{x})$. We compute it with Open CV [Bradski and Kaehler, 2008] function calcOpticalFlowFarneback.

3.2.3 Computing the interpolations

Now let $\mathscr{F}_{\mathbf{A},\mathbf{B}}(r, \cdot)$ be the distortion operator of dense optical flow between images $\mathbf{A}$ and $\mathbf{B}$. Here is $0 \leq r \leq 1$ the magnitude of the distortion. The operator $\mathscr{F}_{\mathbf{A},\mathbf{B}}(r, \mathbf{X})$ computes the optical flow $\mathscr{D}$ and spatially displaces the input image $\mathbf{X}$ by $r\mathscr{D}(\mathbf{A}, \mathbf{B})$. The displacement is applied to the input image with Open CV remap function. In most cases, however, $\mathbf{X}$ is either $\mathbf{A}$ or $\mathbf{B}$. Basically, $\mathscr{F}_{\mathbf{A},\mathbf{B}}(1/2, \mathbf{A})$ distorts $\mathbf{A}$ to correspond to the optical flow from $\mathbf{A}$ to $\mathbf{B}$ at a half magnitude. Naturally, $\mathscr{F}_{\mathbf{A},\mathbf{B}}(0, \mathbf{A}) = \mathbf{A}$ and $\mathscr{F}_{\mathbf{B},\mathbf{A}}(0, \mathbf{B}) = \mathbf{B}$. However, $\mathscr{F}_{\mathbf{A},\mathbf{B}}(1, \mathbf{A})$ is not $\mathbf{B}$, but $\mathbf{A}$, distorted in a way to be similar to $\mathbf{B}$. We compute a sequence of $n + 1$ images $\mathbf{I}_k$, $0 \leq k \leq n$ as

$$\mathbf{I}_k = \frac{k}{n} \mathscr{F}_{\mathbf{A},\mathbf{B}} \left(\frac{k}{n}, \mathbf{A} \right) + \frac{n-k}{n} \mathscr{F}_{\mathbf{B},\mathbf{A}} \left(\frac{n-k}{n}, \mathbf{B} \right). \tag{3.2}$$

Figure 3.4 illustrates this process. In our particular application we compute intermediate images such that the anisotropy is no more $0.5 : 0.5 : 7$ μm, but merely $0.5 : 0.5 : 1$ μm. All figures showing our method in this chapter (except parts of Figure 3.4) feature images $\mathbf{I}_k$.

| (a) First image, $\mathbf{A}$ | (b) Second image, $\mathbf{B}$ | (c) Interpolation from $\mathbf{A}$ to $\mathbf{B}$ | (d) Interpolation from $\mathbf{B}$ to $\mathbf{A}$ | (e) Final result, sum of both interpolations |
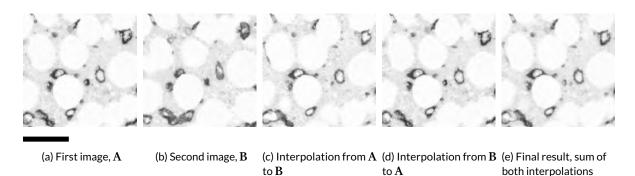
Figure 3.4: Illustrating interpolation with dense optical flow on "bone marrow" data set (Chapter 6), HSV saturation channel, negative image shown. We interpolate with $n = 7$, showing images for $k = 3$. From left to right: first and second input images, $\mathbf{A}$ (a) and $\mathbf{B}$ (b), interpolation from $\mathbf{A}$ to $\mathbf{B}$ (c), i.e., $\mathscr{F}_{\mathbf{A},\mathbf{B}}(3/7, \mathbf{A})$, interpolation from $\mathbf{B}$ to $\mathbf{A}$ (d), i.e., $\mathscr{F}_{\mathbf{B},\mathbf{A}}(3/7, \mathbf{B})$, the final result, i.e., the image $\mathbf{I}_3$ (e). Scale bar is 100 μm.

1. $\mathbf{A} \leftarrow$ read image 1, $\mathbf{B} \leftarrow$ read image 2
2. compute $D_1 = \mathscr{D}(\mathbf{A}, \mathbf{B})$ and $D_2 = \mathscr{D}(\mathbf{B}, \mathbf{A})$ using calcOpticalFlowFarneback
3. **for** $k \in \{0, \ldots, n\}$ **do**
4. $\quad D_1^{(k)} \leftarrow$ scale D_1 with factor k/n
5. $\quad D_2^{(k)} \leftarrow$ scale D_2 with factor $(n-k)/n$
6. $\quad$ compute $\mathscr{F}_{\mathbf{A},\mathbf{B}}(k/n, \mathbf{A})$: apply the map $D_1^{(k)}$ to image $\mathbf{A}$ with remap, yielding $\mathbf{A}^{(k)}$
7. $\quad$ compute $\mathscr{F}_{\mathbf{B},\mathbf{A}}((n-k)/n, \mathbf{B})$: apply the map $D_2^{(k)}$ to image $\mathbf{B}$ with remap, yielding $\mathbf{B}^{(k)}$
8. $\quad \mathbf{I}_k \leftarrow$ alpha-blend $\mathbf{A}^{(k)}$ and $\mathbf{B}^{(k)}$ with weights k/n and $(n-k)/n$
9. $\quad$ write image $\mathbf{I}_k$.

Figure 3.5: Pseudo code for our method.

Pseudo code Figure 3.5 details on our method in pseudo code. We compute the optical flow twice. Notice that due to different displacement representations in calcOpticalFlowFarneback and remap, the optical flows D_1 and D_2 need to be converted to another representation, the x- and y-maps. This operation is, however, trivial at large. Only after the conversion the flow can be scaled. To summarise: we gradually apply the dense optical flow to input images, blending the results to final images $\mathbf{I}_k$.

3.3 Results

First we present the specimens to which we apply our method. Then we showcase how our optical flow-based interpolation improves 3D reconstruction. Finally, we compare our method to the state of the art.

3.3.1 Specimens and the processing pipeline

We report our finding based on multiple data sets. The specimens are listed in Table 3.1. All specimens are human and have been stained for transmitted light microscopy:

Table 3.1: Medical data sets used. All specimens are human. Resolution is stated in μm/pixel.

Organ	Antigen	Stain colour	Resolution	Figure	Chapter
Spleen 1	CD34	brown	0.6	3.6	
Spleen 2	CD34 CD271	brown violet-blue	0.416	3.2, 3.1	Chap. 7
Spleen 3	CD34 SMA CD271	brown brown violet-blue	0.5	3.8–3.11	
Bone marrow	CD34 CD141	brown brown	0.5	3.4	Chap. 6
Tonsil	CD38 SMA	brown brown	0.208	3.7	

- Spleen, single staining: 24 serial sections, immunostained for the glycoprotein CD34 (brown), thus showing primarily capillary walls. Data were acquired by VMscope GmbH with a scanning microscope at 0.3 μm/pixel resolution. The current ROI was resized during processing. The specimen is called here 'spleen 1'. (See also Chapter 7.)
- Spleen, double staining: same specimen as above (24 serial sections), additionally stained for CD271 (violet-blue) after first acquisition. It shows primarily capillaries (brown) and special structures called 'capillary sheaths' (violet-blue). Data were acquired using a Canon 60D camera on a Zeiss Axiophot microscope. We refer to these data as 'spleen 2'. (See also Chapter 7.)
- Spleen, double staining: 24 serial sections, immunostained for the glycoprotein CD34, smooth muscle cells (both brown), and for CD271 (violet-blue). This staining shows primarily capillaries, capillary sheaths, and smooth muscle alpha actin (SMA, targeted with asm-1) in the walls of larger blood vessels, such as arteries and arterioles [Steiniger et al., 2014a,b]. The interior of follicles is also weakly stained in blue. We call these data 'spleen 3'.
- Bone marrow: 21 serial sections, sectioned using a hard plastic embedding technique [Steiniger et al., 2013], immunostained for the glycoproteins CD34 and CD141 (both brown) [Steiniger et al., 2016], thus visualising microvessels (both capillaries and sinuses) and larger blood vessels. We refer to these data as 'bone marrow'. (See also Chapter 6.)
- Tonsil pilot experiment: 5 consecutive sections at 4 μm, immunostained for glycoprotein CD38 and SMA (both brown). Data were acquired using a Canon 60D camera on a Zeiss Axiophot microscope. We refer to these data here as 'tonsil'.

Large images (all except 'spleen 2' and 'tonsil') were pre-registered using a whole-slice registration method [Ulrich et al., 2014b, Chapter 1], then regions of interest (ROI) were selected. These regions were registered again using a fine-grain registration [Lobachev et al., 2017a, Chapter 2]. The 'spleen 2' data set was finely registered only with the same method. The data set 'tonsil' was registered using a rigid-only method. For further processing a single channel was selected to yield grayscale images. Then anisotropy was reduced with interpolation, the actual focus of this chapter. Further processing was done with *3D Slicer* [Fedorov et al., 2012]. We used *Fiji* to produce z-stacks and as a reference implementation for standard interpolation methods.

3.3.2   Interpolation and 3D reconstruction

Figure 3.2 shows the 'spleen 2' data set evolving from registered serial sections to final rendering. The z-profiles of the input stack and interpolated brown staining (after segmentation) showcase our method (Figure 3.2, (b) and (c)). The volume filtering after interpolation is subtle (d), but it makes some difference, as shown by a PSNR difference image at 5% threshold (e) shows. Figure 3.1 demonstrates the final rendering of both brown (CD34, capillary walls) and violet-blue (CD271, capillary sheaths) stainings, interpolated with our method. The weaker staining of some blood vessels outside of the follicle in Figure 3.2 is deliberate. It represents further blood vessel types, irrelevant for present investigation. Only structures that are very intensively stained in blue are represented in Fig. 3.1 (viz. Chapter 7), these correspond to capillary sheaths.

Figs. 3.6 and 3.7 show 'spleen 1' and 'tonsil' data sets respectively. The anisotropy is reduced using an interpolation with our method. Border effects are present in Fig. 3.6. This could be avoided when the interpolation would be performed on a slightly larger data set that would be consequently cropped down.

Figure 3.9 shows the reconstructed meshes from the same input data set 'spleen 3', but for varying interpolation methods. We omit all volume and mesh filtering (aside from mesh repair [Ju, 2004]) for the demonstration purpose. Figure 3.9, (b), shows the final result with our method and all the filtering. Notice that the blood vessel in the foreground features a discontinuity with all methods used, but not with ours. This is due to the optical flow-based interpolation. Our registration method undistorts the sections and aligns them to each other. Nevertheless, rarely microvessels may shift for more than the diameter of small capillaries in the input images. This may happen because of artefacts or registration errors. Through 'shifts' in the flow our method avoids discontinuities, but introduces additional waviness. Our method also features more constant intensity values in intermediate images. A further reason for discontinuities of microscopic blood vessels in other methods is *blending*. Varying intensity in bilinear or bicubic interpolation produces images more pleasant to the eyes, but tightly chosen iso-values in mesh construction yield holes where two structures are blended over (Fig. 3.8).

To give more examples from this thesis, bone marrow reconstructions in Chapter 6 (Fig. 6.4, page 104) and spleen vasculature meshes in Chapter 7 (Figs. 7.1, 7.2, pages 128–129) also use the interpolation method presented here.

3.3.3   Comparing with the state of the art

We compare our method to nearest neighbour interpolation and to bicubic interpolation with *Fiji* [Schindelin et al., 2012] in Figures 3.8–3.11.

The main attribute of our method is that two spatially close, but not necessarily overlapping objects on two consecutive sections are 'morphed' into each other and not 'blended', i.e., are not 'faded out' from one object while 'faded in' to another. The latter is typical of interpolation, but results in smaller intensities during 'blending' that contribute to discontinuities in reconstructed meshes (Fig. 3.9). We *do* use alpha-blending in our computation, but the 'fading' effect is much less with our method. For our medical application retained connectivity is more important than minor shape changes. This behaviour is visible in Fig. 3.11.

Figure 3.8 shows all intermediate images between two consecutive frames, bicubicly interpolated and generated with our method.
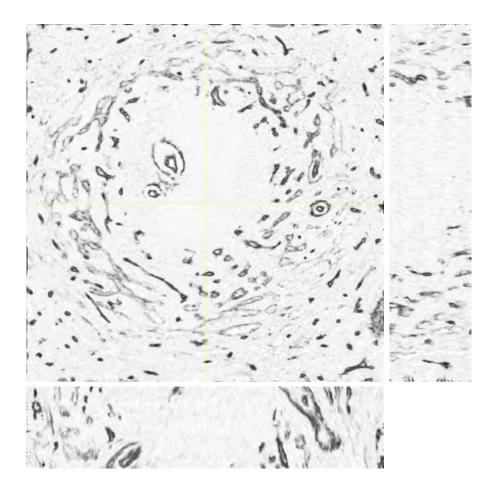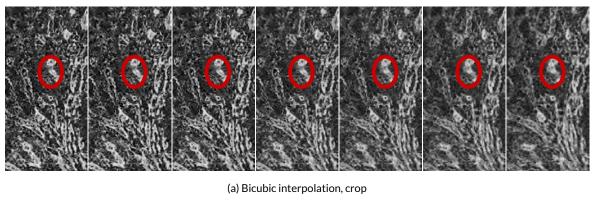
Figure 3.6: Interpolation with our method for 'spleen 1'. We interpolated across 24 serial sections.
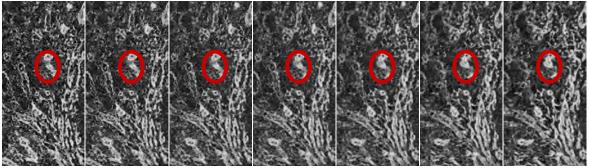
Figure 3.7: A z-profile after interpolation of 'tonsil' data set with our method. The data was generated from 5 initial sections.

3.3.4 Application to generic data

To showcase that our method is applicable not only in medical imaging, we applied our method to images from USC-SIPI image database. We utilised the images 'moon' (5.1.09, upscaled to 512×512 pixels, Fig. 3.12) and 'bridge' (5.2.10, Fig. 3.13). In both cases we applied different distortions to the original image, yielding two input images. Then we generated 12 interpolations between these distorted images with our method. Figure 3.13 shows multiple output images $\mathbf{I}_k$ from our method along with the original image and input data. Here, we have cropped the images to the centre to reduce border effects. We also crop after interpolation in the production use of our method in medical imaging. We deem the results of our method on test images as usable. Thus our method is applicable not only to medical images.

(a) Bicubic interpolation, crop

(b) Our method, crop

Figure 3.8: Interpolation sequence between two consecutive slices of 'spleen 3', a 250×500 crop from centre. (a): bicubic interpolation, (b): our method, $n = 6$. Observe the movement in the marked area.

3.4 Conclusions

We present a special non-linear interpolation method to compensate anisotropy in histological serial sections. Using dense optical flow we determine the "movement" in the consecutive images. This movement is then gradually applied, generating intermediate images. This way our method retains more connectivity in the final meshes than standard interpolation methods. The benefit of our method consists in less intensity changes in interpolated images, and thus less danger that marching cubes algorithms interpret decreasing intensity as a discontinuity.

Our method has been applied to serial sections of specimens from varying human organs. It facilitated advances in fundamental medical research concerning bone marrow (Chapter ~ 6) and spleen (Chapter 7) microvasculature. Our method allowed to improve 3D representations, leading to new insights in microanatomy.

3.4.1 Future work

As we pre-register the data anyway, it is possible to retain the distortions from non-rigid registration phase. One might use these data and not dense optical flow for interpolation. In any case, it would make sense to generalise the distortion data to more than two images and to use e.g., spline-based interpolation. Applications of our method to serial sections from further organs and to data acquired

(a) Overview   (b) Detail, our method, final result with volume and mesh filtering

(c) Detail, nearest neighbour   (d) Detail, bicubic interpolation   (e) Detail, our method
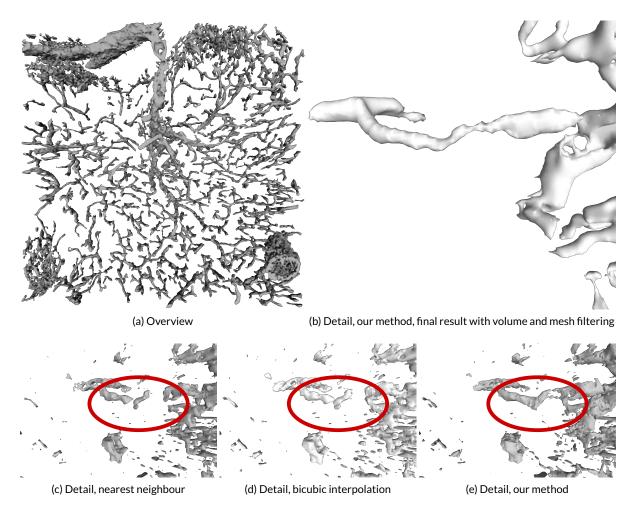
Figure 3.9: 3D reconstructions of the same 'spleen 3' data set after varying interpolation methods. To put more emphasis on interpolation, we used no volume filtering and only mesh healing with PolyMender [Ju, 2004] in (c)–(e), while (b) shows the final result of our pipeline. Notice the discontinuity of blood vessel in (c)–(d), it is avoided in our method (b), (e).

from other modalities (e.g., electron microscopy or MRT) are, of course, possible.

Acknowledgements

(a) Original data, slice 11 of 24

(b) Nearest neighbour

(c) Bicubic interpolation

(d) Our method

Figure 3.10: Comparing various interpolation methods for a 'spleen 3' data set, size $2k \times 2k$ pixels, 24 initial sections. All interpolation methods (b)–(b) operate on CMYK key (black) channel of original data (a). We show the centre of the 24 slice data set, the z-profiles are also cut in the centre of the data set, at $1k$ pixels. We show the negative of the interpolated data. Notice that though our method (d) produces some 'waviness', interpolated data keep similar intensity for varying spatial positions. This is not the case with usual interpolation methods (b), (c). Scale bar = 500 µm.

(a) Bicubic interpolation   (b) Our method
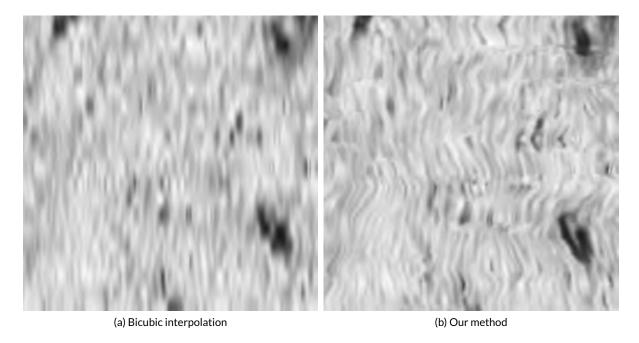
Figure 3.11: Crops from xz-plane stacks from Figure 3.10 at a larger magnification. Bicubic interpolation (a) does not connect slightly misaligned dark blobs (we assume these are capillaries). Our method (b) is able to do so, but produces some waviness.

(a)   (b) $\mathbf{A}, k = 0$ (c) $k = 6$ (d) $\mathbf{B}, k = 12$
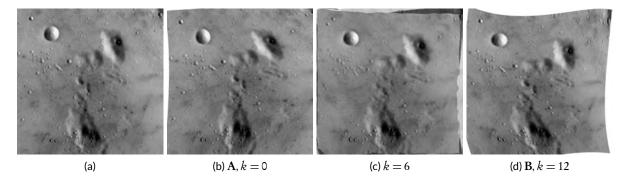
Figure 3.12: Our method on distorted 'moon' images. We show full images, notice the border effects. (a) Upscaled original image. (b) First input image. (c) Interpolation result, image from the middle of the interpolation stack. (d) Second input image.

(a) original   (b) not cropped **A**   (c) not cropped **B**

(d) **A**, $k = 0$ (e) $k = 2$ (f) $k = 5$

(g) $k = 8$ (h) $k = 10$ (i) **B**, $k = 12$
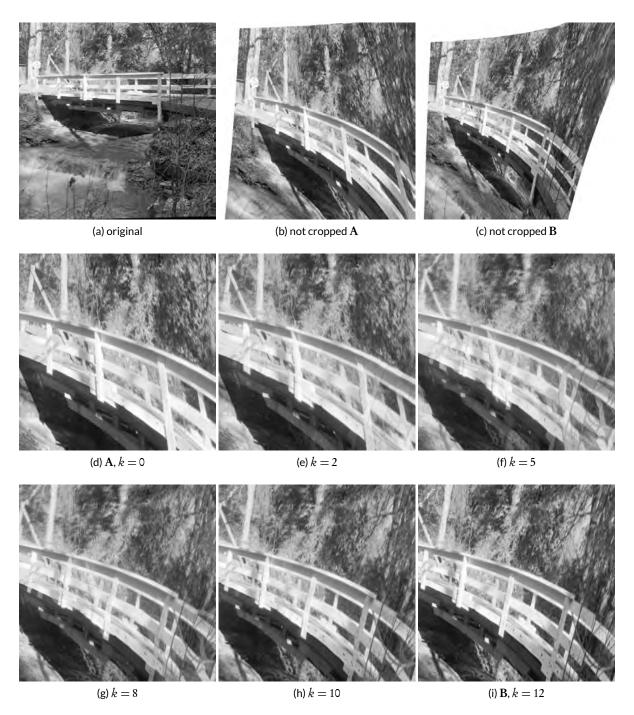
Figure 3.13: Our method on distorted 'bridge' images. Top row shows the original image and both distorted input images. Middle and bottom rows show the cropped input images for reference and the in-between images generated with our method. (a) shows original image, (b) is the full first input image, (c) is the full second input image; (d) is the cropped first input image **A**; (e)–(h) are the interpolated images, $k = 2, 5, 8, 10$ respectively, $n = 12$; (i) is the cropped second input image **B**.

PARALLEL ISO-SURFACE EXTRACTION AND SIMPLIFICATION

In which we discuss how to perform marching cubes on large data sets. This is of relevance also for our medical applications.

This is joint work [Ulrich et al., 2014a] with Christine Ulrich, Nico Grund, Evgenij Derzapf, and Michael Guthe. **Author Contributions:** *Conceptualisation: MG. Data curation: CU, ED. Formal analysis: CU, ED. Funding acquisition: MG. Investigation: CU, NG, ED, OL, MG. Methodology: MG, CU, ED. Project administration: CU, MG. Resources: CU, NG, ED, MG. Software: CU, ED, NG. Supervision: MG. Validation: CU, MG, ED, OL. Visualisation: CU. Writing—original draft: CU, MG. Writing—review & editing: CU, MG, OL, ED, NG. Corresponding author: CU. Presentation: OL.*

When extracting iso-surfaces from large volume data sets, long processing times are required and a high number of polygons is generated. We propose a massively parallel iso-surface extraction and simplification algorithm. The extraction is based on the marching cubes algorithm. In order to process large volume data sets, we perform the extraction with an interleaved simplification step using parallel edge collapses and the quadric error metrics.

Interleaving extraction and simplification is based on locally postponing collapse operations close to the processing front. In contrast to previous methods, we do not need an explicit simplification error fall-off close to the front. Thus we can produce meshes with the same quality as if we would simplify the complete mesh after extraction. By implementing both extraction and simplification on the GPU, we can reconstruct high quality iso-surfaces from large data sets within a few seconds.
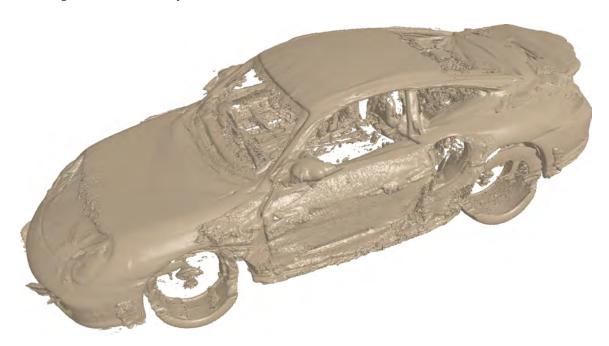
Figure 4.1: Iso-surface extraction and simplification result for the "Porsche" data set with an iso-value of 14. *(Continued on page 69.)*

N THE last decades, the resolution of volume data sets has constantly grown and memory requirements increased. When extracting iso-surfaces from such large volume data sets, a high number of polygons is generated, which also results in long processing times. To efficiently generate meshes from such volumes (especially medical data sets like high resolution histological scans), a fast and efficient algorithm is needed. We combine a massively parallel iso-surface reconstruction algorithm with efficient, high quality parallel simplification.

The marching cubes algorithm is commonly used for iso-surface extraction. It allows producing meshes from large data sets in reasonable time, but—like almost all of its variants—produces a high number of triangles.

Hierarchical approaches [e.g., Müller and Stark, 1993] were proposed to reduce the triangle count during iso-surface extraction. These are however less efficient for tubular structures because they can only uniformly adapt the resolution. To further reduce the mesh, simplification algorithms can be used *after* the extraction or directly *during* this process. The latter can be implemented using a plane sweep algorithm (compare Fig. 4.2) or by partitioning the volume into a regular grid or a space partitioning hierarchy. The processing front—i.e., the outer surface of the already processed volume—needs to be constrained during simplification. Vertices on the front must not collapse, which leads to artefacts, like sliver triangles. The so-called tandem algorithm [Attali et al., 2005] alleviates this problem using a time-lag. The simplification error gradually falls off closer to the front and thus the vertex density increases. While the mesh quality is better, the error falloff constitutes a trade-off between memory consumption and mesh quality.

The main contribution of this chapter is a novel formulation of the tandem algorithm that does not require a time-lag. The key advantage is that the resulting mesh is equivalent to that produced by a complete extraction followed by a sequential simplification. Our implementation is based on a parallel edge collapse algorithm using the quadric error metrics to optimise the vertex positions and normals. Using the GPU for both extraction and simplification allows us to quickly produce high quality meshes from large volume data sets.

4.1   Related work

The main idea of our approach is to interleave iso-surface extraction and simplification on the GPU. We therefore review related approaches on iso-surface extraction and simplification, as well as combinations of both.

4.1.1   Iso-surface extraction

The marching cubes (MC) algorithm [Lorensen and Cline, 1987] divides a voxel grid into cubes and processes each of them independently. Several improvements, including adaptive and dual algorithms, were proposed [Wyvill et al., 1986, Heiden et al., 1993, Montani et al., 1994, Schaefer and Warren, 2004], mainly to prevent holes in the mesh or to improve the mesh quality. Using tetrahedrons [Treece et al., 1999] instead of cubes is also possible. It prevents holes but drastically increases the amount of triangles. Chernyaev [1995] proposed a topologically correct iso-surface generation based on tri-linear interpolation. Other approaches directly reconstruct the surface from sets of orthogonal slices [e.g., Sviták and Skala, 2002, 2004] using contour matching instead of interpolation. Newman and Yi [2006] provide an extensive survey on different MC algorithms and variants.

Performance optimizations include using a modified branch-on-need-octree with min-max decisions [Wilhelms and Gelder, 1990] and span spaces [Shen et al., 1996, Livnat, 1999] which represent
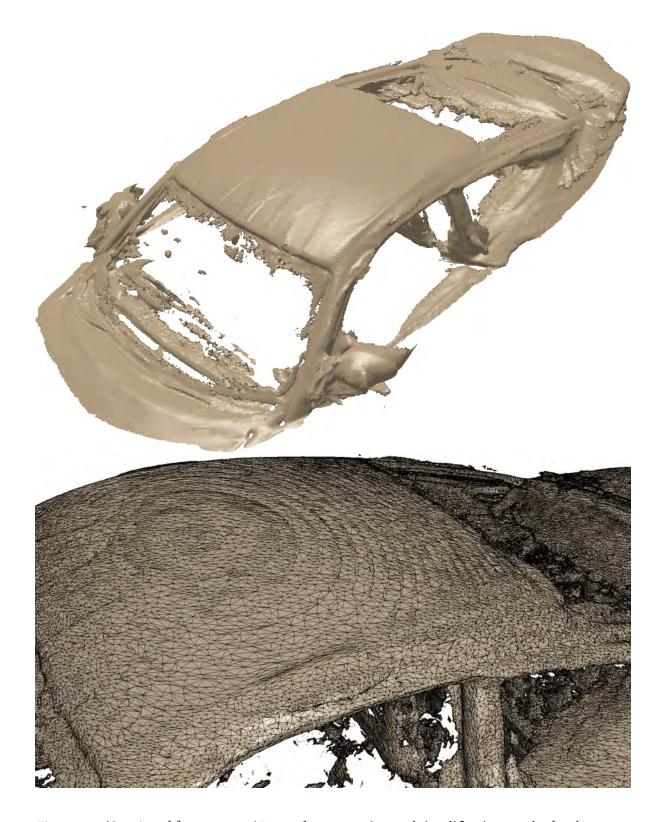
Figure 4.1: (Continued from page 67.) Iso-surface extraction and simplification results for the "Porsche" data set with an iso-value of 14. Top image shows an intermediate result and bottom image shows the triangles of the generated mesh.
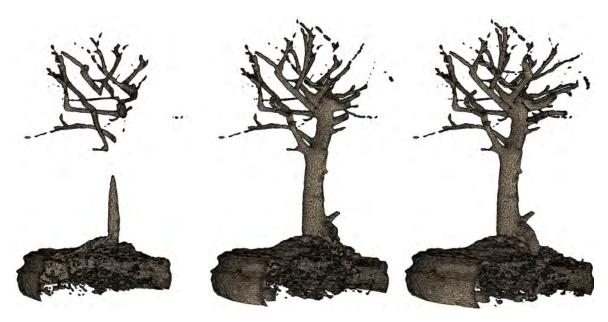
Figure 4.2: These three pictures show the sweep-line. Notice the automatic increase of the triangle density towards the processing front of the extraction.

the cubes as two-dimensional points of their min-max values. For a fast extraction of the iso-surface, partitions (lattice elements) or k-d-trees are used.

The splitting boxes algorithm [Müller and Stark, 1993] was one of the first adaptive methods. A cube intersected by the iso-surface is recursively split and simply checked for sign-changes in every edge. The dual marching cubes algorithm [Schaefer and Warren, 2004] improves the mesh quality and is able to also generate quad meshes. It is based on using the dual of the volume grid, i.e., it places a vertex in each cell that is crossed by the iso-surface. The quadric error function combined with the method of Lindstrom [2000] for positioning dual vertices generate a mesh with better quality and less triangles. This algorithm is also very good in reconstructing sharp features (e.g., edges or corners) but the surface does not accurately approximate the tri-linearly interpolated iso-value.

Recent approaches exploit the processing power of massively parallel graphics processors (GPUs). Reck et al. [2004] proposed an algorithm to extract iso-surfaces from tetrahedral volumes. They pre-select the intersections of surface and voxel grid on the CPU and generate the mesh on the GPU using an interval tree. For rectilinear grids, a tetrahedrisation is required, which leads to a higher number of triangles and also introduces artefacts. Johansson and Carr [2006] also use span spaces for pre-selection and pre-classification. They utilise GPU for interpolation and their approach is not restricted to tetrahedral grids any more. Tatarchuk et al. [2007] propose a hybrid implementation of marching cubes and marching tetrahedra running on the GPU. The iso-surface is generated on the GPU but again contains a higher number of triangles than the original marching cubes algorithm due to the tetrahedrisation. An improvement to marching cubes on the GPU uses the so-called HistoPyramids [Dyken et al., 2008].

4.1.2   Simplification

Mesh simplification is one of the most common techniques for real-time rendering of complex polygonal models and has been an active field of research over the last two decades. A detailed

review of simplification algorithms is given by Luebke [2001]. As we aim at efficient iso-surface reconstruction and simplification from large volumes, we focus on real-time capable simplification algorithms.

Uniform vertex clustering [Rossignac and Borrel, 1993] subdivides the bounding box of the model into cells using a regular grid. All vertices inside the same grid cell are collapsed to their mean. An improved variant is weighted vertex clustering [Low and Tan, 1997]. It better preserves features that are not aligned with the grid. Uniform clustering is relatively fast and gives a precise upper bound for the simplification error. However, a further reduction in flat regions is still possible without increasing the simplification error.

The vertex pair contraction [Garland and Heckbert, 1997, Popović and Hoppe, 1997] has become the most common technique for the simplification of mid-sized triangle meshes. In combination with the introduced quadric error metric, it allows a flexible control over the geometric error and can be used to calculate optimal vertex positions. This approach was also extended to handle an arbitrary number of vertex attributes [Garland and Heckbert, 1998]. A combination of vertex clustering with error quadrics [Lindstrom, 2000] improves the placement of the clustered vertices, but still uses a high number of triangles in flat regions. Shaffer and Garland [2001] proposed to overcome this problem by using a BSP tree instead of a uniform grid. This increased the run time significantly compared to uniform clustering, but the method is still faster than edge collapse simplification. An adaptive vertex clustering using octrees was also proposed by Schaefer and Warren [2003]. Here the run time is even higher than using a BSP tree, but the quality of the simplified mesh can almost compete with edge collapse simplification.

DeCoro and Tatarchuk [2007] proposed a parallel GPU implementation of vertex clustering [Schaefer and Warren, 2003] by implementing an efficient GPU-based data structure. While the performance is very high, it still has the same quality problems as uniform vertex clustering. Recently, a parallel GPU implementation of edge collapse simplification using quadric error metrics [Grund et al., 2011] has been proposed, which we extend in this chapter.

4.1.3   Hybrid algorithms

The hybrid algorithms of Attali et al. [2005] and Dupuy et al. [2010] directly simplify the iso-surface mesh during extraction. The first one uses a sequential marching cubes in a tandem with a simplification algorithm. The extraction and simplification steps alternate layer-by-layer. A *time-lag* is introduced to delay collapses until the extraction front is further away, resulting in a better quality of the simplified mesh. The algorithm of Dupuy et al. [2010] improves this approach by using a load-balanced cluster to parallelize extraction and simplification. In addition, they do not use a plane sweep any more but partition the volume using an octree. Finally, those parts of the mesh that cannot be further simplified are stored on disk to reduce the memory consumption. A similar approach has been proposed for the reconstruction of surfaces from point clouds has been proposed by Cuccuru et al. [2009]. It combines streaming, iso-surface extraction, and simplification by applying local vertex clustering with topology preservation to produce good quality meshes.

4.2   Overview

The core idea of our approach is to interleave a parallel marching cubes with a parallel stream simplifier. As both algorithms run on the GPU, we also minimise communication between host and device as we only transfer the reduced mesh. We chose a plane sweep partitioning due to simplicity, i.e., we process the data in layers, although other partitionings would easily be possible.

Our method consists of two basic modules. The first one is the actual surface extraction based on the marching cubes algorithm. The second module is the mesh simplification that receives the output of the first one as input. The simplification is based on edge collapse operations that contract edges by collapsing two connected vertices (Fig. 4.3). The position and normal of the collapse vertex are computed by minimising the quadric error metric [Garland and Heckbert, 1997] which also defines the collapse cost.
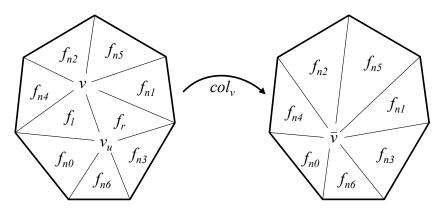
Figure 4.3: Edge collapse. The edge defined by vertex v and v_u is collapsed into the vertex $\overline{v}$.

4.3   Parallel marching cubes

Similar to Attali et al. [2005], we partition the volume into layers for processing. We do however split the volume into partitions containing multiple layers if they fit into memory and process them in parallel. The first kernel calculates the *cube codes,* the eight-bit values encoding if the vertices are inside or outside. This kernel also calculates the intersections between the iso-surface and the cube edges if there are any. Assuming that the four corners and four edges of the cube's bottom are already processed, every thread processes vertex v_4 and the edges e_4, e_7 and e_8 only (shown in blue in Figure 4.4). In addition, the kernel computes the gradient at vertex v_4. Note, that we use a specific kernel for the first slice of each partition to improve performance. The second kernel produces the triangles and thus generates the mesh contained in the layers. This kernel also removes degenerate triangles and feeds the mesh to the second module, the simplification.

We divide a volume of dimension $dim_X \times dim_Y \times dim_Z$ into *layers* and *slices* like Attali et al. [2005] did. The k^{th} *slice* contains all vertices with the same y-coordinate. The k^{th} *layer* is the set of all vertices, edges and patches between or on the k^{th} and the $(k+1)^{\text{st}}$ slice. Figure 4.5 shows the relationship between slices and layers. So the volume comprises the slices from 0 to $dim_Y - 1$ and layers from 0 to $dim_Y - 2$. We then process the layers in ascending order.

We group the slices and layers into partitions during surface extraction. Every partition is composed of N slices and $N-1$ layers respectively. These are processed in two loops and mainly two kernels (see Figure 4.6). Note that two additional slices are required to compute the surface normals from the gradients. Consequently, the partitions overlap by one slice in each direction. As the vertices of the first slice were already calculated in the previous partition, we do however only need one additional slice in each partition. This means that a partition with N slices only contains $N-2$ layers that can be used for the extraction (Figure 4.6). In addition, we use separate kernels for the first and last slices for this reason.
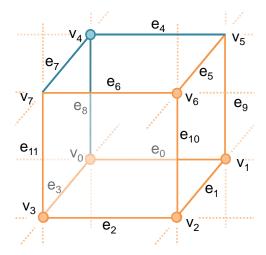
Figure 4.4: Local edge and vertex indices used during the extraction. Each thread processes the blue elements only.
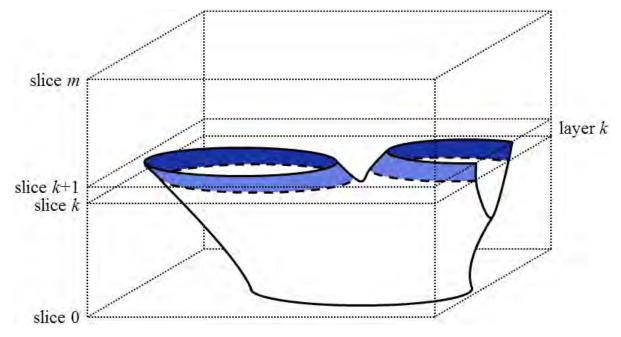
Figure 4.5: Relationship between slices and layers. The k^{th} layer is composed of slices k and $k+1$.

The cube code kernel is executed for each cube (see Figure 4.7) using a thread block dimension of 16×16 or 32×32 depending on the GPU. Assuming that the four corners and four edges of the cube's bottom are already processed, every single thread just processes vertex v_4 and the edges e_4, e_7 and e_8 (see Figure 4.4). An exception are threads on the 'right' and/or 'front' border of a grid that need to process additional vertices.

Every thread also calculates the gradient of vertex v_4. The gradients and the cube codes always have to be calculated for the previous layer as well. A thread sets the first four bits of the cube code and then shifts the code four bits to the right at the end. The parallelization enforces to enumerate cube edges not just locally, as depicted in Figure 4.4, but also globally. This prevents a repeated

Partitions, Remainder = gen_partitions(dim_Y)
kernel_cubecode_init()
for each *partition_index* in *Partitions* **do**
    kernel_cubecode(*partition_index*, $N - 2$)
    kernel_generate(*partition_index*, $N - 2$)
    call simplification_module()
kernel_cubecode_other(*Remainder*)
kernel_generate_other(*Remainder*)
call simplification_module()

Figure 4.6: Parallel extraction and simplification.

k = *partition_index*($N - 2$)   *// compute the global index k*
for each layer in a *partition* **do**   *// a partition consists of N − 2 layers*
    for each *cube* $\in k^{\text{th}}$-layer **do in parallel**
        calculate_gradients()
        generate_cubecode()
        if $0 < cubecode < 255$
            calculate_intersections()
        shift_cubecode()
    $k = k + 1$

Figure 4.7: Cube code kernel.

Table 4.1: Memory consumption and data structures required for the iso-surface extraction. The size of the input grid is $dim_X \times dim_Y \times dim_Z$, while N is the number of slices in each partition.

Buffers	Memory (bytes)
slices	$N dim_X dim_Z$
gradients	$12N dim_X dim_Z$
voxel edges	$24((3N - 1)dim_X dim_Z - N(dim_X + dim_Z))$
cube code	$(N - 1)(dim_X - 1)(dim_Z - 1)$
triangles	$60(N - 1)(dim_X - 1)(dim_Z - 1)$

calculation of intersection points between the iso-surface and cube edges. Figure 4.7 gives an overview of the cube code kernel. Note that Figures 4.6 and 4.7 omit special cases that have to be handled explicitly, e.g., when the volume contains only a single partition.

After the first kernel has calculated the intersections, gradients and cube codes, *kernel_generate* produces the mesh using the classical marching cubes look up table. The look up table stores the topology of the surface. This means that and only intersections of cube edges and the iso-surface have to be calculated and the corresponding mesh is fetched from the look up table. The result is a level set $I_\rho := \{v \in \mathbb{R}^3 : f(v) - \rho = 0\}$, where ρ is the iso-value. Table 4.1 lists all data structures required during iso-surface extraction. Finally, this kernel also removes degenerate triangles and feeds the mesh to the second module, the simplification.
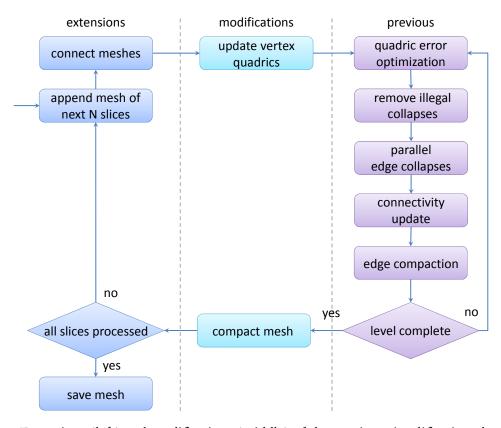
Figure 4.8: Extensions (left) and modifications (middle) of the previous simplification algorithm (right).

4.4 Parallel simplification

The simplification module is based on a parallel GPU simplification [Grund et al., 2011] originally developed for in-core processing of a single mesh. We implemented the following changes to combine it with the stream extraction: First, the mesh generated from N slices is transferred from extraction module. At this point, it does not matter anymore how the extraction module operates, i.e., we could simply replace the marching cubes algorithm with marching tetrahedrons or dual marching cubes. Then the vertex quadrics are computed for all new vertices and vertices on the previous processing front. The vertices along the processing front are marked in the extracted mesh such that they can be identified for the constrained simplification and to later combine the mesh with the next partition. The edge data structure is filled as in the original simplification algorithm [Grund et al., 2011] and the parallel simplification loop starts. The quadric error is optimised and illegal collapses are removed. An edge can only collapse if it is not connected to the processing front, i.e., none of its vertices are marked. In addition, collapses of direct neighbours are also not possible since we do not know their local ordering yet. This can be easily achieved by setting the error of edges connected to the processing front to -1. After removing all illegal collapses, the operations are applied. Finally, the collapsed edges are removed after updating the face and edge connectivity. If no further collapses are possible, the next partition is added to the partially simplified mesh. During compaction, the new positions of the processing front vertices inside the vertex buffer are stored in a look up table. Figure 4.8 shows an overview of the simplification process.

Instead of a time-lag [Attali et al., 2005], our simplification algorithm locally postpones only

Figure 4.9: Interleaved extraction and simplification (bonsai #2, iso 50). The mesh is gradually simplified as the processing front moves without using an explicit time-lag.

those collapse operations that cannot be performed yet. As the collapse operations are local, a mesh with same quality can be produced by any global operation ordering, as long as the local order remains fixed. We exploit this by simply blocking all operations in the direct neighbourhood of the processing front. When additionally enforcing the correct local ordering, the simplification error automatically decreases for vertices close to the current boundary (Figure 4.9). This way, the result will be identical to a simplification of the complete mesh and the operations are performed as soon as possible minimising memory consumption.

The claim above can be proven by analysing which operations are performed by a sequential simplification and by our algorithm. A sequential simplification collapses the vertex pair with the smallest error until a given threshold is reached. As the error of the neigh boring edges increases during the collapse, their error will always be higher after collapsing them, if it was higher before. This implies that the edge is also collapsed by our method at some time during the simplification. If an edge is collapsed by our method, then all neighbouring edges have a higher error. This means that a sequential algorithm also has to collapse exactly this edge before neighbouring edges could be collapsed. Finally, setting the error of edges connected to boundary vertices temporally to −1 and preventing their collapse, only causes a local delay of the simplification, because the neighbouring edges also cannot collapse yet.

When the next partition is added, we simply need to check the marked vertices—i.e., those on the processing front—and find them in the next partition. Using a plane sweep algorithm, these will simply be the first extracted vertices. For other partitioning schemes we simply need to assure that the vertex order is the same in the partitions to combine both meshes in linear time.

4.5   Results

We performed our evaluations on a PC with Intel Core i7 CPU (3.33 GHz), 6 GB of main memory, and an NVidia GTX 580 graphics card. We used CUDA 5.0 to implement the parallel reconstruction and simplification. The volumes originate from Stefan Roettger's volume library[1]. Table 4.2 gives an overview of the volume data sets we used for evaluation. The largest volume is the "Porsche" car model. It also produces the highest number of faces and thus has the highest resource consumption. Both "bonsai #2" and "CTA head" are medium sized. The bonsai contains many ramifications, leading to a high number of faces in the simplified model. Table 4.3 shows the relative and absolute number of cubes crossed by the iso-surface, where the number of generated triangles is roughly twice of that. As the generated mesh also depends on the iso-value, we denote it together with the model in the following.

Table 4.2: Dimension and file size of the models [Roettger, 2013] used for evaluation. The size denotes MBytes in RAW format.

Model	Dimensions	Size
"bonsai #2"	$512 \times 189 \times 512$	48 384
"CTA head"	$512 \times 120 \times 512$	30 720
"Porsche"	$559 \times 347 \times 1023$	193 784

---

[1]http://www9.informatik.uni-erlangen.de/External/vollib/
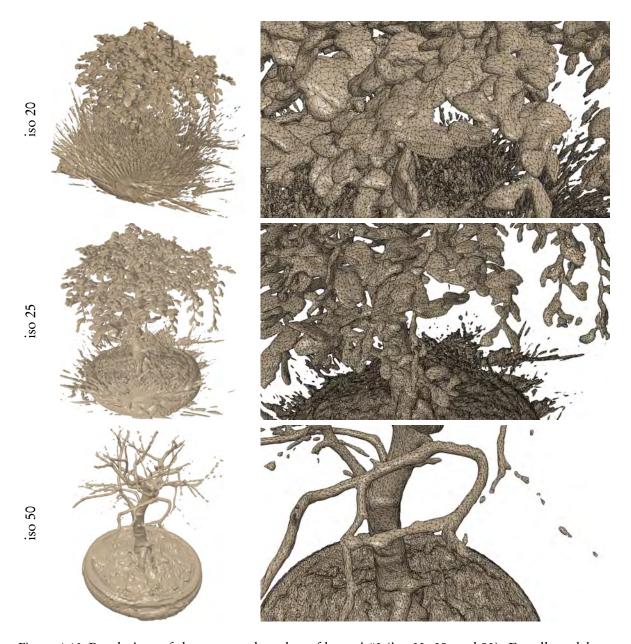
Figure 4.10: Renderings of the generated meshes of bonsai #2 (iso 20, 25, and 50). For all models, closeups with wire frame overlay are shown to the right.

4.5.1   Performance

We used 12 layers per iteration of the algorithm as partition size. The extraction time is almost linear in the number of cubes and slightly increases with the number of generated faces (see Table 4.4) due to the embarrassingly parallel nature of the marching cubes algorithm. The extraction performance ranges from 12.0M ("CTA head", 60) to 12.9M ("Porsche", 14) cubes per second and the number of triangles from 174k ("bonsai #2", iso 25) to 2.24M ("CTA head", iso 60) per second. The memory consumption is dominated by the volume and gradient data in the current partition with a small overhead for the simplified mesh. Note that there is no significant difference in processing times or
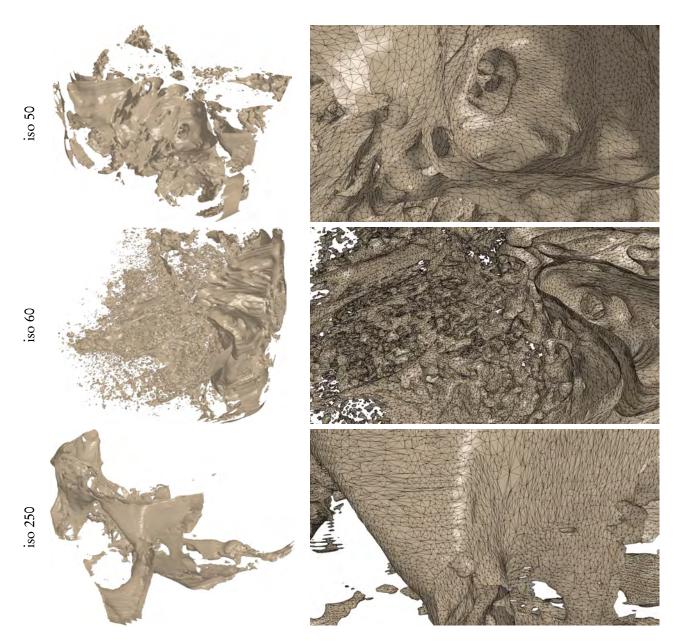
Figure 4.11: Renderings of the generated meshes of the CTA head (iso 50, 60, and 250). For all models, closeups with wire frame overlay are shown besides of them.

memory consumption between medical and other data sets.

The generated meshes are shown in Figures 4.1 ("Porsche"), 4.10 ("bonsai #2"), and 4.11 ("CTA head"). The lowest simplification performance is achieved for the "bonsai #2" due to the highly curved surfaces and the highest performance is achieved for the "Porsche" data set which contains many almost flat regions. The ratio of triangles before and after simplification ranges from 20.7% (CTA head, iso 60) to 100% (bonsai #2, iso 50). The meshes are simplified at up to 258k collapses per second. The bonsai #2 data set (iso 25) has similar characteristics as the "old bone" [Attali et al., 2005] but almost three the size. With an iso-value of 20 it resembles the "young bone" data set, again with about three times the size. Unfortunately, the original data sets used by Attali et al. [2005]

Table 4.3: Relative and absolute number of crossed cubes depending on the iso-value, given in parenthesis. In addition, the number of generated faces before simplification is shown.

Model	Crossed cubes		Number of faces
	Per cent	Number	
"bonsai #2" (20)	4.5	2 203 645	4 405 952
"bonsai #2" (25)	2.3	1 125 808	2 252 046
"bonsai #2" (50)	0.7	329 858	658 158
"CTA head" (50)	3.1	956 441	1 913 256
"CTA head" (60)	9.5	2 961 041	5 878 764
"CTA head" (250)	2.2	694 745	1 392 432
"Porsche" (14)	2.4	4 744 499	9 580 084

Table 4.4: Computation time in seconds for surface extraction and simplification, number of faces after simplification and maximum memory consumption (in MBytes). Cf. Table 4.3 for number of faces before simplification.

Model	Time, seconds		Number of faces	Memory
	Extraction	Simplification		
bonsai #2 (20)	4.1	5.7	4 396 060	480.5
bonsai #2 (25)	4.0	2.9	2 245 412	430.6
bonsai #2 (50)	3.8	1.0	658 158	394.7
CTA head (50)	2.6	1.2	1 845 976	423.4
CTA head (60)	2.6	5.3	1 216 734	514.2
CTA head (250)	2.6	1.4	1 392 432	411.5
Porsche (14)	15.3	21.1	4 140 690	923.2

were not available for a direct comparison. On our system, the algorithm of Attali et al. needs about 60 seconds for the bonsai #2 (iso 25) and approximately 125 seconds for bonsai #2 (iso 20), so our algorithm is 8.8 times and 12.7 times faster respectively. Our approach also compares favourably with Dupuy et al. [2010] although the latter method scales almost linear with the number of CPU cores (4 in our system) and is thus almost four times faster than Attali et al. [2005].

4.5.2   Memory consumption

Figure 4.12 shows the total memory consumption when processing the "Porsche" data set in detail. Note that the GTX 580 does not have enough memory to store the complete model, so we had to use the partitioning. For each partition, we plot the maximum memory consumption during extraction and simplification along with the number of faces after simplification. While the memory gradually increases with the size of the generated mesh, it is dominated by the data required to process the current partition.
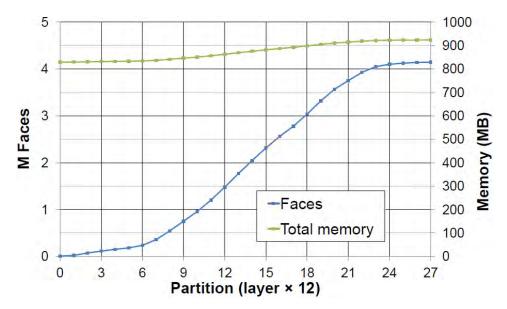
Figure 4.12: Total memory consumption and number of faces in the extracted and simplified mesh after complete processing of the "Porsche" data set.

4.5.3 Mesh Quality

As shown above, our method produces meshes of the same quality as the underlying simplification algorithm. We compared our simplification with simplifying the complete extracted mesh using the algorithm of Grund et al. [2011]. For the smaller models, the runtime and the simplified mesh were the same (up to floating point round-off errors). The larger models like "CTA head" at iso-value 60 and "Porsche" do not fit into the graphics memory of the GTX 580 and cannot be simplified with that algorithm.

We experimented with the error-threshold and analysed the directional bias. In contrast to Attali et al. [2005], the directional bias is—as expected—almost zero because we do not require any time lag.

4.6 Conclusion and future work

Removing the explicit time-lag by locally blocking/postponing simplification operations enables us to reduce memory consumption and processing time. Due to the massively parallel implementation on the GPU, our approach can process large volumes within a few seconds. In addition, it directly benefits from the future improvements of graphics hardware or other parallel systems.

By guaranteeing the same local ordering of collapse operations as a sequential simplification of the complete mesh, we achieve the same quality as the underlying simplification algorithm. This also implies that no artefacts at partition boundaries are introduced. Our improved simplification algorithm works with *any* mesh generation algorithm. So we could easily replace the marching cubes algorithm with dual marching cubes, marching tetrahedrons, or others.

Our current implementation is limited to models for which at least three slices—i.e., one layer—fit into memory because of the plane sweep partitioning. For very large volume data sets, we could use a regular grid instead without changing the core algorithm. The only part that needs to be modified is the fusion of the next partition's mesh with the already simplified one. However, the run time

complexity of this step will still be linear as the vertex ordering of the iso-surface extraction is fixed within each partition.

QUALITY CONTROL OF HISTOLOGICAL 3D RECONSTRUCTIONS IN VIRTUAL REALITY

In which we present a method for quality control of our reconstructions. We utilise an approach in virtual reality. This is joint work with Moritz Berthold, Birte S. Steiniger, and Michael Guthe. The paper is submitted to "IEEE Computer Graphics and Applications".

Author Contributions: *Conceptualisation: OL. Data curation: OL, MB, BSS. Formal analysis: OL. Funding acquisition: MG. Investigation: OL, MB, BSS. Methodology: OL, BSS. Project administration: OL. Resources: MG, BSS, OL, MB. Software: MB. Supervision: MG. Validation: OL, BSS, MG. Visualisation: OL, MB. Writing—original draft: OL. Writing—review & editing: BSS, OL, MB, MG. Corresponding author: OL.*

3D reconstruction is a challenging current topic in medical research. We perform reconstructions from serial sections stained by immunohistological methods. This chapter presents an immersive solution to quality control (QC) such reconstructions. QC is essential to establish correct digital processing methodologies. We propose visualisation in virtual reality (VR) for QC of 3D reconstructions. Domain experts are also able to obtain new insights with our approach. In this manner, we advance research on the microanatomy of human bone marrow and spleen. Both 3D reconstructions and original data are available in VR. QC is streamlined by subtle implementation details and general immersion in VR.

(a)   (b)
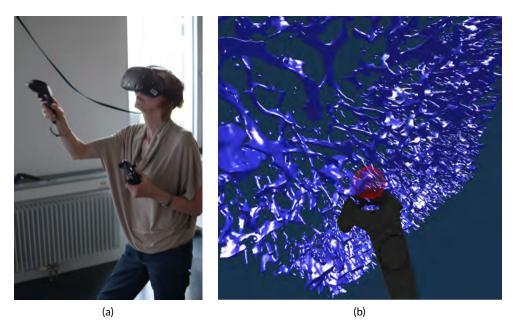
Figure 5.1: Annotating a region with HTC Vive.   a) The user, operating the controller. b) A screenshot of our application showing the view of the left eye of the head-mounted display. The QC visualisation shows a mesh for a human bone marrow specimen, stained for CD34 plus CD141 (Chapter 6).

ISUALISATION is an important part of modern research. Among other fields, it gained a lot of importance in medicine and life sciences. It is important not only to obtain images, but to be able to gasp and correctly interpret them. Further, there is always the question, whether the obtained model is close enough to reality. The method we present in this chapter is a typical example of this very generic scheme, but it is quite novel in the details.

3D reconstruction from histological serial sections closes a gap in medical research methodology, as Appendix A details. The specimen are biologically processed: fixated, embedded, sectioned, stained. Then the sections are scanned and further digitally processed. We normalise, register (Chapters 1, 2), interpolate (Chapter 3) the volumes. The volume data is further filtered, as are the meshes generated from volumes. Chapter 4 details on mesh construction, the filtering is detailed on in Part ˜ II. Both kinds of filtering operations and the actual mesh construction feature some parameters, most prominently, the iso-value. How can we ensure the reconstruction is correct?

We now present an up-to-date method for controlling that 3D reconstructions tightly correspond to the original immunostained sections by directly comparing the surface reconstructions to the original serial sections. This method accelerates QC by showing single sections in the visualised mesh, without volume rendering. In this manner we can verify that the reconstruction features everything needed and nothing unneeded. The decisive innovation is that we inspect the 3D model and the sections in virtual reality (VR). Figure 5.2 documents a QC session from 'outside'.

Figures 5.1, 5.3–5.7 feature screenshots from the 'inside' of VR. Our domain experts are much better trained in distinguishing and analysing details in stained histological sections than in reconstructed meshes. However, only 3D reconstructions provide an overview of continuous structures spanning multiple sections, e.g., blood vessels.

Further, the reconstructed mesh permits novel findings. Domain experts often encounter problems when trying to understand 3D arrangements using conventional mesh viewers. For this reason, we previously showed pre-rendered videos to the experts to communicate the reconstruction results and to enable detailed inspection. Video data are, however, limited by the fixed direction of movement and the fixed camera angle. In VR the user can freely and intuitively control her angle of view and her movement through the model. Immersion allows for much easier and more thorough inspection of visualised data. Occlusion of decisive structures in the reconstruction does no longer pose a problem.

5.1   Background

5.1.1   Related work on immersive visualisation

Immersive visualisation is by far not a new idea, Brooks Jr. [1999] quotes a vision of Sutherland from 1965. However, not a multi-million-dollar training simulation, but an immersive scientific visualisation was quite hard to obtain in earlier years [van Dam et al., 2000]. What has massively changed the game recently is the availability of inexpensive hardware such as Oculus Rift or HTC Vive head-mounted displays (HMD). This fact (and the progress in GPU performance) allows for good VR experiences on commodity hardware. Immersive visualisation has been previously suggested for molecular visualisation [Stone et al., 2010], medical volumetric data [Shen et al., 2008], for dentistry [Shimabukuro and Minghim, 1998], and computational fluid dynamics [Quam et al., 2015]. More relevant to our approach is the visualisation of the inside of large arterial blood vessels [Forsberg et al., 2000]. A recent trend is to utilise VR in medical education and training [Mathur, 2015, Chan et al., 2013]. The availability of head-mounted displays has sparked some new research [Chen et al., 2015, Choi et al., 2016, Inoue et al., 2016] in addition to already covered fields.

Figure 5.2: A user in VR. The large display mirrors the headset image.

A visualisation support for HTC Vive in a popular medical imaging toolkit has been presented before [Egger et al., 2017]. Unlike our approach, their method is tied into existing visualisation libraries. Our method is a stand-alone application, even if easily usable in our tool pipeline. Further, we need to mandatorily visualise both reconstructed meshes *and* original input data, as we focus on QC.

5.1.2 Features of VR visualisation

Our domain experts provided feedback on the earlier versions of the software in order to shape our application. The following features were deemed necessary by medical researchers:

- Load multiple meshes corresponding to parts of the model and to switch between them.
- Load the original data as a texture on a plane and blend it in VR at will. The experts need to discriminate all details in the original stained sections.
- Remove the reconstructed mesh to see the original section underneath.
- Provide a possibility to annotate a 3D position in VR.
- Provide a possibility for voice recording in annotations (work in progress).
- Design a method for sharing the view point and current highlight with partners outside VR (trivial with Steam VR), and for communicating findings in VR as still 2D images in research papers (an open question).

In this chapter we establish the method that allows the experts to compare processed meshes (optionally coloured) with registered ROI images in an immersive setting using a VR head-mounted display.

5.2   System architecture and features

Here we present an overview of the application we developed for VR-based visualisation and QC of reconstructions from histological data.

The following text elaborates on existing software and hardware components we are using, on design decisions in application control, on creation and saving of annotations, on anti-aliasing, and, finally, on the creative usage of front face culling.

5.2.1   Components

Our application makes use of existing software libraries. We load meshes with the VCG library. Multiple meshes with vertex colours are supported. We utilise Open GL 4.2. We enable back face culling, multisampling, and mipmaps. The section textures are loaded with the FreeImage library. The Steam VR library is used for interaction with Vive controllers and the headset.

With respect to hardware, the system consists of a desktop computer with a Steam VR-capable GPU, HTC Vive with controllers, and a XBox One controller. The setup is inexpensive, posing a sharp contrast to the earlier years of VR.

5.2.2   Controls

For control, a simple keyboard and mouse interface (for debugging outside VR), XBox One controller, and Vive native controllers can all be used. Our initial idea was to use a XBox 360 or a XBox One controller, as such controllers provide a simple control metaphor. A XBox controller has a multitude of buttons, triggers, and sticks that can be bound to the various features to control. The controller was slated for manipulating the meshes in orientation, magnification, and position, as well as controlling the number and combination of meshes and slices shown at a time. Sticks were very suitable to rotate and move the objects in space. We attempted to maximise the number of inputs by using the letter keys ('A', 'B', 'X', and 'Y') as function modifiers. However, the expert users were not acquainted to gaming controllers and could not see the XBox One controller in VR. Thus, initial error rates were high when they tried to simultaneously use an 'X' key and a 'D-Pad' (for original section functions) in blind.

A more intuitive approach with the native Vive controllers was targeted. But switching controllers in the immersion was quite cumbersome (compare Figs. 5.2, 5.1). In most cases it turned out that the immersed user did not need to scale, translate, and rotate the mesh with controller input.We hence have kept the keyboard-and-mouse and the XBox controller options, but duplicated most of the input actions with Vive controllers. Native Vive controllers proved their benefits. Although the metaphors were much more complicated, the intuitive control payed off immediately. Further, the visibility of the tracked controllers in VR helped a lot.

Our current usage model is to employ two people on two sides of the immersion, the operator outside and the immersed domain expert in VR. The operator is responsible for initial launch of the software and can correct the zoom and position of the mesh in the rare occasion it is needed. The operator can follow the expert. The HMD feed is mirrored on a monitor (Fig. 5.2). Further, annotations significantly improve communication. The main controller actions of the domain expert are:

- Blending the original data in or out;
- Blending the mesh(es) in or out;
- Advancing the currently displayed section.

The most significant user interaction happens with intuitive movements of the immersed user around (and through) the displayed entities in VR.

5.2.3 Communicating from immersion to the outside and afterwards

We have two or more users. One, typically the domain expert, is immersed into VR. Others, including the operator, are outside, but follow the immersed user on a mirrored screen. The immersed expert can verbally communicate with the users outside. With a Vive controller and our annotation tool interesting areas in the visualisation can be shown to the users outside in real time. The annotations are saved for further analysis. An example use case: after the domain expert has marked suspicious areas, the 3D reconstruction expert can inspect them in a later VR session. Reconstruction improvements can be deduced from this information.

5.2.4 Annotation markers

We designed a spherical selection tool for marking points in space. The sphere is located at the top front of a Vive controller and can be seen in the virtual space (and, by proxy, also in the mirrored display, Fig. 5.1). We need to note, however, that the annotation sphere appears much more vivid to the VR user than it appears on screenshots. For the convenience of the developer, the sphere always has a radius of one in the internal representation, even if sphere radius is freely changeable by the user. Thus, only the full transformation matrix from the controller space to the mesh space and the colour of the marker need to be saved. The actual radius of the sphere can be derived from multiplication of the transformation matrix and a vector of length one.

We have reconstructed the 3D shape of smaller and larger microvessels in hard serial sections of human bone marrow [Steiniger et al., 2016, Chapter 6]. Figure 5.3 shows (a): a single section (part of the input data to reconstruction), (b): a volume rendering of all 21 sections, and (c): our 3D reconstruction. The process of annotation is in Fig. 5.3, (d). Next subfigures show further investigation of the annotated area in VR in combined mesh-section view (e), and when showing only the corresponding section (f). To discriminate between capillaries (smaller, coloured red) and sinuses (larger, coloured green), we computed shape diameter function on the reconstructed meshes and colour-coded the result values on the mesh, as shown in (c)–(e). Chapter 6 provides more details on this. The handling of the reconstruction and serial section data in VR showcases the annotation process.

5.2.5 Annotation protocol

The current state of the visualisation is immediately written into a protocol file. This protocol includes the markers' positions, the list of meshes, and the list of serial sections. The markers set during the immersive annotation session are meant to be processed afterwards. The protocol file can be loaded again into the application to exactly restore the same meshes, original data, and markers as before. In this manner the users can further investigate the issues that were marked earlier.

5.2.6 Anti-aliasing

If a direct rendering approach is used, there is a very dominant aliasing effect at certain points. We used multisampling on meshes and mipmaps on textures to alleviate this problem, as detailed in the next section.

In VR, the 'screen' is extremely close to the viewer. Hence, aliasing becomes troublesome, despite of the relatively big resolution of the Vive HMD: 1080×1200 pixel per eye at 90 fps. Without anti-aliasing the immersed user perceives a lot of glittering and rapid intensity changes. Aliasing mostly occurs on the edges of objects and on the surface of the detailed section textures. The probable
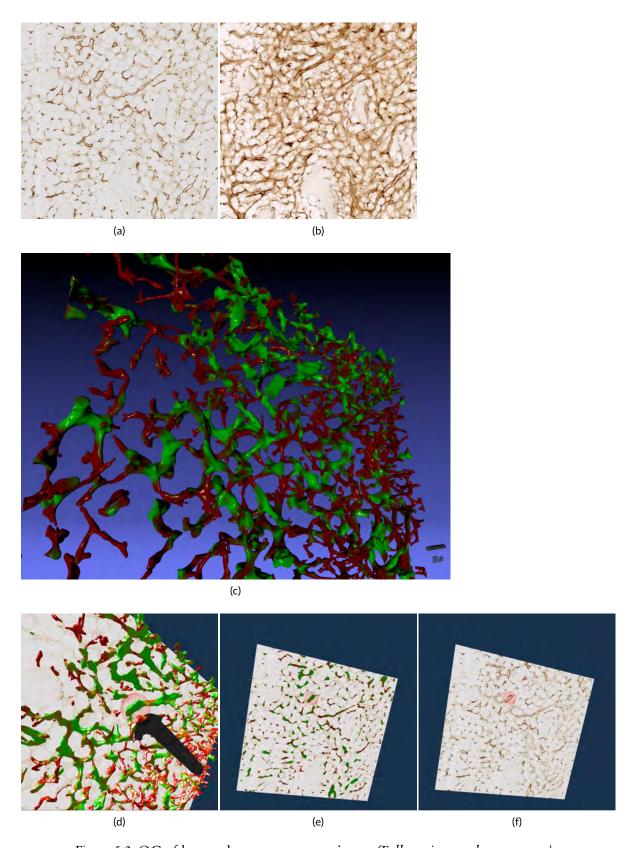
Figure 5.3: QC of human bone marrow specimen. *(Full caption on the next page.)*

Figure 5.3: (Continued from the previous page.) Images, renderings, and VR screenshots showing mesh reconstructions of blood vessels in a human bone marrow specimen, stained with anti-CD34 plus anti-CD141. (a): A single section. The staining colour is brown for both molecules in the original data. (b): A volume rendering of 21 consecutive serial sections. (c): The reconstructed mesh (Chapter 6). It shows shape diameter function values, colour-coded from red to green. (d): We annotate a position of interest in the mesh in VR. An original section is seen in the background. (e): We have found the section containing the annotation, the mesh is still visible. (f) Only the section with the annotation is shown in VR. Domain experts can now reason on the stained tissue at the marked position.

culprit is the tracking system of the HMD. Even very small changes in position can result in a change of perspective and therefore an intensity shift at some pixels. Humans mostly do not hold *completely* still for long. Minor measurement errors might also occur in the light house tracking. These minor shifts in position influence very short frame sequences. Position change is mostly unnoticed, but intensity changes in some pixels are disturbing. Visually, pixel values change most where the aliasing effect is observable. This flicker—or temporal aliasing, so to say—can happen very frequently and hurts the user experience a lot. To counteract rendering aliasing, we use multisampling for general anti-aliasing and mipmapping for better display of textures, as explained below.

Multisampled anti-aliasing (MSAA), is technically just rendering the scene at a higher resolution. The result is then resized to the desired smaller resolution by interpolating between the samples for each pixel. Mipmapping is used to display a texture at a lower resolution. If a texture is displayed directly at $1/n$ its resolution, only every n^{th} pixel is shown, which is undesirable. Mipmapping is build-in into Open GL. Basically, an image pyramid is created for mipmapped textures. At a lower display resolution, the smaller pyramid levels are used. These anti-aliasing measures mostly undid the negative temporal aliasing effect.

5.2.7 Front face culling

Consider the interplay between the model and original serial sections. A section is not an infinitely thin plane. We show the original data as an opaque parallelepiped that is one section thick in the z direction and spawns over the full surface in xy plane. The actual data points of the mesh, corresponding to the displayed section, are *inside* the opaque block. Decisive parts of the mesh are occluded by the front face of the parallelepiped. To give an example, it may be a connection or an artefact in the model. On the one hand, this is, of course, not desired, and requires correction. On the other hand, the behaviour of the model spanning for multiple sections *in front* of the current section is best studied when it ends at the front face of the parallelepiped. The solution is to enable or disable front face culling of the original data display at will.

With front face culling enabled, the user can look inside the opaque block with original section texture. This is well suited for the inspection of lesser details and small artefacts. (Figure 5.7 features a real life example, observe the edge of the section shown.) The general behaviour of the model across multiple sections can be tracked more easily with front faces of the original section on display. The ability to switch between these two representations accelerates QC. It allows the domain experts to QC the details of 3D reconstruction faster.

5.2.8   Communication

The domain expert immersed in VR needs a way to communicate her findings to the operator and further members of the group outside. The mirrored view from the VR headset is available to the audience. However, without a beacon visible in VR it is almost impossible to understand what the expert tries to show. Having a Vive controller in VR and an annotation tool solves this problem.

5.3   Results

5.3.1   Hardware

We conducted our tests on a 64-bit Intel machine with i7-6700K CPU at 4 GHz, 16 GB RAM, and Windows 10. We used NVidia GTX 1070 with 8 GB GDDR5 memory, and HTC Vive. Our VR application was focused on Vive; however it should perform well also with other headsets such as Oculus Rift. It was possible to use a previous-generation GPU, we also tested our application with NVidia GTX 960.

5.3.2   Immunohistological stainings, an overview

Following immunohistological stainings have been used to stain specimen in this chapter:

- *Anti-CD34*: primarily stains capillary walls in human spleen and bone marrow. Typically used colour is brown. Some stem cells in bone marrow are also stained.
- *Anti-CD141*: stains sinus walls in human bone marrow. Typically used colour is brown.
- *Anti-SMA*: stains smooth muscle alpha-actin. It is present, e.g., in walls of larger blood vessels on the 'input' arterial side, the so-called 'arterioles'. Typically used staining colour is brown.
- *Anti-CD271*: stains capillary sheath cells and additional fibroblast-like cells in human spleen. The sheaths are multi-cellular structures around the initial segment of human splenic capillaries. Sheath cells obviously represent the sessile fibroblast-derived part of capillary sheaths. It may be that they filter macromolecular substances from the incoming blood. Typically used staining colour is blue. Specialised fibroblasts inside the follicles are more weakly stained with this antibody.

5.3.3   Input data

The input data was generated from the registered stack of serial sections. Typical data volume was $2.3k \times 2.3k \times 161$ voxels, with z-axis interpolation. The original data featured 21 to 24 sections. The final meshes had typically 1.7M–2.3M vertices. (Most of the GPU memory used by the application was occupied by textures anyway.) Real time rendering was possible with Vive-native resolution at 90 fps. Original data were projected as textures, typically at $2.3k \times 2.3k$. While it was possible to show all sections at once, most often only one section was shown at a time.

   We QC the following data sets derived from the bone marrow of a 53-year-old male and from the spleen of a 22-year-old male.

1. 'Bone marrow': stained with anti-CD34 plus anti-CD141 (both brown), 4 ROIs 3500×3500 pixel at 0.5 µm/pixel, 21 serial sections (Fig. 5.3, Chapter 6).
2. 'Follicle-double': stained with anti-CD34 (brown) followed by anti-CD271 (blue), three ROIs 2300×2300 pixel at 0.416 µm/pixel, 24 serial sections (Fig. 5.4, Chapter 7);

3. 'Follicle-single': stained with anti-CD34 (brown), a ROI $4k \times 4k$ pixel, 0.3 µm/pixel, 24 serial sections (Fig. 5.5, Chapter 7);
4. 'Red pulp': spleen stained with anti-CD34 plus anti-SMA (both brown), followed by anti-CD271 (violet-blue, different pigment than above), 12 ROIs $2k \times 2k$ pixel at 0.5 µm/pixel, 24 serial sections (Figs. 0.3, 5.6, 5.7);

5.3.4 'Bone marrow'

Shape diameter function on the reconstructed mesh allows to distinguish capillaries from sinuses (Fig. 5.3). Domain experts could investigate the mesh, mark interesting positions, and return to the section view to reason about them.

5.3.5 'Follicle-double' and 'follicle-single'

The human spleen contains accumulations of special migratory lymphocytes, the so-termed follicles. We reconstructed the capillaries inside and outside the follicles. Before our efforts (Chapter 7), only one study existed on this topic [Kusumi et al., 2015]. Fig. 5.4 shows one of three ROIs that were used for QC.

Our reconstruction demonstrates that follicles are embedded in a superficial capillary meshwork resembling a basketball basket. Figure 5.4 shows that our VR tool enables easy annotation and projection of the original data leading to further results. In Figs. 5.4, (e), (f), some brown dots have been marked inside a follicle. The 3D model shows, that the dots indeed represent capillaries cut orthogonally to their long axis. Thus, we additionally find that infrequent capillaries also occur inside the follicles. The superficial capillary network of the follicles is thus connected to an internal and to an external network of capillaries. The latter two networks have a shape which is totally different from the superficial follicular network. The external network is partially covered by capillary sheaths stained in blue colour. In total, we examined three 'follicle-double' data sets in VR.

Fig. 5.5 also shows that the annotated elongated structures in the follicles and in the T-cell zone at least partially belong to long capillaries, which accompany the outside of larger arteries, so-termed *vasa vasorum*. With our VR-based method, we investigated this $4k \times 4k$ ROI at 0.3 µm/pixel and three further (non-shown) ROIs with 1600×1600 pixels at 0.6 µm/pixel.

5.3.6 'Red pulp'

The location of capillary sheaths in human spleens has not been clarified in detail up to now. Our 3D reconstructions indicate that sheaths primarily occur in a post-arteriolar position in the part of the organ, which does not contain lymphocyte accumulations (so-termed red pulp), although length and diameter of the sheaths are variable. Many sheaths are interrupted by the boundaries of the model. For this reason it makes sense to collect only sheaths which are completely included in the reconstruction. In Figure 7 we have annotated the ends of some sheaths with our VR tool.

Figure 5.7, (a) shows an overview of the annotations. In Figs. 5.7, (b)–(d) it becomes clear, that the sheaths indeed end at the marked positions. Notice the enabled front face culling on the section parallelepiped in the closeups. Figure 5.7, (e) additionally shows the reconstructed meshes for the sheaths. We show a single ROI at $2k \times 2k$ pixels. This work was largely motivated by the urge to inspect 'red pulp' data sets. We show a single ROI at $2k \times 2k$ pixels. We have inspected twelve such ROIs with our VR application.
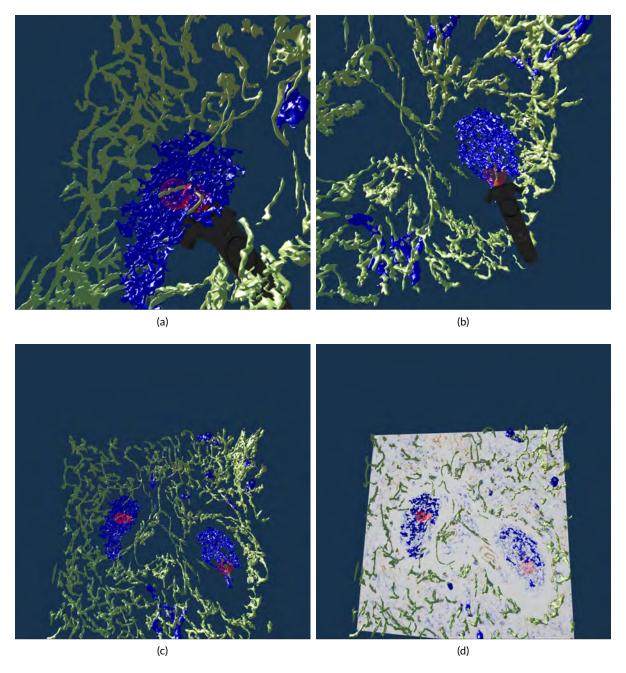
Figure 5.4: Real or artefact? In VR, we have spotted some blood vessels going through the inside of the follicle (large blue blob) in this reconstruction from human spleen data, stained anti-CD34 (brown) and anti-CD271 (blue). The blood vessels in questions are annotated (a), (b). We can look at the meshes only (c) or also show the original data (d). All images in Fig. 5.4 are screenshots from our application. *(Continued on the next page.)*

(e)   (f)
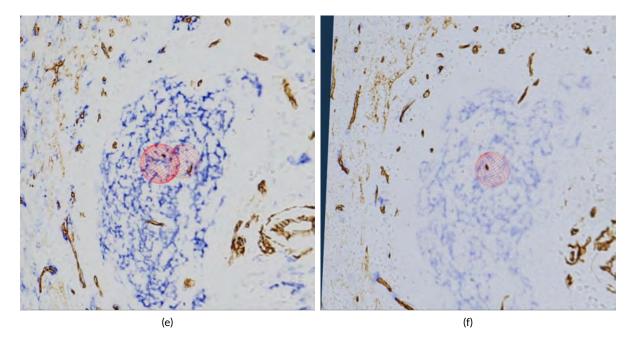
Figure 5.4: (Continued from previous page.) Real or artefact? In VR, we have spotted some blood vessels going through the inside of the follicle (large blue blob) in this reconstruction from human spleen data, stained anti-CD34 (brown) and anti-CD271 (blue). A closer view (e), (f) confirms: the reconstruction is correct, these structures are CD34-positive objects inside the follicle. As the structures in question continue through multiple sections, they do not represent single CD34-positive cells. Hence the objects in question must be blood vessels. All images in this figure are screenshots from our application.

5.3.7 Impact

The presented QC tool plays an important role in our pipeline. 3D reconstructions from serial histological sections exist, but are quite rare. Reasons are both technical difficulties to create them and struggles to QC the reconstruction. To create a good reconstruction, a proper registration, correct reconstruction options, and possibly also inter-slice interpolation, are all sought for. For QC we need to ensure the correctness of processing, identify what the reconstruction actually shows and keep artefacts at bay by creating a better reconstruction if needed. Finding a good reconstruction from previously unexplored data with a new staining is an iterative process. While we create the reconstructions quite efficiently, QC was a lot of work in the past. With an immersive VR application, QC is much easier and faster.

Our experience emphasises the importance of VR-based QC. Our previous 3D reconstruction study [Steiniger et al., 2016, Chapter 6] featured $3500 \times 3500 \times 21$ voxels in four regions. From each reconstruction a further version was derived. They did not need to be quality controlled again, but their inspection was crucial to produce a diagnosis for medical research. We used both a non-VR visualisation tool for QC and pre-rendered videos for inspection. It took us multiple meetings to QC these reconstructions with the non-VR tool. (This tool is briefly described in Section 6.3.12). Deducing the anatomical findings from pre-rendered videos was also not easy for the domain experts. The total data amount was roughly $1.03 \cdot 10^9$ voxels. Unlike, the number of voxels was slightly larger, and QC was much faster with the new method. We quality controlled *twelve* regions with

(a)   (b)
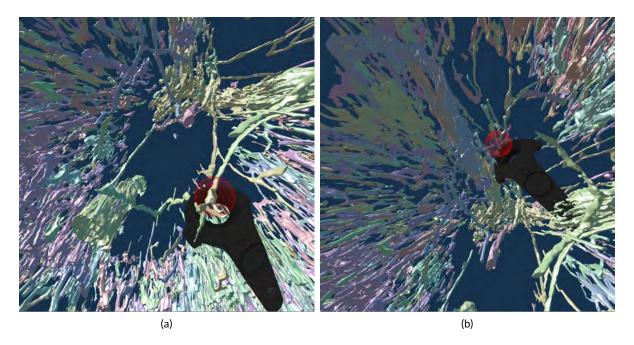
Figure 5.5: VR screenshots showing mesh reconstructions of blood vessels in human spleen specimen, anti-CD34 staining. Unconnected mesh components were set to distinct colours. We have marked some smaller blood vessels that follow larger ones. A part of these are *vasa vasorum*.

(a)   (b)   (c)
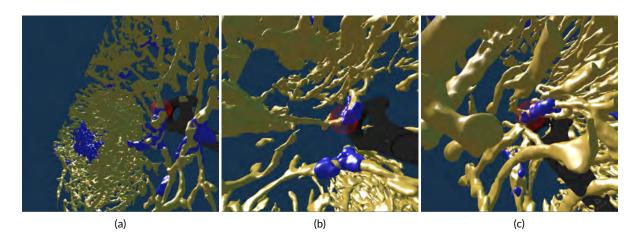
Figure 5.6: Working with our annotation tool. We show VR screenshots of our application, in a human spleen specimen. We have annotated some ends of capillary sheaths. This is the same ROI as in Figure 0.3. Anti-CD34 plus anti-SMA stainings (brown) are shown as a gold-brown mesh. Anti-CD271 (violet-blue) staining corresponds to the blue mesh, the capillary sheaths. We do not show the heuristics for arterioles.
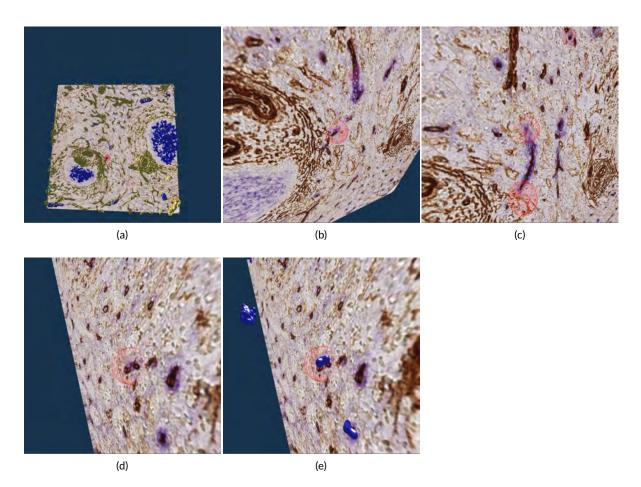
Figure 5.7: Investigating the annotated regions, VR screenshots of our application. It is the same VR session as in Figure 5.6; we have annotated some ends of capillary sheaths in meshes reconstructed from human spleen data. (a): Overview. (b)–(d): Original data and an annotation. Experts can more easily reason on such visualisations. (e): The same annotation as in (d), showing additionally the mesh for sheaths.

$2000 \times 2000 \times 24$ voxels per day with our VR-based method. At the present state, roughly $1.15 \cdot 10^9$ voxels may be QC'ed by a single person per day.

5.4 Conclusions

3D reconstructions from histological serial sections require quality control (QC). Domain experts have up to now not been satisfied by previously existing QC methods. We present a VR-based solution to explore mesh data. Our application also permits to superimpose the original serial sections. Such display is essential for QC. Immersion accelerates QC by a magnitude. Our users can annotate areas of interest and communicate the annotations. This greatly enhances productivity and allows for more precise reconstructions. Our quality controlled 3D reconstructions enable new insights in microanatomical research.

5.4.1   Future work

Standalone usage by the immersed domain expert is planned. The application evolves to be usable with the Vive controllers only. Minor improvements at user input handling, e.g., more input combinations and gestures, are always viable. A planned feature is to record spoken annotations for every annotation marker. Recorded memos would facilitate to better understand the markings at their revision.

There are further ways to improve the anti-aliasing and thus even further improve the immersive user experience. Possibilities to consider are interpolation (for higher internal frame rates) and a small averaging window or Kalman filter on few past positions (to remove outliers).

A promising idea is to learn better view angles from the view transformation matrices saved as parts of annotations. View-dependent rendering approach [e.g., Burns et al., 2007, but done for volume rendering there] might help to find most interesting views for better pre-rendered videos. Combining the view matrix with the original mesh and location of annotations we hope to spawn enough data for machine learning.

Texture compression in general and volume compression techniques in particular [e.g., Guthe et al., 2002, Guthe and Goesele, 2016] would help to reduce the GPU memory consumption caused by data for original slices.

Of course, viable future work includes applications of this method to reconstructions of further organs and tissues (e.g., future bone marrow or tonsil probes) and expansion to further modalities of medical (such as MRI or CT) or non-medical data. Possible examples of the latter include materials science, computational fluid dynamics, and, most surely, computer graphics.

Part II

APPLICATIONS IN MEDICAL 3D RECONSTRUCTION

Beautiful objects are wrought by
study through effort.

———————————————————————
Democritus

THREE-DIMENSIONAL ARRANGEMENT OF HUMAN BONE MARROW MICROVESSELS FROM IMMUNOHISTOLOGY IN UNDECALCIFIED SERIAL SECTIONS

In which we explore the applications of our 3D reconstruction pipeline applied to the bone marrow specimens. The previews of videos are shown in Appendix C.

The arrangement of microvessels in human bone marrow is so far unknown. We combined monoclonal antibodies against CD34 and against CD141 to visualise all microvessel endothelia in 21 serial sections of about 1 cm² size derived from a human iliac crest. The specimen was not decalcified and embedded in Technovit®9100. In different regions of interest, the microvasculature was reconstructed in three dimensions using automatic methods. The three-dimensional models were subject to a rigid semiautomatic and manual quality control. In iliac crest bone marrow, the adipose tissue harbours irregularly distributed haematopoietic areas. These are fed by networks of large sinuses, which are loosely connected to networks of small capillaries prevailing in areas of pure adipose tissue. Our findings are compatible with the hypothesis that capillaries and sinuses in human iliac crest bone marrow are partially arranged in parallel.

This research [Steiniger et al., 2016] appeared in collaboration with Birte S. Steiniger, Vitus Stachniss, Verena Wilhelmi, Anja Seiler, Katrin Lampp, Andreas Neff, and Michael Guthe.
Author Contributions: *Conceptualization: BSS. Data curation: OL, BSS. Formal analysis: OL. Funding acquisition: BSS, OL, VS. Investigation: BSS, VW, AS, KL. Methodology: BSS, VS, OL, MG. Project administration: BSS, OL. Resources: BSS, VS, AN, MG, OL. Software: MG, OL. Supervision: BSS, OL, MG. Validation: BSS, VS, VW, OL, MG. Visualization: OL. Writing—original draft: BSS. Writing—review & editing: BSS, VS, VW, OL. Corresponding authors: BSS, OL.*

ICROANATOMY of the bone marrow has met with increasing interest in recent years. Haematopoietic stem and progenitor cells (HSPCs, [Kusumbe et al., 2016, Hooper et al., 2009, Celso et al., 2009]), plasma cells [Zehentmeier et al., 2014] and T memory lymphocytes [Tokoyoda et al., 2009, 2010] have come into focus, because all these cells have been found to reside in special bone marrow microenvironments, so-termed niches. These niches are regarded to be primarily formed by sessile cells of local mesenchymal origin providing essential mediators for survival of the respective cell type. Motile cells such as macrophages may also contribute. The microvasculature of the bone marrow is relatively dense and the relation of niches to this vasculature is regarded to be important. Special niches have been described for HSPCs and for more mature red and white blood cells, located near the endosteum and near the blood sinuses, respectively [Nombela-Arrieta et al., 2013]. However, it is not entirely clear, whether there are distinct micro-compartments for different stages of blood cell development [Kiel and Morrison, 2008, Celso et al., 2009, Mercier et al., 2012]. Especially the location of haematopoietic stem cells in mouse femora is disputed [Nombela-Arrieta et al., 2013, Acar et al., 2015]. Niches are, however, also relevant for mature cells, which immigrate into the bone marrow and stay there, such as plasma cells. Exit from the bone marrow also plays a role not only in the course of haematopoiesis [Cohen et al., 2015], but also during recirculation of mature lymphocytes, which at least partially travel through the bone marrow [Cariappa et al., 2005, Tokoyoda et al., 2009, 2010, Itkin et al., 2016]. Bone marrow sinuses

form lymphocyte entry and egress portals of decisive function in the physiological setting [Itkin et al., 2016]. This is also important after bone marrow or haematopoietic stem cell transplantation and in spreading of multiple myeloma.

Astonishingly few histological facts are known about rodent and even fewer about human bone marrow microvasculature. Most investigations on bone marrow vessels were done in the femoral growth plates [Itkin et al., 2016, Kusumbe et al., 2016] or the femoral diaphysis [Nombela-Arrieta et al., 2013] and the tibia [Li et al., 2009] of mice. Flat bones were rarely investigated in this species [Sipkins et al., 2005, Takaku et al., 2010] and may exhibit a unique vasculature. Even in mice the nomenclature applied for different-sized bone marrow microvessels is not consistent. Thus, bone marrow arterioles were described to split up into sinuses [Itkin et al., 2016] or arteries were described to give off capillaries connected to sinuses [Celso et al., 2009]. Again, other researchers mentioned 'sinusoidal capillaries' [Kusumbe et al., 2014] or arteries, arterioles, capillaries and sinuses [Li et al., 2009]. In addition, arteries and two sequential types of microvessel endothelial cells, which most likely correspond to capillary and sinus endothelia, were visualised in mice [Kusumbe et al., 2016]. Interestingly, it was shown in mice that endothelia located at the arterial side of the microvasculature were associated with less activated haematopoietic stem and progenitor cells (HSPCs) in comparison to sinus endothelia [Itkin et al., 2016, Kunisaki et al., 2013]. Arterioles were described to be primarily situated near the endosteum in mouse femora [Kunisaki et al., 2013]. Several publications found that arteriolar and sinusoidal endothelial cells in mouse femoral bone marrow differ by phenotype [Hooper et al., 2009, Kopp et al., 2009, Tokoyoda et al., 2010]. Thus, arteriolar endothelia were described to express Sca-1, but not VEGFR3 or CD201, while sinus endothelia were positive for VEGFR3 and CD201, but not for Sca-1. Tie-2 (CD202b) was preferentially detected in arterial, but not in sinusoidal endothelial cells [Li et al., 2009]. Phenotypical differences and localisation of bone marrow microvascular endothelial cells may not only influence haematopoietic precursors but also mature immunocompetent lymphocytes and a large number of plasma cells. In spite of this fact, a correct histological classification of the single components in bone marrow microvascular networks is lacking in most mouse studies.

It is very likely that bone marrow microvessels in human flat bones, such as the iliac crest, are arranged differently from the microvessels of mouse femoral growth plates. We suppose that especially the amount and distribution of adipose tissue and haematopoietic areas differ according to species, anatomical location and age. In humans, bone marrow microvessels have up to now been primarily demonstrated in small paraffin-embedded biopsies using monoclonal antibodies (mAbs) to CD34 [Kvasnicka and Thiele, 2004, Panteli et al., 2005, Boveri et al., 2008, Negaard et al., 2009, Takaku et al., 2010]. We have now used undecalcified serial sections of a representative iliac crest specimen spanning about $1\,\text{cm}^2$ and a combination of antibodies against CD34 and CD141 to analyse the 3D arrangement of microvessel endothelium in human bone marrow. Our findings indicate that small-sized and large-sized microvessels, which most probably correspond to capillaries and sinuses, occupy different locations. Both vessel types are at least partially arranged in parallel.

6.1   Results

6.1.1   Histology

We show that high-quality sections of undecalcified bone specimens measuring up to $1\,\text{cm}^2$ can be cut with a modern motorised rotary microtome, if a suitable plastic embedding technique is used (Fig. 6.1). We chose to work with Technovit® 9100, but needed to substantially modify the manufacturer's protocol for handling this methacrylate formulation. The modifications consist of reducing the
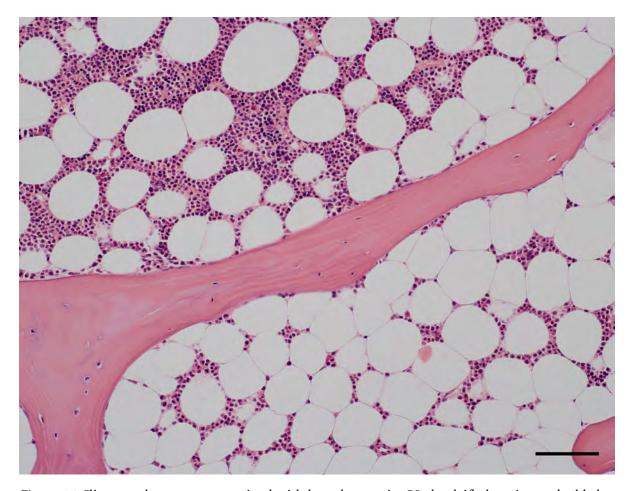
Figure 6.1: Iliac crest bone marrow stained with hemalum-eosin. Undecalcified section embedded in Technovit® 9100. The quality of the section is equivalent to a decalcified paraffin section. Note the variable density of haematopoiesis within directly adjacent adipose tissue areas. Scale bar = 100 µm.

viscosity of the methacrylate solution used for infiltration and of prohibiting access of oxygen during polymerisation. In addition, a type C tungsten carbide blade is essential for cutting large numbers of non-compressed bone sections. Only aqueous solutions, but not alcohol, should be used as cutting fluids. Unless high-temperature antigen retrieval is mandatory, Technovit sections can be applied to aminosilane-coated glass slides for de-plasticising and staining.

6.1.2 Immunohistological detection of bone marrow microvessels

Applying a widely used mAb against CD34 to undecalcified human iliac crest sections from a 53-year-old healthy male patient revealed that CD34 was present in endothelia of small microvessels in the adipose tissue, but was only weakly expressed in large microvessels in haematopoietic areas. We thus additionally tested a mAb against CD141 (thrombomodulin). This mAb stained endothelia in large microvessels, while small microvessels reacted more weakly. Double staining showed that there was a patchy distribution of CD34 or CD141 in large microvessel endothelium (Fig. 6.2). We thus decided to use a mixture of both antibodies for detecting all microvessels. Further analysis of vessels
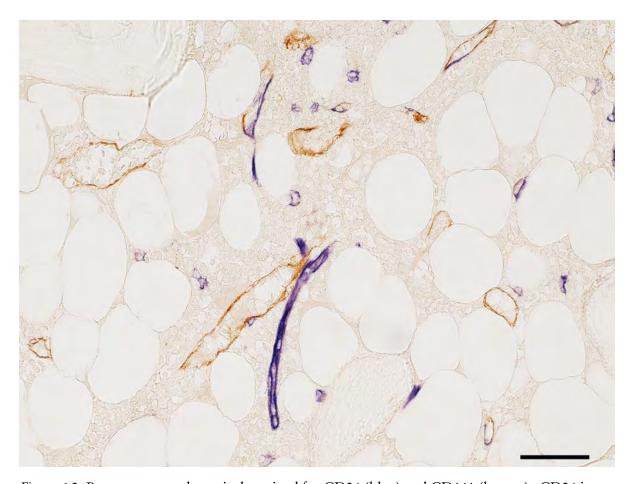
Figure 6.2: Bone marrow subtractively stained for CD34 (blue) and CD141 (brown). CD34 is visualised in small microvessels (capillaries) and is less intensely expressed in large microvessels (sinuses) which exhibit patchy blue and brown staining. Scale bar = 50 μm.

was performed by staining for smooth muscle alpha-actin to detect arterioles.

CD34 was not only expressed in endothelia, but also in round extravascular cells which were interpreted as HSPCs. CD141 also occurred in cells outside microvessels, which might correspond to dendritic cells [Meixlsperger et al., 2013].

6.1.3   Three-dimensional models of microvessel networks

Twenty-one serial sections of the specimen were cut and stained with a mixture of anti-CD34 and anti-CD141. Each stained section was scanned with a Leica SCN 400 automated slide scanning microscope. The bone specimen chosen contained evenly distributed fat cells with intercalated areas of haematopoiesis increasing in size towards the periphery of the specimen. In the centre of the sections almost pure adipose tissue prevailed (Fig. 6.3). Bone trabeculae were also more densely arranged at the periphery, especially at one edge of the specimen which seemed to be connected to a cortical bone region. The specimen chosen was compared to five additional iliac crest specimens and found to be representative with respect to the fact that the haematopoietic tissue was focally distributed in the omnipresent adipose tissue. Our study is based on regions with an intermediate density of haematopoietic cells. Four regions of interest (ROIs), R1–R4, (Fig. 6.3) were chosen. R1
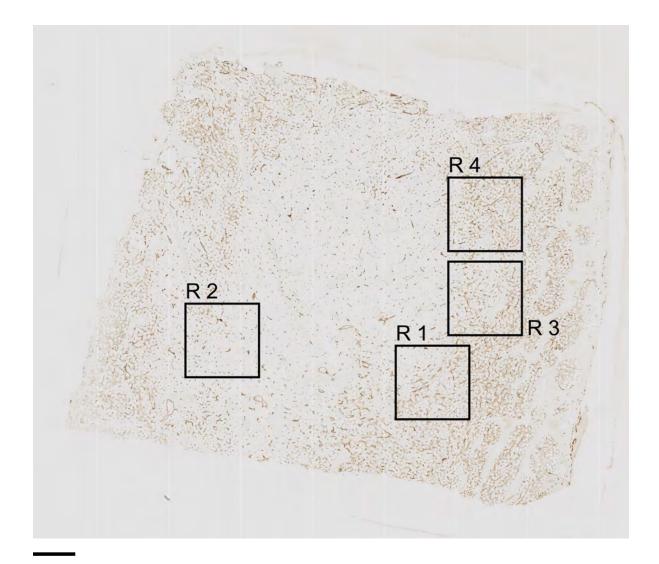
Figure 6.3: The first of 21 serial sections immunostained for CD34 plus CD141 with regions of interest R1 to R4. Scale bar = 1 mm.

and R2 contained adipose tissue with and without haematopoietic tissue, while R3 and R4 contained only adipose tissue with haematopoiesis. Arteries and typical arterioles were absent in these ROIs as judged from inspection of the immunostained sections.

We applied a new multi-resolution non-rigid registration [Lobachev et al., 2017a, Chapter 2], implicating a hierarchy of feature sizes. Subsequently, volume filtering, 3D reconstruction, mesh filtering, and rendering were carried out. The results were visualised in video format using two types of 3D models (Fig. 6.4).

In the first type of models (Fig. 6.4b,e,h,k), vessels with a shape diameter smaller than 20 µm were depicted in red, vessels with more than 40 µm shape diameter were coloured green and transitional vessels had a brown colour. All non-connected components with a maximum diameter less than 125 µm were excluded. This corresponds to less than 18% of the model surface.

These models demonstrated that networks of small and large microvessels predominantly occupied different locations in the ROIs. The innermost area of R4 (Fig. 6.4k) did, however, form an

(a) R1: The first section

(d) R2: The first section

Figure 6.4: First scan of R1, R2 immunostaining in combination with front views of the 3D models coloured for shape diameter and connectivity visualisation. Scale bar for all images = 200 µm. *(Continued on the next page, full caption on page 106.)*

(b) R1: The shape diameter

(e) R2: The shape diameter

(c) R1: The connectivity study

(f) R2: The connectivity study

Figure 6.4: (Continued from previous page.) First scan of R3, R4 immunostaining and front views of the 3D models coloured for shape diameter and connectivity visualisation. *(Full caption on the next page.)*

(g) R3: The first section

(j) R4: The first section

(h) R3: The shape diameter

(k) R4: The shape diameter

(i) R3: The connectivity study

(l) R4: The connectivity study

Figure 6.4: (On two previous pages) First scan of R1 (a), R2 (d), R3 (g) and R4 (j) immunostaining in combination with front views of the 3D models coloured for shape diameter (b, e, h, k) and connectivity visualisation (c, f, i, l). In (b), (e), (h), and (k) a shape diameter below 20 µm is depicted in red and above 40 µm in green. A transition from red to green occurs in microvessels of shape diameters from 20 µm to 40 µm with middle at 23.5 µm. In (c), (i), and (l) the largest continuous microvessel network is depicted in light blue and all smaller networks are green. In (f) there are several large networks, all coloured in blue. Small structures sized between 50 µm and 12.5 µm are red in (c), (f), (i), and (l), and all structures with a diameter less than 12.5 µm in any dimension are omitted. (b, c) R1; (e, f) R2; (h, i) R3; (k, l) R4. Scale bar for all images: 200 µm.

exception. Single small vessels were sometimes observed inside networks of large vessels, but these small vessels were very short and did not form networks. Connections of small vessel networks to large vessel networks were present, but relatively rare. In such cases, the connecting small vessel often anastomosed to a large vessel network end-to-side with an abrupt change of calibre. Comparison to the immunostaining (Fig. 6.4a,d,g,j) revealed, that networks of small microvessels were primarily associated with areas of pure adipose tissue, while networks of large microvessels were associated with adipose tissue harbouring haematopoietic areas. We thus interpret small microvessels as capillaries and large microvessels as sinuses. R1 contained microvessels of more than 40 µm diameter in a haematopoietic area, which were, however, accompanied by a non-connected network of small microvessels (Fig. 6.4b). We assume that the large vessels represent collecting sinuses or post-sinusoidal venules and the small vessels correspond to capillaries.

Besides capillaries, pure adipose tissue outside the ROIs additionally contained single very large vessels, presumably arterioles and venules. Visual control of the immunostained sections and staining for smooth muscle alpha-actin in directly adjacent sections showed that only very few arterioles with a continuous layer of smooth muscle cells were present within the four ROIs. Thus, almost all vessels visualised corresponded to capillaries and/or sinuses.

The second type of 3D models (Fig. 6.4c,f,i,l) was created by showing the largest network of connected microvessels in blue. All other smaller networks were visualised in green, and structures smaller than 50 µm, but larger than 12.5 µm in any dimension, were attributed a red colour. Structures with a diameter less than 12.5 µm in all dimensions were removed. In three of the four ROIs (Fig. 6.4c,i,l) there was one major network of large vessels (blue) and several unconnected smaller networks (green). R2 (Fig. 6.4f) was an exception, because there were ten unconnected large microvessel networks of similar size. We visualised all ten in blue. The red structures primarily represented fused single cells, which appeared elongated in direction of the z axis.

This second type of models showed that small microvessels formed networks with few connections to networks of large microvessels in the 21 serial sections. The centre of R4 (Fig. 6.4k,l) was an exception, because it contained an arrangement of small vessels connected to a network of larger vessels. This arrangement was also detectable by shape diameter function (Fig. 6.4k).

The 3D models were subject to a thorough quality control (QC, Figs. 6.5, 6.6a,b, 6.8; Fig. C.2, doi: 10.5281/zenodo.128861) consisting of two semiautomatic and one manual procedure. The only step that could not be controlled was the initial registration. Quality assessment of this step was described by Lobachev et al. [2017a]. All QC methods revealed that blind ends of vessels occurred in the 3D models. Most of these ends were either due to registration errors because of tears in sections or to rather short interruptions of small capillaries after mesh construction (for details of QC see materials and methods, discussion; Fig. 6.8, Fig. C.2). Few 'real' ends were found when the models
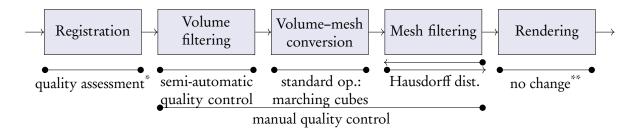
Figure 6.5: Pipeline of image processing and QC methods, both semi-automatic and manual.
* The quality of registration is commented on in Chapter 2 and Lobachev et al. [2017a].
** 3D rendering does not affect the data.

were compared to the registered scans. Thus, we are sure that the models presented come very close to reality and loss or addition of major microvessel network components do not occur.

In summary, the microvessels of human iliac crest bone marrow appeared to form a closed vasculature, because erythrocytes were not found outside vessels. All 3D models showed, that the microvessel network at the surface of bone trabeculae did not differ in shape from the network at other locations. Networks of small microvessels, i.e., capillaries, were preferentially associated with pure adipose tissue or adipose tissue with few haematopoietic cells. Networks of large microvessels, i.e., sinuses, were prevailing in areas of more haematopoietically active adipose tissue. Both types of networks had only few connections, i.e., large microvessels did mostly not co-exist together with small microvessels, but they replaced them. Our 3D models provoke a new hypothesis on the arrangement of capillaries and sinuses in human iliac crest bone marrow.

6.2 Discussion

Our study establishes new immunohistological and visualisation procedures for the 3D analysis of human bone marrow morphology in flat bones using large undecalcified sections. The results demonstrate the 3D structure of bone marrow microvasculature in 21 serial sections of the iliac crest and form a basis for further investigations in larger series of sections.

We show that small and large microvessels in human iliac crest bone marrow do not only differ by size but also by preferential expression of CD34 or CD141, although both features do overlap. It is very likely that the small vessels correspond to capillaries and the large vessels represent sinuses. 3D reconstruction of microvessels after combined immunostaining for CD34 and CD141 reveals that both types of microvessels tend to form mutually exclusive networks. This provokes the hypothesis, that capillaries and sinuses are at least partially arranged in parallel. Our findings are thus different from the arrangement published for microvessels of the growth plates or the diaphysis in mouse femora [Nombela-Arrieta et al., 2013, Kusumbe et al., 2014, Kunisaki et al., 2013]. The video presentations of the 3D models (cf. Fig. 6.4) indicate that in most areas of the four ROIs smaller microvessels, i.e., capillaries, have only few connections to larger microvessels, i.e., sinuses. Capillaries do, however, accompany larger sinuses or initial venules (Figs. 6.4b,c). The serial sections investigated comprise areas of pure adipose tissue without haematopoiesis containing only small microvessels and adipose tissue with different amounts of haematopoietic cells associated with networks of large microvessels. Both types of tissue are included in two of the ROIs (Fig. 6.4a–f).

We have only stained endothelial cells and thus we cannot unequivocally differentiate very small arterioles (metarterioles) from capillaries in the sequence of the 21 serial sections by morphological
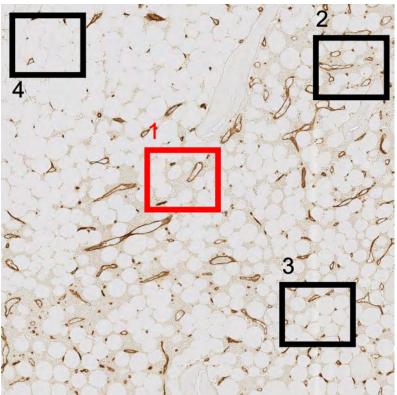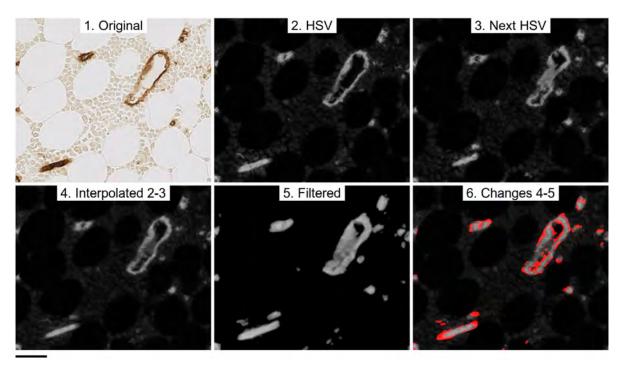
Figure 6.6: Semi-automatic QC.

(a) R1 containing 4 small regions for semi-automatic QC of volume data processing. The region visualised in (b) is coloured red. Scale bar = 200 μm.

(b) Sequence of QC visualised from left to right in upper and lower part of panel: original scan (1), HSV conversion of scan of current serial section (2), HSV conversion of next serial section (3), interpolation between both serial sections: in this case the fourth of seven images is shown (4), final result of all filtering operations of interpolated images (5), areas different between interpolated images and result of filtering coloured red (6). Scale bar = 50 μm.

inspection. For this purpose double-staining for endothelial and smooth muscle cells would have been necessary. Such a procedure might, however, introduce additional problems into in the models. Thus, staining for smooth muscle cells was performed in the same specimen, but not in the section series.

We define all vessels containing smooth muscle cells in their walls (diagnosed by morphology or immunohistology) as arterioles and hypothesise that these vessels feed capillaries (small microvessels) which only connect to sinuses (large microvessels) near haematopoietic areas. Capillaries and sinuses most probably both drain into initial venules (very large microvessels without smooth muscle cells). In aged human adult flat and long bones (such as the femoral head), adipose tissue is the prevailing bone marrow component. Haematopoietic areas are focally distributed in the adipose tissue with a tendency to increase towards the compact bone surface. Pure bone marrow adipose tissue only contains arterioles, capillaries (small microvessels) and venules. In directly adjacent focal haematopoietic areas sinuses (large microvessels) and initial venules (extremely large microvessels) prevail, while capillaries are almost absent. Thus, capillaries and sinuses tend to occur in a parallel arrangement in human bone marrow adipose tissue with and without haematopoiesis. Because arterioles are scarce, it is likely, that sinuses are fed by capillaries primarily at the surface of haematopoietic areas, while the major part of the sinus network inside the areas does not coexist with capillaries. We can, however, not directly demonstrate, which vessels feed the sinus network, because 21 sections are not sufficient to follow branching arterioles and because small arterioles were also not detected by immunohistology.

Our hypothesis is derived from a single specimen and thus needs to be regarded with some reservation. We did, however, investigate iliac crest bone marrow specimens from five additional individuals aged 30 to 60 years without performing serial sections or 3D reconstructions. The 2D arrangement of microvessels inside and outside haematopoietic areas was similar in all specimens.

A very special aspect of our method is that we work with undecalcified bone tissue embedded in methacrylate for cutting serial sections of large bone marrow specimens. This has not been reported for Technovit® 9100 before. We avoid the drawbacks of paraffin embedding, such as tedious incubation in EDTA solutions and the uncertainty of analysing the residual calcium content of the specimen. Thus, the overall processing time of the specimens is accelerated, even if the embedding process itself is somewhat more time-consuming than in routine paraffin histology. All steps of our method may be performed with normal laboratory and precision engineering equipment, except polymerisation and sawing. The pressure polymerisation device can, however, be easily replicated (Fig. 6.7). Modern rotary microtomes are able to produce large high quality sections of undecalcified bone tissue embedded in methacrylate, if a tungsten carbide type C knife and an adequate cutting fluid is used. The distortion is comparable to that of paraffin sections. However, bone tissue is brittle and small cracks in bone trabeculae cannot be entirely avoided. Our embedding method can also be used for sawing and grinding techniques to investigate the integration of metal implants into bone tissue.

Technovit® 9100 is entirely removed from the sections before immunostaining, but this process needs to be carefully monitored. Our method is associated with good preservation of different antigens in the sections. Other investigations have already demonstrated that embedding in Technovit® 9100 is suitable for immunohistology [Yang et al., 2003, Willbold and Witte, 2010, Steiniger et al., 2013]. In general, we recommend using a high amplification method for immunostaining, such as a tyramide-enhanced immunoperoxidase procedure. Antibodies need to be individually tested, because preservation of epitopes may be better or worse in comparison to paraffin sections. If a two-component epoxide gluing method is used, high temperature antigen retrieval is also possible in undecalcified bone sections.

The visual analysis of microvessel networks in the models is demanding, although only a rather small number of serial sections is used. We thus try to alleviate the interpretation by introducing

two different types of 3D models, one based on vessel shape diameter and the other based on the connectivity of the networks. Both visualisations assist the observer, but the information needs to be carefully interpreted. In Figs. 6.4b,e,h,k a vessel shape diameter less than 20 μm in any direction is represented in red colour. Small capillaries are thus nicely highlighted, but broad and extremely flat sinuses also appear red, if viewed from their broad side. There is a rather high number of flat sinuses, extending in direction of the z axis. Thus, these models cannot be used to distinguish individual capillaries and sinuses by colour, i.e., by size. They are only intended to help locating the different networks. The second type of models (Figs. 6.4c,f,i,l) reveals the size of unconnected vessel networks. With exception of R2, the largest network is marked blue and all other unconnected networks of smaller size are green. This visualisation also shows that capillaries are largely separate from sinuses, but it has to be kept in mind that random interruptions of collapsed capillaries and other errors occur. We assume that there need to be several artificial interruptions until networks fall apart, but it cannot be excluded that the real connectivity of the networks in the series is greater than demonstrated. In addition, connections may also be present outside the 21 sections investigated.

Proper thresholding is decisive for the models demonstrated in Figs. 6.4b and 6.4c. Removing non-connected structures smaller than 125 μm (i.e., smaller than 5% of the main diagonal) is a prerequisite for computing shape diameter function, because the interior of vessels needs to be cleared of precipitates and blood cells. The threshold was determined optically. It also permits better recognition of the overall microvessel network, because it reduces the complexity of the model. Single and merged cells, parts of small non-connected capillaries and parts of vessels isolated by the cut surface are removed, leaving connections between microvessels intact. To demonstrate this fact, we specially controlled the model shown in Fig. 6.4b by selectively visualising the removed structures in 3D. The threshold for Fig. 6.4c was 12.5 μm (i.e., only 0.5% of the main diagonal). It primarily excludes single or fused cells, i.e., structures most probably occurring in only one of the 21 sections.

The image processing pipeline for the 3D models is adjusted to accommodate several problems associated with immunostaining. First, the sectioning procedure may lead to occasional defects in vessel endothelium. Second, staining intensities of endothelia within one and the same vessel may be extremely variable, depending on the angle of cutting. In tangential sections, endothelia leave less immunoreactive material as compared to cross-sections or longitudinal sections and the staining is thus much lighter. These problems are solved by several volume filtering steps.

Background staining does not form a major problem with exception of the background artificially introduced by the scanning microscope. This might have been avoided by fine-tuning the digital processing during data acquisition. Procedures that help with the consistency of the model, such as colour conversion to HSV colour space and the closing operation (see materials and methods), may connect small separated stained and unstained cells to a close-by blood vessel. The chosen values for closing and further operations are thought to improve the quality of the model while keeping the amount of 'fused' cells low. In some sections megakaryocytes have a light brown colour. It is possible to exclude these cells by proper thresholding.

Data processing as performed does not produce smooth vessel surfaces due to cells fused to these surfaces or due to some persisting holes in the walls of larger blood vessels caused by the sectioning or filtering procedures. Further improvements of the processing pipeline may help. However, the aggressive removal of artefacts partially present in the original data might disturb the model and steer it away from reality, for example by increasing the number of 'fused' single small cells.

Errors of registration due to sectioning artefacts can neither be avoided nor controlled, because the arrangement of the microvessels and trabeculae before sectioning of the specimen is unknown. We do, however, regard the novel registration method applied [Lobachev et al., 2017a, Chapter 2] as optimal, because it uses automatic detection of landmarks of various scale allowing to register both

small and large image elements. In addition, the method functions for features with a high degree of self-similarity, typical of haematopoietic bone marrow. Lobachev et al. [2017a] corroborate this by a sequence of quality assessments.

Fusion of cells partially result from the colour space conversion procedure where the saturation does not decrease as fast as the immunohistological staining, i.e., in the narrow spaces between blood vessels and uncoloured cells. The colour intensity of 'unstained' cells is sometimes artificially increased by the scanner so that their saturation is above the segmentation threshold. A higher segmentation threshold is not possible during volume filtering, as this might produce even more holes and blind ends in the vasculature. Another cause for cell fusion was due to the fact that two points are connected by the closing filter if the distance between them is less than 10 voxel 5 µm × 5 µm × 10 µm.

Immersion fixation of the specimens cannot be avoided. In the bone marrow, this procedure obviously provokes the collapse of capillaries, but not sinuses. Thus, the diameter of such vessels may be near 3 µm to 4 µm in adipose tissue. Overall, the diameter differences of blood vessels in the sections are huge. Our volume filtering approach leads to a slightly increased diameter of all blood vessels. This is not conspicuous in large blood vessels, but in capillaries. Nevertheless, the calibre differences between capillaries and sinuses are still clearly visible. In addition, the increase in capillary diameter due to image processing is not really relevant, because it counterbalances the effect of immersion fixation.

Sinuses and fused single cells tend to be elongated in the direction of the z axis. This phenomenon is most likely explained by volume anisotropy. Volume filtering may have introduced a directional bias.

In detail, the semi-automatic and manual QC methods revealed that blind ends or interruptions of microvessels were due to one of the following problems listed in the order of the image processing pipeline.

1. Errors of registration occurred because of disruption or compression artefacts during sectioning of the specimen, especially near bone trabeculae. This error sometimes provoked further errors because of the extreme distortion of adjacent tissue. It was detected by manual QC and occurred infrequently.

2. Single stained and unstained cells were often fused to the surface of vessels forming artificial buds or short branches. This error was associated with conversion of the original data to HSV and was detected semi-automatically. With respect to unstained cells it was partially due to surface contrast.

3. Errors primarily associated with filtering of HSV data frequently occurred and were detected semi-automatically. They led to holes or partial loss of the walls of larger vessels. Occasionally this problem provoked interruption of a larger vessel.

4. Extremely small capillaries were sometimes discontinuous for a short distance and then re-appeared. This problem was relatively frequent and occurred during mesh processing. It was made evident by manual QC (Fig. 6.8) and by the visualisation of Hausdorff distances (non-shown).

5. In very few cases did vessels really end inside the series of 21 sections. This phenomenon was detected manually and may represent angiogenesis.

Figs. 6.4c,f,i,l show, that in addition to blind ends, short elongated red structures were present between the microvessel network, which extended primarily in direction of the z axis. These

structures were also analysed by manual QC. Most of them were found to be due to volume processing and fusion of single cells located outside vessels. Alternatively, small accumulations of granular extracellular material with a light brown colour also caused short elongated structures. This extracellular material was visible in scans, but not in the original sections.

Collapsed capillaries pose the only reconstruction-associated systematic problem that could not be totally remedied. Some of these capillaries are very thin so that short interruptions occur during processing of the surface mesh in spite of high staining intensity. This problem is responsible for most of the interrupted vessels in the final model. The overall structure of the microvessel networks is, however, not altered by this drawback.

Our quality control methods permit direct monitoring of the entire processing pipeline by semi-automatic procedures. In addition, we can manually control for suspicious structures in the final model (Figs. 6.5, 6.8). In order to achieve models of high reliability, we need to use a large number of data processing procedures. Manual quality control ensures that the reconstruction process does not lead to major artefacts, i.e., selective loss of connections between microvessels of different calibre or accumulation of errors. The microvessels shown in the model are highly congruent to the stained vessels if superimposed on the registered original section (Fig. 6.8). We are not aware of any other models of microvessel arrangement that have been controlled with such stringency.

Others have previously tried to reconstruct microvessel networks from serial immunostained sections of human soft tissues stained for CD34 [Gijtenbeek et al., 2005, Gilhuis et al., 2006, Cassoni et al., 2008]. 3D reconstructions of microvessel endothelia in large normal human bone marrow specimens have not been published before. Previous reports are limited to small human bone marrow biopsies stained for CD34 [Kvasnicka and Thiele, 2004, Takaku et al., 2010]. We are also not aware of any other 3D reconstruction methods for serial sections that use feature-based non-rigid registration, volume and mesh filtering as well as mesh-based analysis in combination with a comprehensive QC.

Overall, the models indicate that human iliac crests do not only harbour 'sinusoidal capillaries' or sinuses, but that capillaries as well as sinuses exist, but preferentially occupy different locations. In addition, staining for endothelia and smooth muscle alpha-actin showed that the distribution of arterioles in the human iliac crest differs from that in mouse femora [Kunisaki et al., 2013]. In the 21 serial sections we investigated, arterioles were rare. They were not associated with the endosteum and tended to occur in groups, preferentially in areas of adipose tissue without haematopoiesis.

We aim at reconstructing entire sections instead of ROIs and at investigating larger series of sections to find the connections between smooth muscle alpha-actin-positive arterioles and microvessels. However, the information to be handled is extremely large. Reconstructing 21 entire sections will amount to 8.7×10^{10} voxel and will necessitate special measures to reduce a z-axis bias. We regard this endeavour justified, because human bone marrow does is not only associated with haematopoiesis, but also represents a special type of secondary lymphoid organ and a reservoir for mature granulocytes. Visualising the morphological correlates of these functions in 3D may be achieved by double-staining for microvessels together with lymphocyte populations, granulocytes, macrophages or stromal cells.

6.3   Materials and methods

6.3.1   Histological procedures

Bone specimen and fixation

The iliac crest specimen came from a 53-year-old healthy male patient and was left over after retrieval of bone for facial reconstructive surgery. This study was approved by the ethics committee of the

Medical Faculty of Marburg University (AZ 35/12). Written informed consent was obtained. The bone pieces were immediately immersed in 3.7% formol in tap water for 24 hours at 4 °C and then transferred via 50% isopropanol to 70% isopropanol for storage. Where possible, larger bone particles were reduced to square tiles of a volume smaller than $1\,cm^3$ with less than 3 mm thickness using a diamond band saw. Most particles were, however, irregular in shape. The material was removed from the iliac crest approximately 5 to 10 cm dorsal of the anterior superior iliac spine.

Dehydration and defatting

Two tiles each are wrapped in hanging Technovit-resistant mesh tissue in a custom-made support frame designed for small 200 ml screw-cap glass containers. Dehydration is accomplished by immersing 6 to 8 tiles in isopropanol/water 80% once for one hour, then in isopropanol/water 90% twice for one hour, and finally in isopropanol 100% three times for one hour. For removal of fat, Xylol is applied once for one hour and subsequently overnight.

Infiltration

The bone specimens are embedded in Technovit® 9100 obtained as a kit from Kulzer, Wehrheim, Germany. The infiltration protocol specified in the kit proved inadequate for reproducible embedding results and had to be optimised. The key alterations to the Kulzer procedure are the use of an infiltration solution with reduced viscosity obtained by preparing solution A with only 25 g of polymethylmethacrylate powder instead of 80 g and by introducing an acetone step before infiltration by solution A to optimise penetration into the specimen. Using 80 g of PMMA powder, as recommended by the manufacturer, leads to long stirring times for dissolving the powder and yields a fluid of high viscosity. High viscosity prevents the uniform distribution of Technovit® 9100 in and around hard as well as soft tissues and provokes bubble formation in the centre of the specimen during polymerisation.

Destabilisation of the base solution

A 250 ml chromatography column with a glass filter of porosity P4 (pore size 10–16 µm) is filled with 50 g of basic Al_2O_3 (Roth, Karlsruhe, Germany, No. X908.1). A maximum of 3 l to 4 l of Technovit base solution is slowly passed (at about 300 ml/h through the Al_2O_3 and subsequently stored in brown glass bottles at −20 °C. Porosity P4 is chosen to retain Al_2O_3 particles, which may damage the tungsten carbide knife of the microtome.

Solution A Technovit® 9100 is prepared in advance as stock solution A containing 3 g dibenzoylperoxide and 25 g polymethylmethacrylate (PMMA) powder per 500 ml destabilized Technovit base solution. In detail, 250 ml base solution is stirred, 25 g PMMA powder is added and the base solution is filled up to 400 ml. When the solution is clear after about 3 hours of stirring, 3 g dibenzoylperoxide is added and base solution is filled up to 500 ml.

Infiltration procedure The low viscosity of solution A improves and simplifies preinfiltration and infiltration. Specimens are transferred from xylol to 100% acetone for 6 hours with one change of acetone and then infiltrated with solution A for 40 hours at 4 °C.

(a) External aspect   (b) Technical drawing
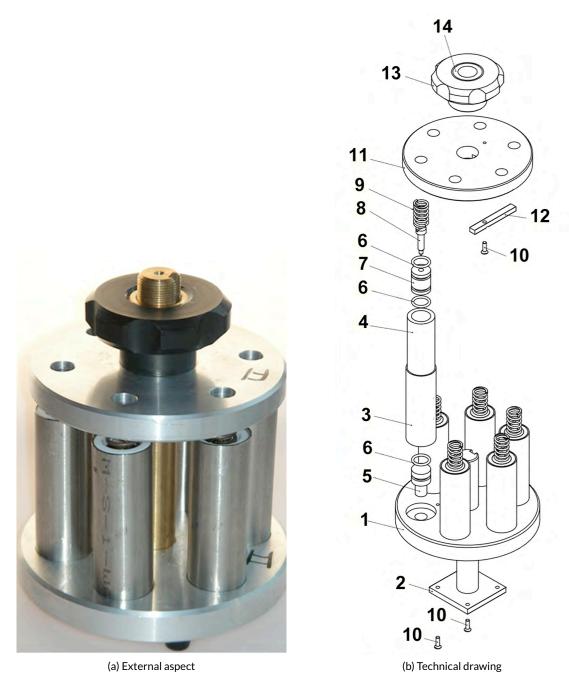
Figure 6.7: Pressurising device for polymerisation of Technovit® 9100. External aspect (a) and technical drawing (b) of the newly designed device for polymerizing Technovit® 9100. Components: 1: bottom pressure plate, 2: drawbar, 3: steel jacket, 4: pressure cylinder, 5: piston, 6: O ring, ISO 3601, 12 mm × 2 mm, 7: vented piston, 8: venting screw, 9: compression spring, ISO 10243, 1.6 mm × 12.5 mm × 21.5 mm, 10: M3 countersunk slot head screw, ISO 2009, 10 mm, 11: upper pressure plate, 12: pilot pin, 13: screw for upper pressure plate, 14: thread insert. Items 4, 5, and 7 consist of teflon.

6.3.2 Polymerisation under pressure

The polymerisation moulds provided by the manufacturer proved totally insufficient, because they were not sealed so that air intruded due to the vacuum arising from polymerisation shrinkage, which can reach 15% by volume. Oxygen slows or even prevents polymerisation of Technovit® 9100. We thus constructed a pressurising device (Fig. 6.7) with six metal-sleeved Teflon tubes of 16.0 mm internal diameter and about 100 mm length as specimen containers. These containers have sealed mobile bottom and top stoppers to prevent vacuum formation inside. The top stopper of each container always adapts to the surface of the shrinking Technovit solution because it is forced downwards by a spring fixed to a screwed-in metal pressure plate. This principle also prevents bubble formation within the specimen.

Solution B

Stock solution B is prepared according to the manufacturer's recommended procedure by first adding 4 ml of N,N,3,5-tetramethylaniline (corresponding to 4.3 g) to 30 ml base solution, then adding 2 ml of 1-decanthiol (corresponding to 1.7 g) to the mixture and finally filling base solution up to 50 ml. Solution A and B can be stored at $-20\,°C$.

6.3.3 Positioning of specimen tiles for embedding and polymerisation

The specimen tiles are embedded in an upright position to prevent bubbles caused by polymerisation shrinkage of Technovit® 9100. For this purpose, they need to be held in place by a carrier consisting of the same Technovit used for embedding. Technovit is polymerised in the pressurising device as described below and the resulting methacrylate cylinders of 16 mm diameter and 30 mm length are cut into three shorter cylinders of 10 mm length. On one end face of each cylinder a groove of 3 mm width and 2.5 mm depth is milled into the plastic using a twin-fluted end mill. The groove passes through the centre of the end face. The specimen tile is inserted into the groove thus avoiding the need for glue. The cylinder with the tile is then put onto the bottom stopper and inserted into the polymerisation container. This procedure prevents the specimen from becoming dislodged during the filling process.

6.3.4 Embedding protocol

Shortly before starting the embedding procedure, stock solutions A and B are removed from the refrigerator. One part of stock solution B is added to nine parts of stock solution A and stirred with a glass rod for exactly 1 min. Care is taken to quickly proceed with the following embedding procedure. After insertion of the specimen and the bottom stopper, the polymerisation solution is poured into the tube up to a mark and the top stopper is inserted. The stopper contains a screw-hole for the escape of air and for observing whether it touches the surface of the solution during lowering. When this position is reached, a screw is inserted into the hole and tightened. After filling and closing all six tubes in this way, a pressure plate with an affixed spring for each tube is screwed in and tightened until a pressure of about 6 bar (6×10^5 Pa) has built up.

    The pressurising device is then placed into an air-tight plastic bag and put into a refrigerator for 4 days at $-14\,°C$ to prolong the polymerisation process. After 4 days the device is transferred to an incubator at $37\,°C$ for one day to guarantee complete polymerisation. The metal sleeves are then removed from the Teflon tubes and the specimens are pushed out of the tubes with a special lever press.

6.3.5   Preparation of embedded specimen tiles for sectioning

The specimens are embedded in an upright position and thus need to be cut from the methacrylate cylinders, turned through 90° and re-fixed in a horizontal position for sectioning with a microtome. The upper face of each cylinder is first made transparent by hand-polishing with increasingly fine grinding paper (grit 1200, 2500 and 3000).

The planes of two parallel cuts and one rectangular cut to be performed with a contact point diamond band saw are then marked on the specimen cylinder using a stereo loupe. Two of these cuts are parallel to the axis of the cylinder, one cut running directly along the specimen surface. A further cut, perpendicular to the axis, separates the specimen from the cylinder. The exposed specimen is then rotated through 90° and glued to the remaining cylinder using isocyanate glue (Loctite 4305, Henkel, Düsseldorf, Germany) for 1 min with subsequent polymerisation in a UV-A polymerisation device (Ivoclar, Schaan, Liechtenstein) for 90 s. To assist cutting with a microtome, the methacrylate embedding material surrounding the specimen is chamfered at an angle of 10°.

6.3.6   Quality control of embedded specimens

Before cutting, the specimen surface is polished as mentioned above and the quality of the embedding process, i.e., the absence of bubbles and other artefacts in the specimen, is photo-documented using a modified conventional transmitted light microscope. The modification consists of exchanging the revolving nosepiece below the eyepiece beam splitter for an $f = 50\,\text{mm}$ objective (Rodagon, Schneider, Kreuznach, Germany) and mounting an SLR camera with a 4/3 chip onto the tube. This arrangement permits the observation of the surface and also the interior of the embedded specimen by using either light reflected at an angle of 30° or transmitted light. Using reflected light necessitates a specimen holder, which may be tilted. For transmitted light the depth of focus is about 2 mm with aperture $f/16$. Thus, the interior of the specimen may be visualised by two photographs.

6.3.7   Cutting of hard serial sections

Hard specimens containing bone can only be cut in series, if an up-to-date motorised rotary microtome, such as Leica RM2255, is used. Older motorised microtomes do not provide sufficient mechanical precision and stability. In addition, it is essential that the knife has a type C tungsten carbide blade. This blade has a cutting angle of 27° and thus leads to less compression and distortion of the sections than the conventional type D blade with an angle of 42°. Type D knives inevitably lead to sections shortened by one third in the direction of cutting due to compression. Type C knives are, however, more easily worn than type D knives as they are much less resistant to mechanical lesions during cutting.

The Technovit cylinder containing the specimen is fixed to the microtome using a holder for square specimens and a custom-made split brass adapter for cylindrical specimens to prevent the embedding plastic from damage caused by repeated mounting and cutting. Alternatively, a proprietary segmented holder for round specimens may be used, but this damages the plastic.

Hard plastic sections need to be cut using a fluid of low surface tension to prevent the sections from rolling up during cutting. The solution is distributed on the section surface and on the blade with a brush before each cut. The lower part of the section, which only consists of embedding plastic is held with fine forceps during cutting and the section is then transferred to a drop of cutting fluid on a slide. In our experience an aqueous solution of 1% BSA (bovine serum albumin) proved optimal for cutting and attaching the sections to silanised glass slides. The manufacturer of Technovit® 9100

recommends alcohol solutions for cutting, but we found that—even if used at high dilution—alcohol provokes cracks in Technovit® 9100 and thus should be totally avoided.

The section on the slide is then covered with a polyethylene film of 25 μm thickness provided by the manufacturer. Starting from one end of the slide the film is pressed onto the slide and attached uniformly by squeezing out the fluid underneath the section. The slides carrying the films are finally placed between two sheets of slide-sized blotting paper to prevent them from damage. They are inserted into a frame-shaped slide press and dried at 50 °C overnight under the pressure of 2 N m, developed by a tightening torque.

6.3.8 Immunohistology

Primary antibodies

The following primary monoclonal antibodies (mAbs) were used: MAb QBend10 against CD34 (Dianova, Hamburg, Germany, No. DLN-09135), mAb TM1009 against CD141 (DAKO, Hamburg, Germany, No. M 0617), and mAb asm-1 (Progen, Heidelberg, Germany, No. 61001) against smooth muscle alpha-actin.

The antibodies had been pre-tested on sections of human spleen specimens embedded in paraffin and in Technovit® 9100. They can be used without antigen retrieval.

Staining Procedures

Methylmethacrylate is removed from the sections by three 20 min treatments with methoxyethylacetate (MEA). The slides are then transferred to MEA:water solutions of 80%, 50%, and 20% for 7 min each followed by two additional 10 min rinses in distilled water.

Single staining, peroxidase technique Endogenous peroxidatic activity is removed by treating the sections with 0.4 U/ml glucose oxidase/10 mM beta-D-glucose/1 mM NaN_3 in PBS for 60 min at 37 °C. After washing in PBS, MAb QBend10 and mAb TM1009 are mixed in PBS/1% BSA/0.1% NaN_3 containing 0.003 mg/ml avidin at a final dilution of 1 : 3000 and 1 : 60, respectively. Both reagents are applied overnight at 4 °C. Then the sections are incubated with the biotinylated secondary antibody of the Vectastain® Elite Kit for Peroxidase (Vector Laboratories, Burlingame, USA, No. BA-9200, via Alexis, Grünberg, Germany) for 30 min at room temperature as recommended by the supplier with a final concentration of 0.02 mg/ml biotin in the solution. Subsequently, the avidin-biotinylated peroxidase complex is prepared according to the instructions and also applied for 30 min at RT. MAb TM1009 needs amplification by tyramide enhancement, a procedure, which leads to increased deposition of biotin by peroxidase-catalysed binding of biotinylated tyramide to antigen-containing areas. For this purpose, biotinylated tyramide is prepared by incubating 50 mg NHS-LC-biotin (Pierce No. 21335) in 20 ml 0.025 M borate buffer pH 8.5 with 15 mg tyramine-HCl (Sigma No. T 2879) on a stirrer overnight. After filtration, this solution is either stored frozen for further use or diluted 1 : 500 in TBS pH 7.6 containing 0.09 mM H_2O_2 and immediately applied to the sections for 10 min at RT. Finally, after washing, the avidin-biotinylated peroxidase complex of the Vectastain kit is applied again as described above and peroxidase reactivity is revealed in brown colour by a diaminobenzidine (DAB) reaction in TBS. The sections are then dehydrated and coverslipped in Eukitt®(Sigma-Aldrich, Steinheim, Germany, No. 03989). A negative control using the Vectastain kit without a primary antibody is included in each experiment. Sections of Technovit-embedded human spleen tissue are used as positive controls. MAb asm-1 against smooth muscle alpha actin was diluted

1 : 1000. This reagent was used with tyramide amplification and with or without high temperature antigen retrieval.

Single staining, alkaline phosphatase technique   MAbs QBend10 or TM1009 are used at 1 : 500 or 1 : 60 dilutions in the same buffer as mentioned above. Instead of the biotinylated secondary antibody and the ABC complex, the ultravision labelled polymer kit (Thermo Fisher Scientific, Schwerte, Germany, No. TL-060-AL) for alkaline phosphatase (AP) or the Vectastain AP elite kit (Vector Laboratories, Burlingame, USA, No. AK-5000, via Alexis, Grünberg, Germany) are applied according to the instructions of the manufacturers. The presence of alkaline phosphatase is detected in blue colour by a standard reaction with naphthol-AS-MX-phosphate in Tris/HCl pH 8.2 and Fast Blue BB salt containing 240 µg/ml levamisole. The slides are coverslipped in Mowiol (polyvinyl alcohol, Sigma-Aldrich, Steinheim, Germany, No. 32.459-0).

Double staining   MAb QBend10 is applied first at 1 : 500 and visualised with Fast Blue as described above. After washing, the sections are then incubated with mAb TM1009 at 1 : 60 with detection by tyramide amplification and DAB. Double-stained specimens are coverslipped with Mowiol. Two controls omitting either the first or the second primary antibody are always included in each experiment. This is a subtractive method. Cells stained for the first primary antibody do no longer stain for the second primary antibody.

Photodocumentation of staining results

The slides are photographed using a Zeiss Axiophot microscope with a Canon 60D digital SLR camera using the life view function for focusing on a screen.

Videodocumentation of working procedures

The methods video (Fig. C.1) is produced using a Panasonic Lumix DMC-GH4 camera and a Lumix G Vario 14–140 mm/3.5-5.6 lens as well as a H-ES045 Leica DG Makro Elmarit 45 mm/2.8 lens. *Blender* 2.7.6 (Blender Foundation, Amsterdam, Netherlands) open source software is used for video editing.

6.3.9   Imaging methods

Registration and selection of regions of interest (ROIs)

The serial bone sections are scanned using a Leica SCN 400 microscope with a ×20 lens which produces an output resolution of 0.5 µm/pixel. The single sections of the series are not aligned. The computer-based processing procedure consists of five consecutive coarse steps: registration, volume filtering and segmentation, volume-to-mesh conversion, mesh filtering including shape diameter function (SDF) processing or visualisation of connected components and rendering (Fig. 6.5).

   The complete scanned images of each section are very large, typically $23k \times 27k$ pixels. We assumed the thickness of a single section to be 7 µm, which is the average in a thickness range 5 µm to 10 µm. The images need to be registered before any further processing can take place. Registration was done using a sparse single-resolution non-rigid registration method (Chapter 1). Non-rigid registration was introduced to mend the distortion by the sectioning procedure. We deem this registration a coarse one.

The ROIs can only be selected after coarse registration (Fig. 6.3). The aim was to find interesting regions partially containing adipose tissue as well as haematopoietic areas and lacking artefacts through the whole stack of 21 sections. The size of ROIs selected was $4k \times 4k$ pixel. After coarse registration we applied a refined registration method. This method represents a sparse multi-resolution non-rigid registration [Lobachev et al., 2017a, Chapter 2]. A hierarchy of feature sizes was used to successfully register not only the large features (corresponding to large blood vessels or other large structures such as trabeculae), but also smaller features that correspond to microvessels. After the registration we trimmed the edges of the image stack to obtain cut surfaces on the sides. The final resolution was $3.5k \times 3.5k$ pixels in 21 images. The registration software was custom-written. It is available as part of the supplemental material published by Lobachev et al. [2017a].

Interpolation and conversion of colour data

At this point we still work with the original image data corrected for distortions. For further processing, including the interpolation, we convert the RGB (red, green, blue) data of the scanned serial sections stained in brown to the saturation channel of the hue saturation value (HSV) colour space. Basically, in the HSV saturation channel, the image is light where the saturation is high in stained cells and dark in non-stained areas. We use single-channel images to reduce data volume. This reduction is required for subsequent compensation of anisotropy by interpolation that increases the number of images from 21 to 140. We basically reduce data size from 4.79 GB to 1.60 GB in uncompressed form by using one colour channel with $3500 \times 3500 \times 140$ voxels. Alternative and potentially simpler transformations are theoretically possible. For example, the inverted red channel of the RGB data alone would also produce a good result.

The resolution of the scan in the xy plane (i.e., in the scanning plane) is $0.5\,\mu m$/pixel. However, the resolution in the direction of the z axis (i.e., the thickness of the section) is only $7\,\mu m$/section. This high anisotropy would prevent us from obtaining optimal results. Hence, we interpolate between each pair of consecutive sections (Chapter 3). We use dense optical flow [Farnebäck, 2003] to obtain the 'movement' between section i and section $i + 1$ for all i from 1 to 20. This motion information is then used to generate intermediate images between two sections (Chapter 3). The resolution in the direction of the z axis thus becomes $1\,\mu m$/image. The resulting 'stack' of 140 images is the input volume data for further processing. This operation is performed with a custom-written software based on the open source library Open CV (version 3.1.0, [Bradski and Kaehler, 2008]).

Filtering

To enhance the final result, we apply a set of filters to the volume data. Basically, we use a sequence of simple image processing operations aimed at obtaining better input data for the 3D reconstruction. In this step, we remedy intensity changes still visible in the images due to tangential sectioning of stained vessels or cells or to small interruptions in the contours of blood vessel walls.

The image processing operations consist of the following steps. First, we threshold the volume: all pixels with intensity below 70 (of 255) are set to black. This is a coarse form of segmentation. Next we apply a 3D greyscale closing operator with kernel radius of 10 voxels. Then the closing filter is used to mend at least some missing parts of blood vessel walls. We need to eliminate the inner contours of the vessel walls, to permit correct computation of large vessel diameters in later steps. The hole fill method is one procedure to approach this problem. Further procedures to remove remnants of double contours are executed at the mesh processing stage. After applying a close hole filter we perform a greyscale 3D dilating operation with a radius of 3 voxels. Next, another closing

filter with a radius of 5 voxels is applied. Finally, the image is blurred with a Gaussian blur procedure with sigma value 0.33 to smooth the surface of the reconstructed blood vessels. In summary, six filters, namely Threshold 70, Close 10, Fill hole, Dilate 3, Close 5, and Blur 0.33, are applied to the initial interpolated volume.

All these operations are performed in the open source software *3D Slicer* (version 4.5.0-1 r24735, [Fedorov et al., 2012]). The changes which the original data undergo are shown for 4 small regions within each of the four large ROIs (Fig. 6.6a,b; Fig. C.2, doi: 10.5281/zenodo.128861). This quality control procedure permits visual comparison of the original aligned section (Fig. 6.6b1), the HSV conversion (Fig. 6.6b2) and (Fig. 6.6b3), the interpolation (Fig. 6.6b4), and the final result of filtering (Fig. 6.6b5). Differences between Fig. 6.6b4 and Fig. 6.6b5 are highlighted in Fig. 6.6b6.

Mesh construction

Next, we apply the marching cubes algorithm [Lorensen and Cline, 1987] to convert the volume representation to a mesh. We start with a volume data set, similar to a usual 2D image. We have voxels, '3D pixels', for each point in the direction of the x, y, and z axis. After performing the marching cubes algorithm, we obtain a 'mesh', a network of triangles that basically corresponds to surfaces of constant intensity (so-termed iso-value) in the volume. Basically, our desired result is a network of triangles corresponding to the surface of blood vessels we aim to depict. Meshes are the standard data representation form in computer graphics. The marching cubes operation was also performed using *3D Slicer*.

The output mesh is large (typically about 1.5 GB) and direct processing is troublesome. Application of smoothing and decimation methods build-in in *3D Slicer* did not produce the desired result as they greatly impacted the quality of the output mesh. We thus apply a series of filtering operations to reduce data size and enhance the quality of the final visualisation. First, we remesh and heal the initial mesh using *PolyMender* (version 1.7.1; [Ju, 2004]), then we apply 10 iterations (the default setting) of the Taubin smoothing filter [Taubin, 1995]. The goal of 'Taubin smooth' is to reduce the unevenness of the blood vessel surface without disturbing object size. It is essential to use Taubin and not Laplace smoothing filter, as Laplace smoothing shrinks the lumina of the blood vessels. This step happens in *MeshLab* (version 1.3.3, [Cignoni et al., 2008]). Next we proceed to 'solidify' and remesh the data again. We aim for a representation of blood vessels as solid forms even though we use the mesh, a surface-based representation. For this purpose, we apply an automatic processing operation available in *MeshMixer* (version 11.0.544, [Schmidt and Ratto, 2013]). We also use this software to improve our presentation of cut surface planes.

6.3.10   3D models

3D models are presented using two different colour codes. In the first type of models (Fig. 6.4b,e,h,k) we need to totally remove all inner contours inside the mesh network. We map each vertex with its volumetric obscurance factor [Iones et al., 2003], basically noting down how visible the particular point is. Then we remove points (actually: vertices with connected edges and faces) that have low visibility. In order not to disturb the actual mesh surface by this procedure we choose the selection threshold in a way preventing deletion of vertices on the mesh surface. Hence, we delete only vertices inside the closed structure. Subsequently, we eliminate all non-connected components inside and outside blood vessels the size of which is less than 5% of the space diagonal of the model, i.e., components with a largest diameter less than 125 μm.

The resulting mesh is processed using the so-termed shape diameter function (SDF, [Shapira et al., 2008]). The basic idea of SDF is simple: from each point (vertex) in the mesh, rays are cast in varying directions. When a ray encounters a surface in the mesh, it stops. The shortest of all rays is recorded at the vertex. If we additionally utilise the normals, i.e., if we know the 'inside' and 'outside' of the mesh, and if rays are only shot inside the mesh, we obtain the diameter information at each vertex, the final result of SDF. We use the implementation of SDF from *MeshLab*.

To visualise the SDF data, we colour all vertices red if the SDF value is under 20 µm, and green, if it is above 40 µm. Intermediate values are shown as a gradient from red to green (Fig. 6.4b,e,h,k). The gradient is centred at 23.5 µm. In order to establish sectioned surfaces of the entire model, we replace the colour information at the sectioned surface by light grey using a custom-written mesh segmentation script.

The second type of models (Fig. 6.4c,f,i,l) demonstrates the largest non-connected microvessel network in blue colour, smaller non-connected networks in green and small components sized 12.5 µm to 50 µm in their largest diameter in red. The smallest components (smaller than 12.5 µm in largest diameter) are discarded using a filter in *MeshLab*. The cut surface of the ROI is coloured light blue. The connectivity test is a simple mesh-based analysis technique, which basically detects connected triangles.

6.3.11 Quality control of the three-dimensional models (Fig. 6.5)

In order to control the quality of the 3D models generated from the four ROIs, we established two semi-automatic and one manual quality control (QC) methods. In detail, we devised two semi-automatic QCs for volume filtering and mesh filtering and a manual QC covering all operations from the result of registration to the input into the video rendering program.

For semi-automatic QC of all volume-related procedures after the registration step we designed a sequence of optical visualisations in three to four small regions per ROI. In these small regions we tried to detect errors caused by processing of the volume data, consisting of conversion to the saturation channel of the HSV colour space, interpolation between two consecutive sections in HSV and application of several types of filters (see also 'filtering', Fig. 6.6a,b; Fig. C.2).

The second semi-automatic QC method consisted of measuring Hausdorff distances between meshes. It was also used for repair purposes. For the processing of cut surfaces we need to resort to algorithms that may significantly alter further parts of the mesh surface. Among other operations, we remesh and solidify the mesh representation of the vasculature to produce filled, flat and segmented cut surfaces. A control step is mandatory afterwards to ensure consistency. We automatically measure Hausdorff distances to detect parts of the marching cubes output mesh that are lost during the processing and then add the lost parts to the processed mesh.

The Hausdorff distance computation is a unidirectional representation of distances from one mesh to another [Alt and Guibas, 1999]. At each vertex of the first mesh rays are shot in all directions. As soon as a ray intersects the surface of the second mesh, the ray is terminated. The length of the shortest ray is noted at the vertex. This procedure defines the distance between two meshes. However, if a component is not present in the first mesh, a distance cannot be determined at that point. For this reason, we need to compute Hausdorff distances in both directions. A large distance from the first mesh to the second identifies structures present in the first mesh and lacking or extremely different in the second mesh and vice versa.

In detail, we compute Hausdorff distances between the result of *PolyMender* on the marching cubes mesh (first mesh) and the result of the last step of our pipeline (second mesh) and vice versa. All structures from the first mesh that are more than 18 µm away from the second mesh are added

to the second mesh. After this procedure short discontinuities may still be present in the repaired blood vessels, but larger parts of microvessels are no longer missing, unless registration or volume processing errors have occurred. The computation of Hausdorff distances provides a reliable tool to control the second part of the reconstruction procedure. We performed manual QC (Fig. C.2, see below) of blind vessel ends in the four models visualising Hausdorff distances in forward direction and found that any missing structures were either correctly detected by the method or were due to the first and not to the second part of the processing pipeline.

With respect to manual QC, we hypothesised that bone marrow microvasculature represents a closed system and sampled 10 vessels per ROI, which had blind ends in the 3D-models and thus did not reach a second surface. In addition, we controlled single small elongated structures located between the microvessels (shown in red in Fig. 6.4c,f,i,l) In order to obtain more information about such unexpected structures, we devised a special program (described next section) which permitted capturing selected structures in the mesh representation of each ROI and projecting these structures either onto single registered scans or onto the entire series of scans. A three-dimensional impression emerges when the view vector is tilted by the user.

6.3.12   The volume renderer

The volume renderer program loads the (registered) stack of serial sections with the reconstructed mesh (or parts thereof) and enables the 3D representation of both simultaneously. Basically, the user can view the whole volume or page through the sections in series and view the corresponding parts of the 3D reconstruction—from different angles and with arbitrary zoom factor (Figure 6.8). This tool enables us
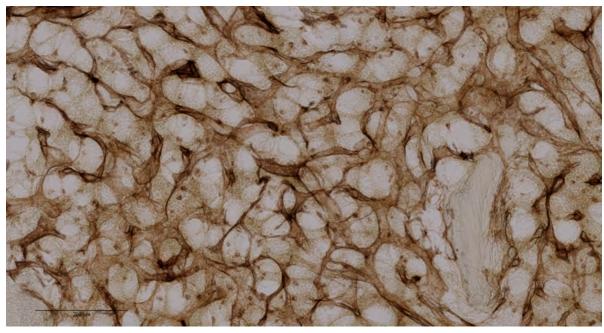
- To blend in the reference slices or their cumulated volume rendering at will. Domain experts are much better trained at reading stainings in slices.

- To visualise and verify ruptures in early reconstructions, e.g., an interruption in an actually continuous microvessel due to registration errors or wrong settings—in other words: to perform a quality control of our reconstructions, as detailed above.

During quality control we found some blindly ending blood vessels that were not artefacts. We attributed them to agniogenesis. We enhanced the volume renderer with 3D stereo support for a more immersive presentation and improved user experience. An even further leap in user experience was obtained in later experiments through VR immersion. Chapter 5 details on the VR application.

Acknowledgements

(a) Volume data representation.

(b) Visualisation of final, filtered mesh vs. corresponding single slice.

Figure 6.8: Showcasing our volume renderer. Endothelia of blood vessels are stained brown. The blended-in mesh is blue.

CAPILLARY NETWORKS AND FOLLICULAR MARGINAL ZONES IN HUMAN SPLEENS. THREE-DIMENSIONAL MODELS BASED ON IMMUNOSTAINED SERIAL SECTIONS

In which we discuss medical results obtained from our 3D reconstructions. This chapter is written after our paper Steiniger et al. [2018]. It emerged in collaboration with Birte S. Steiniger, Christine Ulrich, Moritz Berthold, and Michael Guthe.

Author Contributions: *BSS designed the study and wrote the manuscript. CU helped establishing the methods for registration and 3D rendering of follicular capillaries. MB developed an application for quality control in VR. MG supervised digital processing, and OL designed and performed the details of all digital processing steps for 3D reconstruction.*

We have reconstructed small parts of capillary networks in the human splenic white pulp using serial sections immunostained for CD34 alone or for CD34 and CD271. The three-dimensional (3D) models show three types of interconnected networks: a network with very few long capillaries inside the white pulp originating from central arteries, a denser network surrounding follicles plus periarterial T-cell regions and a network in the red pulp. Capillaries of the perifollicular network and the red pulp network have open ends. Perifollicular capillaries form an arrangement similar to a basketball net located in the outer marginal zone. The marginal zone is defined by MAdCAM-1$^+$ marginal reticular stromal cells. Perifollicular capillaries are connected to red pulp capillaries surrounded by CD271$^+$ stromal capillary sheath cells. The scarcity of capillaries inside the splenic white pulp is astonishing, as non-polarised germinal centres with proliferating B-cells occur in adult human spleens. We suggest that specialised stromal marginal reticular cells form a barrier inside the splenic marginal zone, which together with the scarcity of capillaries guarantees the maintenance of gradients necessary for positioning of migratory B- and T-lymphocytes in the human splenic white pulp.

Videos 7.1a–7.4 referenced in this chapter are available in data repository under doi: 10.5281/zenodo.845670, as detailed in Appendix D.

HE SPLEEN is the organ with the most enigmatic microvasculature, both in humans and in rats or mice [Snook, 1964, Balázs et al., 2001, Kusumi et al., 2015]. In all three species the main part of the organ, the so-termed red pulp, harbours two types of microvessels, namely capillaries and sinuses, which form the end of the arterial tree and, respectively, the beginning of the venous part of the circulation. The exceptional feature is that both types of microvessels are not connected to one another in humans [Steiniger et al., 2011a]. The capillaries have open ends delivering the blood to strands of reticular connective tissue. Blood plasma and all blood cells finally enter venous sinuses from the outside via openings in their walls. Thus, the spleen is the only organ exhibiting an open circulation where blood flows in spaces not delimited by endothelial cells, but by fibroblasts. Apart from different types of fibroblasts, specialised macrophages, plasma cells and potentially also mast cells form more sessile cell populations in the splenic red pulp.

Besides this very special structure of the red pulp, the spleen also hosts dense accumulations of more or less migratory lymphocytes, which are named the white pulp. The white pulp is populated by T- and B-lymphocytes, which migrate in special compartments located in the vicinity of branches of the splenic artery termed central arteries. Details of splenic white pulp compartments were first described in rats and mice [Snook, 1964, Nieuwenhuis and Ford, 1976, Veerman and van Ewijk, 1975].

The white pulp consists of periarterial lymphatic sheaths (PALSs) predominantly occupied by T-lymphocytes and of hemispherical follicles attached to the sheaths where B-lymphocytes predominate. Both lymphocyte types are attracted by compartment-specific resident connective tissue cells termed fibroblastic reticulum cells (FRCs) in case of PALSs and follicular dendritic cells (FDCs) in case of follicles. During B-cell immune reactions secondary follicles are formed, which comprise a germinal centre with proliferating antigen-stimulated B-lymphocytes and a surrounding mantle zone formed by naive recirculating B-lymphocytes. An additional compartment is the marginal zone (MZ), which surrounds both PALSs and follicles in rats and mice. The MZ is separated from these compartments by a blood-filled cleft called the marginal sinus and by marginal metallophilic macrophages (MMMs, Kraal and Mebius [2006], Dijkstra et al. [1985]). The rat MZ is difficult to classify because it belongs to the open circulation system and harbours a minor number of red blood cells, but at the same time it is densely populated by MZ B-lymphocytes. Thus, it may be regarded as part of the red pulp or of the white pulp. In rats and mice B-lymphocytes in the MZ are a heterogeneous population comprising memory B-cells [Gray et al., 1982, Liu et al., 1988], but also B-lymphocytes of an independent MZ developmental pathway [Hammad et al., 2017].

In adult human spleens T- and B-cell areas of the white pulp and also the MZ are arranged differently from rats and mice [Steiniger et al., 2005, 2003, 2014b, Steiniger, 2015]. In contrast to both rodent species, the PALSs are of limited length and spherical follicles form the major compartment of the white pulp. Human CD27$^+$ memory B-cells [Bagnara et al., 2015] are only present at the surface of follicles and are absent from the surface of the PALSs. A marginal sinus and MMMs are also lacking in humans. The human MZ cannot be defined by B-cells, because CD27$^+$ memory B cells both occur among the most superficial FDCs in the follicular mantle zone and among marginal reticular cells (MRCs, Steiniger [2015]). These cells form a third stromal cell type superficial to FRCs and FDCs. They have been first described in mouse lymph nodes [Katakai et al., 2008]. It is not clear, whether the superficial follicular stromal cells in human spleens are totally identical to mouse MRCs. MRCs are the only organ-specific cells defining a MZ-equivalent in the human spleen.

In previous publications we have analysed the composition of an additional species-specific trait of human spleens, the capillary sheaths [Steiniger et al., 2014a]. The most unusual cell type in these structures are strongly CD271 positive (CD271$^+$) capillary sheath cells. These cells represent the most strongly CD271$^+$ cell type in human spleens. CD271 is also present in FDCs, FRCs, and with less intensity in ubiquitous splenic fibroblasts [Steiniger et al., 2014b]. It is highly likely that sheath cells represent the only sessile cells in human capillary sheaths. They are surrounded by macrophages and recirculating B-cells [Steiniger et al., 2014a]. Capillary sheaths are absent in mouse or rat spleens. This is exceptional, because capillary sheaths or equivalent structures termed ellipsoids occur in most vertebrate species.

The microvasculature of human spleens has been investigated in 3D before [Kusumi et al., 2015, Steiniger et al., 2011a, 2003]. However, up to now the exact position of splenic capillaries in the white pulp has not been localised. To obtain a more comprehensive overview of all capillaries in a human spleen, we now reconstructed serial sections stained for CD34 alone and for CD34 followed by visualisation of CD271. We describe three different interconnected capillary networks. In addition, we definitively prove that humans do not exhibit a marginal sinus. The 3D models developed are quality controlled by direct comparison to the original immunostained serial sections in virtual reality (VR).

7.1 Results

7.1.1 Stromal cells, B lymphocytes, and microvessels at the surface of splenic follicles

The distribution of CD34 and CD271 in human spleens has already been described by us in 2D [Steiniger et al., 2007, 2014b, Steiniger, 2015], but in the present study some additional details were found. CD34 is present in endothelial cells of larger splenic arterial and venous vessels and in capillaries. In addition, fibroblasts located periarterially, within trabeculae and in the splenic capsule are stained. Perifollicular sinus endothelia are also faintly positive, while the majority of the splenic red pulp sinuses appear to be CD34 negative (Fig. 7.1a,b,d,e,g,h, Fig. 7.2a,b,d,e,g,h). Single round CD34$^+$ cells are also present in the red pulp, which may represent haematopoietic precursor cells in the open circulation.

When single sections or overlays were inspected in the present investigation, large parts of the white pulp did not contain any CD34$^+$ endothelial cells except in central arteries and their branches (Fig. 7.1a,b,d,e,g,h, Fig. 7.3). Few single faintly CD34 positive cells were seen in PALSs, which remain to be further analysed (Fig. 7.1h, Fig. 7.3). As described previously [Steiniger et al., 2007, 2011a], there was a network of capillaries in the red pulp. Careful analysis of the follicular surfaces for CD34 revealed that a network of strongly positive capillaries was located inside a ring of faintly stained larger sized sinuses (Fig. 7.1a,b,d,e,g,h, Fig. 7.3, Fig. 7.4d). The individual perifollicular capillaries appeared more irregular and the network exhibited a different branching pattern than that found in red pulp capillaries (Fig. 7.3). The perifollicular capillary network appeared to give off branches to red pulp capillaries (Fig. 7.4a), which could not be followed further in single sections.

We first tried to localise the perifollicular CD34$^+$ capillary endothelial cells in relation to the different cell types described to occur at the surface of human splenic follicles. The network was embedded in MRCs expressing MAdCAM-1, smooth muscle alpha-actin (SMA), and CD141 (Fig. 7.4a–c). In the specimen investigated most of the perifollicular CD27$^+$ B-lymphocytes were located centrally to the perifollicular capillary network with scattered faintly CD27$^+$ B-cells occurring between the capillaries (Fig. 7.4d).

CD271, the low affinity nerve growth factor receptor p75, is most strongly expressed in capillary sheath cells of isoprismatic shape in the red pulp as well as in FDCs [Steiniger et al., 2014b,a]. FRCs including adventitial cells around central arteries are less strongly stained. Fibroblasts at the surfaces of splenic trabeculae, but not in their interior, are also positive. In addition, there is a faintly positive ubiquitous stromal cell population in red pulp cords (Fig. 7.2a,d,g).

These stromal cells appeared to express more CD271 in the vicinity of the perifollicular capillary network mentioned above (Fig. 7.2a,b,d,e,g,h). The CD271$^+$ fibroblasts surrounding the perifollicular capillary network were sometimes as strongly stained as capillary sheath cells, although they were morphologically different. Interestingly, there was a conspicuous CD271$^-$ area at the follicular surface between the CD271$^+$ FDCs and the CD271$^+$ fibroblasts surrounding the capillary network (Fig. 7.2a,d,g). This area continued around the PALSs. The CD271$^-$ area around the follicles seemed to host those CD27$^+$ memory B-cells (Fig. 7.4d, Steiniger [2015]), which expressed CD27 most strongly. In addition, few scattered IgD$^+$ cells, many CD4$^+$ T-cells (Fig. 7.4e), and—in one specimen— also a few plasmablasts positive for Ki-67 and intracellular IgM occurred in this location (Fig. 7.4f). The CD271$^-$ area was primarily supported by MAdCAM-1$^+$ fibroblasts (Fig. 7.1b, Steiniger et al. [2014b]) and might represent a special migration compartment. The innermost strongly CD27$^+$ B-cells, however, clearly overlapped with typical CD271$^+$ mantle zone FDCs [Steiniger et al., 2014b]. Thus, CD27$^+$ B-cells were found in two different stromal compartments defined either by FDCs or by MRCs at the follicular surface. This fact excludes that CD27$^+$ B-cells can be used for defining a
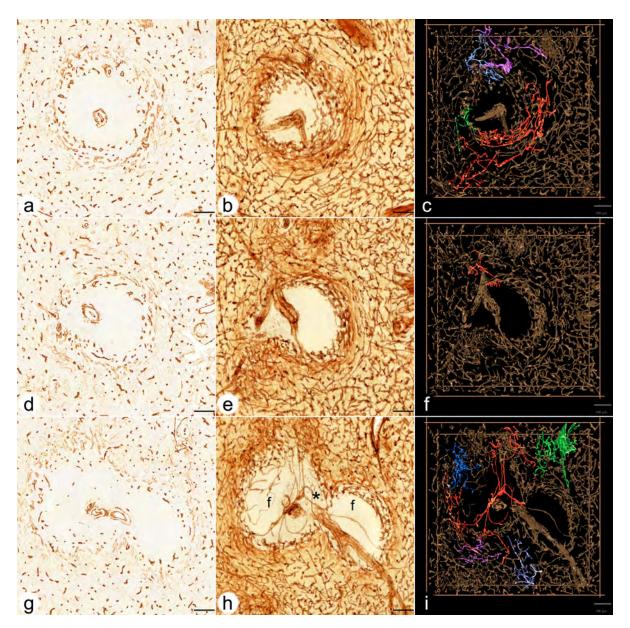
Figure 7.1: Spleen follicles, single-staining. **(a, b, c) ROI 1** (a) first section stained for CD34, (b) overlay of all 24 registered sections, (c) first frame of video showing a 3D model with highlighted connections among perifollicular and red pulp capillaries. **(d, e, f) ROI 2** (d) first section stained for CD34, (e) overlay of all 24 registered sections, (f) first frame of video showing a 3D model with highlighted connections among perifollicular and red pulp capillaries. **(g, h, i) ROI 3** (g) first section stained for CD34, (h) overlay of all 24 registered sections, (i) first frame of video showing a 3D model with highlighted connections among perifollicular and red pulp capillaries. (g, h, i) show two follicles (f in (h)) and a part of a PALS (asterisk * in (h)) in between. Internal capillaries arise from a central artery in the PALS. Staining of red pulp sinuses is eliminated as far as possible in c, f, i. CD34 is present in capillary endothelia, adventitial fibroblasts of arteries, perifollicular sinus endothelia (weak) and in fibroblasts in trabeculae.
All scale bars = 100 μm.
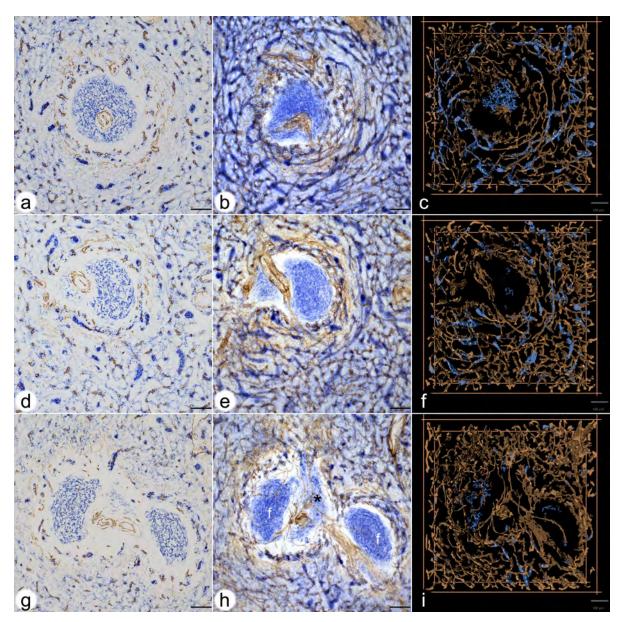
Figure 7.2: Spleen follicles, double staining. **(a, b, c) ROI 1** (a) first section stained for CD34 (brown) and for CD271 (blue), (b) overlay of all 24 registered sections, (c) first frame of video of the 3D model. **(d, e, f) ROI 2** (d) first section stained for CD34 (brown) and for CD271 (blue), (e) overlay of all 24 registered sections, (f) first frame of video of the 3D model. **(g, h, i) ROI 3** (g) first section stained for CD34 (brown) and for CD271 (blue), (h) overlay of all 24 registered sections, (i) first frame of video of the 3D model.
CD271 is most strongly expressed in stromal capillary sheath cells. FDCs in follicles, FRCs in a PALS and ubiquitous interstitial fibroblasts are also positive. Staining of red pulp sinuses for CD34 and of ubiquitous fibroblasts for CD271 is eliminated as far as possible. (g), (h) and (i) show two follicles (f in (h)) and a part of a PALS in between. Internal capillaries arise from a central artery in the PALS (asterisk * in (h)).
All scale bars = 100 µm.
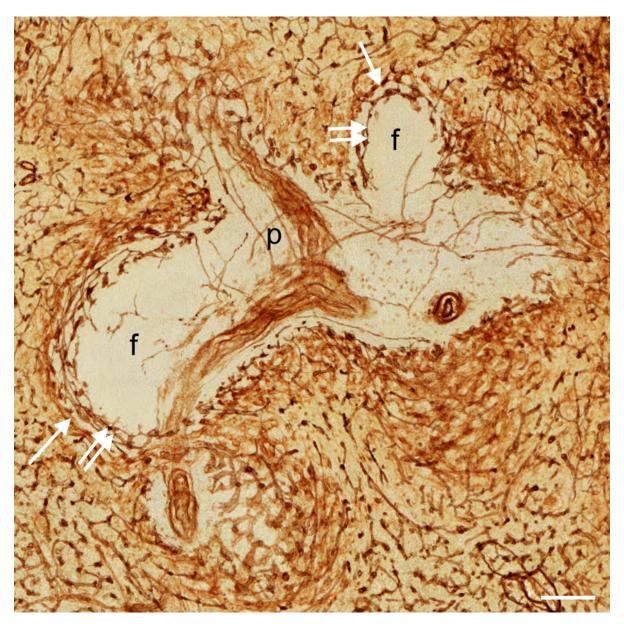
Figure 7.3: Distribution of CD34 in an overlay of 24 registered sections of ROI 4 showing a PALS (p) with a branching central artery and two sectioned follicles (f). There is a clear difference in the localisation of the perifollicular capillary network (double arrows) and the perifollicular sinus network (arrows). Internal capillaries of the PALS appear to run in the adventitia of the central artery forming vasa vasorum.
Scale bar = 100 μm.

MZ equivalent in human spleens. Detailed inspection of the perifollicular capillaries revealed that some of them had open ends (Fig. 7.4c).

7.1.2   The perifollicular capillary network in 3D models

A 3D model of the perifollicular capillary network was constructed from 24 serial sections taken from a typical adult spleen. Four regions of interest (ROIs) were chosen for a detailed analysis (Fig. 7.5). The serial sections were first stained for CD34 in brown colour. After scanning, they were stained for CD271 in blue and photodocumented again. Thus, ROI 1–3 were represented by a single-coloured and by a double-coloured data set. ROI 4 was only captured in brown. Both data sets were visualised as sequences of sections (Videos 7.1a,b, 7.2a,b, 7.3a,b), overlays (Fig. 7.1b,e,h; Fig. 7.2b,e,h) and as 3D models (Videos 7.1c,d, 7.2c,d, 7.3c,d). The single-coloured sets were used to visualise the perifollicular capillary network after optimal exclusion of the weakly CD34$^+$ perifollicular sinuses. Exclusion of sinus staining was mandatory to find connections among the perifollicular and the red pulp capillary networks. Figs. 7.1c,f,i and Videos 7.1c–7.3c demonstrate that the perifollicular capillary network is connected to the network of red pulp capillaries. The connections highlighted represent an underestimation of their real number, because of the rigorous filtering methods to exclude sinuses. It is highly likely that this procedure interrupted a large number of connections between perifollicular and red pulp capillaries.

Overlays of all sections (Fig. 7.1b,e,h), sequences (Videos 7.1a–7.3a), and 3D models (Fig. 7.1c,f,i; Videos 7.1c–7.3c) revealed astonishingly few capillaries inside the white pulp. Thus, the perifollicular capillary network appears to be primarily fed by red pulp microvessels. The follicles were located within their perifollicular capillaries similar to balls within basketball nets. The capillaries were clearly distinguishable from perifollicular sinuses (Fig. 7.3). They appeared to continue at the surface of the PALSs, but this could not be investigated in detail because of the limited number of serial sections.

Staining for CD34 in brown with subsequent visualisation of CD271 in blue revealed the arrangement of capillaries and capillary sheath cells in the surroundings of follicles. For 3D reconstruction, the iso-values for mesh construction were chosen to eliminate the ubiquitous light blue CD271$^+$ fibroblasts in the red pulp, which would have obscured the model. Thus, we accepted losing the majority of the CD271$^+$ FDCs and all FRCs to improve visualisation of capillary sheaths.

The frequency of capillary sheaths around follicles did not seem to differ from that in the red pulp. Interestingly, some of the perifollicular capillaries that had been demonstrated to be connected to the capillary network of the red pulp by single-staining for CD34 were shown to be connected to capillaries bearing sheaths (Fig. 7.1c versus Fig. 7.2c). Capillary sheaths were elongated and often branched structures of variable length. The CD271$^+$ stromal sheath cells were isoprismatic and resembled epithelia. Their unstained nuclei often left gaps inside the cells, which sometimes provoked dissociation of sheaths during 3D reconstruction. The perifollicular capillary net was accompanied by branched CD271$^+$ fibroblasts, which occasionally were so strongly stained that they could not be eliminated from the 3D model shown in Fig. 7.2c,f,i and in Video 7.1d–7.3d. These cells were, however, morphologically different from typical sheath cells.

7.1.3   Internal capillaries in the white pulp in 3D models

In two single follicles investigated (Fig. 7.1a–f; Fig. 7.2a–f; Video 7.1a–c, 7.2a–c) internal capillaries were not observed in 24 serial sections. To further analyse this phenomenon, we chose two additional regions, both containing two follicles and an associated PALS at low (Fig. 7.1g–i; Fig. 7.2g–i,
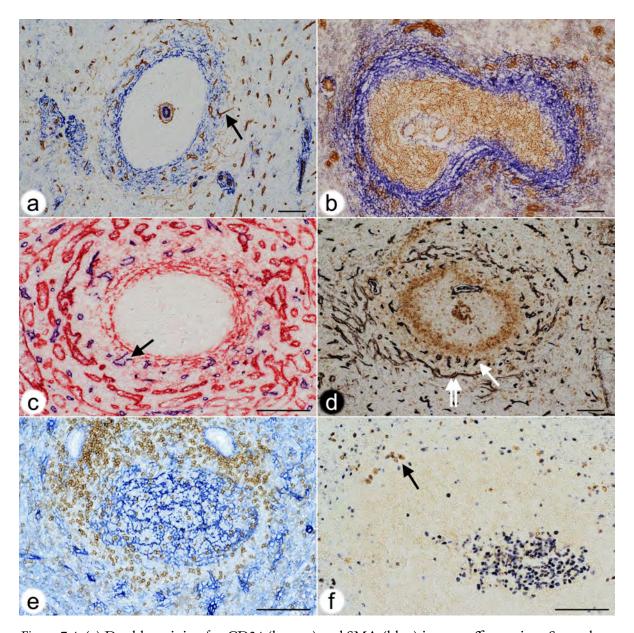
Figure 7.4: (a) Double staining for CD34 (brown) and SMA (blue) in a paraffin section. Strongly CD34$^+$ capillary endothelia form a shell around the follicle. Sinus endothelia are weakly stained. A capillary (arrow) in the outer MZ forms a connection to red pulp. Same individual as Figs. 7.1 and 7.2.

(b) Double staining for CD271 (brown) and MAdCAM-1 (blue) in a cryosection. All FDCs are CD271$^+$. MAdCAM-1$^+$ perifollicular MRCs define the MZ and harbour capillaries surrounded by light brown fibroblasts. CD271$^+$ capillary sheaths occur superficially in the MZ. Same individual as Figs. 7.1 and 7.2.

(c) Double staining for CD34 (blue) and CD141 (red) in a paraffin section. CD141 demonstrates venous sinuses in the red pulp and MRCs. Arrow indicates the open end of a perifollicular capillary. Same individual as Figs. 7.1 and 7.2. *(Caption continued on the next page.)*

Figure 7.4: (Continued from previous page.)
(d) Double staining for CD34 (black) and CD27 (brown) in a paraffin section. The CD34$^+$ capillary network is located in a superficial area of the follicle with scattered CD27$^+$ B-cells. Double arrow shows perifollicular capillaries, arrow indicates sinus. Same individual as Figs. 7.1 and 7.2.
(e) Double staining for CD4$^+$ T-cells (brown) and CD271$^+$ FDCs and fibroblasts (blue) in a paraffin section. CD4$^+$ T-cells originating from a PALS (upper part of image) occupy the CD271$^-$ area underlying the perifollicular capillary network accompanied by CD271$^+$ fibroblasts. Same individual as Figs. 7.1 and 7.2.
(f) Double staining for Ki-67 (blue) and intracellular IgM (brown) in a paraffin section. A germinal centre is located in the lower right part of the image. IgM$^+$ plasmablasts with blue Ki-67$^+$ nuclei (arrow) are located at the surface of the unstained follicle.
Scale bar for all figures = 100 μm.

Videos 7.3a,c) and high magnification (Fig. 7.1). In these regions it was evident that a few capillaries were present in the white pulp, which arose from side branches of the central arteries. The capillaries ran alone and had an extremely long and often rather straight course without branching. Some of these vessels appeared to form vasa vasorum for the central artery. Other capillaries, however, randomly traversed the white pulp without association with any histologically defined regions (Fig. 7.1h,i; Video 7.3a,c). Interestingly, several of the internal white pulp capillaries seemed to end in the perifollicular capillary network (Fig. 7.3).

7.2 Discussion

Our results show that three different but interconnected capillary networks occur in human spleens, a red pulp network, a perifollicular network, and an internal network in the white pulp. It is likely that the perifollicular network is also present around the PALSs, but this was not investigated in detail. The red pulp network contains sheathed capillaries some of which are directly connected to perifollicular capillaries. We show that perifollicular capillaries form part of the open splenic circulation, because unequivocal open ends are observed in single sections.

We have chosen a sample of a representative adult spleen for investigation. Samples of this spleen have been immunostained with a large number of antibodies directed against different T- and B-lymphocyte populations, macrophages and stromal cells in the past. We consider the organ as representative, because the findings obtained with these antibodies correspond to those seen in the majority of adult spleens with secondary follicles we have investigated during the last two decades. Up to now about 80 spleens have been studied altogether.

3D reconstruction is a tedious, time and data volume consuming process, which presently precludes investigating a larger number of individuals. A single scanned section (of 24) amounts to 2.4 GB in the original quality. Each single-stained ROI consisted of 4.1×10^8 voxels after inter-slice interpolation, a double-stained ROI amounted to 8.5×10^8 voxels. ROI 4 had 2.57×10^9 voxels. The size of the initial meshes after the marching cubes procedure ranged from 345 to 1626 MB for brown staining. The total data of the 3D reconstruction project reached 517 GB. Because of the sheer data volume, the specimen and the ROIs for 3D reconstruction need to be very carefully selected.

Our investigation tries to solve the question whether there is an equivalent of the rodent MZ and the marginal sinus in human spleens. The answer is complicated by the fact that the MZ was
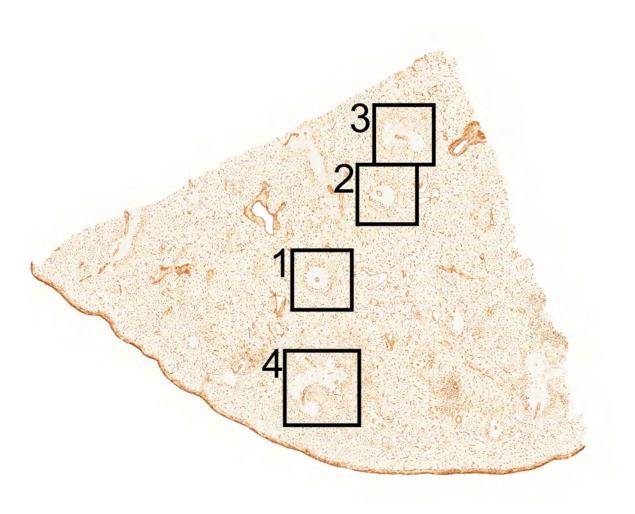
Figure 7.5: Overview of the section stained for CD34 showing localisation of ROI 1–4. The length of the cut edges is about 1 cm.

first described in rats based on routine histology, which does not reveal all relevant cell types. In rats, it was first detected that the MZ is populated by memory B-lymphocytes [Liu et al., 1988]. In humans, the antigen most often used to detect memory B-lymphocytes is CD27, although this antigen is not present in all memory B-cells [Bagnara et al., 2015]. CD27$^+$ B-cells can only be visualised after first staining T-cells, which also express this antigen [Steiniger et al., 2003, 2005]. Human CD27$^+$ memory B-lymphocytes behave differently from rodent MZ B-cells. Most rat and mouse MZ B-cells are located superficial to all parts of the white pulp, i.e., to follicles, PALSs, and the marginal sinus. The majority of human splenic CD27$^+$ B-cells is, however, only found around follicles and is situated below the perifollicular capillary network, i.e., located towards the interior of the follicles. In apparent contrast to rodent MZ B-cells, human CD27$^+$ B-cells appear in two different stromal compartments, characterised by CD271$^+$ FDCs and by MAdCAM-1$^+$ MRCs. Thus, the localisation of MZ B-cells appears to be species-specific. In rats and mice the MZ has long been regarded as a 'static' compartment [Gray et al., 1982]. In contrast, it has also been reported, that rat and mouse MZ B-cells enter follicles in case of activation by LPS and other substances and perhaps even in the steady state [Gray et al., 1984, Liu et al., 1991, Ferguson et al., 2004, Arnon et al., 2013, Cinamon et al., 2008, Turner and Mabbott, 2017]. Interestingly, mAbs detecting rat MZ B-cells also react with

B-cells in follicles [Dammers et al., 1999]. A final solution to this conundrum is still lacking.

3D reconstructions of serial sections have been attempted before [Steiniger et al., 2003, 2011a, Ma et al., 2008b, Saalfeld et al., 2012, Kusumi et al., 2015]. Our registration method [Lobachev et al., 2017a, Chapter 2] is novel in the way it uses computer vision methods to better align the sections. In contrast, most other methods, such as *elastix* [Klein et al., 2010], are image-based. *TrackEM2* [Cardona et al., 2012] from *Fiji* [Schindelin et al., 2012] also targets serial sections. The inter-slice interpolation we use [Lobachev et al., 2017b, Chapter 3] is a dense method. It utilises dense optical flow [Farnebäck, 2003]. A sparse interpolation for electron microscopy has been suggested before [Carata et al., 2013]. Recently, a VR visualisation tool was published for medical applications [Egger et al., 2017], but we used our own custom-written VR visualisation software (Chapter 5).

The monoclonal antibodies (mAbs) used are not totally specific to single cell types. This problem needed to be circumvented for 3D reconstruction. Thus, data processing of sections single-stained for CD34 was optimised to visualise only capillary and large vessel endothelial cells. We tried to eliminate the light brown staining of perifollicular sinuses as far as possible. This method inevitably led to the loss of connections among capillaries. The connections among the perifollicular and the red pulp capillary networks shown in Figs. 7.1c,f,i and Videos 7.1c–7.3c thus only represent examples of all connections which were present. The reconstructions visualise qualitative, but not quantitative facts. A similar condition also applies to the reconstructions of double-stained sections (Figs. 7.2c,f,i; Videos 7.1d–7.3d), where the weakly CD271$^+$ ubiquitous fibroblasts of the red pulp were eliminated to permit better recognition of capillary stromal sheath cells. This procedure was associated with the loss of a variably large number of CD271$^+$ FDCs in follicles (Fig. 7.2c,f,i; Videos 7.1d–7.3d).

Stringent quality control thus played a decisive role for interpretation of the 3D models. We achieved this by using a VR headset to walk through the reconstruction and inspect it from the interior and from all directions. VR permitted positioning each single registered section into the model so that all 3D structures could be directly compared with the staining results by viewing either the model (or parts of it) or the section or both (Video 7.4).

Up to now there is only one comparable investigation on the 3D structure of human spleen microvessels [Kusumi et al., 2015]. These authors used 50 to 100 serial sections to reveal CD34 and thus they were able to follow larger arterial vessels and their branches, which was not possible in our study. Somewhat similar to our results, Kusumi et al. [2015] did not detect capillaries in the human splenic white pulp and postulated that blood only comes from a superficial capillary plexus supplied by arterioles from the red pulp. Our results indicate that central arteries may have side branches supplying PALSs and follicles, but these vessels are astonishingly rare. We do not agree with Kusumi et al. [2015] with respect to the location of the perifollicular capillary network as visualised in their schematic drawings. The authors postulate the existence of a MZ in human spleens without defining this region and they suppose that the capillary plexus exists at the inner border of the MZ. The analysis of cell phenotypes in our present and previous studies does, however, indicate that the CD34$^+$ perifollicular capillaries are located within the area of MRCs located superficial to the majority of strongly CD27$^+$ B-cells, i.e., they exist in the most superficial part of the MZ.

The scarcity of capillaries in the human splenic white pulp and especially in the follicles, may be related to the fact that adult human spleens are immunologically quiescent. Full-blown germinal centres seldom occur in healthy adults. If germinal centres are detected, they are rather small and non-polarised, although B-cell proliferation is still present (Fig. 7.4f). Thus, spleens should also be investigated in immune reactions to see whether new white pulp capillaries appear on demand. However, the lack of internal capillaries may also be related to the necessity of maintaining a gradient of blood-borne mediators across follicles and PALSs to permit proper B-cell and T-cell recirculation. This may also be true for the access of antigen to follicles. In addition, the fact that most of the

capillaries at the surface of follicles are supplied with blood from the red pulp might secure that the blood has previously passed sheathed capillaries. The function of capillary sheaths in humans is unknown. Many human sheaths—perhaps even all—are located in an immediate post-arteriolar position [Steiniger, 2015]. There are indications that immune complexes and carbon particles are retained in capillary sheaths termed ellipsoids in birds and other animals [Oláh and Glick, 1982, Oláh et al., 1984, 2013]. Our previous results have shown that human splenic capillary sheaths are composed of endothelia and pericytes, CXCL13$^+$ stromal sheath cells, macrophages, and recirculating B-lymphocytes [Steiniger et al., 2014a, Steiniger, 2015]. Thus, sheaths may hypothetically extract large size antigens from the blood and participate in initial activation and guidance of naive B-cells immigrating into the spleen.

In mice and rats the marginal sinus does not only form the inner border of the MZ to provide direct access of blood to the MZ, but cells in its inner wall also establish a size barrier for substances approaching the white pulp from the open red pulp circulation [Nolte et al., 2003]. Humans neither possess a marginal sinus nor marginal metallophilic macrophages. It is impossible to functionally investigate whether a barrier restricting size, charge or other properties of blood-borne substances also occurs in human spleens and which cell type is involved. The present investigation suggests that the innermost SMA$^+$ MRCs (Fig. 7.4a) might cooperate with MRCs of additional phenotypes to form a barrier between the open circulation of the red pulp and the white pulp. These fibroblasts and the thick fibres they produce are sometimes visible in conventional paraffin sections or in sections stained for CD27 (Fig. 7.6a,b). They form the only barrier-like structure that exists at the follicular surface. SMA$^+$ fibroblasts also occur in the inner wall of the marginal sinus in rats [Steiniger and Barth, 1999] and may also be responsible for establishing a barrier around the white pulp in this species. In humans, CD27$^+$ memory B-cells and other lymphocytes may be easily able to cross such a barrier. Free erythrocytes are, however, only present outside the innermost SMA$^+$ fibroblasts. Erythrocytes stay in an area, which may be termed the 'outer MZ', which also harbours the perifollicular capillary network with surrounding weakly CD271$^+$ fibroblasts. Switched IgM$^-$IgD$^-$CD27$^+$ B-cells appear to also sojourn in the outer MZ [Steiniger et al., 2014b]. This distribution may correspond to the distinct recirculation behaviour of non-switched and switched human memory B-cell populations noted by others [Bagnara et al., 2015]. Interestingly, the region immediately inside the putative barrier is a CD271$^-$ region with strongly CD27$^+$ B-cells. It may be termed the 'inner MZ'. The characteristic feature of both MZ compartments is the presence of MAdCAM-1$^+$ stromal cells.

7.2.1   Perifollicular capillary sheaths in 3D models

3D visualisation of capillaries combined with immunohistological staining for different cell types in 2D allows the following speculation: CD27$^+$ B-cells in healthy adult humans enter the follicular surface from the red pulp in close association with MAdCAM-1$^+$ perifollicular stromal cells by immigration from open-ended perifollicular capillaries or from the splenic cords. CD27$^+$ B-cells then probe the CD271$^+$ FDCs in the follicular mantle zone and leave again in direction to the red pulp if cognate antigen is not encountered. Probing may be especially important for non-switched CD27$^+$ B-cells. This scenario is derived from the findings that—in contrast to IgD$^+$ recirculating naive B-cells—CD27$^+$ B-cells do not accumulate at any other place in human spleens. Only CD27$^+$ plasmablasts and plasma cells, but not CD27$^+$ B-cells, form groups of cells in the red pulp. CD27$^+$ B-cells are also absent from capillary sheaths. Emigration back to the circulation most probably leads recirculating cells into the perifollicular sinuses. If the expression of CD34 in perifollicular sinuses is not a mere epiphenomenon, it may somehow help CD27$^+$ B-cells and other lymphocytes to re-locate into the venous vasculature.
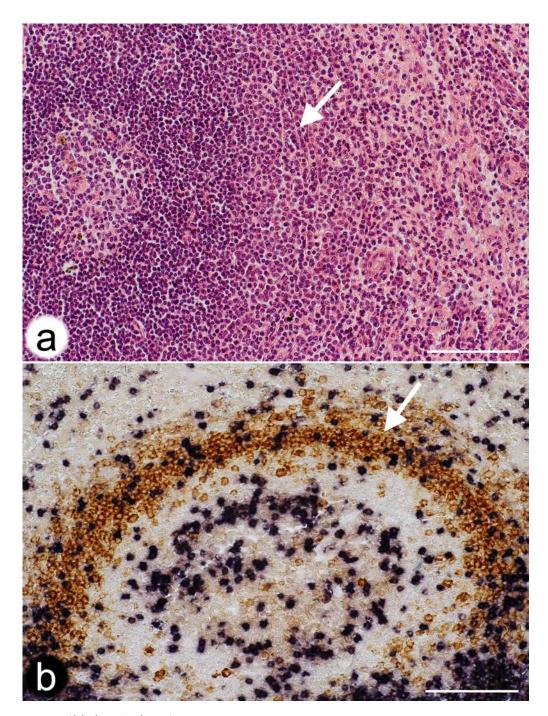
Figure 7.6: Possible barrier locations.
(a) Hemalum-Eosin-stained paraffin section of a typical secondary follicle in an adult human spleen. Arrow indicates a layer of thick fibres and fibroblasts potentially forming a barrier against the open circulation. Same individual as Figs. 7.1 and 7.2.
(b) Double staining for CD3[+] T-cells (blue-black) and CD27[+] B-cells (brown) in a paraffin section. Arrow indicates fibres and fibroblasts outside the most strongly CD27[+] B-cells potentially forming a barrier against the open circulation.
Scale bars = 100 μm.
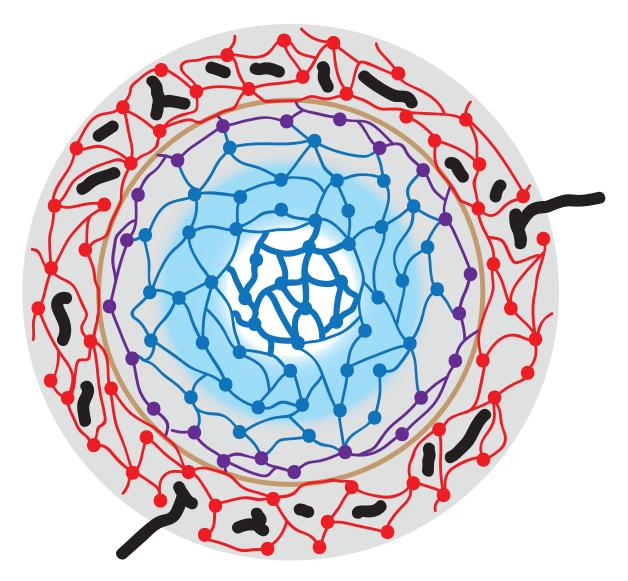
Figure 7.7: Schematic drawing visualising stromal cells, capillaries and the location of CD27$^+$
B-lymphocytes at the surface of a secondary follicle in a human spleen.
SMA$^+$MAdCAM-1$^+$ MRCs are depicted in red colour, SMA$^-$MAdCAM-1$^+$ MRCs in violet
and CD271$^+$ FDCs in blue. The innermost FDCs (thick blue lines) belong to the germinal
centre. Perifollicular capillaries are black and the putative size barrier (formed by the red cells)
is light brown. Gray background colour represents the area with CD27$^+$ memory B-cells and
blue background shows the follicular mantle zone. Note that the gray area overlaps the outermost
blue FDCs and thus the mantle zone. There is no border between the mantle zone and the area
of CD27$^+$ B-cells. The follicle-associated T-cell area is not depicted.

The most strongly stained CD27$^+$ B-cells at the follicular surface are located in a ring-like (2D)
or shell-like (3D) region containing MAdCAM-1$^+$ fibroblasts, but totally lacking any expression of
CD271 in paraffin and cryosections (Fig. 7.7) [Steiniger et al., 2014b, Steiniger, 2015]. This region also
harbours CD4$^+$ T-cells. In some spleens it is populated by a small number of Ki-67$^+$ cells containing
intracellular IgM and coexpressing Ki-67, i.e., by plasmablasts, which continue at the surface of the

PALS. These cells are supposed to be strongly CD27$^+$ and may be partially responsible for the variable expression strength of CD27 in the inner MZ. It is likely that the CD271$^-$ region represents a migration compartment, which is passed by cells moving either in parallel (plasmablasts and CD4$^+$ T-cells) or at a rectangular direction (CD27$^+$ B-cells) to the surface of the follicle. With exception of the innermost follicular CD27$^+$ B-cells, MAdCAM-1$^+$ fibroblasts and CD27$^+$ B-cells are always closely associated at the follicular surface. It cannot be entirely excluded, that the expression of CD27 is upregulated in memory B-cells in the vicinity of follicles. This might explain the reduced staining intensity in the most superficial CD27$^+$ B-cells. On the other hand, the cells may somehow be sorted for strength of CD27 expression by MRCs and FDCs.

Our present and past investigations [Steiniger et al., 2014b, Steiniger, 2015] clearly show, that the human splenic marginal zone can only be defined by resident stromal cells, the MRCs. The most comprehensive phenotypical trait of human MRCs is the expression of MAdCAM-1, while SMA reveals a somewhat smaller MRC population [Steiniger et al., 2014b]. Unfortunately, up to now human MAdCAM-1 cannot be reliably detected in paraffin sections, if cryostaining is used as a standard. Perifollicular MAdCAM-1$^+$ fibroblasts have been found in each human spleen investigated so far, but these cells were lacking in tonsils (Fig. D.1). This is astonishing, because MAdCAM-1 is regarded as a typical adhesion molecule of mucosa-associated lymphatic tissues. Theoretically, it should be abundant in tonsils and absent from spleens, but in reality the reverse is true. Further antigens of MRCs, such as Notch ligands [Fasnacht et al., 2014], remain to be reliably detected by immunohistology in humans.

In comparison to rats and mice, the surface of splenic follicles has a more complicated structure in humans, because an easily detectable barrier to the open splenic circulation is lacking. Blood is primarily transported to the follicles by a superficial network of fine capillaries located inside a network of large perifollicular/red pulp venous sinuses. This capillary network is part of the open splenic circulation and is primarily connected to red pulp capillaries. It may form an equivalent of the rodent marginal sinus, but its microanatomical location is different. If a size barrier for entry of blood-borne molecules into the white pulp exists in humans, it is most probably located internal to the capillary network and associated with MAdCAM-1$^+$ and/or SMA$^+$ MRCs. CD27$^+$ memory B-lymphocytes occur in two different stromal compartments at the surface of follicles, defined by MRCs or by FDCs. Thus, only the perifollicular spleen-specific stromal cells, but not B-lymphocytes, define the human splenic MZ. A refined analysis of microvessels and stromal cells at the surface of mouse and rat splenic follicles is still lacking. Such a study is likely to reveal more similarities to humans than so far recognised.

7.3 Materials and methods

7.3.1 Specimens

The spleen specimen used for 3D reconstructions came from a 22-year-old male patient. It was fixed in 3.7% formaldehyde/tap water for 24 h, washed in tap water for 12 h and processed to paraffin via graded solutions of isopropanol and xylol. Double staining in single sections involved this specimen and additionally a specimen of a 17-year-old male patient (Fig. 7.4f, 7.6b). Both specimens were representative of a large number of adult spleens with small secondary follicles investigated during recent years.

7.3.2   Single staining for CD34 in serial sections

The sections were de-paraffinised and stained for CD34 applying mAb QBend10 (Dianova, Hamburg, Germany No. DLN-09135) diluted 1 : 1500 in PBS/1% BSA/0.1% NaN$_3$ overnight after treating the sections with glucose oxidase (GO, Sigma-Aldrich Chemie Gmbh, Munich, Germany, No. G-6641) at 100 U/ml in PBS, pH 7.2, containing 20 mM beta-D-glucose and 2 mM NaN$_3$) for 1 hr at 37 °C to remove endogenous peroxidase. Antigen retrieval was not performed. Binding of the mAb was visualised using the Vectastain Elite Kit for peroxidase (VEP, Vector Labs, Burlingame, USA, No. PK-6100) for mouse IgG (Vector Labs No. BA-9200) and diaminobenzidine (DAB) as chromogen. The sections were coverslipped in Eukitt.

7.3.3   Double-staining for CD34 and CD271 in serial sections

After scanning, the coverslips were removed in xylol, the sections were autoclaved in citrate buffer pH 6.0 and incubated with mAb EP1039Y (GeneTex No. GTX61425, via Biozol, Eching, Germany) overnight to reveal CD271. EP1039Y was diluted 1 : 150 and revealed with VEP for rabbit IgG (Vector Labs, No. BA-1000). Bound antibody was revealed in blue using High Def Blue Peroxidase (HDBP, Enzo Life Sciences, Lörrach, Germany). The sections were coverslipped in polyvinyl alcohol (Mowiol).

7.3.4   Double staining for additional antigens in single sections

Unless indicated otherwise, immunohistology was performed in paraffin sections fixed in 3,7% formaldehyde for 24 h. With exception of double staining for CD34/SMA, where antigens were retrieved after visualisation of CD34, antigen retrieval was accomplished by autoclaving the deparaffinised sections in citrate buffer at pH 6.0 before the first antibody incubation. Endogenous peroxidase was always inactivated with glucose oxidase. Double staining for CD3 and CD27 was performed as described in Steiniger et al. [2005]. Briefly, CD3 was revealed first using an alkaline phosphatase (AP) system and NBT/BCIP as blue/black chromogen, followed by staining for CD27 in brown using DAB. Double staining for CD271 and MAdCAM-1 had to be done in cryosections and has been described in Steiniger et al. [2014b]. Staining was first performed for CD271 in brown using DAB as chromogen followed by revealing MAdCAM-1 by an AP system and the Fast Blue chromogen. Double staining for CD34 and CD141 was performed according to Steiniger et al. [2011a] by first detecting CD34 in blue and subsequently CD141 in red using an AP system.

CD4 was detected by mAb 4B12 (DAKO, Hamburg, Germany, No. M7310) diluted 1 : 200 using VEP and DAB, followed by CD271 using mAb EP1039Y at 1 : 100 and the UltraVision System (LabVision, Fremont, USA via Thermo Fisher Scientific, Schwerte, Germany, No. TL-060-AL) for AP. AP was revealed by High Def Blue for AP (Enzo Life Sciences, Lörach, Germany). For visualisation of CD34 and SMA, mAb QBend10 (Dianova, Hamburg, Germany, No. DLN-09135) was diluted 1 : 1500 and revealed with VEP (mouse) and DAB, followed by mAb asm-1 (Progen, Heidelberg, Germany No. 61001) at 1 : 200 and the same detection system using HDBP as chromogen. For CD34 and CD27 mAb QBend10 was used at 1 : 1500 with VEP (mouse) and HDBP followed by mAb 137B4 (Quartett, Berlin No. 030410901) at 1 : 20, VEP (mouse), tyramide amplification and DAB. Ki-67 was shown using mAb MIB-5 (kindly donated by J. Gerdes, Borstel, Germany) at 1 : 100 with VEP (mouse) and HDBP, combined with anti-IgM (DAKO, Hamburg, Germany, No. A425) diluted 1 : 2000 for intracellular staining of plasma cells and DAB. All double-stained sections were mounted in Mowiol.

7.3.5 Image acquisition and processing for 3D reconstruction

Image acquisition

The single-stained serial sections were acquired by VMscope GmbH (Berlin, Germany) with a Zeiss Mirax scanning microscope and a ×20 lens. Double-stained ROIs of serial sections were acquired using a Canon 60D camera on a Zeiss Axiophot microscope with a ×10 lens. The same camera was also used for documentation of all double-stained single sections.

General outline of processing

The acquired data were normalised [Reinhard et al., 2001, Khan et al., 2014] and registered [Lobachev et al., 2017a, Chapter 2]. We defined ROIs that were processed and visualised. We show the resized registered sections as sequences (Video 7.1a,b–7.3a,b) and as overlays, i.e., frontal volume renderings (Fig. 7.1b,e,h; Fig. 7.2b,e,h). We used inter-section interpolation [Lobachev et al., 2017b, Chapter 3] to reduce anisotropy. Colour processing reduced the data to a single channel. In case of double staining, the colours were separated and processed individually. We converted the volume data to a surface representation (a mesh) using marching cubes algorithm [Lorensen and Cline, 1987]. The version we used included a simultaneous simplification [Ulrich et al., 2014a, Chapter 4]. The mesh was corrected for minor inconsistencies [Ju, 2004] and smoothed [Taubin, 1995]. We also removed small unconnected components, typically those that were smaller than 2% of the main diagonal.

Our registration [Lobachev et al., 2017a] uses a sparse approach. This approach matches and moves detected image features instead of pixels and allows for processing of the whole sections [Ulrich et al., 2014b, Chapter 1]. Our registration features a non-rigid undistortion module to cope with individual distortions of the sections during cutting. The module moves individual areas of the sections to minimise their distortion while maximising the number of features matched across the whole series. In detail, we overlay a control point grid over the image. The control points are then moved driven by feature correspondence. The position of control points defines a spline-based undistortion applied to the image. We at once compute all undistortions in the series. Our registration is a multi-resolution method. After the first undistortion, the number of control points is doubled in each direction and the control points are moved again towards finer features, which results in a more refined correction of distortion. We typically use four iterations of our non-linear undistortion process during registration.

The inter-section interpolation [Lobachev et al., 2017b, Chapter 3] uses optical flow [Farnebäck, 2003] to reduce anisotropy. Basically, further intermediate images are inserted between two sections. This provides for much smoother reconstructions. Using this interpolation the z resolution was improved from 7 µm to 1 µm.

We developed a special software for registration as well as for interpolation. Mesh generation was done with our own software for single staining data and with *3D Slicer* (version 4.6.0, Fedorov et al. [2012]) for double staining. Volume filtering for double staining was performed with *3D Slicer*. We used *Fiji* (version 1.51n, Schindelin et al. [2012]) for data conversion and for colour deconvolution. Mesh filtering was performed with *MeshLab* (version 1.33, [Cignoni et al., 2008]). Rendering was done with *Cinema 4D* (version R14.0429, MAXON Computer GmbH, Friedrichsdorf, Germany). The results are shown in Fig 7.1c,f,i, in Fig. 7.2c,f,i, and in Videos 7.1c,d–7.3c,d. The models were quality-controlled in VR with our own specially designed software (Video 7.4, Chapter ˜ 5). Videos were encoded with *FFmpeg* (version 3.2.2, https://ffmpeg.org).

Single staining

For single-stained sections, we converted the registered stack to grayscale using Open CV [Bradski and Kaehler, 2008]. The volume data were directly converted to a mesh without further mesh processing. This eliminated theoretically possible artefacts. Thus, some connections in the vasculature might have been lost, but artificial connections could be avoided during processing. For ROI 1–3 the data had $1600 \times 1600 \times 24$ voxels at $0.6\,\mu m/pixel$ in the xy plane and $7\,\mu m/pixel$ in direction of the z axis. ROI 4 had $4000 \times 4000 \times 24$ voxels at $0.3\,\mu m/pixel$ in the xy plane and the same z resolution. The z resolution of all ROIs was increased afterwards to $1\,\mu m/pixel$ by interpolation (Chapter 3). Interpolation and mesh construction operated on grayscale images. The iso-value for ROI 1 and 2 was 119. An iso-value of 110 was used for ROI 3. The iso-value of 119 was initially optimised for best removal of weakly CD34-positive perifollicular sinuses. However, this value led to a lot of artefacts in ROI 3. This problem may be due to slightly variable staining intensities of capillaries in different follicles caused by subtle variations in the quality of fixation. Fixation may be better at the surface of the specimen and slightly reduced deeper inside the tissue.

Double staining

3D reconstruction of double-stained sections required the separation of both staining colours and hence further processing. The brown staining of CD34 was separated as the black (K) channel of the image stack in the CMYK colour space. In our experience, this yielded best and most straightforward results compared to colour deconvolution. However, we had to use colour deconvolution [Ruifrok and Johnston, 2001, Onder et al., 2014] for the blue staining of CD271. It was not possible to obtain a reliable separation of the blue chromogen with simple colour space conversions, as was the case with the brown chromogen. Optical flow interpolation was performed after colour deconvolution. Before colour deconvolution, we removed background in *Fiji* [Schindelin et al., 2012].

Each channel had $2300 \times 2300 \times 24$ voxels at $0.4\,\mu m/pixel$ in the xy plane and $7\,\mu m/section$ in the z plane before interpolation. We interpolated to $1\,\mu m/pixel$ in direction of the z axis, similar to other data sets.

The interpolated volume data were further processed before mesh construction. For the brown stain, we applied a grayscale closing filter with radius 7–7–3 and Gaussian blur with $\sigma = 1$. This ensured smoother surfaces. For the blue stain, we changed the radius of the closing filter to 11–11–4. Otherwise, the processing was the same as for the brown stain.

After mesh construction for both brown and blue stains, we applied the same mesh processing steps as outlined above. An exception was the size of removed components for the blue stain, which was 10% of the main diagonal of the mesh. This action radically removed small unconnected structures in the reconstructions, while allowing us to maintain the sensitivity of a relatively low iso-value. The removal of small components was justified, as our objects of interest, the capillary sheaths, are rather large objects. Apart from capillary sheaths the most striking features of double staining for CD271 were FDCs in the follicles. The capillary sheath cells formed large connected structures in the reconstructions, which—in contrast to FDCs—were unaffected by the removal of small components. We also removed unconnected components smaller than 5% of the main diagonal for brown staining of ROI 1.

We typically used an iso-value of 80 for the construction of meshes from the volume data of both blue and brown stains. For the brown stain in ROI 1 we used a value of 130. The range of values in both cases is 0–255, as we used 8 bit per channel.

Visualisation and quality control The final meshes were rendered using *Cinema 4D* software. Minor processing improved the rendering of cut surfaces. We highlighted microvessels connecting the perifollicular capillary network with red pulp capillaries. Quality control of our reconstructions was performed using a custom-written VR tool (Chapter 5). Hardware used for quality control included a HTC Vive headset and a PC with NVidia GTX 1070 graphics processing unit. Our tool allowed an intuitive inspection of the reconstructed meshes. It was also possible to visualise the image of the original registered section at the correct position in the reconstruction. This allowed us to inspect suspicious and/or interesting areas of the meshes with immediate reference to the original immunohistological staining (Video 7.4).

Ethics

Human specimens were obtained before the year 2000. Acquisition conformed to the regulations of Marburg University Hospital at that time, when a formal ethics vote was not required. Verbal informed consent that a sample of the organ was to be used for basic anatomical research and was to be handed over to the first author of the paper was obtained from the patient by the attending surgeon in the Department of General Surgery of Marburg University Hospital. The first author of the paper did not know the patients nor any data except age, sex and the fact that the patients had been healthy before splenectomy became necessary.

Acknowledgement

Part III

FURTHER APPLICATIONS OF COMPUTER GRAPHICS

Paradox though it may seem—and paradoxes are always dangerous things—it is none the less true that Life imitates art far more than Art imitates life.

O. Wilde, *The Decay of Lying*

OPTIMISING MULTIPLE CAMERA POSITIONS FOR THE DEFLECTOMETRIC MEASUREMENT OF MULTIPLE VARYING TARGETS

Which is based on our paper on camera position optimisation for the SAICA project [Lobachev et al., 2013], written in collaboration with Martin Schmidt and Michael Guthe.

Author Contributions: *Conceptualisation: MG. Data curation: MG, OL, MS. Formal analysis: MG. Funding acquisition: MG. Investigation: OL, MG, MS. Methodology: MG, OL. Project administration: MG. Resources: OL, MG, MS. Software: MG, OL, MS. Supervision: MG. Validation: OL, MG, MS. Visualisation: MG, OL, MS. Writing—original draft: OL, MG, MS. Writing—review & editing: MG, OL, MS. Corresponding author: OL. Presentation: MS.*

We present a device for detection of hail dents in passenger cars. For this purpose we have constructed a new multi-camera deflectometric setup for large specular objects. Deflectometric measurements have strict constraints how cameras can be placed—for instance: angular restrictions and distance limitations. An important trait of our system is the static setup—we use a single setup for camera configuration for all objects to be scanned.

We render the camera images and analyse them for the deflectometric needs to optimise the camera placement w.r.t. multiple parameters. Important ones are the positions of the cameras—reflections of the patterns should be clearly visible. Camera parameters are computed using a global optimisation procedure for which we efficiently generate a good starting configuration. We introduce an empiric quality measure of a particular camera configuration and present both visual and quantitative results for the generated camera placement. This configuration was then used to build the actual device.

(a) An overview of the setup with colour coding to show different projection screens

(b) View through one of the cameras where the borders between screens are clearly visible

Figure 8.1: We optimise the camera positions for optimal inspection of specular surfaces using reflections of projected patterns. These are synthetic images, Figure 8.2 shows the actual device.

AILSTORMS are one of the most problematic events for car insurance companies as numerous claims need to be processed in the shortest possible time. Currently, each car needs to be assessed and documented by a human expert using a camera and a pattern that is reflected by the damaged surface. As this process is very time consuming, insurance companies are highly interested in automating assessment and documentation. The main challenge is reducing the recording time to process all damaged cars shortly after the event. Analysis and assessment may take longer since the claims settlement—i.e., repair—is the limiting factor anyways.

Traditional methods like laser measurement can only reconstruct an object if it consists of a diffuse material [e.g., Levoy et al., 2000]. For the inspection of specular surfaces, these methods are not applicable. One possible solution to this problem *deflectometry*. Cameras record reflections from a specific pattern on the surface. Once the pattern and the positions of the camera and the pattern generator (e.g., a fixed pattern or a display) are known, the surface normals can be calculated by comparing the distorted pattern reflection with the expected one (Figure 8.1).

In this joint project with i-Lumica AG, PHIcom GmbH, OGP Messtechnik GmbH, ALTRAN Group, AXA Auto Competence Center AG and ExactVision Bildverabeitungssysteme GmbH, we developed a specific deflectometric measurement system. The goal was to detect hail damage in cars, shown in Figure 8.2, (d). The whole measurement procedure, including driving the car in and out should take no more than three minutes. Our setup significantly differs from the traditional approach where a human expert uses a single camera and pattern to document the hail damage:

- We consider complete *large* objects: passenger cars.

- A large *variety* of different objects—any car, not a particular brand and model—should be measured with a completely static setup.

- We do not move a single camera and pattern projector above the surface, but use a multi-camera setup with multiple projectors.

Figure 8.2, (a) shows the actual measurement device. It is built using 13 back projection screens, in between cameras can be mounted on aluminium beams.

For objects with a complex form, one camera alone can not measure the complete surface. However, setting up more than one camera can mitigate this problem and lead to an almost complete coverage—we will show this fact practically. Positioning the cameras and maximising the coverage is performed using global optimisation. (The goal is to find not a local maximum, but a global one.) An important issue in our setting is the optimality of camera positioning for a *broad set* of scanned objects. We do not optimise the setup for a single particular target, but for a set of different targets.

In this paper we address the optimisation of the camera placement. We use optimisation metaheuristics to maximise the *coverage*, i.e., the amount of the surface that the cameras can inspect. Car insurance companies require at most 3% tolerance in the number of hail dents. If we reach a sufficient coverage, we can presume that the remaining area has the same density of hail dents or less. This assumption is feasible as the uncovered area is both visually and physically less accessible and thus a lower density of hail dents can be assumed. If we obtain a coverage of 94% and estimate $\sim 3.2\%$ more hail dents than found, we have 3% too many if there are none and 3% too few if the density is the same as on the covered area.

The main contributions of this chapter are:
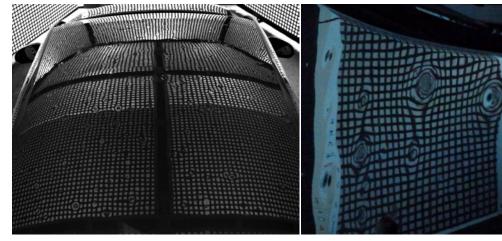
- We present and discuss our novel setup and its optimisation.

- Using ray tracing we obtain qualitative measures for the *coverage* of a given camera configuration.

- We propose an algorithm for the initial camera placement that produces good starting values, essential for the subsequent optimisation.

- Using quantum annealing for further optimisation, we have achieved almost complete coverage ($\geq 94\%$ of the visible surface).

- We verified our method on a broad set of car models, including a car that was not used during optimisation.

- We use a *static* setup, i.e., a single configuration is used for all models.

(a) A picture of the actual device

(b) Testing the projection

(c) Hail damage under deflectometry, an overview

(d) Hail damage on a (white) car where the dents are clearly visible as distortions of the reflected pattern

Figure 8.2: We optimise the camera positions for optimal inspection of specular surfaces using reflections of projected patterns. These are the images of the physically build device and pattern projection in action.

8.1   Related work

There are many applications for which it is necessary to reconstruct the surface of an arbitrary object. Often, only contactless measurement methods are possible [Balzer et al., 2011]. Existing approaches that measure the quality and characteristics of a surface mainly depend on diffuse materials. Highly specular materials and objects with complex form are harder to measure. Two different methods form the so called 'traditional techniques'. These include, but are not limited to, stereoscopic vision and triangulation.

Stereoscopic vision uses two images that are recorded from different positions to create a 3D representation of a given object [Faugeras and Keriven, 1998]. The two images can be matched by a set of feature points or by dense matching. The distance to the camera can then be estimated for each pixel based on the parallax. This is also used by almost all primates for spatial vision [Qian, 1997]. Stereoscopic vision fails in case of specular surfaces because the points can not be matched in the exact surface plane, but only in the mirror plane, which lies behind the actual surface.

The second approach is the triangulation of incident light on a surface [Levoy et al., 2000]. A laser scanline is projected onto the surface and recorded by a camera above the laser. Then the distance to the surface can be calculated from the vertical displacement using triangulation. This process depends on diffuse or partially diffuse surfaces and does not work for highly specular objects. If the diffuse reflection is too dim, the surface needs to be coated, which is not possible for all types of objects. These techniques can be used as a fallback for surfaces that are not specular enough and therefore not suitable for deflectometry.

8.1.1   Deflectometry

The deflectometric inspection of specular surfaces [Sanderson et al., 1988, Ihrke et al., 2010] simulates the way how humans recognise the shape of such surfaces: we look at the distortion of the reflected patterns and infer the shape of the surface from the reflection and the original form of the pattern. A stereo setup makes it possible to infer the point cloud of the surface from a set of specular reflections [Longuet-Higgins, 1981, Balzer et al., 2011]. The reflected patterns are projected on a back light screen using a pattern generator. A camera records the reflected image of the pattern on the surface. Because the pattern has a known structure and can additionally be modified when taking several consecutive images, it clearly shows the distortions induced by the shape of the surface.

Although deflectometry has been known for 25 years, practical methods became feasible only recently [Denkena et al., 2002, Knauer et al., 2004, Kammel and León, 2008, Hong et al., 2009]. Our method differs from these in several aspects, although others also compare with a reference [Kammel and León, 2008, Balzer et al., 2011]:

- *Multiple cameras.* Most approaches use a single camera and move it along the object. We use many fixed cameras.

- *Static setup.* Our setup is not repeatedly optimised for each particular model, but the camera positions and orientations remain fixed in productive use.

- *Multiple objects.* Our system scans multiple different cars without modifications of the setup in-between.

- *Optimisation.* We optimise the camera positions, which is not considered in most papers on deflectometry. We discuss related work for optimal camera placement below.

- *Stereo.* We could also use deflectometric stereo, i.e., the reconstruction of 3D image with multiple sensors [Schultz, 1994, Knauer et al., 2004, Balzer and Werling, 2010, Balzer et al., 2011]. In an interesting alternative approach, Zheng and Murata [2000] rotate their objects in order to spare the number of cameras and projection screens.

Directly related to our work are the papers by Balzer et al. [2011] and Hong et al. [2009]. The first uses a deflectometric setup to find a dent in a part of a passenger car. In contrast to our work, a single sensor and a single projection screen are moved on a robot arm. The second utilise two cameras and five screens for deflectometry of solder joints. They were not concerned with the optimisation of camera positions; further we use much more cameras.

8.1.2 Optimisation

For deflectometric measurements, the camera must be placed in such a way that the whole surface of the object is visible and entirely covered by the reflected pattern. For objects with increased complexity, such as car bodies, this is not feasible using a single camera and pattern only. Multiple cameras must record the surface. The necessary condition is that every point on the surface can be seen by at least one camera. This 'art gallery problem' has received profound research attention. While this is necessary for any measurement, it is not sufficient. The cameras must view the surfaces in such an angle that every point on the surface is also covered by a reflection of the pattern [Fleishman et al., 2000, Erdem and Sclaroff, 2006, Murray and Kim, 2007].

In mechanical measurement apparatuses, the placement of cameras can not be arbitrary. Certain conditions must be met, such as traverses and beams that can be fitted with cameras and there must be a minimum distance between two cameras. As the installation of cameras is tedious and time-consuming, a preceding optimisation of camera positioning is crucial to build the device. Camera position optimisation [Halper and Olivier, 2000] was most often researched for a single camera. The random walk method [Spitzer, 2001] is often used for camera positioning, while some other approaches, like genetic algorithms [Olague and Mohr, 2002], are viable.

Random walk.   This optimisation method [Chandrasekhar, 1943, Spitzer, 2001], abbreviated here with RW, is a typical approach to camera placement [Zhao et al., 2009]. It is often used in computer vision [Bovet and Benhamou, 1988, Meila and Shi, 2000, Grady, 2006]. The camera positions and target directions are 'jittered' a bit with random values, then the new position is evaluated. If it is better than the previous one it is kept and otherwise discarded. This method often gets stuck in a local optimum if the step size is too small.

Adaptive simulated annealing.   The adaptive simulated annealing method, in short ASA, is a global optimisation metaheuristic [Kirkpatrick et al., 1983]; it is adaptive because the step size is adjusted to the current state of the problem. ASA has an important property of being a *global* optimisation method. 'Usual' optimisation algorithms are local, i.e., as soon as the local optimum is found, no better solution is reached. In contrast, the ASA can 'move away' from the local optimum in a search for a global optimum.

Quantum annealing.   This method [Apolloni et al., 1989, Finnila et al., 1994] (abbreviated QA) slightly differs from ASA. To achieve global maximum, ASA allows certain 'setbacks': the result might be worse temporarily. In contrast, QA 'tunnels' through regions between local optima, which also allows QA to find a global optimum. More formally, in QA the tunnelling distance for the

current step determines possible candidates for the next state; in ASA this depends on the notion of temperature in the annealing analogy. QA can be seen as a quantum Monte-Carlo method. QA is also similar to RW, the key difference lies in how the 'jitter' radius is defined, depending on the step count. In RW this radius is constant. In QA the radius is chosen based on a normal distribution. The variance of this distribution is steadily decreasing, but very large, 'tunnelling' jumps are still possible, although less and less probable with increasing number of steps.

While both ASA and QA could be used in our case, the latter has several advantages. The possible camera positions do not form a connected manifold and thus jumps are inevitable. While this is natural in QA, it needs to be specially handled in ASA. In addition, QA has been evidenced to need relatively few iterations compared to ASA in order to find a good solution [Martoňák et al., 2004]; in our case iterations are quite costly for setups with many cameras and objects.

8.2   Measurement system

Our setup includes 13 projection screens placed in a 'greenhouse' structure as shown in Figures 8.1, 8.2. Note the spaces between the screens. First of all, this is a construction-induced limitation; however, we also use these spaces to place cameras without obstructing the projection screens. The placement of the cameras on the supporting beams induces some limitations to the camera placement, as discussed below.

Each camera records several images using different patterns on the projection screens. Based on that data, the hail dents are detected as shown in Figure 8.3. To register the dents on the car and remove duplicate ones, we need to determine the exact position of the car. For this purpose, we mount four wide-angle cameras: to the left, to the right, behind and above the car. Then we use background subtraction [Piccardi, 2004] to detect the car. We compare the captured images to the views of 3D models, minimising the difference between the covered pixels. This gives us the position and actual reference model of the car. Note that the reference car models are segmented, as we need to assign the hail damage to construction parts. Projecting the images back onto the reference model gives us the exact positions of the hail dents. Subsequently, duplicates are removed. The final list per part is the result of the deflectometric measurement.

8.2.1   Constraints

Our approach considers a number of constraints regarding the installation and then optimises the complete camera set. All parameters are subject to a global optimisation procedure that will generate the parameter values with the highest coverage possible with respect to the specified constraints. We change both position and orientation of the camera, as well as the focal length of the lens. Hence, each camera theoretically possesses 7 degrees of freedom: 6 DOFs for the camera position and orientation and one for the focal length, which we choose from a fixed set. The position is limited by the construction constraints: the cameras are placed on beams between the projection screens (Figure 8.4) and the orientation is also limited to 2 DOFs. So in total there are only $3 + 1$ DOFs per camera. A further constraint is the *distance* between cameras. Due to construction limitations two cameras need to be almost 10 cm apart. We set this distance to exactly 10 cm in our optimisation.

Summarising, some constraints are implied by the physical properties of the cameras: focal distance, depth of field, resolution of the sensor. Other constraints emerge from the physical construction: we place cameras not everywhere, but only where it is mechanically feasible.

We define several criteria that need to be fulfilled for each point on the surface of the car body in order to be classified as covered. These include basic visibility (visible by at least one camera),

Figure 8.3: A camera image showing detected hail dents. In addition, the dents are classified into those that can be repaired without much effort (cyan circles) and those that cannot (red circles).

Figure 8.4: Camera mounted between two projections screens. Note that while the view direction can be modified, the up vector of the camera is fixed.

extended visibility (visible by at least two cameras if 3D reconstruction is desired), visible reflection of pattern display, and reflection adequate for deflectometry. We will discuss these criteria in detail. Notably, surface parts with interreflections (e.g., concave parts of the car surface) are marked as not covered.

8.2.2 Visibility

The most basic requirement for coverage is the visibility of each point. This is easily determined: If a point can be seen by at least one camera, it is visible. To compute the visibility of a certain point, several steps are performed. The images resulting from all of the cameras in the scene are processed. Covered pixels are transformed into covered patches on the surface of the object. Any area not covered by a camera is not visible.
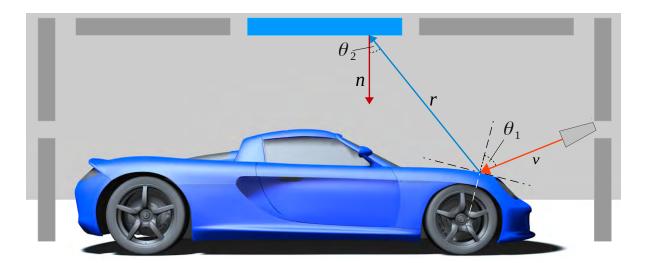
Figure 8.5: Fresnel reflection angles on the surface and display.

8.2.3   Reflections of pattern display

For each point, the eye vector describes the direction from which the camera sees this point. A reflection can occur when the reflection vector **r** belonging to the eye vector **v** hits one of the pattern displays. Figure 8.5 illustrates this concept.

One problem during the optimisation with respect to the visibility of reflected patterns are the gaps between the display segments. They result from technical and mechanical constraints during construction of the system. Solving this problem is easy, as the affected points have to be covered by a different camera.

The reflection quality is the most interesting and difficult requirement. An adequate reflection is one that fulfils the deflectometric requirements [Kammel and León, 2008]. For our system, we can reduce these requirements to the following:

- Possible reflection

- Actually visible reflection

- Reflected image must not hide details.

The possible reflection is already guaranteed by the preceding paragraph. It is necessary for any further calculations. Whether the reflection is actually visible depends on several factors. The limitations depend on the characteristics of the surface material. In our case, we observe reflections on car paint. The viewing angle influences the strength of reflections on the surface. This is known as the Fresnel effect. The angle between surface normal and viewing direction θ_1 differs between points on the surface. Any points that are viewed from an angle lower than the threshold cannot be marked as covered. This is the first part of the conditions that lead to sufficient reflection quality. The glass panels of the projection screens induce another limitation of the viewing angles. The angle θ_2 between the normal of the screen surface and the vector **r** must not be too large, otherwise the Fresnel reflection creates an overlay on the pattern. In our experiments, we found that up to 60° the reflection does not cause measurement artefacts.

Furthermore, we want to maintain a good quality of the reflected image. Strictly speaking, there are two kinds of lateral shrinking of the pattern: if the surface curvature is too large and if the

projected pattern is viewed under a too sharp angle. Both are important for our application, but in the following we discuss the one that is actually controllable with camera placement. Now, shrinking the image in one direction leads to severe loss of resolution. We do not want the shrinking in lateral direction to exceed $1:2$. This leads to the second limitation: the viewing angle θ_1 between surface normal and view vector $\mathbf{v}$ must be less than $\theta_{\max}$, the angle where the lateral shrinking is 0.5. Any angles larger than $\theta_{\max} = 60°$ are still visible, but lead to distortions due to the extreme shrinking of the reflected pattern. Since we want to avoid such errors, $\theta_{\max}$ is the upper bound for the acceptable reflection angle.

The final restriction is the minimum resolution on the surface. The projected pixel should be smaller than a given threshold in both directions. Note that the projected size is not only distance and focus dependent, but also changes with θ_1.

8.3 Optimisation

The goal of our simulation approach is to optimise the camera placement in a scene of multiple cameras and an object. The more parts of the object's surface are visible, the better the placement is. Therefore the coverage, a measure for covered surface, is an important objective function. We aim to optimise the input parameters (position and view direction of every camera in the setup) leading to the best possible coverage. The simulation data was used to determine the number of required cameras and their optimal position and orientation for the construction of the actual physical device, see Figure 8.2.

The *coverage* is the value we are maximising, i.e., the quality measure of the camera placement. We define it as the relative car body surface area, which is seen by a camera and covered with a suitable pattern reflection. We combine the individual coverages for all cameras to obtain the total covered area per model. Then we combine the relative coverages of all models using

$$x_{\max} = \arg\max_{x \in D} \left(\min_{1 \leq i \leq N} \left\{ \frac{A_{\text{cov},i}(x)}{A_{\text{total},i}} \right\} + \frac{\lambda}{N} \sum_{i=1}^{N} \frac{A_{\text{cov},i}(x)}{A_{\text{total},i}} \right),$$

where D is the search space, as discussed in Section 8.2.1, and i iterates through the N car models. We use the penalty weight λ to increase the convergence rate. In this manner not only the worst case is optimised, but also the others. We use quantum annealing to compute $x_{\max}$ and thus maximise the coverage.

8.3.1 Computing the coverage

We use ray-tracing on commodity graphics hardware to simulate the recorded images for a set of cameras and then compute the covered area $A_{\text{cov},i}$ of each car. During this process, we estimate the total area of the car body surface in which the pattern reflection is clearly visible and which meets all constraints discussed in Section 8.2.1. Basically, we add up all pixels of the specular car surface for that the above constraints are satisfied. This value is then divided by the total externally visible area that is computed before starting the optimisation.

Each car model is prepared in order to efficiently measure the covered area. The model is first split into reflective parts that need to be measured—i.e., the car body—and the remaining parts that may occlude the reflective ones or the screens. In addition to calculating per face and per vertex normals for the reflections, a texture map is created by unwrapping the car body using simple box mapping [Bier and Sloan, 1986]. The Figure 8.6, (a), shows the generated texture atlas of a car body. We use

the texture map to determine if a pixel is externally visible: we simply trace several rays outward and check them for occlusion.

(a) unwrapped texture atlas   (b) the model

(c) example atlas mapped onto a model
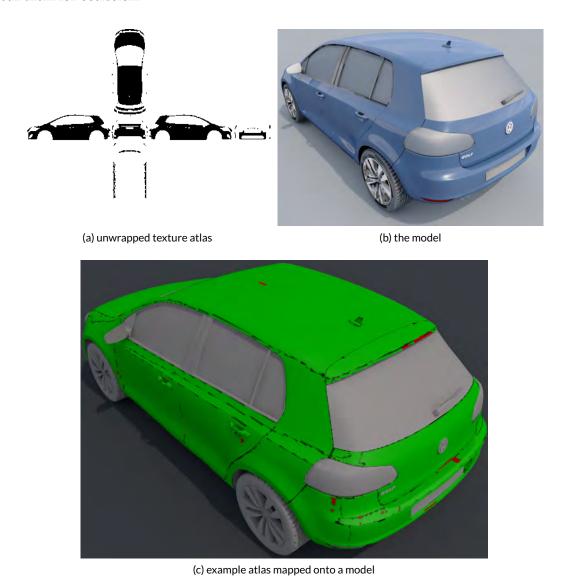
Figure 8.6: Texture atlas of a sample car (a), the model used (b), and computed coverage mapped onto the model (c). Almost everything is green, meaning sufficient coverage.

The computation of the specific reflection of each eye vector is necessary in each optimisation step. The form of the object does not change over time and iterations, therefore, the face-normals are computed only once. Apart from computing the normals, we need to unwrap the texture map of the mesh. Unwrapping describes the procedure of mapping the complete surface onto a plane. For complex objects, the surface mapping must be cut in pieces to prevent errors and reduce distortion. During the optimisation process, we create a texture atlas for the coverage [Bier and Sloan, 1986]. The coverage can be seen directly there. The process of unwrapping assigns the matching texture coordinates to each vertex of the object. This gives us the possibility to map a texture onto the object that fits exactly to the texture atlas.

For each optimisation pass the view of every camera is rendered to determine which parts are visible and contain usable reflections. Then each pixel of the camera view is mapped into the texture atlas. As the only information we need for this mapping is the texture coordinate of the pixel, we simply store that in the render buffer. For pixels that do not meet the constraints, we simply store invalid coordinates (e.g., $[-1,-1]$). Figure 8.7 shows the content of this render buffer for one of the cameras.

Figure 8.7: Intermediate buffer storing the texture coordinates of each pixel containing a suitable reflection of a pattern. These images are generated for every camera and mapped into the texture atlas.

8.3.2 Initial placement

All random-walk based methods, including QA, are *local* search methods. Good start values are crucially important for the quality of the result and the runtime. The first phase generates an initial placement of all cameras.

We place a camera, to cover the largest 'hole'. These are detected by low pass filtering the texture atlas using a Gaussian filter of radius 0.1 of the texture size. Then we search for the local maxima and keep the largest 10. To place a camera, we choose randomly one of these and determine position and normal at the closest texel containing a part of the surface. Based on position and normal, we first place the camera 1 m above that point. Then we choose a random position within a distance of 1 m from that point and project the camera to the closest free location on a beam. This process is repeated 100 times for each camera and the placement with the highest coverage is chosen. Note that the filtering and search are performed once for each camera, so the additional overhead is relatively low compared to QA iterations.

Note that the initialisation search phase is not required, as QA would have found good solutions anyway, but greatly improves the speed of the computation.

8.3.3 Output

To visualise the quality of the currently found solution and build the actual device, we produce several outputs. These are a texture depicting the current coverage, a text output of the minimum coverage, and the camera placement itself. Using colour-coding, we can show each point's characteristics.

Covered pixels are green and uncovered ones are red. In addition, pixels that are not visible from above or beside the car are black. The latter is, e.g., the case for pixels covered by the license plate, in joints, or in the bottom-facing parts of the car. Figure 8.6, (c), shows such a visualisation.

The textual output of the simulation is the list of cameras with their position in the construction co-ordinates and their look-at point. We also implemented a script to import these in 3D-modeling/CAD software to visualise the cameras and for the construction of the actual device.

Table 8.1: Left columns: run time for three optimisation methods, right: mean run time for the performance-tuned QA method with initial placement phase, designated 'IP+QA*'. All approaches use 100 cameras, we measure the time for 1000 iterations. Time is in seconds.

100 cameras, 1000 iterations				
Method	RW	SA	QA	IP+QA*
Time, s	8344.34	7851.55	6995.35	3256.38

8.4   Results

We optimised the camera positions for *twelve different* car models. These models were chosen to be as variative as possible to ensure covering maximal shape space. The models were: Audi R8, BMW X5, Chevrolet Lumina, Fiat Grande Punto, Ford Focus ST, Volkswagen Golf 6, Mazda MX5, Mercedes C, Peugeot 107, Porsche Carrera GT, Renault Clio 2, Toyota Auris. The cameras had a 6.9×5.5 mm CCD sensor with a resolution of 1280×1024 pixels. The cameras used were capable of producing images at 100 fps. These parameters were chosen because of the sharp capture time limitations. The focal lengths of the lenses were 9 mm, 12 mm, 16 mm, 25 mm, and 35 mm. As desired resolution on the car body surface, we set a minimum of one pixel per millimetre. Relatively small resolution of the cameras was compensated by their speed; more cameras were added to achieve the desired resolution on the car surface. We found the minimum number of cameras required to achieve the desired coverage of 94% to be 99 for our test set. While the minimum coverage was 94.0006% (Chevrolet Lumina), the peak coverage was 97.0595% (Volkswagen Golf 6). Still remaining red spots (see Figure 8.8) were outliers outside the desired coverage. The high number of cameras stemmed from the desired accuracy of ∼0.5 mm for the detection.

The optimisation was performed on an Intel Core-i7-3770K at 3.5 GHz with 16 GB RAM and NVidia GeForce GTX 680, running Windows 7. For the quantum annealing we used $\lambda = 10^{-3}$. The number of iterations is 10000, disregarding heat-up. In each iteration we either changed the position of a single camera or moved all cameras. In the first case the distance of the camera 'jitter' ranged from 50 cm to 0.1 mm, and in the second case from 5 m to 1 mm. The cool-down factor was $\varepsilon^- \approx 0.99915$, the heat-up factor was $\varepsilon^+ \approx 1.1857$. Each time, we found a better solution, the temperature was increased by that factor. In total, the optimisation took 25.5 hours. The optimisation time is less than the three days required to set up the rest of the device before the cameras could be mounted.

We compared the number of required cameras for all cars against the number required when optimising for each car separately. Figure 8.8 shows the coverage–number of cameras relation. For all 12 cars in a static setup, the number of cameras (99) roughly doubles the number required for each car alone (35–67). Figure 8.9 shows the coverage visualised on some of the car models. Most important here is the top view as hail dents are much more common on horizontal surfaces. The last image (Fig. 8.9, (f), page 163) shows an experiment to test our approach. After optimising the
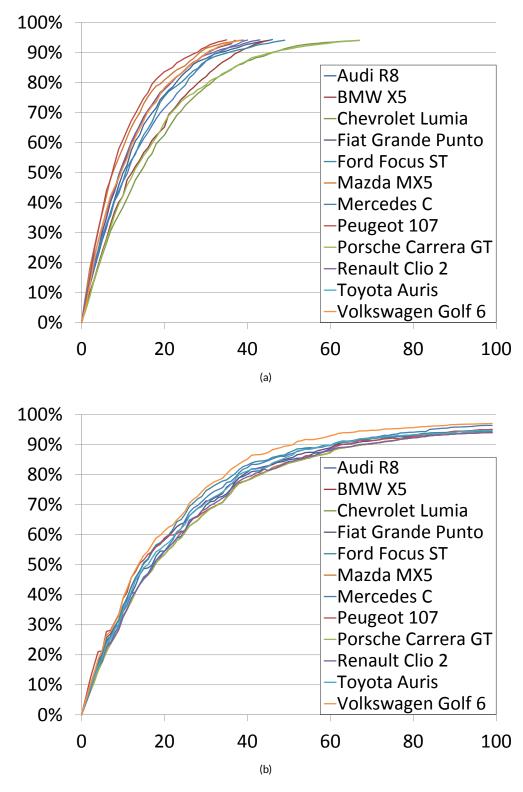
Figure 8.8: Comparison of coverage when optimising for each car independently (a) vs. optimising for all cars together (b). The horizontal axis represents the number of cameras, the vertical axis shows the coverage.

(a) Ford Focus ST

(b) Renault Clio 2

Figure 8.9: Visualisation of the coverage for the models used during the optimisation. Green is covered, red is not covered, black is not externally visible. *(Continued on page 162.)*

cameras for the 12 cars, we used this setup to compute the coverage of the first generation Pontiac Firebird, which was *not* included in the optimisation set. We reach a coverage of 93.87%. Hence, an acceptable coverage is also possible for cars that were not considered during the optimisation of the camera placement. This makes it easy to use the camera setup for cars that are newly introduced to the market, because the shape space changes only marginally in short time. Still, the setup needs to be revised at some point as those chances in shape space can add up and need to be covered. The bigger the shape space is during optimisation, the less revisions are due. This has also a positive effect on custom car alterations, e.g., spoilers. Surface parts occluded by these are less prone to dents.

8.5 Conclusion and future work

We have presented an approach to build a deflectometric measurement system for large specular objects using a multi-camera setup. The objects measurable by our setup are complete cars—a quite large field of work. The device we describe here was actually built and is now operational at AXA, Bern. We use quantum annealing combined with an effective initial camera placement to determine the optimal camera placement. We consistently reached high coverage rates for multiple simulated vehicles using a single configuration—the set of camera positions, orientations and focal lengths, shared for all vehicles. Approximately 48% more cameras are required for all cars than for the most difficult one alone.
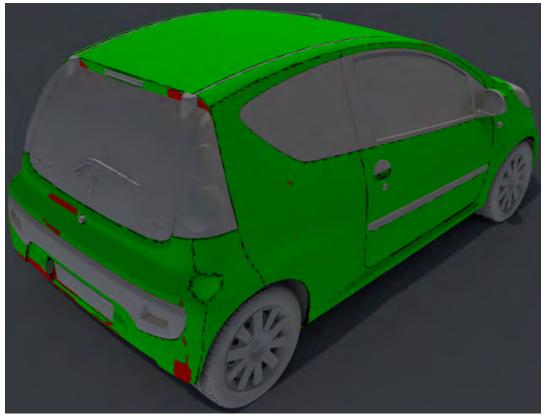
There are some issues to investigate in our future work. Other build forms for the supportive constructions for projection panels and camera mounting could be viable. It could make sense to try some further optimisation methods on the whole setup, not only for the cameras. The cameras also have further advanced parameters one could change during the optimisation, like the tilt for the selective focus. An interesting option is to dynamically group the cameras for 3D reconstruction without reference models. While 3D reconstruction is not necessary in the current system, it could be of interest for future projects.

Acknowledgements
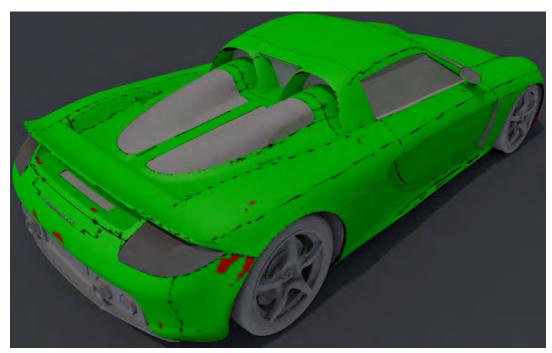
(c) Mercedes-Benz C-class

(d) Peugeot 107

Figure 8.9: (Continued from page 160.) Visualisation of the coverage for the models used during the optimisation. Green is covered, red is not covered, black is not externally visible. *(Continued on the next page.)*

(e) Porsche Carrera GT

(f) Pontiac Firebird

Figure 8.9: (Continued from previous pages.) Visualisation of the coverage. Green is covered, red is not covered, black is not externally visible. The bottom image shows the coverage of a car that was not part of this set. Nevertheless, we still achieve a good coverage.

DIRECT VISUALISATION OF CRYPTOGRAPHIC KEYS FOR ENHANCED SECURITY

In which we present a method for visualisation of small binary data. This chapter follows our publication Lobachev [2017], accepted for publication in The Visual Computer.

Author Contributions: *Complete work was done by OL alone.*

PGP public keys are relatively small binary data. Their hashes are used and also visualised for comparison and validation purposes. We pursue a direct, but previously unused approach. We produce colourful images of public keys and other binary data by generating drawing primitives from binary input. Optionally, we also include the hashes in the visualisation. The visualisation of raw data together with its hash provides a further security benefit. With it we can visually detect hash collisions.

The primary focus of this chapter is a direct visualisation of public keys. We tune the transparency heuristics for better results. Our method visually detects key spoofing on real SHA1 collision data.

```
99 02 0d 04 55 5b 4b 79
01 10 00 a7 8e 75 e2 13
e6 65 86 a4 36 66 8e 76
9f 53 d3 51 fe d5 eb 58
3c ca 41 a9 db db 07 84
c8 ff 06 19 bc d1 be 2f
8f 67 2c 10 2e e4 0c 33
d5 fd fd 1f da 18 d6 75
35 6c 9e b0 aa 28 87 91
79 db f3 34 cc 73 a2 71
08 e9 3a 17 6c e1 e3 c0
...
```
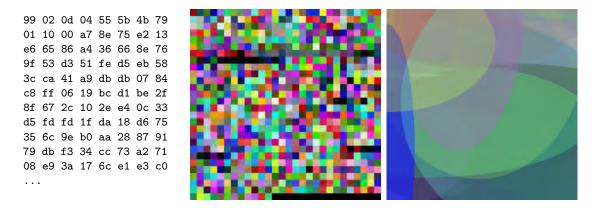

Figure 9.1: Interpreting author's public *PGP* key (first few bytes are shown left) as a sequence of pixel colours is not aesthetically pleasant (middle), while interpreting the same data as a sequence of drawing statements (right) produces nicer images.

```
+---[RSA 2048]----+
|XBoo+.           |
|X o=.            |
|.= .o            |
|+ o.o            |
| E =oo. S        |
|o *.B+..         |
|.+oOo+           |
|.==.+ .          |
|o=+.             |
+----[SHA256]-----+
```
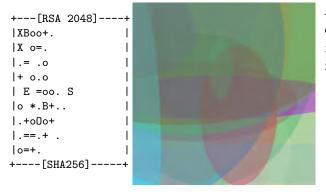

Figure 9.2: Visualising *ssh* keys. Left: an existing hash-based key visualisation in *ssh*, right: direct visualisation of the key with our method.

UBLIC KEY cryptography is an important application in the modern world. Almost all communication channels in 2017 are encrypted using session keys that emerge from public-key-based key exchange mechanisms. *PGP* [Zimmermann, 1995] was one of the front runners of using encryption in email exchange and pioneered the usage of military-grade encryption standards by the masses. Modern *PGP* implementation is *GnuPG* [GPG2].

The typical way of obtaining a *PGP* public key is a key server or an initial message with the key attached. But how can one be sure that the key is correct and originates from intended source? The initial approach was to verify a *fingerprint* of a key [Schneier, 2007]. Fingerprint is a cryptographic hash of a key with some metainformation added before hashing. Ideally, a fingerprint should be obtained from an alternative source (such as a phone call) and verified with the fingerprint of insecurely obtained key. However, fingerprints are quite tedious to transmit and to compare for a human. Efforts to improve this [BSD General Commands Manual, 2017, Perrig and Song, 1999] produce either a small ASCII-art image or a 'random art' of the fingerprint. For these visualisations following holds:

- They are visualisations of a *hash*,
- They are not quite aesthetic,
- They can serve for a fast visual validation, as varying keys (even by few bits) should have very different hashes,
- They are prone to hash collisions if their hash function is compromised.

A very straightforward alternative mostly fails. We might try to interpret a key as a binary sequence of pixel colours. As the cryptographic data is mostly randomly distributed over the value range, the result looks mostly like random noise and is ill-suited for a visual comparison by a human, as Figure 9.1 shows.

For the best presentation of our method we provide a short introduction to public key cryptography and a notation of cryptographic protocols first and only then list the features of our contribution.

9.0.1 Asymmetric cryptography

In contrast to the 'conventional', symmetric cryptography, the asymmetric cryptography, also called 'public key cryptography', utilises *two* keys [Schneier, 2007]. Assume, Alice wants to send an encrypted message to Bob. Alice's private key is known only to Alice, Alice's public key is—as the name says—publicly available, same for Bob. Now, to encrypt a message to Bob, Alice fetches his public key. With it, the message can be encrypted in such a manner that only the owner of a corresponding private key—namely, Bob—can decrypt it. Section 9.0.2 presents such an interaction in a more formal manner.

An attacker, Eve, could intercept Alice's request for Bob's public key, supply her own public key instead, and continue with the man-in-the-middle attack. In the case, Alice's message is encrypted with Eve's public key, Eve can decrypt it, fetch Bob's genuine public key, and encrypt it again. Neither Alice, nor Bob would notice anything in this simple protocol, but their communication is compromised. Section 9.0.2 gives a formal description of the above attack, Section 9.2.5 showcases attacks that can be prevented with our method.

Hash functions reduce an arbitrary long sequence to a string of fixed size. Naturally, there are collisions: two different inputs that result in the same hash value. However, cryptographic hash functions are designed in such a way that:

- Minor changes in the input result in large changes in the hash value;
- It is very hard to find a hash collision.

Previously safe hash functions are regularly compromised with advances in cryptography and with increasing computational power. This already happened to MD5 [Wang and Yu, 2005, Stevens, 2006] and SHA1 [Wang et al., 2005, Stevens et al., 2017], to name relatively recent and resonant hash function compromises.

To validate a public key, *fingerprinting* is used. It is basically a hash value of the public key with some metainformation. Alice could have obtained the fingerprint from an alternative source, such as a phone call to Bob. Now it is possible for Alice to verify that the public key, she has fetched before, indeed belongs to Bob. The problem is: almost nobody compares fingerprints, the process is tedious and too complicated. Further, if the hash function is compromised, the attacker could design her, 'fake' public key in such a way that the fingerprint is still correct. This is called 'key spoofing'.

Theoretically, this is remedied with constant maintenance: if we always switch to a better, yet uncompromised hash function, everything could be fine. Here, we suggest a different approach to solve this problem: we *visualise* the cryptographic data. In a distinction from all other related work, we visualise the public keys directly. (We can add hashes afterwards for even more security.) In this manner we can detected a spoofed key even if the fingerprints correspond. To fool our visualisation, the attacker needs to find a key that *looks* similar to the genuine key—and then the subsequent hash visualisation should also look similar enough. A different attack vector is required for this, this requirement makes typical key spoofing more complicated.

9.0.2 Formal notation for cryptographic protocols

To formally describe the interactions between multiple parties we use a simple notation similar to one by Cervesato et al. [1999]. When Alice (A) sends Bob (B) a message, we write

$$A \rightarrow B : \text{message contents.}$$

An encrypted message is denoted with $\{\text{plain-text}\}_{\text{key}}$.

A typical *PGP* workflow

Alice (A) and Bob (B) want to communicate securely. Bob's public key is available at a public keyserver (S). Alice requests the key, the keyserver supplies Bob's key K_B and its fingerprint $f(K_B)$:

$$A \rightarrow S : \text{request public key}(B)$$
$$S \rightarrow A : K_B, f(K_B).$$

Alice should have obtained the fingerprint $\hat{f}(\hat{K}_B)$ via a different channel, such as calling Bob per phone. If $\hat{f}(\hat{K}_B)$ and $f(K_B)$ are identical, the keyserver supplied a genuine key K_B. This step is, however, tedious and is often (and mistakenly!) omitted. Now Alice can use the key K_B to encrypt her plaintext message m to Bob:

$$A \rightarrow B : \{m\}_{K_B}.$$

Only Bob can decrypt $\{m\}_{K_B}$. The workflow is similar in many other cryptography applications.

A simple attack

Now we also have the attacker Eve (E), who manages to tamper with Alice's internet connection so Alice never reaches the genuine public keyserver S:

$$A \rightarrow E : \text{request public key}(B)$$
$$E \rightarrow A : K_E, f(K_E).$$

Alice receives from Eve a malicious key K_E with a matching fingerprint instead of correct Bob's key K_B. If Alice never bothers to verify the fingerprint with Bob, she now sends a message that only Eve (still tampering with Alice's internet connection) can read. Eve intercepts the message:

$$A \rightarrow X : \{m\}_{K_E}.$$

Alice thinks X is Bob, but X is Eve. Eve can decrypt the message, as she has created the maliciously used public key K_E and knows the corresponding private key. Having the plaintext message m, Eve can even obtain the genuine Bob's public key K_B (in the manner described above) and play the man in the middle. Eve tampers the message to be m' and sends it:

$$Y \rightarrow B : \{m'\}_{K_B}.$$

Bob may think Y is Alice, but Y is in fact Eve. If Alice would verify the fingerprint, this would not happen. Our approach allows for more visual verification of public keys. Of course, more advanced techniques such as public key based signature would have prevented this very simple attack and would request more complicated approaches. However, the core idea is the same.

9.0.3   Contribution

In this chapter we propose a visualisation of short binary sequences (mostly public keys) that:

- Makes key collisions are visually detectable (Sections 9.2.4 and 9.3.1), thus improving security. Basically, the visualisations are *different* for two different keys with the same hash;
- Is aesthetically pleasant (as, e.g., Figs. 9.1 and 9.10 show);
- Is the same for same keys;
- Works for arbitrary binary data. Our method is independent from cryptographic algorithm that utilises the keys we visualise.

9.1   Related work

To our knowledge such visualisations have not been attempted before. The closest approaches known to us are the hash visualisations: the work of Perrig and Song [1999] and a the visualisation used in *ssh* [BSD General Commands Manual, 2017] (called `VisualHostKey`, Fig. 9.2, left). Our approach in same figure, right, visualises the *key* directly. Perrig and Song [1999] visualise only the *hashes* using the so-called 'random art' technique. We deem our method more simple. In our opinion it also produces nicer images. Further, `VisualHostKey` is an ASCII-based visualisation of hashes. We convert input bytes to a (potentially scalable) graphical visualisation. Basically, where `VisualHostKey` has a printable subset of ASCII characters and their position, we have coordinates and colours on the canvas. In a contrast to both Perrig and Song and `VisualHostKey`, we visualise the keys directly. It is possible to combine the key and the hash visualisations in our method. Awni [2017] shows an
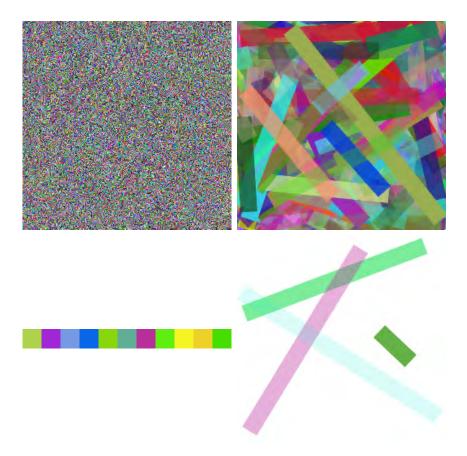
Figure 9.3: Visualising 192 KB of random data directly (top line) and its SHA256 hash (bottom) as RGB pixels (left) and with our "line" method (right). Notice how our method shows the size of input data, while remaining in the same gist of visualisation for data range from 32 to 192K bytes.

interesting approach. There, faces are generated from hashes of the keys to provide better recognition. While one might argue, that people can better differentiate faces than abstract patterns, Awni still utilises hash functions, while we use the public key and add, optionally, its hash.

Hash function attacks [Wang and Yu, 2005, Wang et al., 2005, Liang and Lai, 2007, Stevens et al., 2017] are the reason why hash-only visualisations are vulnerable. In this chapter we use SHA2 hashes for which no collisions are yet known. Generally, our approach can be used with any hash algorithm, but also completely without it. Our method facilitates a detection of a changed key even if the hash value is the same. We show this on a mock example (Fig. 9.8), but also on real SHA1 collision data (Fig. 9.12, data courtesy of Stevens et al. [2017]). Counter-cryptanalysis [Stevens, 2013] also detects possible collisions, as they typically rely on very special circumstances.

Loosely related to our method are graphical passwords [Dhamija and Perrig, 2000, Suo et al., 2005]. Visual cryptography [Naor and Shamir, 1995, Hou, 2003] is a completely different technique, where cryptographic methods are used to convey and securely share visual information. We, however, visualise cryptographic data.

Visualisations in other areas of computer security include [Teoh et al., 1998, Conti et al., 2005, Nataraj et al., 2011]. Generally speaking, both Teoh et al. [1998] and Conti et al. [2005] visualise network traffic on a 2D canvas. Conti et al. [2005] focus on data analysis. Features like typical packet size and other characteristics can be seen in their visualisations. Their goal is the detecting of

reoccurring patterns and overall structure. We focus on generating memorable shapes—at least our user should notice if the visualisation is suddenly different from how it always was, as detailed below. Teoh et al. [1998] showcase an activity visualisation in a network as a 2D plot. In their visualisations, e.g., an ongoing attack can be easily identified. They also mention further techniques and 3D visualisations, such as classical information visualisations of network activities. Our visualisations do not look at the network at large, but focus on a single most critical information during a concrete network session, the identity of the counterpart. We are not concerned with network connections at all in this chapter and visualise the public keys as such.

Nataraj et al. [2011] visualise malware as grayscale images. Such visualisations allow for similarity classifications. Basically, we allow our users to do the same with cryptographic keys. In contrast to our approach, they use a more sophisticated visualisation technique. Further, Nataraj et al. use GIST features [Oliva and Torralba, 2001] in a post-mortem analysis, while in our application non-similarities are striking enough to be noticed with a naked eye. Our application strategy is to visualise cryptographic data during each session. The user would notice a quite different image (Fig. 9.8) if an attack is being executed. A yet another difference is that Nataraj et al. concentrate on malware payloads, somewhat large data, while cryptographic keys and accompanying data typically stay well under 10 KB. Precisely their small size allows us to visualise the data directly.

A long-time companion of cryptography in practical applications is steganography. It is not concerned with how to make a message unreadable to third parties, but with hiding the fact the message is present [Subhedar and Mankar, 2014]. A connected issue is digital watermarking [Cox et al., 2007]. Typical targets are images, although there are methods to hide information in meshes [Cheng and Wang, 2006, 2007].

9.2   Methods

9.2.1   Informal description

We interpret the input bytes as (parts of) vector graphics statements. For example, two bytes state the starting coordinates, two bytes state the end coordinates and four bytes state RGB colour and alpha channel (Fig. 9.3). While it is possible to use an additional byte for the line thickness (Fig. 9.4), we decided against it. The final result of our 'line' visualisation is shown in Figures 9.3 and 9.9. This approach is much more aesthetically pleasant than the direct depiction of values (Figs. 9.1, 9.3). The visualisation also shows how large the input is (Figs. 9.3 and 9.9), as its size in bytes influences the number of line strokes. Key size is a crucial information that is often needed, as it defines the grade of security. Keys with shorter key size (in the same cryptographic algorithm) might be compromised earlier, so they need to be replaced sooner.

Still, we were not yet satisfied with this visualisation and changed a major detail. Instead of painting lines, we paint ellipses. In this way the first four bytes specify the parameters of the ellipse. The next three bytes define the colour and the last, eighth byte defines the transparency. Basically this approach produces our visualisations in Figures 9.1 and 9.10. We provide more visualisations in the supplementary material (doi: 10.5281/zenodo.817656).

9.2.2   Formal description

For our 'line' method, in an 8-byte sequence $a, \ldots, h$ we define

$$from = [a, g] \quad \text{and} \quad to = [b, h],$$

Figure 9.4: Here we also include line thickness into the set of varying parameters. For best results, here the transparency also should be adjusted, however this was not the case with these images.

the line of constant thickness is drawn on a 256×256 canvas between points *from* and *to*. The bytes $c, \ldots, f$ define directly RGBA colour of said line. This method produces images in Figure 9.9—simple but already more visually pleasing than a direct pixel-based visualisation (Fig. 9.3) or hash-based ASCII visualisation (Fig. 9.2). We have also experimented with varying line thickness, but were less satisfied with these results (Fig. 9.4).

For our 'ellipse' method, in an 8-byte sequence $a, \ldots, h$ we define on a 256×256 canvas an ellipse position and size using $a, \ldots, d$, per

$$C = [a, b], \quad D_1 = c, \quad D_2 = d,$$

where C is the centre of the ellipse, D_1 and D_2 are both diameters. The RGB colour of the current ellipse is defined as e, f, g. Using the 'real' transparency byte h as the input value t in the below procedure, we obtain the adjusted transparency value $\hat{t}$ (Section 9.2.3). This results in visualisations shown in Fig. 9.1, bottom right, Fig. 9.2, right, Fig. 9.5, right, Fig. 9.6, right, Figs. 9.10, and 9.11.

Possible future adjustments may include varying the drawing primitive or using segments of ellipses instead of permanent full 360° primitive.

9.2.3   Transparency adjustments

When we directly use transparency values from eighth byte of the block, it produces less nice images. Basically, it happens when the transparency is low and the ellipse size is high, as Figure 9.5, left shows. To adjust this, the final method uses a distance measure obtained from a not quite straightforward heuristics.

A direct approach would be to take the area of the ellipse, or a heuristics for that, and to increase the transparency for large area ellipses. We still want to factor in the last, eighth byte from the input into the transparency value. However, we found that such approach produces less visually appealing images (Fig. 9.6 compares with our actual method). The reason is clear: now absent larger opaque features yield a bleak, not very memorable images (Fig. 9.7). Hence, we follow a different, more involved, heuristics.
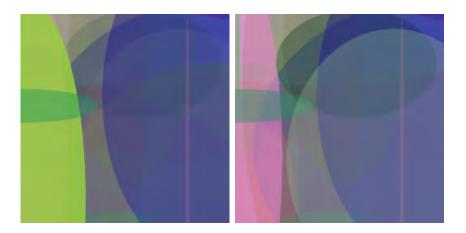
Figure 9.5: The transparency problem. The visualisation left uses direct transparency values from the visualised key. On the right hand side the same key is visualised using adjusted transparency values.
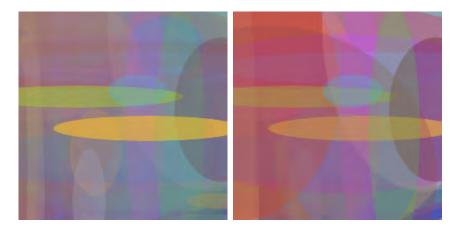
Figure 9.6: Adjusting transparency values. Left: a straightforward way with area heuristics, right: our final approach.

Our transparency heuristics uses the first four input bytes ('coordinates') and the 'real' transparency value t (i.e., value of the eighth byte). It computes the adjusted transparency $\hat{t}$ from D:

$$\hat{t} = \left\lfloor \delta + \frac{s\,t}{D} \right\rfloor \bmod T.$$

In our implementation we used the constants $\delta = 10$, $s = 100$, and $T = 200$. We compute D as $|a - c| + |b - d|$ for first four bytes $a, \ldots, d$. This is in a sense a Manhattan norm. To give an example, for $D = 500$ and $t = 200$ the adjusted transparency value is $\hat{t} = 50$. Figure 9.5, right, Figure 9.6, right, and Figure 9.10 show visualisations with transparency values adjusted. Notice that this means a non-even distribution of transparency across the image, but the particular configuration presented was empirically found to be quite aesthetically pleasant. The more 'linear' adjustment of transparency was visually less viable, as Figs. 9.6 and 9.7 show.

For the visualisation the input is processed in eight byte sequences. The trailing number of bytes that is smaller as eight is discarded from visualisations shown here.
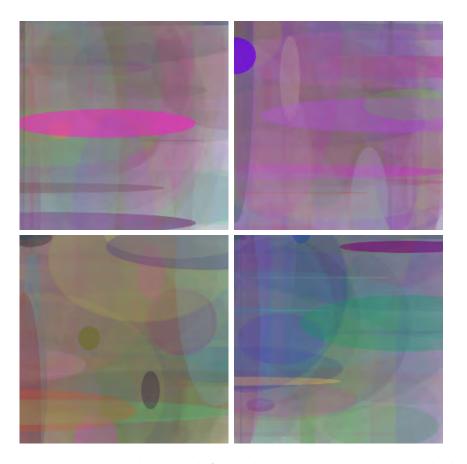
Figure 9.7: Examples when straightforward transparency is aesthetically bad.

9.2.4 Benefit for security

Our method is suited to visualise *any* kind of binary data. Visualising a key alone makes the visualisation prone to subtle key changes that would be immediately detected with a cryptographic hash function, but might go undetected with a key-only visualisation. Further, if a hash function is compromised, and collisions are easy to find, then key validation and visualisation with these hashes is also compromised.

To combat this we visualise *both* the raw key and its SHA256 hash—a 256 bit version of the SHA2 family of hash functions [NIST, 2015]. Figure 9.8 showcases the result. If a key is changed in a manner undetected by our key visualisation, the hash will be different (Fig. 9.8, top). For demonstration purposes we decoupled the key and hash visualisations and used the genuine key with a different hash. If the hash function is compromised and allows to create hash collisions easily, the key would be very different, even if the hash is the same. Fig. 9.8, left, shows this. Even with the same hash the actual keys differ a lot, which can be visually detected.

To circumvent such a combined visualisation, the attacker needs to find hash collisions in a cryptographic hash function at will, while minimally changing the key bits (such that direct key visualisations look similar), and freely changing the last < 8 bits in the key that are currently omitted in our visualisations. The latter is circumvented easily by padding the key to a length divisible by eight before visualising. Finding a collision that only minimally changes the the visualisation output might be quite hard, as Figs. 9.9 and 9.11 showcase quite different images all the time. These figures
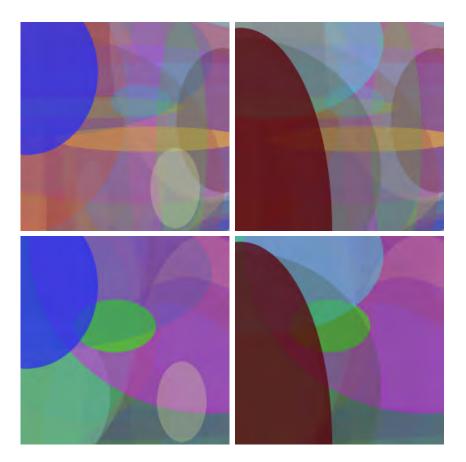
Figure 9.8: Showcasing how key spoofing can be detected with our method, a synthetic image. (Figure 9.12 shows real collision detection in SHA1.) All visualisations use the same method. Top left: original *PGP* key with its genuine SHA256 hash. This is the initial reference for the user. Top right: we assume the key is rigged in a way that is not detectable in our visualisation, but the hash has changed. This is strikingly visible. To produce this image, we have basically changed the hash, but not the key. Bottom left: assume it was possible to find a hash collision, the hash is the same as the genuine hash. But the key itself has drastically changes to allow the collision, which is detected with our visualisation. In other words, in case of key spoofing, a key with the same hash can be found, but the *key* visualisation of the spoofed key differs. Bottom right: for completeness, a fake key with a hash that does not match the reference hash. (In this image we merely visualised a different key with *its* hash.)

show key-only visualisations of public keys from author's key chain.

   We deem the attack a quite challenging task as mapping minor changes in the input to large changes in the output is the essence of a good hash function. The attacker would need to find a hash collision (i.e., a very different 'fake' key with the same hash as genuine key) while maintaining a low bit-wise difference between the genuine key and the fake key.

   Section 9.3.1 shows a real-world example using SHA1 collision data.

9.2.5 Protocol descriptions of prevented attacks

Next, let us regard two scenarios in which our visualisations would be of an immediate utility. We use the formalism from Section 9.0.2.

Weak hashes

Imagine, the hash function used for fingerprinting is compromised, but Alice never took notice of it. Eve tricks Alice into using a wrong key; Eve can intercept Alice's communication to Bob.

A has Bob's key K_B and is aware of its fingerprint $f(K_B)$

E creates a fake key K_E with matching fingerprint.

The key K_E is maliciously generated in such a manner that fingerprint $f(K_E)$ is the same as $f(K_B)$. This is only possible because f is compromised.

E tricks A into using K_E instead of K_B.

A computes the fingerprint of K_X, it is valid.

The key K_X is what Alice believes is K_B, but it is really K_E. The fingerprint did not change. Now Alice sends a message using this key:

$$A \rightarrow X : \{m\}_{K_X}.$$

Eve can intercept the encrypted message and read it, because the key K_X is actually K_E. All key visualisation systems based on the same fingerprint method f would fail, as the hash function at the heart of f is compromised. With our approach, the key is visualised directly and Alice can immediately see the difference (Fig. 9.8). Next, regard a more real-life example.

Online banking example

As of now, too many people are tricked into entering their online banking credentials at a wrong website. Illegitimate acquisition of credentials is called 'phishing'. Thus the question 'How can I allow the user to quickly verify that the public key of this online banking site is genuine and the same as it was on the real online banking site yesterday?' has, unfortunately, enough motivation. We regard this case next.

In a HTTPS session [Rescorla, 2000], a session key $K_{A,B}$ is negotiated. A *certificate c*, a combination of server public key and metainformation, such as domain name, is cryptographically signed during the initial setup of the service. To alleviate the validation procedure, a hierarchy of certificate authorities is used. Basically, a chain of 'who trusts whom' is established, from the server certificate c up to a root certificate that is universally trusted. At a HTTPS session, the certificate of the website is checked whether it is still trusted. Here, we replace the hierarchy of certificate authorities with a stronger version: a direct query to the trusted third party. The problem is that in both cases the validity is established for the link website–entity that obtained the certificate. Bluntly, if a bank has a legitimate certificate for bankname.com, the attacker can obtain a *legitimate* certificate for mybankname.net. Both certificates are valid and legitimate. It is the user's responsibility to discern if the domain name is correct and whether the certificate was issued to the correct legal entity. The users, however, often make mistakes and trust wrong websites.

Alice (A) is going to visit her online banking site (B). Their communication is encrypted by a session key $K_{A,B}$. The correctness of the communication can be verified by checking the certificate c

of the website with a trusted third party (T). The certificate is basically a public key with metadata.
The correct behaviour is:

$$A \rightarrow B : \{\text{access website}\}_{K_{A,B}}$$

$$B \rightarrow A : \{\text{login form}, c, f(c)\}_{K_{A,B}}$$

$$A \rightarrow T : \{\text{is } c \text{ valid?}\}_{K_{A,T}}$$

$$T \rightarrow A : \{\text{our fingerprint for } \hat{c} \text{ is } \hat{f}(\hat{c})\}_{K_{A,T}}$$

$$A \text{ compares } f(c) \text{ and } \hat{f}(\hat{c}) \text{ and they are identical}$$

$$A \rightarrow B : \{\text{credentials for money transfer}\}_{K_{A,B}}.$$

Notice that we do not automatically verify if the certificate c belongs to the bank B, but if it is
a correct, valid certificate. An attacker might also own such a certificate, of course for a different
website. She can lure Alice to it, while pretending to be the bank, as we will see below.

$$A \text{ is lured to a website } E, \text{ thinking it is } B$$

$$A \rightarrow E : \{\text{access website}\}_{K_{A,E}}$$

$$E \rightarrow A : \{\text{login form}, c', f(c')\}_{K_{A,E}}$$

$$A \rightarrow T : \{\text{is } c' \text{ valid?}\}_{K_{A,T}}$$

$$T \rightarrow A : \{\text{our fingerprint for } \hat{c}' \text{ is } \hat{f}(\hat{c}')\}.$$

$$A \text{ compares } f(c') \text{ and } \hat{f}(\hat{c}') \text{ and they are identical.}$$

Notice that all Alice knows is that the certificate is valid. It is her responsibility to check that c' is
actually a certificate of her bank B. If Alice has failed to notice the counterpart E she is talking to is
not the bank:

$$A \rightarrow E : \{\text{credentials for money transfer}\}_{K_{A,E}}.$$

$$A \text{ looses all her money to } E.$$

If Alice would have compared our visualisation of the certificate c with the visualisation of malicious c',
she would have noticed the fact of an attack. The point is to recognise that the malicious website uses
a different certificate c', as valid as it might be. Now, online banking is a quite routine operation.
If the original visualisation is displayed every time Alice uses online banking, we speculate that she
would remember it after few sessions. Our visualisations are quite distinctive to be easily remembered
and discerned by humans (Figs. 9.9–9.11).

The bank might even supply the visualisation on a different channel, so Alice does not need to
remember the genuine certificate image. For example, the visualisation of genuine online banking
certificate could be placed on the banking card. Of course, Alice would also notice every good-willed
change of the certificate, and hence her card would need to be re-issued every time the certificate
changes. This is a distinct issue, however.

9.3   Results and Discussion

We have visualised keys from author's key ring with lines in Figure 9.9. Figure 9.10 shows the same
keys, visualised using ellipses with adjusted transparency. In this case we add the SHA256 [NIST,

2015] hash to the visualised data. We utilised the SHA256 implementation from *Open SSL* [Open SSL]. All images are quite different, it is easy for a human to distinguish them. In most cases the visualisations are even quite artistic.

9.3.1 SHA1 Collision

The weaknesses of SHA1 hash algorithm have been known for more than a decade [Wang et al., 2005]. Recently, the first actual collision became publicly available [Stevens et al., 2017]. As lined out in Section 9.2.4, our visualisation should reveal key spoofing based on hash collisions. Hence, we have tried our approach on the collision example with SHA1 hashes. Notice that everywhere else in this work we use SHA2 hashes. For this hash function no practical collisions are known. Available files[1] with collision are small PDF files. We have produced a series of images from 512 byte chunks at same offsets in the respective files. We additionally compute and visualise in the same image a SHA1 hash of each chunk. (This does not help, as the hashes suffer from the collision and are identical.) Both files are mostly identical, with a smaller sequence with the collision that is different in respective files. These sequences have however the same SHA1 hash [Stevens et al., 2017]. Figure 9.12 shows that in our visualisation said differences are detected despite the collision. This distinguishes our system from *all* other methods used for visualisation of public keys, as all these methods utilise hashes.

9.3.2 Discussion

Should our visualisation provide *similar*, but subtly different images for varying keys, these keys are also similar in their byte-wise values. This is a corollary from using a direct visualisation and requiring same visualisations for same keys. For production use, however, we augment the key visualisation with a subsequent hash visualisation. Now we obtain different visualisations also for similar keys. We are able to visually detect the key spoofing, as the visualisations are *different* for two different keys with the same hash.

The method is specifically geared toward small files, such as public keys. However, it makes also little sense to visualise a large file directly, rather interesting is its hash. Au contraire, in public keys both hash and direct visualisations are viable. Hash visualisation contributes the cryptographic information: keys that differ only in few bits should be considered as radically different. Hash collisions, though, would deem two radically different keys as the same. Our method (in key + hash mode) shows similar (but varying, because of hash) visualisations for similar keys. Our method cannot fall into the pitfall of an eventual hash collision as not *only* hashing is used. In the actual key + hash visualisation, two kinds of attacks can be detected. A 'visual' key spoofing attack—that does not exist yet—can be detected via depiction of varying hash values of an uncompromised hash function. A hash collision (stemming from usage of a compromised hash function) can also be detected via direct key visualisation. We model these two situations in Figure 9.8.

Compared to Perrig and Song [1999] we provide a nice and potentially scalable graphics by *directly* painting the key (and, eventually, additional data). Our method is more direct, this conveys more security. The `VisualHostKey` visualisation from *Open SSH* [BSD General Commands Manual, 2017] is used in production. It visualises hashes as an ASCII-art. Displaying an image and not an ASCII-art should be a non-issue nowadays. However, it has the same weakness as other hash visualisations. All cryptographic security visualisations known to us, visualise hashes of the keys. This makes them prone to hash collisions. Our method does not have this weakness. We enhance security by visualising the keys directly and optionally adding the hash visualisation.

---

[1]SHA1 collision files were obtained from http://shattered.it/.
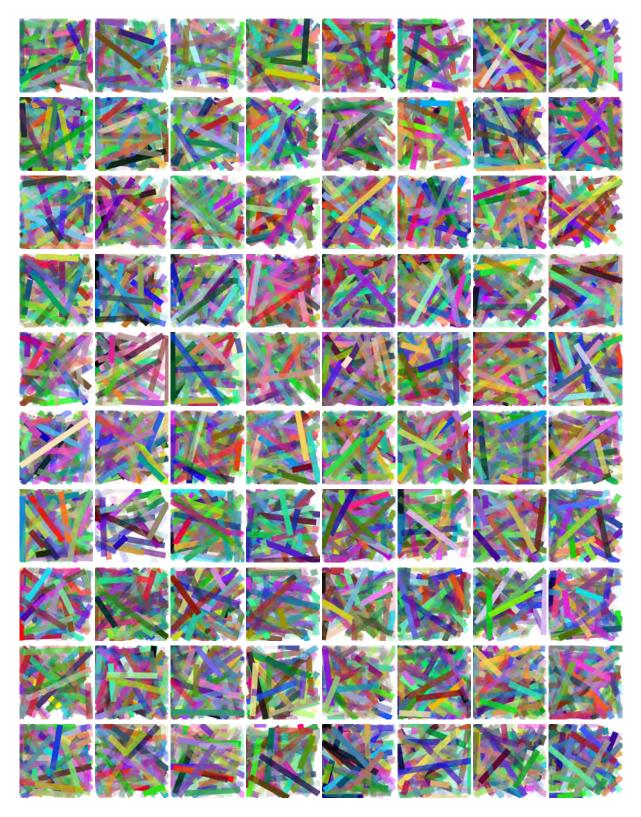
Figure 9.9: Lines visualisation of public keys from author's *PGP* keyring. This is a key-only visualisation.
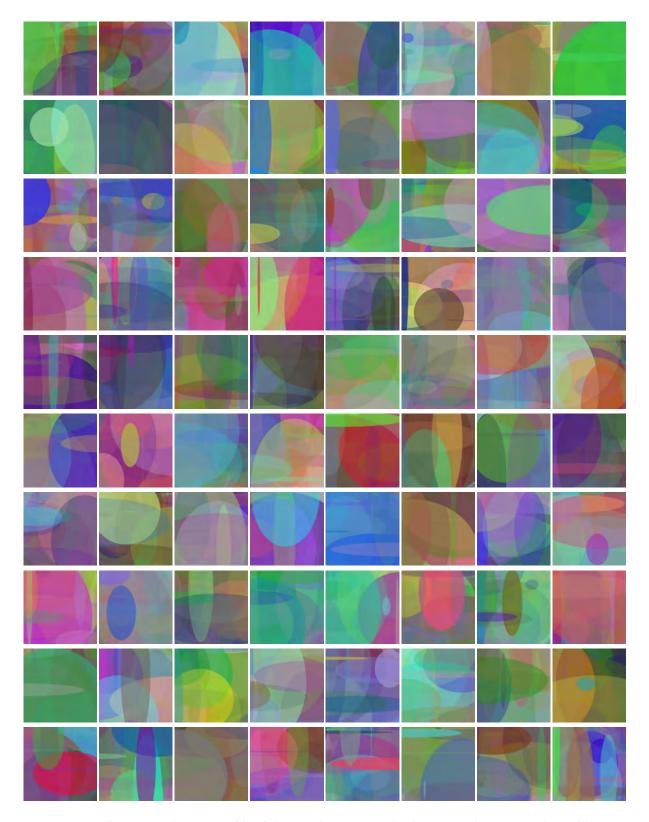
Figure 9.10: Ellipse visualisations of both key and a SHA256 hash in a single image. The public keys originate from author's *PGP* keyring.
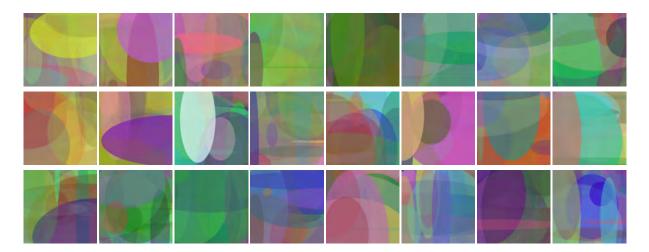
Figure 9.11: Ellipse visualisation, keys only.

Even if the current implementation of our method omits the last few (< 8) bytes from the input, it might be possible to reconstruct or at least approximate parts of the key from the visualisation. Hence, only public keys (which are available to anyone per design) should be visualised with our method.

9.4   Conclusions and future work

We present an approach to visualise cryptographic public keys with their hashes. Unlike previous approaches, our visualisation is direct. Our method produces aesthetically pleasant visualisations (Fig. 9.10). Further, when both the key itself and its hash are visualised, our approach makes key spoofing and similar man-in-the-middle attacks much more complicated for the attacker. We made two binary data blocks with the same SHA1 hash visually distinguishable (Fig. 9.12).

One possible idea for the future work is to spend more bytes for further decisions during image generation. Possible options include using multiple visualisation methods in one image and using gradients instead of solid colours. Line thickness could be also utilised as an input parameter, though for such visualisations (Fig. 9.4), a similar kind of transparency heuristics is needed as we introduced for our ellipse visualisations.

While visualising binary data and its hash, we treat both as equal binary data. It makes sense to visualise data and hash in the same image, but with different methods, e.g., 'ellipse' for data and 'lines' for hash. Because presented method provides additional security, it would make sense to integrate it into cryptographic software and sensible communication protocols, e.g., online banking.

Acknowledgment
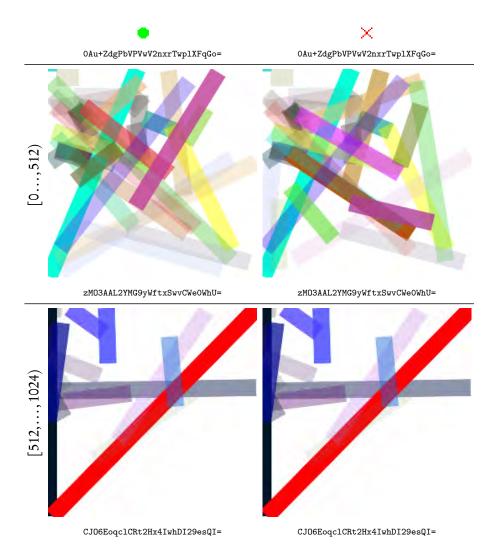
Figure 9.12: Detecting differences in real SHA1 collision data. We process 512 bytes at once as a single block, visualised together with SHA1 hash for the block. We also show the SHA1 hash for each block below its visualisation. Left and right shows two different files, as top line shows. Their SHA1 hashes (shown in Base64) are, however, identical. The files are mostly identical, as visualised further 512 B block in the bottom row shows. The block $[512,\ldots,1024)$ features a lot of zeros and close to zero values, hence it looks more sparse in our visualisation. The blocks in the middle row $[0,\ldots,512)$ are noticeably different in the visualisation; this is clearly visible. However, their SHA1 hashes are the same. Thus, key spoofing based on hash collisions can be detected with our method.

Supplementary material

Supplementary material showing more visualisations is available in the Zenodo repository under doi: 10.5281/zenodo.817656.

Part IV

CONCLUSIONS

There is no real ending. It's just the place where you stop the story.

$$\overline{\hspace{4cm}}$$

F. Herbert

CONCLUSIONS

E CONCLUDE this thesis with a few closing remarks, a summary, and an outlook. Our methodology pipeline (Part I) appears quite finished. Indeed, it has been used—and will also be used in the future—for production-grade 3D reconstructions in medical research. We showcase the currently obtained results in Part II. (Part III shows our tangential works in computer graphics.) However, our research area is by far not closed. On one hand, 3D reconstructions of further specimens and from further organs are viable and partially already in the works. The most prominent example are follicles in tonsils. On the other hand, there is always room for the improvement of methodology. We elaborate on possible future directions of our work in Section 10.2. Next, we summarise our findings.

10.1 Summary

We developed a chain of tools for 3D reconstruction from medical 2D data, originating from invasive probes, the 'serial sections'. These sections were acquired as very large images (up to $30k \times 30k$ pixel), obtained with special optical microscopes. The images did not spatially match each other, due to inevitable position changes during the biological production phase. First, we presented a method to register the full sections. We used a computer vision approach, instead of image-based methods, which are usual in this field. We developed an innovative feature-based registration method. It was *sparse*, since between the feature detection and application of transformation we worked only with features and not with underlying images.

However, in most cases the medical researchers were interested in smaller special areas of a series, 'regions of interest' (ROIs). To better register these areas we developed a new approach. To register ROIs down to the size of capillaries, we applied our multi-resolution, sparse, non-rigid registration. In a hierarchy of feature sizes, we iteratively distorted the images based on correspondence of larger features first, then on smaller and smaller ones. A thorough quality evaluation of our registration method compared it with alternatives. Our method performed very well w.r.t. seven quality measures.

After the registration the volume data suffered from strong anisotropy. The reason was clear: while the resolution of optical scanning microscopes was in the range of 0.3–0.5 µm/pixel, the typical thickness of a single section in our processing pipeline was 7 µm. Thus, we developed an interpolation method based on optical flow. With it we reduced the anisotropy from $14:1$ to $2:1$.

Next in our pipeline we applied some domain-specific volume filtering. After that, the mesh could be constructed. The meshes were also be filtered for better presentation. We developed a method to interleave mesh construction and simplification on GPU. Having obtained the meshes, we could finish with rendering. But were our reconstructions correct? The question of the correctness of the reconstructions arose quite early, during the development of methods for the reconstructions from human bone marrow. We developed an out-of-core volume renderer, that could also overlay the reconstruction over the original data.

The emergence of inexpensive virtual reality (VR) head-mounted displays and tracked controllers was a game changer in multiple areas of computer graphics. In a novel development, such hardware allowed domain experts to inspect our 3D reconstructions intuitively. Merely displaying the meshes
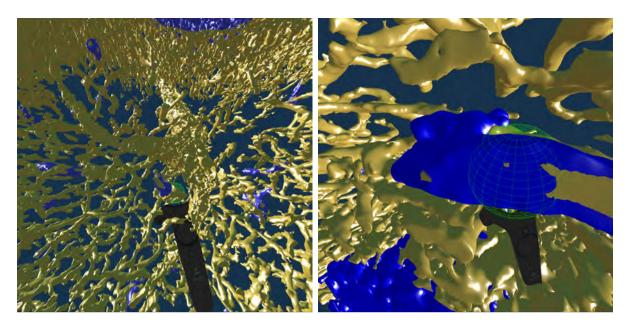
Figure 10.1: A VR-based inspection and annotation of spleen specimens stained against CD34 and SMA (golden mesh) and against CD271 (blue mesh).

was the most basic feature, we did more. Domain experts were used to working with the original section images. So, the question 'What is this structure?' was better solved with original (but registered) data. However, the question 'What is the form of this structure?' could only be answered by inspections of our 3D reconstruction.

Annotations were crucial for communication (Fig. 10.1). Our innovative VR-based visualisation allowed the user to mark structures of interest. These markings could, of course, be displayed again, communicated, and shown outside of the VR in mesh space. Our VR-based tool accelerated the quality control.

We applied our 3D reconstruction pipeline to medical research data. With our processing pipeline, we reconstructed meshes from serial sections of human bone marrow and spleen. We were the first to visualise microscopic blood vessels in human bone marrow specimen in 3D and to find that two types of microvessels there (known from other probes and species) were arranged mostly parallel.

We also studied another organ, the spleen. Its tissues differ a lot from bone marrow. Figure 10.2 shows a connectivity analysis of human spleen microvasculature around a follicle—a compartment for more or less migratory lymphocytes. To produce this visual, we performed a 3D reconstruction using our method and colour-coded components of the mesh. Such visualisations enabled further insights in medical research. Perifollicular capillary network in human spleens appears to be primarily fed by red pulp microvessels. This is a new result in microanatomy.

Finally, the last two chapters were dealing with additional topics from computer graphics. We optimised camera positions in a special deflectometry device for better and faster assessment of hail damage on passenger vehicles. With our camera placements, different vehicles could be inspected without reconfiguration of the device. We also achieved high coverage values (> 94%).

We used image processing techniques to visualise cryptographic keys. Our visualisations are very direct, but at the same time aesthetically pleasant. We do not use hashes or fingerprints to visualise public keys, which singles out our method among all other approaches. We can detect some kinds of cryptographic attacks with our visualisations.
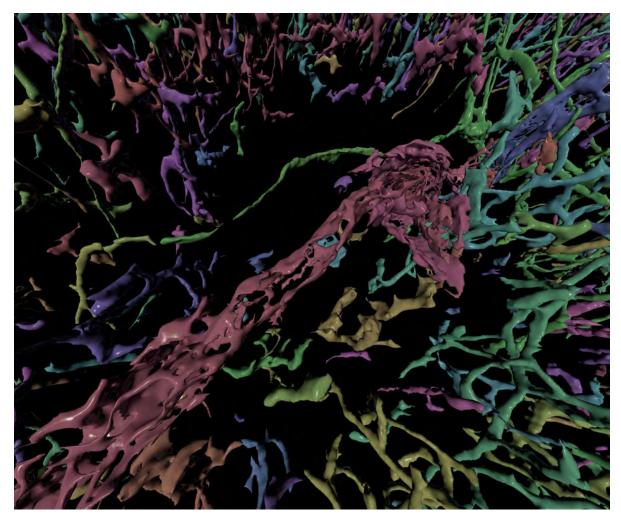
Figure 10.2: Connected component analysis in spleenic white pulp. ROI 3 from Chapter 7.

10.2 Future work

The uttermost quality criteria of a 3D reconstruction from serial sections is quality of the registration. While we are quite pleased with our registration method [Lobachev et al., 2017a, Chapter 2], there is always room for improvement. Our current method works well on regions of interest. In the light of acquisition of hundreds of sections in a single series, two questions arise.

- *Performance of the registration.* Though we already improved the performance from Chapter 1 to Chapter 2 with a faster Krylov subspace solver, a GPU-based solver and a GPU-based feature detector would speed up the registration.

- *Processing of larger sections.* Is it possible to process the whole section at once with the fine-grain method? An idea for the future work is *mosaicking,* as used for rigid registration of electron microscopy images by Saalfeld et al. [2010]. Combining mosaicking and our feature-based multi-resolution method is a promising research idea.

It might be also viable to use the correspondences between the sections from the registration for generation of intermediate images. Chapter 3 currently uses dense optical flow for this purpose.
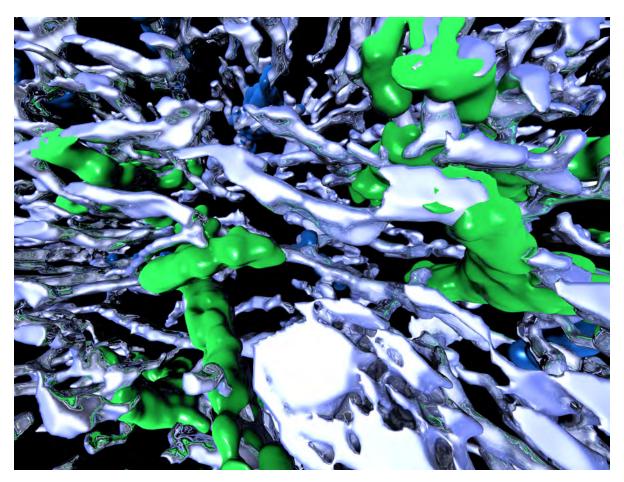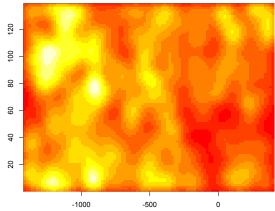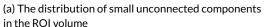
Figure 10.3: Blood vessels and fibroblasts (glossy) and capillary sheaths (blue and green) in a human spleen. The capillary sheaths are classified into ones surely following arterioles (shown in green) and others (blue).
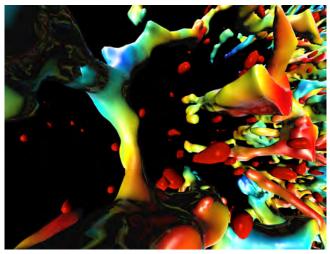
A lot of further work is in sight in VR-based quality control methods. Firstly, some technical improvements to the application are viable. Overlays and elaborate menu structure would allow the immersed user to control the application in absence of the operator 'outside' the VR. It might be of interest to use progressive meshes [Hoppe, 1996, Derzapf and Guthe, 2012] to display the reconstructions. One further improvement could be a possibility to record vocal remarks to the annotations the user places in VR.

Secondly, extensions to the existing VR-based verification paradigm are viable. The most prominent of them is the multi-user VR. The idea is to share the same 'virtual space' among multiple users, each with her own headset. Such extension would allow, e.g., *direct* collaboration of the 3D reconstruction expert and the domain expert in VR. In order to make this possible, the application has to be divided into a server part and a client part. Network protocols need to be developed for the transmission of the initial scene and the current positions of headsets and controllers. Once again, the usage of progressive meshes, but also texture compression techniques, would help to transmit the initial scene over the Internet in a more responsive manner.

As for the medical part of this work, we look forward to further investigating both already explored and new organs and tissues. To give a few examples:

(a) The distribution of small unconnected components in the ROI volume

(b) The distance of small unconnected components to blood vessels in R1

Figure 10.4: The distribution of small unconnected components (that would be stem and progenitor cells) in bone marrow regions R1 (left) and R3 (right) from Chapter 6. In (a) we show the probability to see a vertex of a small component in y vs. z coordinate of the volume. The axes show internal mesh dimensions. In (b) we visualise the distance between small components and microscopic blood vessels. The Hausdorff distance is encoded as rainbow, dark red in 0 µm, dark blue is 100 µm.

- In the red pulp of the spleen, capillaries have a special surrounding endothelia, the so-called 'capillary sheaths'. Although they were first observed very early, the questions of their exact function and even of their location remain open. Our 3D reconstruction aims to answer the latter question. With VR-based visualisation domain experts can assess the location of the capillary sheaths relative to corresponding arterioles. We will also be able to produce some numerical results w.r.t. to a number and location of the sheaths in different regions. We show one such region in Figure 10.3.

- The multitude of single loose CD34$^+$ cells in the bone marrow specimen lead to an interesting medical research question. These cells are haematopoietic stem and progenitor cells. Their location w.r.t. to the different microscopic blood vessel types has been investigated in rodent models, but not in humans. Today we can only say, that they are definitely not evenly distributed in the probe volume, as Figure 10.4, (a), shows. Small components appear to be close to sinus (b), but this issue needs to be investigated in a more detail.

- An application of the hard embedding technique to other bones, such as femoral head in thigh, and a study of larger serial sections of paraffin embeddings of the spleen specimens would both be a very challenging, but interesting projects. Requirements would include large section series, automated scanning, and, of course performance adjustments to our processing pipeline to facilitate the processing of very large data volumes in a timely manner. With these prerequisites we could reconstruct the vasculature (or other tissue structure) in said specimens for multiple millimetres, possibly even in single centimetre range.

  This would be an amazing and a very exhaustive project: a centimetre doesn't seem to be a lot, but in the z-axis direction processed at, e.g., 1 µm/image, a centimetre means 10k images.

(a) First image from a z-stack   (b) Second image from a z-stack   (c) Major differences between (a) and (b)
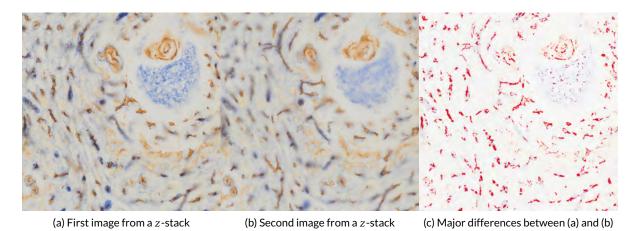
Figure 10.5: Images from a thick serial section. The section was 25 µm thick, acquired with an Olympus VS120 microscope at 0.3 µm/pixel.

A technical issue to explore in the future is the computer-based reconstruction of data from *thick* serial sections. We have processed the sections with a typical thickness of 7 µm. This thickness has been established for visual inspection under a microscope. How can we inspect, e.g., 25 µm thick sections? Successful processing of thicker sections would facilitate less cutting and thus increase the performance of biological processing. Larger acquired volumes lead to larger 3D reconstructions. Figure 10.5 shows an example of a thick section. We are concerned with *where* in the image stack would be a structure in focus. We also need to remove from all images all the instances of structures not in focus.

This leads us to a problem statement in the field of computational photography: a deconvolution of a focus stack to yield a 3D volume. In optical transmitted light microscopy such deconvolutions would be much harder than in other fields, such as computed tomography and fluorescent microscopy. More careful acquisition of initial data, possibly with coded aperture, might be necessary to obtain good results. In a more generic sense, we would like to apply the repertoire of computer vision and computational photography to optical transmitted light microscopy in the context of automated full-sized acquisition of serial sections. The final goal would be, as already the case in this work, a full-fledged 3D reconstruction. The novelty would stem from drastically increase data size, even when compared to our, quite large, volumes.

Our methods have been used in microanatomy research. However, they are not application-specific. It appears quite promising to apply our methodology to the 3D reconstructions in other research areas. To give a few examples:

- *Materials science.* Inspection of materials by means of grinding produces a series of images, but also destroys the material in the process. We imagine that our method might be quite applicable there, as materials science is concerned with 3D structure of various materials.

  Possible applications also include research dentistry, botany, and bionics.

- *Chemistry.* A particular area of chemistry, chromatography, is concerned with filtering and travel duration of fluids through a chromatographic column. Conversely, the shape and inner structure of real columns matters and is a subject of research.

With our 3D reconstructions we could facilitate better understanding of chromatographic columns from real-life data.

- *Oceanography*. Our registration methods were designed for series of larger 2D images. However, our 'sparse' approach might be applicable for 'stitching' of sea floor data acquired from sonar sessions.

APPENDICES

L'avenir, tu n'as point à le prévoir
mais à le permettre.

Your task is not to foresee the future, but to enable it.

<div style="text-align: right">

A. de Saint Exupéry,
Citadel

</div>

HISTOLOGY, SERIAL SECTIONS, AND IMMUNOSTAINING

Which provides a brief overview of the medical processing of the specimen for the inspection under a microscope. We utilise the same techniques to obtain data for 3D reconstructions.

ISTOLOGY is one of the oldest medical imaging techniques. Removed biological specimens are embedded in paraffin or hard plastic to facilitate sectioning. They are sectioned into thin layers ('sections'), colour is added via 'staining', and a stained section can be inspected under a microscope. The sectioning procedure, however, may damage and definitely distorts the source material. Improvements from using better cutting blades, motorised sectioning devices ('microtomes'), and immunohistology (most notable for us: monoclonal antibodies, detailed below) allowed for more selective discrimination of varying cell types and more robust and stable sectioning. With the advance of immunohistological stainings, specific antigens in the specimen can be targeted. With modern microtomes, consecutive, 'serial' sections of the specimen can be obtained and processed. The emergence of optical scanning microscopes allowed to acquire whole sections in a very high resolution and to deal with region selection and further processing later—and digitally. In this manner the microscopy session and image acquisition were reversed and separated from each other, the so-called 'virtual microscopy' became viable. With image pyramids and proper UI users can navigate in an overview image and zoom in at will. Modern optical scanning microscopes and overall advances in the sectioning and staining procedures allow for the digital acquisition of serial sections. In digitised series of sections even more spectacular possibilities emerge.

A.1 Serial sections

Long structures, like blood vessels, are impossible to capture on a single section. Serial sections provide the data required to track the flow of vasculature, but tracking a single blood vessel in a series is impossible without computer-based aid. Only a thin section of the specimen—typically about 7 µm thick—is immunostained to enhance focusing in transmitted light microscopy. If larger structures (such as a microvessel network) are to be observed, serial sections need to be produced.

With further processing (most notably: registration), a stack of serial sections can be treated as volume data. In this thesis we do not stop there, but perform a full-fledged *3D reconstruction* from immunohistologically stained serial sections. General histological procedures before digital processing of our sections include embedding, sectioning, immunohistological staining, and image acquisition. Our toolsuite (Part I) facilitates digital processing of serial sections from registration to rendering and quality control.

3D reconstruction from histological serial sections is a very useful tool that enables insights otherwise impossible. While with, e.g., rodent models, some other techniques are viable, it is not the case with human specimens. Conventional MRI, CT, and ultrasound methods do not have the desired resolution. This is also true for micro-CT and similar methods. Modern optical scanning microscopes produce images at 0.3–0.5 µm/pixel in the xy plane for probes that easily reach $1\,cm^2$ area. Micro- and nano-CT methods either have orders of magnitude worse resolution or cannot deal with that large probes. (A recent effort used micro-CT data as a guide for registration of serial sections [Nagara et al., 2017].) Even worse, there is no way to *identify* the structures of interest in human specimen under non-invasive imaging techniques. Conversely, immunohistological staining

provides a reliable method to mark specific kinds of molecules in the cells as long as these molecules can be adequately fixed after specimen removal. It is possible to unequivocally identify, e.g., the cells forming the walls of small blood vessels, the so-termed endothelial cells, using modern histological methods. We need to stain the specimen to see the cells of interest. We section the specimen for further processing and acquisition. However, sectioning of more than $\approx 30\,\mu m$ of tissue is impractical for paraffin embeddings. An alternative might be fluorescence microscopy. In non-human probes it is even possible to genetically enhance cells with fluorescence, thus improving the separation of cell types. However, fluorescence labelling degrades very fast under illumination, while stainings for transmitted light microscopy are known to survive for decades. An interesting alternative is tissue clearing [Richardson and Lichtman, 2015]: the refraction index of the tissue is adapted to the index of embedding fluid, but this approach is still highly experimental and has not been tried profoundly on human specimens and with transmitted light microscopy. Basically, for microanatomical research on human specimens no other technique except serial sectioning is practical and solidly established.

A.2   Immunohistology

Immunohistological staining is a procedure to detect proteins or other antigens in thin sections of tissue that have been typically fixed and embedded in paraffin or a similar material for cutting with a microtome. Alternatively, frozen sections may be used. There are different detection systems available. The system used for our experiments is called the 'ABC'-system (avidin-biotinylated peroxidase complex system). For processing of sections, the embedding material needs to be removed in a first step. Then appropriate antigen retrieval is performed and endogenous biotin and peroxidase activity is blocked. We first use a specific antibody, which binds to the surface and/or the interior of the cut cell. Subsequently, a second antibody binding to the first one is applied. This secondary antibody has been conjugated to the small molecule biotin. Biotin is used, because it is bound by the protein avidin at one of the highest affinities known in biological systems. To detect the biotin-labelled secondary antibody, we use a molecular complex of avidin and biotinylated peroxidase, containing many unoccupied biotin-binding sites in the avidin molecules. When the molecular complex is applied to the sections, it binds to biotin of the secondary antibodies. This means, that many peroxidase molecules are present at the site of the antigen. Peroxidase is an enzyme that catalyses an oxidative polymerisation of soluble colour substances such as diaminobenzidine to an insoluble product in the presence of hydrogen peroxide. Thus, a permanent brown precipitate is deposited, where the enzyme has exerted its function at the location of the first antibody. Another enzyme that can be used for immunohistology is alkaline phosphatase, which produces red or blue colour substance precipitates at the location of its action. Both types of enzyme reactions (peroxidase and alkaline phosphatase) can be sequentially combined to result in a subtractive staining procedure. The number of applicable enzymes and colours is limited so that in most cases not more than three colours are used for conventional light microscopy.

FEATURE-BASED MULTI-RESOLUTION REGISTRATION OF IMMUNOSTAINED SERIAL SECTIONS. SUPPLEMENTARY MATERIAL

Which serves as a supplementary text to Chapter 2. The main goal is to provide more information. Some details not mentioned there are shown. This chapter also shows additional figures and gives an exact definition of quality measures used. The pinnacle of this chapter is the full quality evaluation of our registration method. It was too large to be shown in main text. Here, we prefix all the references to the main text with "M."

ONSIDER for a second a *direct* match of key points in an all-against-all manner. It is neither feasible nor does it yield good results, as Figure B.1 shows. Thus we rematch the key point pairs basing on a weighted implementation of radius match. Then least-squares optimisation follows to yield the positions of B-spline control points for the actual non-rigid distortion. Compare the matches of a rigid stage with the matches of the non-rigid stage: in Figure B.1, (a) shows decimated direct matches, (b) is the rigid match we find in the first iteration, (c) shows matches from the non-rigid registration. M2.2.2

We present a side view on serial sections of spleen, double-stained, in Figure B.2 (a)–(c) show the horizontal the *z*-stacks and demonstrate that our method is superior not only to the single-resolution registration ('A'), but also to the repeated application of single-resolution registration to the ROI ('A–A'). Compare Figure B.2 a) to Figure B.2 b): we see some positive effects of a second application of non-multi-resolution registration. However, in Figure B.2 c) we see a more significant effect of our method ('A–B'). The shapes of the blood vessels and other objects in this *z*-stack are better, esp. in marked positions. This is even more pronounced in the volume renderings in Figure M2.12 in the main text. M2.2.3

In Figure B.3, (a) shows the crop zone and the ROI (whole image shown) used for the fine-grain registration. Fine capillaries are not properly registered using the slice-wide, coarse registration in (b). We used the method 'A' here. It served as an input for (c) and (d). Compare the fine-grain registration using (c) single-resolution registration 'A' again (yielding 'A–A') with (d) our method 'B', the multi-resolution registration (resulting in 'A–B'). Small capillaries should profit from our method as they are too small to produce features that are large enough to be regarded by a single-resolution method, esp. in a slice-wide application. We compare these two methods in detail in main paper.

Figure B.4 shows the full volume rendering of a single-stained spleen specimen. Figure M2.8 (in main text) shows a region from this image that was only cut out, without any re-registering. Figure B.5 shows the processed bone marrow specimen as a volume rendering of the alignment with our method. We used two levels of multi-resolution registration. M2.3.3
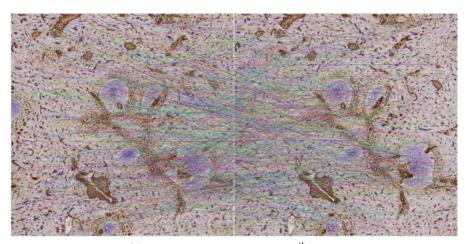
To further highlight our method we show *z*-stacks and additional volume renderings of another bone marrow ROI in Figures B.6–B.7. Observe the precision of section matching in the *z*-stacks (Figure B.6). These images are not directly comparable to each other, because different registrations have distorted them differently. Instead, compare the quality of the separate shapes without comparing the shapes directly. We also show complete volume renderings accompanied by more detailed crops (Figure B.7). The crops in Figure B.6 b), d), f) show visually the benefit from using our method. M2.4
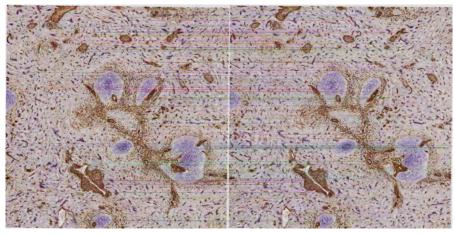
(a) naive matching, showing each 10th match

(b) matches during rigid matching

(c) matches during the non-rigid registration

Figure B.1: Matches in a ROI from spleen data, double staining. Fig. (c) shows matches at the end of the first non-rigid registration phase, i.e., those with largest search radius and size. We add further (smaller) matches within a smaller search radius in the next phases. The registered data set is in Figure M2.9 in the main text.

(a) Only section-scaled single-resolution registration.

(b) Single-resolution registration applied again to the ROI only.

(c) Our approach, 2-level multi-resolution registration applied to the ROI.
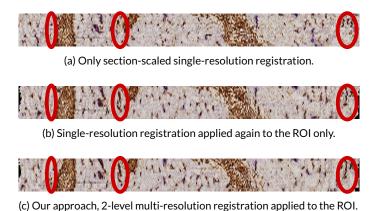
Figure B.2: Stack of 24 serial human spleen sections double-stained for CD34 plus SMA (brown) and CD271 (blue). Compare the z-profiles of this specimen. All the sections underwent a single-resolution registration and were then cropped to ROI. Note the differences in the vascular network at the marked positions.

(a) Overview   (b) Coarse ('A')   (c) 'A' again   (d) 'A–B', ours
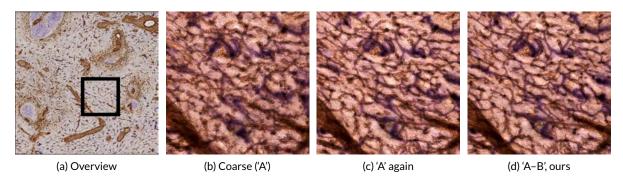
Figure B.3: Stack of 24 serial human spleen sections double-stained for CD34 plus SMA (brown) and CD271 (blue). We showcase how well small capillaries are represented, this is pivotal for histological research basing on this images. Capillary sheaths are visible as blue areas around brown capillaries. Compare the crops from the volume rendering, from left to right: (a) legend, a single slice showing the registered ROI and post-registration crop location. (b) coarse, section-wide registration only using single-resolution registration of Ulrich et al., (c) the series from b) was registered *again* using the single-resolution method, (d) the series from b) was registered using our method. We used four levels of multi-resolution registration.

B.1   Applying our method to standard images (Section M2.4.5)

To show that our method is applicable not only to histological images, we generated distorted series from standard test images and registered them with our method.

We took two standard images from signal processing: 'Lena' and 'mandrill'. To use them for the image registration, we randomly distorted the images. We moved random points of the input in random direction and distance using Shepard's distortion. Both the offsets and the initial positions were generated using a normal distribution with $\sigma = 204$ for positions; with $\mu = 2.5$ and $\sigma = 4$ for offsets. (Fig. B.8) All the difference images here are negated for the sake of presentation—white is no difference. The image size was 512×512 pixels. We upscaled the input images after the distortion $6\times$.
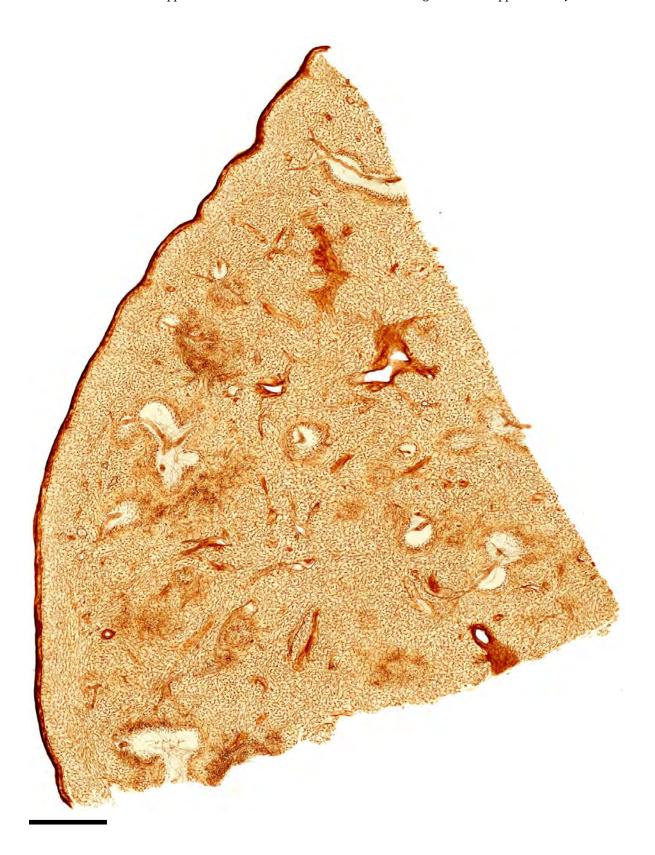
Figure B.4: Volume rendering of spleen, single-stained. *(Full caption on the next page.)*

Figure B.4: (Continued from previous page.) Volume rendering of spleen, single-stained for CD34. We overlay 24 images after 4-levels of our multi-resolution registration. Figure M2.8 in the main text shows a crop of this overview image.
Scale bar is 1 mm.
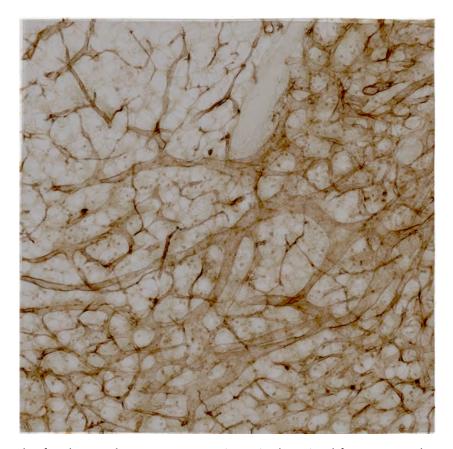
Figure B.5: Stack of 30 human bone marrow sections single-stained for CD34 and CD141 (both in brown). In contrast to Section M2.3 and Figure M2.11, it was registered with *two* levels of our method, after the single-resolution registration was applied both to the full section and to the ROI. We evaluate this ROI in Figure B.17.

Then we performed the registration using our method, see Fig. B.8, page 205 for the results.

B.2   Extended Quantitative Evaluation

B.2.1   Definitions of the Quality Measures (Section M2.4.3)

We discuss seven different quantitative measures of registration quality for each of the specimens (Fig. M2.5–M2.7) below. The paper has omitted the formula for the quality measures used, we fix this here.
    For the notation, consider two images **A** and **B**, two single pixel values at (i, j) are A_{ij} and B_{ij},

<div style="text-align: right">M2.4.3</div>

(a) Original (after section-wide single-resolution registration)

(b) ROI-based single-resolution registration, used as input for our method

(c) Our method improves the overlay

(d) For reference: z-profile using *elastix* on saturation channel, negated for presentation

Figure B.6: Human bone marrow, z-profiles of 30 sections single-stained for CD34 and CD141 (both in brown), showing the same ROI as in Figure B.7

the mean values at (i, j)-centred window are $\overline{A}_{ij}$ and $\overline{B}_{ij}$, i.e., for $n \times m$-sized window,

$$\overline{A}_{ij} = \frac{1}{nm} \sum_{k=i-n/2}^{i+n/2} \sum_{l=j-m/2}^{j+m/2} A_{kl}.$$

The cross-correlation in the same $n \times m$ window is

$$\rho_{ij}(\mathbf{A}, \mathbf{B}) = \frac{1}{nm} \sum_{k=i-n/2}^{i+n/2} \sum_{l=j-m/2}^{j+m/2} (A_{kl} - \overline{A}_{ij})(B_{kl} - \overline{B}_{ij}).$$

With $\partial A_{ij} / \partial x$ we denote the derivative of $\mathbf{A}$ at (i, j) in the direction of x axis.

To obtain the data for quality evaluation we registered the *whole* series with each method, picked a pair of consecutive images and compared them using the following measures over the whole ROI (we show later a detailed view when applicable):

- structural similarity (SSIM) is

$$\frac{(2\overline{A}_{ij}\overline{B}_{ij} + C_1)(2\rho_{ij}(\mathbf{A}, \mathbf{B}) + C_2)}{(\overline{A}_{ij}^2 + \overline{B}_{ij}^2 + C_1)(\rho_{ij}(\mathbf{A}, \mathbf{A}) + \rho_{ij}(\mathbf{B}, \mathbf{B}) + C_2)},$$

with stabilisation constants C_1 and C_2, using 11×11 Gaussian weighting. Final single SSIM value is the mean over the overall SSIM image.

(a) Original (after section-wide single-resolution registration)

(b) A detail from (a)

(c) ROI-based single-resolution registration, shows almost no effect in volume rendering, used as input for our method

(d) Single-resolution registration, a detail from (c)

(e) Our method improves the overlay

(f) Our method, a detail from (e)
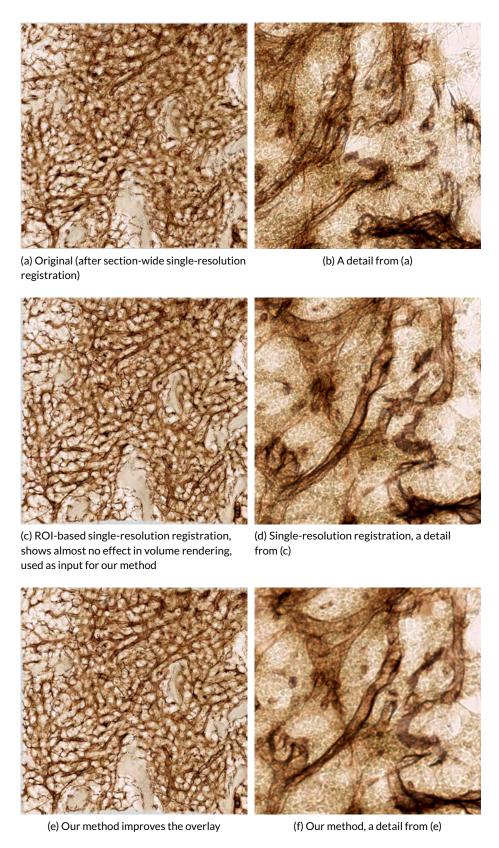
Figure B.7: Stack of 30 human bone marrow sections. *(Caption on the next page)*

Figure B.7: (Figure on the previous page.) Stack of 30 human bone marrow sections single-stained for CD34 and CD141 (both in brown), showing another ROI.

The stack shows capillaries (microvessels with small diameter) and sinuses (microvessels with large diameter). Both microvessels form a network with smooth transitions from capillaries to sinuses. Sinuses may form round or flattened structures. It is impossible to represent blood microvessels in such length without registering a series of sections.

A repeated application of single-resolution registration to a) would not show a better result, as it only operates on larger key points. These, however, were already aligned in (c). We need to align smaller key points with our method to yield a difference (e). These results are clearly visible in selected details (b), (d), and (f): notice the blurry blood vessel contour in d) and how it is more pronounced in f).

Figures b), d), f) where shown in a smaller magnification in the main paper as Fig. M2.4. Fig. B.6 also shows the same ROI. Our registration method here uses *four* levels to produce an even better stack than in Fig. B.5. Fig. M2.11 (main document) is similar to (e), but shows another ROI, a larger view of Fig. B.5.

- pixel-wise sum of squared differences (SSD), i.e., $\sum (A_{ij} - B_{ij})^2$.

- zero-mean sum of absolute differences (ZSAD) with window 16×16,

$$\sum (A_{ij} - \overline{A}_{ij}) - (B_{ij} - \overline{B}_{ij}).$$

  We compute the absolute differences pixel-wise, but beforehand subtract the mean of the values in the window around the pixel from the actual pixel values. The resulting images are quite dark, for the sake of presentation we adjust the brightness and contrast to $+50$ in all presented ZSAD images. The numerical values are not corrected.

- mutual information (MI) increase. We compute the mutual information for an image pair. For the actual value we divide the mutual information in two registered consecutive images of a scan series by the mutual information in the original, unregistered image pair to obtain the increase of mutual information compared to baseline.

- gradient cross-correlation (GCC):

$$\rho_{ij}(\partial A_{ij}/\partial x, \partial B_{ij}/\partial x), \quad \rho_{ij}(\partial A_{ij}/\partial y, \partial B_{ij}/\partial y).$$

  We compute the gradient separately in the direction of the x and y axes using Scharr operator, compute the cross-correlation using window size 16 and then blend both images together. As the result is too dark, we use gamma correction with $\gamma = 6$ before blending. The numerical values are not corrected.

- Farnebäck's dense optical flow that is defined in the following way. We model a signal with quadric polynomials, e.g., $f_1(\mathbf{x}) = \mathbf{x}^T \mathbf{A}_1 \mathbf{x} + \mathbf{b}_1^T \mathbf{x} + c_1$ for a symmetric matrix $\mathbf{A}_1$, a vector $\mathbf{b}_1$ and a scalar c_1. A globally displaced by $\mathbf{d}$ signal f_2 is defined in a similar way in terms of $\mathbf{A}_2, \mathbf{b}_2$ and c_2. We can express $\mathbf{d}$ as $-\mathbf{A}_1^{-1}(\mathbf{b}_2 - \mathbf{b}_1)/2$. Now, to make the computation more practical we use *local* polynomial approximations of images. In the following description of the optical flow we omit the $\mathbf{x}$ argument, e.g., $\mathbf{A}$ always stands for $\mathbf{A}(\mathbf{x})$, a local approximation. Let

$$\mathbf{A} := (\mathbf{A}_1 + \mathbf{A}_2)/2 \quad \text{and} \quad \Delta\mathbf{b} := -(\mathbf{b}_2 - \mathbf{b}_1)/2.$$

(a) Original image

(b) A distorted image used for registration

(c) The difference between (a) and (b)
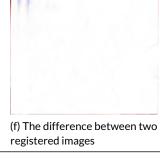
(d) The difference between two differently distorted images used as input for registration

(e) A registered image using our approach

(f) The difference between two registered images

(g) Original image

(h) A distorted image used for registration

(i) The difference between (g) and (h)
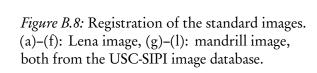
(j) The difference between two differently distorted images used as input for registration

(k) A registered image using our approach

(l) The difference between two registered images

Figure B.8: Registration of the standard images. (a)–(f): Lena image, (g)–(l): mandrill image, both from the USC-SIPI image database.
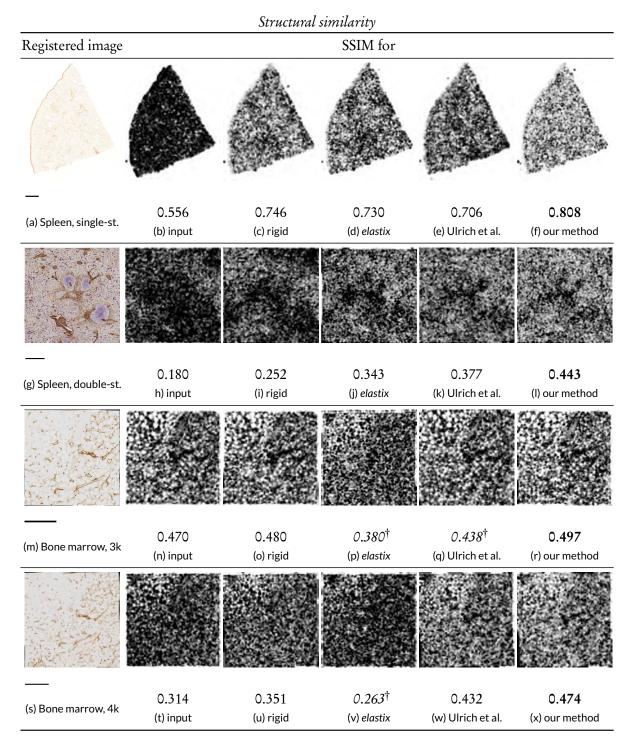
Figure B.9: Structural similarity index (SSIM). This is an extended version of Figure M2.13 in the main text, showing SSIM for spleen, single-stained (top row, (a)–(f), scale bar is 1 mm), spleen, double-stained (second row, (g)–(l), scale bar 500 μm), bone marrow (3k side, two levels of our method, third row, (m)–(r), scale bar 500 μm), bone marrow (4k side, four levels of our method, bottom row, (s)–(x), scale bar 500 μm) specimens. *(Continued on the next page.)*

Figure B.9: Continued from previous page. The left column, (a), (g), (m), (s), shows a registered image, all other columns show SSIM images and SSIM values: the second column (b), (h), (n), (t) shows SSIM for input data, third one, (c), (l), (o), (u), shows SSIM for rigid-only registration, fourth and fifth columns show state of the art: (d), (j), (p), (v) features *elastix*, (e), (k), (q), (w) shows Ulrich et al., sixth and last column, (f), (l), (r), (x), shows our method.
The higher the value and the lighter the image, the better. **Bold** shows the best value, *italics* with † denotes decrease from baseline.

The distance is now a spatially varying displacement field $\mathbf{d}$ with $\mathbf{A}\mathbf{d} = \Delta\mathbf{b}$. Let w be a weight function—the edges of an image have a smaller weight. Now, in a neighbourhood I following holds:

$$\mathbf{d}(\mathbf{x}) = \left(\sum_I w\mathbf{A}^T\mathbf{A}\right)^{-1}\sum_I w\mathbf{A}^T\Delta\mathbf{b},$$

the sums run over all $\mathbf{x}$ in I, as $\mathbf{x}$ is an omitted for brevity argument to $\mathbf{A}$, $\Delta\mathbf{b}$, and w. In practice we compute $\mathbf{A}^T\mathbf{A}$, $\mathbf{A}^T\Delta\mathbf{b}$, and $\Delta\mathbf{b}^T\Delta\mathbf{b}$ (needed for the confidence value) pointwise.

The flow is computed on multiple scales to alleviate the problem with larger displacements. We computed it based on a classical 2^n image pyramid with 3 levels, averaging window size of 15 pixels, three iterations at each level, neighbourhood size for polynomial expansion of 5 pixels. The derivatives for polynomial expansion are smoothed with $\sigma = 1.2$.

- The Jaccard similarity measure for two sets X and Y is defined as

$$J(X,Y) = \frac{\#(X \cap Y)}{\#(X \cup Y)},$$

the valid values lie between zero (no similarity) and one (full similarity). As our inputs are images, say, $\mathbf{A}$ and $\mathbf{B}$ we threshold the data at a given value, yielding a coarse kind of segmentation. Let the binary threshold of an image $\mathbf{A}$ at value x be $t_x(\mathbf{A})$. Then we compute the Jaccard measure for the resulting binary images. For multi-channel images we perform the computation channel-wise and take the maximum. Let r, g, b be the channels, and $\mathbf{A}^r$ a single channel of an image $\mathbf{A}$. Hence, the Jaccard measure we use is

$$J(\mathbf{A},\mathbf{B}) = \max_{c \in \{r,g,b\}} \frac{\#(t_x(\mathbf{A}^c) \cap t_x(\mathbf{B}^c))}{\#(t_x(\mathbf{A}^c) \cup t_x(\mathbf{B}^c))},$$

we use RGB channels when applicable.

The result is presented in Fig. B.9 (SSIM for all specimens), Figs. B.10–B.12 and Table M2.2 in the main text (optical flow), Table 2.3 (Jaccard measure for all specimens), Fig. B.14 for spleen, single-stained, Fig. B.15 for spleen, double-stained, Figs. B.17 and B.18 for bone marrow present results of SSD, ZSAD, mutual information increase, and GCC for respective specimens.

Figs. B.9, B.17 and B.18 show the results both for bone marrow with two levels of registration with our method and with four levels of registration. This facilitates the comparison of different multi-resolution approaches and underlines the importance of iteratively decreasing feature sizes in a registration.

B.2.2   Full evaluation (Section M2.4.3)

Chapter 2 does not show the complete evaluation because of article size limitations. We show it here in complete detail. Images in Figs. B.14–B.17 show crops to emphasise details of the registration. Below the images we show mean values for whole images, not for the crops shown. All differences and improvements in per cent are relative values. M2.4.3

The SSIM and Jaccard measure images and values are better when lighter and higher. SSD, ZSAD, GCC, optical flow images and values are better when darker and lower.

SSIM for all specimens (Fig. B.9)

Figure B.9 shows the structural similarity index for spleen, single-stained (b)–(f), spleen, double-stained (h)–(l), and bone marrow (n)–(r). Figs. B.9 a), g), m) show corresponding specimens registered using our method.

The SSIM images were very convincing for spleen, single-stained specimen. The input overlay Fig. B.9 b) produced a very dark (i.e., bad) image, our method (f) was much lighter, the lightest in the series. The images for *elastix* (d) and the method of Ulrich et al. (e) were quite similar to rigid-only transformation (c), but a bit *darker*, which corresponded to a minor decrease of the quality value. Overall, our method showed 45% improvement over baseline, while *elastix* yielded 31% of improvement. The relative difference between the SSIM for our method Fig. B.9 f) and for *elastix* (d) was over 10%. The difference between SSIM for our method (f) and for the method of Ulrich et al. (e) was over 14%.

Again, SSIM for spleen, double-stained were much darker for the initial overlay (h) than for our method (l). The SSIM value here improved more than by 146%, the best overall result. As for *elastix* (j) the improvement was 90.5%. All three non-rigid methods, (j)–(l), showed SSIM images much lighter than the rigid-only transformation (i).

As for bone marrow at two levels of our method, the difference was not that high. The image for our method Fig. B.9 r) appeared a bit lighter (i.e., better) than the input overlay (n). Our multi-resolution method improved the overlay of two images by 5.7%, however repeated application of the non-multi-resolution method by Ulrich et al. (q) did not improve the situation at all. The SSIM quality decrease was here −6.8%. Same holds for *elastix* (p), it impaired the registration by −19%.

The bone marrow specimen at four levels of our method (x) showed more than 50% of SSIM improvement over the coarse-only registration (t) for a 4k × 4k ROI. The result of *elastix* (v) was visibly darker, indeed it showed a decrease in SSIM value w.r.t. baseline (t). It corresponded to over 44% quality decrease when compared to our method. Ulrich et al. (w) facilitated 37.6% improvement over baseline (t).

Optical flow (Section M2.4.3, Figs. B.10–B.12)

We measure the magnitude of dense optical flow between to consecutive images based on work of Farnebäck. This method basically approximates the both signals locally in neighbourhood with quadric polynomials, then solves for the translation vector. For the practical reasons the neighbourhoods are averaged and the fitting happens iteratively on multiple scales. We use pairs of consecutive registered images as input and evaluate the motion vector field in terms of its magnitude and angle. M2.4.3

Table 2.3 lists the means of the magnitude of translation vector between two consecutive registered images. Here we extend it with Figures B.10, B.11, and B.12 that show angle and magnitude between two slices as a colour coded image for single-stained spleen, double-stained spleen, and bone marrow specimens.
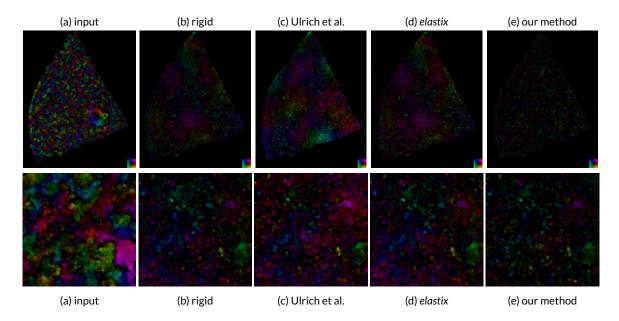
Figure B.10: Colour-coded optical flow as displacement angle and magnitude for single-stained spleen specimen. From left to right: input, Ulrich et al. (Chapter 1), *elastix*, our multi-resolution method (Chapter 2). Top row shows whole image, bottom row shows a crop. The smaller the intensity of the image, the better.
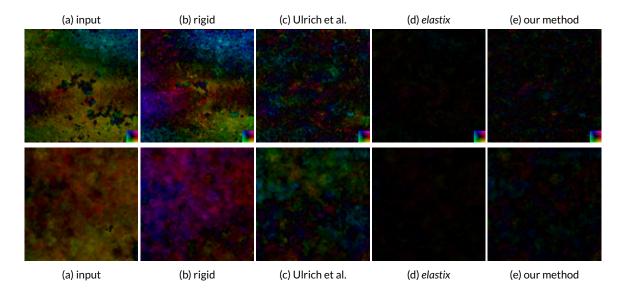
Figure B.11: Colour-coded optical flow as displacement angle and magnitude for double-stained spleen specimen. From left to right: input, Ulrich et al. (Chapter 1), *elastix*, our multi-resolution method (Chapter 2). Top row shows whole image, bottom row shows a crop. The smaller the intensity of the image, the better.
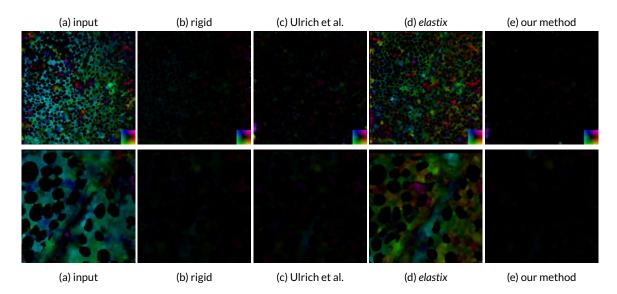
Figure B.12: Colour-coded optical flow as displacement angle and magnitude for bone marrow specimen. From left to right: input, Ulrich et al. (Chapter 1), *elastix*, our multi-resolution method (Chapter 2) Top row shows whole image, bottom row shows a crop. The smaller the intensity of the image, the better.

The colour-coded optical flow for single-stained spleen specimen (Fig. B.10) appeared quite dark for rigid-only transform (b) and our method (e) which is good, Ulrich et al. (c) and *elastix* (d) appeared somewhat lighter, while input overlay was the brightest (i.e., the worst). This corresponded well with Table M2.2.

In double-stained spleen spleen analysis in Figure B.11, *elastix* (d) and our method (e) were dark, where Ulrich et al. (c) showed some coloured spots. These were almost everywhere in input overlay (a) and rigid-only transform (b), though in a different manner. This visual did not quite correspond with the numerical findings of Table M2.2, but the probable culprit was the difference between single-channel and colour data (*elastix* ran on blue channel only), the distribution of motion magnitudes, or even the dimension of the filter kernels in the implementation of dense optical flow we used as detailed in the main document (Sec. M2.4.3).

Bone marrow specimen showed a very bright picture of optical flow between two consecutive images (Fig. B.12) in input overlay (a) and *elastix* (d). Images for rigid-only transform (b), method of Ulrich et al. (c), and our method (e) were quite dark. The numerical evaluation in Table M2.2 stated that Ulrich et al. outperformed our method by a relatively small margin, but it also stated that rigid-only transform was worse than *elastix*, which contradicted with the visuals. Judging from visuals only (i.e., Fig. B.12 c) and e)) we would deem these two methods as approximately equal.

At this point we can only re-state the assumption from the main paper that the averaging window of the dense optical flow method used was larger than the finest-grain improvements in capillaries and similar small structures. Visual inspection and results of 3D reconstruction (Figs. B.6–B.7) clearly indicated than the result of Ulrich et al. was less usable than the result of our registration, which contradicts Table M2.2.

Table B.1: Jaccard measure for all three specimens. The values were computed channel-wise where applicable, we show here the value for the best channel. This is the extended version of Table M3. We denote the best overall value in **bold**, the values worse than baseline are marked with *emphasis* and †, values with possible compatibility problems (single-channel vs. best of three RGB channels) are marked with *. The value marked with ¶ was too low (< 0.06558) due to a too high threshold. The values are scaled between 0 and 1.0, the higher the better.

Specimen	Size	Lev.	Thresh.			Method		
				input	rigid	Chapter 1	*elastix*	Chapter 2
Spleen, single-stained	8000	4	220	0.9701	0.9785	*0.9762*†	*0.9774**,†	**0.9829**
Spleen, double-stained	5032	4	100	0.8622	0.8755	0.8917	*0.8158**,†	**0.9062**
Bone marrow	4032	4	150	0.9583	*0.9562*†	0.9612	*0.8083**,†	**0.9631**
			220	0.6744	0.6905	0.7169	—¶	**0.7224**

Jaccard measure (Section M2.4.3, Table B.1)

Table B.1 extends the Table M3 from the main text. Fig. B.13 shows the results of thresholding for double-stained spleen specimen (top) and bone marrow specimen (bottom).

We chose the threshold value quite high to obtain distinctive values; the Jaccard measure for our method applied to single-stained spleen specimen reaches 1.0 at threshold 100. The stainings we used were quite distinctive, thus we chose the threshold 220 for the same specimen in Table B.1.

In Table 2.3 our method outperformed all other methods used. Both the method Ulrich et al. (Chapter 1) and *elastix* were worse than baseline for spleen, single-stained. In fact, *elastix* performed not so well against this measure, it was consequently below the baseline for all specimens.

This table differs from Table M2.3 in the main text in the last line: we experimented with a larger threshold on bone marrow specimen. Our method from Chapter 2 showed again its superiority w.r.t. this measure. As for *elastix*, we obtained in this setting an absurdly low value which we did not include in the table.

Spleen, single-stained (Fig. B.14)

Visually, the SSD did not change a lot between rigid-only registration, *elastix*, and our method; ours might be a little clearer. The numerical values, however, showed that our method displayed a better mean value. Our method was 14.5% better.

ZSAD for our method in single-stained spleen was visually much clearer and of less intensity. (For the sake of presentation we even had to increase brightness and contrast in all ZSAD images.) The numerical values showed that the ZSAD mean was the best in our method, with 22.7% improvement.

MI showed a slight increase between rigid-only registration and our method. The value for *elastix* was formally better, but it was not really comparable as *elastix* used only the red channel of the images. Thus we evaluated the information increase for this colour channel only. (Section B.2.2 discusses an experiment with grayscale-based *elastix* and grayscale-based MI on our registration to obtain comparable values.)

The GCC values—basically, the 'movement' in the image pair—improved in all cases compared to the original overlay. Our method delivered the best average value overall: 0.64 vs. 0.80, i.e., a 25% improvement.

(a) Spleen, double-stained   (b) Bone marrow
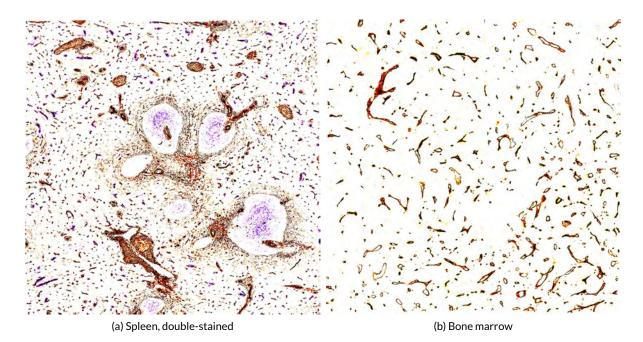
Figure B.13: The pixels occurring in numerator or denominator of the Jaccard measure shown for: (a): a ROI in the double-stained spleen specimen at threshold 100, (b): a ROI in the bone marrow specimen at threshold 150. In (a) the blue staining shows additional information, the brown staining mirrors the microvascular system. In (b) the pixels shown show quite nicely the stained blood vessel walls. Hence the Jaccard measure compares the performance of the registration methods exactly at the point of our interest.

Our method showed best results for all quality measures up to MI increase, where the data was not quite comparable.

Spleen, double-stained (Fig. B.15)

The SSD images showed basically noise, however, in case of *elastix* and our method one might find some stationary points, especially with our method. This corresponded to best mean (and also median, not shown) improvements.

ZSAD was much more informative for this specimen: the "noise"was greatly reduced, the image for our method looked less noisy and also visibly darker. Numerically this meant that our method delivered the best results. The mean value of ZSAD for *elastix* showed roughly 4% quality decrease, while our method had a 21% better mean ZSAD value than rigid-only registration. The median for our method (not-shown) was 24% better.

We obtained an MI increase of 195% which is the best value across the methods tested. However, the value delivered by *elastix* (183%) was not really comparable to our results. The reason for this was that mutual information increase (i.e., a quotient of mutual information in the registered image pair and in original image pair) was computed for RGB images in case of rigid-only transform and our method in Figure B.15, but only the blue channel was used in both image pairs for *elastix*. This was the same image channel *elastix* operated on in Fig. M2.12. If we used only the blue channel for our method, we would obtain 176.48% increase. However, this would not be fair towards our method as

Spleen, single-stained

(a) Overview, detail selection.

(b) Full ZSAD

Sum of squared differences

17.9

(d) rigid

$19.0^{\dagger}$

(e) *elastix*

15.3

(f) our method

(c) detail #13

Zero-mean sum of absolute differences

4.4

(h) rigid

$4.7^{\dagger}$

(i) *elastix*

3.4

(j) our method

(g) detail #14
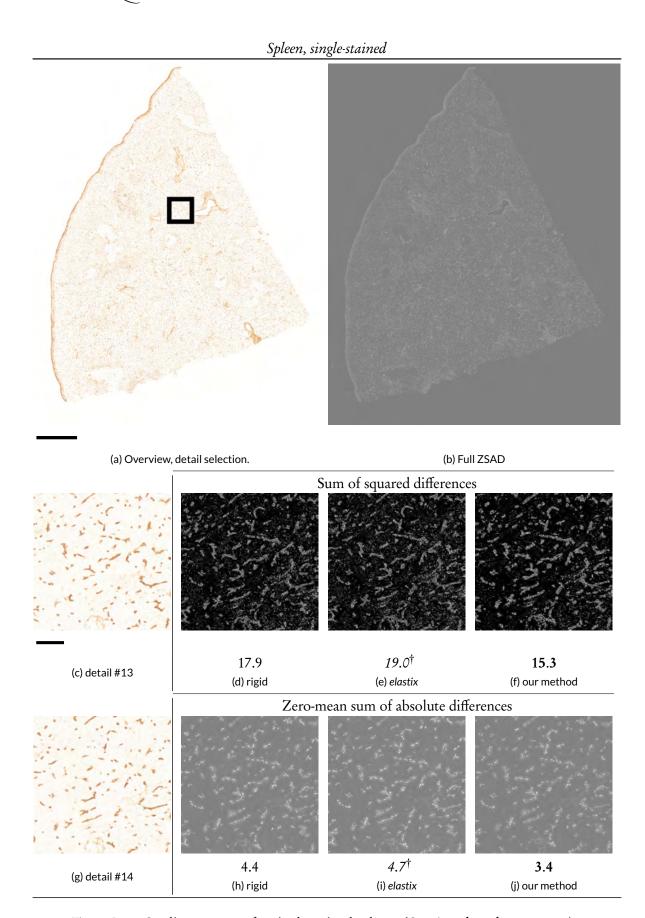
Figure B.14: Quality measures for single-stained spleen. *(Continued on the next page.)*

Spleen, single-stained (cont.)			
original	rigid	*elastix*	our method
100%	154.93%	163.26%*	158.23%
(k) Mutual information increase			

Gradient cross-correlation

0.79	0.68	0.69	**0.64**
0.80	0.68	0.70	**0.64**
(l) original	(m) rigid	(n) *elastix*	(o) our method
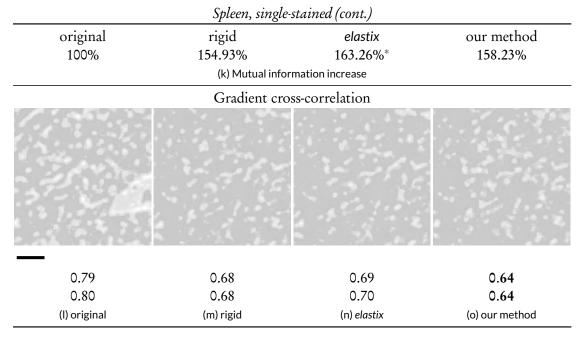
Figure B.14: (Continued from previous page.) Quality measures for single-stained spleen, full section (same as Fig. B.4). (a) Detail selection. Scale bar = 1 mm. (b) Zero-mean sum of absolute differences (ZSAD) measure for the whole image. Brightness and contrast for ZSAD images was corrected with +50. (c) and (g) The detail of the first and second sections. Scale bar = 50 μm. (d)–(f) sum of squared differences (SSD). (h)–(j) Zero-mean sum of absolute differences (ZSAD). (k) Mutual information (MI) increase. (l)–(o) Combined cross-correlation of gradient approximation (GCC) with gamma enhanced to +6. The values below the images show the mean values of full image-wide quality measure. In case of GCC the first (second) line states the mean of derivatives in the direction of the *x* (*y*) axis.
Aside from (k), the smaller the values and the darker the images the better. **Bold** is best, *italics* with † denotes decrease from the baseline, * marks not comparable values.

it did not operate on separate colour channels. If we compared the intensities only (and run *elastix* on intensity images), *elastix* would outperform our method by 3.6%.

The gradient cross-correlation (GCC) showed some darker areas in case of our method, these might guide us to perceive the whole image as more dark, esp. when compared with the original, unregistered image pair or with the rigid-only method. As for mean values, the improvement of our method translated to 12.3%. (The non-shown maximal values improved up to 61%.)

Our method showed better results than *elastix* for *all* measures on double-stained spleen specimen, however the MI values were not quite comparable.

Bone marrow, two levels (Fig. B.17)

Here we compare the method of Ulrich et al., a non-multi-resolution feature-based registration method, to our approach w.r.t. bone marrow images at two levels of multi-resolution registration. We also show the unregistered overlay and rigid-only registration as intended baselines.

The sum of squared differences (SSD) for bone marrow images basically showed noise everywhere besides the interior of large unstained cells (adipocytes, light grey blobs in the microscopic images)

and the non-stained interior of a blood vessel. Still, our method showed the best mean value (by a narrow margin) and was equal to the older method in median and maximum (not shown).

ZSAD quite clearly showed the edges of the blood vessels in a light colour with an apparent matching problem, e.g., in rigid-only registration (h). These values seemed to increase for the method of Ulrich et al. (i) and decrease for our method (j) which led to the best mean value (and also maximum, not shown) in our case and to a ZSAD-quality decrease for the method of Ulrich et al. The difference between mean ZSAD for Ulrich et al. and our method was about 10%.

As we compared MI for RGB images, all values were fully comparable. Unfortunately, *all* registrations of the ROI were worse than the initial coarse-grained registration ('original'). Our method posed a 21% improvement over Ulrich et al., which was, however, still not enough to come over the baseline. We have no plausible explanation for this behaviour.

The GCC images showed quite nicely how the gradient cross-correlation values in our method fade on a larger surface to darker grey, the same colour as the distinctively visible large unstained cells. These images were gamma-corrected with a quite high value of $\gamma = 6$ for presentation sake, so that dark grey was a rather low value. The numerical data for the full and gamma-uncorrected ROIs confirmed this. Our method was best in both directions for the mean value.

If we would simply sum up how often we obtained the best values for each quality measure, we obtain would a score of our method to the method of Ulrich et al. of 5 to 1. We ignored here MI, as both methods performed non-ideally w.r.t. this measure. Still, the MI value for our method was 21% better than that of Ulrich et al. We discussed the non-ideal performance of optical flow as a quality measure above.

Bone marrow, four levels (Fig. B.18)

This discussion is based on four levels of our registration method, the input ROI is $4k \times 4k$, everything else is as in Section B.2.2.

Our method showed the best overall mean SSD value. The result of Ulrich et al. was marginally *worse* than the input overlay. In a contrast, our method managed to improve over the same overlay by 4.7%. The SSD images were not very informative, it is basically noise outside of larger cells. Still we saw some minor improvements (dark areas) in SSD for our method, that were not present in other versions.

ZSAD showed the contours of the blood vessels with an apparent matching problem quite clearly in a light colour, e.g., in (e) and (f). These values seemed to decrease both for Ulrich et al. (g) and for our method (h), however our method showed less bright spots that would correspond to a mismatch. This led to the best mean value (and also maximum and median, not shown) The difference between mean ZSAD for the method of Ulrich et al. and our method was almost 20%.

As we compared MI for RGB images, all values were this time fully comparable. Here, the single-resolution registration was marginally ($< 0.7\%$) better than our method. Both showed improvements over rigid-only registration and input images.

The reason for the not so well performance of our method w.r.t. MI might lie in the specimen, there were a lot small cells that basically constitute background noise. Visual comparisons of reconstructions based on method of Ulrich et al. and on our method (Figs. B.6–B.7) underlined the superiority of our method in the representation of small capillaries.

The GCC images showed how the gradient cross-correlation values in our method faded on a larger surface to darker grey, the same colour as the distinctively visible large unstained cells. We also saw in Fig. (m) some black spots not present in other images. Notably, even dark grey was a quite low value ($\gamma = 6$ was applied). The numerical data on full and gamma-uncorrected ROIs confirmed

Spleen, double-stained

(a) Overview and detail selection   (b) Full ZSAD

Sum of squared differences

| 99.7 | 98.5 | **96.3** |
| (d) rigid | (e) *elastix* | (f) our method |

Zero-mean sum of absolute differences

| 33.1 | *34.5*[†] | **26.2** |
| (h) rigid | (i) *elastix* | (j) our method |
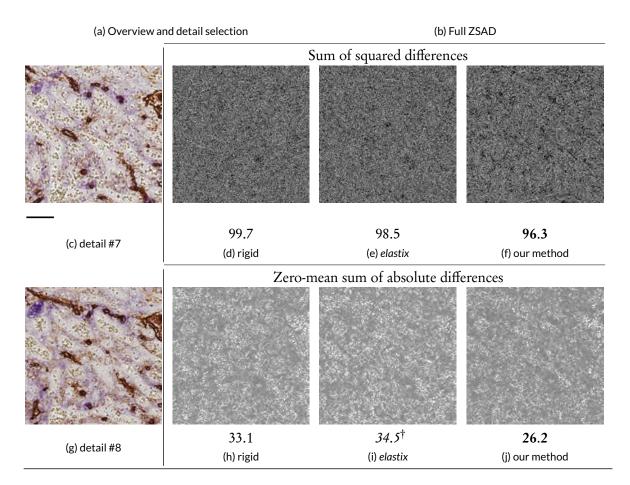
Figure B.15: Quality measures for double-stained spleen. *(Continued on the next page.)*

	Spleen, double-stained (cont.)		
input	rigid	*elastix*	our method
100%	107.08%	183.06%*	**195.37%**

(k) Mutual information increase

Gradient cross-correlation

1.95	1.90	*1.96*†	**1.71**
1.99	1.94	1.98	**1.77**
(l) input	(m) rigid	(n) *elastix*	(o) our method
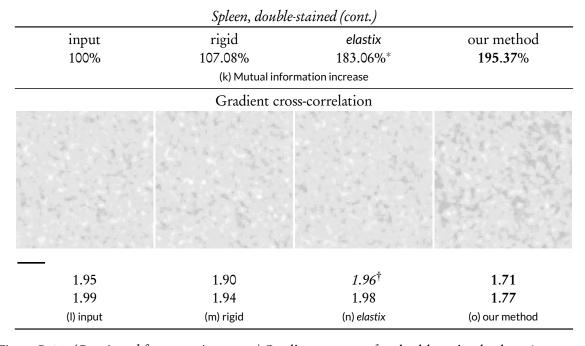
Figure B.16: (Continued from previous page.) Quality measures for double-stained spleen (same ROI as in Fig. M2.9.) (a) Detail selection. Scale bar = 500 µm. (b) Zero-mean sum of absolute differences (ZSAD) for full ROI. (See (h)–(j) for crops.) ZSAD images were enhanced with brightness and contrast +50. (c, g) The details of the first and second sections. Scale bar = 100 µm. (d)–(f) Sum of squared differences (SSD). (h)–(j) Zero-mean sum of absolute differences. (k) mutual information (MI) increase with original image pair as a baseline. Note that *elastix* (marked with asterisk) shows data for the blue channel only, but our method produces RGB images. (l)–(o) Combined cross-correlation of gradient approximation (GCC). Images are +6 gamma-corrected. We show mean values of full image-wide quality measures. For GCC we state both the mean of the derivative in the direction of the x axis (first line) and y axis (second line).
Aside from (k), the smaller the values and the darker the images the better. **Bold** shows the best value, *italics* with † denotes decrease from baseline, asterisk (∗) marks incomparable values.

this. In the direction of x (y) axis, Ulrich et al. (l) improved 10.6% (6.6%) over input overlay, our method (m) showed 17.2% (16.1%) improvement, and *elastix* (not shown) was 9.2% (7.6%) better than its single-channel input. Our method (m) was best in both directions.

If we simply summed up how often we obtained the best values for each quality measure, our method would clearly outperform the method of Ulrich et al.

B.3 Online materials

Further supplementary material, including input data, full-scale volume renderings, full-scale images of quantitative evaluations, Windows executable of our application, and *elastix* parameter files, is available under https://gdv-server.inf.uni-bayreuth.de/gdvcloud/index.php/s/NnSov0O65n9Gp01 and doi: 10.5281/zenodo.164404.

The supplementary video is also available under https://gdv-server.inf.uni-bayreuth.de/gdvcloud/index.php/s/sl3s5Ilyf4BrDe7.

Bone marrow, 2 levels, 3k

(a) Overview and detail selection   (b) Full ZSAD

Sum of squared differences

	49.1	50.5†	**48.5**
(c) detail #12	(d) rigid	(e) Ulrich et al.	(f) our method

Zero-mean sum of absolute differences

	12.3	12.6†	**11.3**
(g) detail #13	(h) rigid	(i) Ulrich et al.	(j) our method
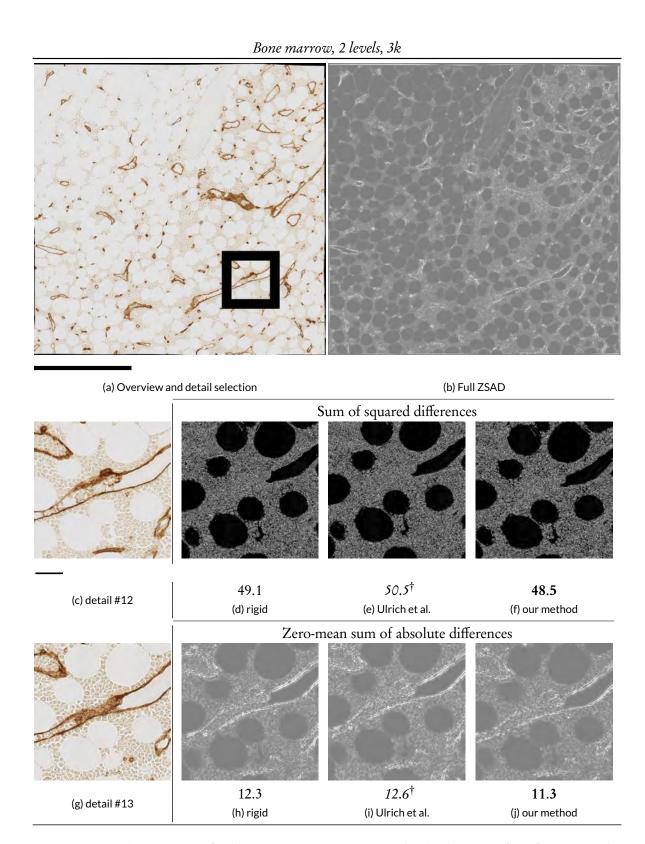
Figure B.17: Quality measures for bone marrow specimen, two levels. *(Continued on the next page.)*

Bone marrow, 2 levels, 3k, cont.

original	rigid	Ulrich et al.	our method
100%	97.52%†	75.46%†	91.32%†

(k) Mutual information increase

Gradient cross-correlation

1.36	1.33	1.36	**1.30**
1.38	1.42†	1.38	**1.34**
(l) original	(m) rigid	(n) Ulrich et al.	(o) our method

Figure B.17: (Continued from previous page.) Quality measures for bone marrow specimen, same ROI as in Fig. B.5. (a) The detail selection, scale bar = 500 µm. (b) Complete ZSAD quality measure. (See (h)–(j) for crops.) (c) The detail of the first section, scale bar = 100 µm. (d)–(f) Sum of squared differences (SSD). (g) The detail of the second section. (h)–(j) Zero-mean sum of absolute differences (ZSAD). We enhanced the brightness and contrast. (k) Mutual information (MI) increase from baseline to the method from the corresponding column. (l)–(o) Combined cross-correlation of gradient approximation (GCC). These images are enhanced with brightness and contrast +50.

We state the mean values for each image-wide quality measures. For GCC the first (second) line states the mean of derivatives in the direction of the x (second) axis.

Aside from (k), the smaller the values and the darker the images the better. **Bold** shows the best value, *italics* with † denotes decrease from baseline.

Bone marrow, 4 levels, 4k

(a) Overview and detail selection

(b) Full ZSAD

Sum of squared differences		
56.4	54.5†	**52.5**
(d) rigid	(e) Ulrich et al.	(f) our method

(c) detail #12

Zero-mean sum of absolute differences		
13.8	13.0	**10.9**
(h) rigid	(i) Ulrich et al.	(j) our method

(g) detail #13

Figure B.18: Quality measures for bone marrow specimen, four levels. *(Continued on the next page.)*

Bone marrow, 4 levels, 4k, cont.

original	rigid	Ulrich et al.	our method
100%	121.4%	**140.8%**	139.8%

(k) Mutual information increase

Gradient cross-correlation

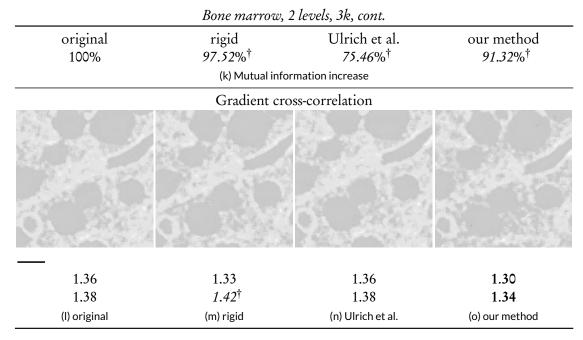| 1.51 | 1.49 | 1.36 | **1.29** |
| 1.50 | 1.43 | 1.41 | **1.30** |
| (l) original | (m) rigid | (n) Ulrich et al. | (o) our method |
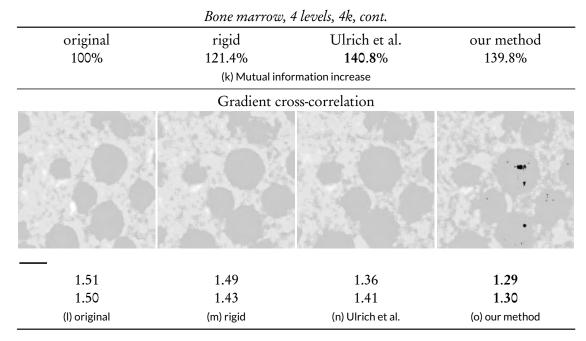
Figure B.18: (Continued from previous page.) Quality measures for bone marrow specimen, extended version of Fig. M2.16. (a) The detail selection. Scale bar $= 500\,\mu$m. (b) Complete ZSAD quality measure. (See (h–j) for crops.) (c) The detail of the first section, scale bar $= 100\,\mu$m. (d–f) sum of squared differences (SSD). (g) The detail of the second section. (h–j) Zero-mean sum of absolute differences (ZSAD). We enhanced the brightness and contrast. (k) Mutual information (MI) increase from baseline to the method from the corresponding column. (l–o) Combined cross-correlation of gradient approximation (GCC). These images are enhanced with brightness and contrast $+50$.

We state the mean values for each image-wide quality measures. For GCC the first (second) line states the mean of derivatives in the direction of the x (y) axis.

Aside from (k), the smaller the values and the darker the images the better. **Bold** shows the best value, *italics* with † denotes decrease from baseline.
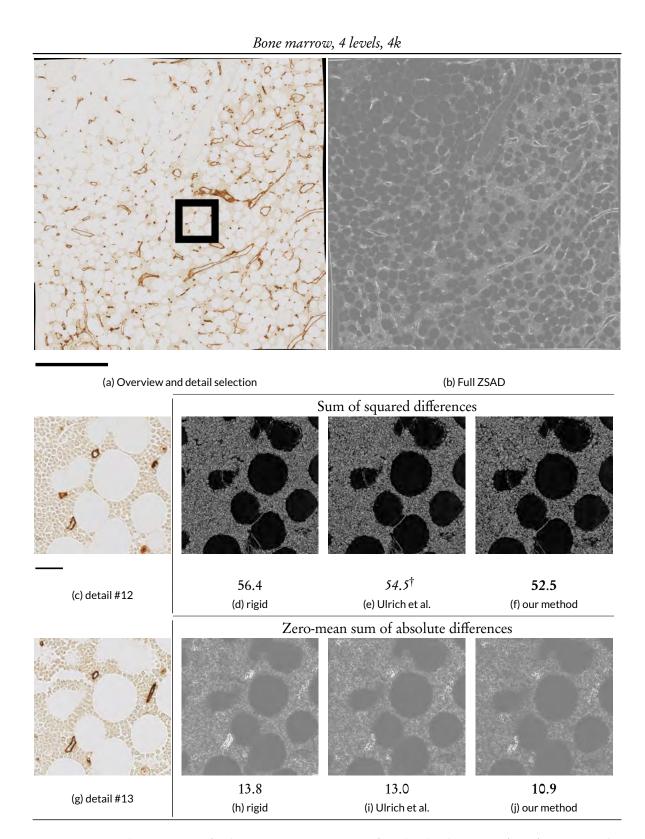
THREE-DIMENSIONAL ARRANGEMENT OF HUMAN BONE MARROW MICROVESSELS FROM IMMUNOHISTOLOGY IN UNDECALCIFIED SERIAL SECTIONS. SUPPLEMENTARY MATERIAL

Which supplements Chapter 6 by showing stills for videos mentioned in main text.
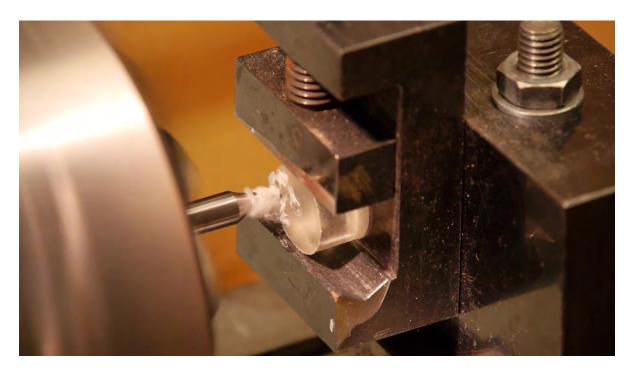
Figure C.1: Preview of Video 6.1. Methods of specimen processing and staining for serial Technovit® 9100 sections, video under doi: 10.1371/journal.pone.0168173.s020. High-resolution video deposited under doi: 10.5281/zenodo.141383.
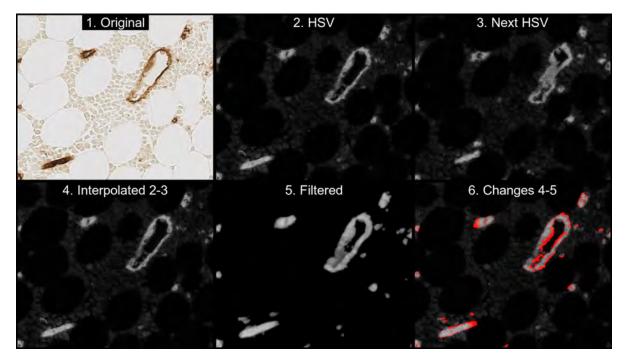
Figure C.2: Preview of Video 6.2. Sequence of QC in the small region depicted in Fig. 6.6a,b. The remaining 13 videos are found in the repository (doi: 10.5281/zenodo.128861). For each small region, the following processing steps are visualised from left to right in the upper and lower part of the panel: original registered data (1), HSV data conversion (2), HSV data conversion of next serial section (3), interpolations between both serial sections resulting in 7 images (4), final result of all volume filtering operations of interpolated images (5), areas differing between interpolated images and the result of filtering visualised in red colour (6).
Video is available under doi: 10.1371/journal.pone.0168173.s021. High-resolution video deposited under doi: 10.5281/zenodo.128861.

CAPILLARY NETWORKS AND FOLLICULAR MARGINAL ZONES IN HUMAN SPLEENS. THREE-DIMENSIONAL MODELS BASED ON IMMUNOSTAINED SERIAL SECTIONS. SUPPLEMENTARY MATERIAL

In which we support the claims of the main text (Chapter 7) and provide an overview of accompanying videos and further supplementary material.

SUPPLEMENTARY material to the main chapter is available in the Zenodo repository (doi: 10.5281/zenodo.845670). The material includes the sequence videos of single (Video 7.1a, 7.2a, 7.3a) and double (Video 7.1b, 7.2b, 7.3b) stainings and the 3D models of single (Video 7.1c, 7.2c, 7.3c) and double stainings (Video 7.1d, 7.2d, 7.3d). Video 7.4 shows quality control using virtual reality. The settings for colour deconvolution are also included in the repository.

Figure D.1 shows a cryosection of a tonsil that is used to support our line of though on page 139. Videos 7.1a,b–7.3a,b page through the section series. We have omitted their previews here. Figures D.2–D.7 show single frames from videos of 3D reconstructions from single (c) and double (d) stainings from ROI 1–3 respectively. Figure D.8 showcases the quality control process for these reconstructions. We detail on it in Chapter 5.
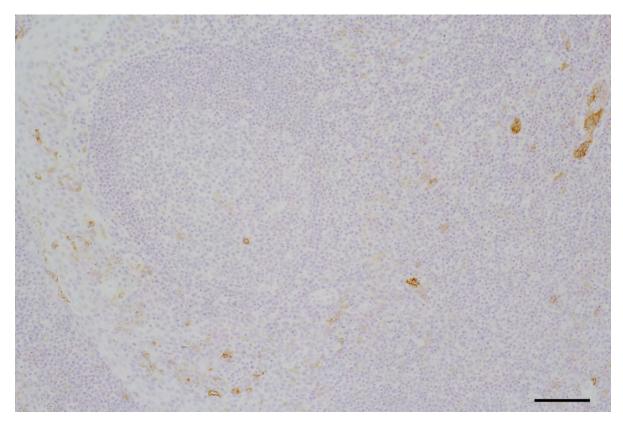
Figure D.1: A secondary follicle in a tonsil. *(Caption continued on the next page.)*

Figure D.1: (Caption for the figure on the previous page.) Cryosection of a secondary follicle in a tonsil stained for MAdCAM-1 (mAb 10A6). Very few faintly stained dendritic-like cells are observed in the tonsillar epithelium (upper left corner) and in the surrounding connective tissue. The specificity of this reaction is not clear. A minority of high endothelial venules (right upper corner) has MAdCAM-1$^+$ endothelium. Perifollicular staining is absent. 12-year-old female patient. The result is representative of 6 tonsils from adults and children removed because of recurrent tonsillitis.

Hemalum counterstain, scale bar = 100 μm.
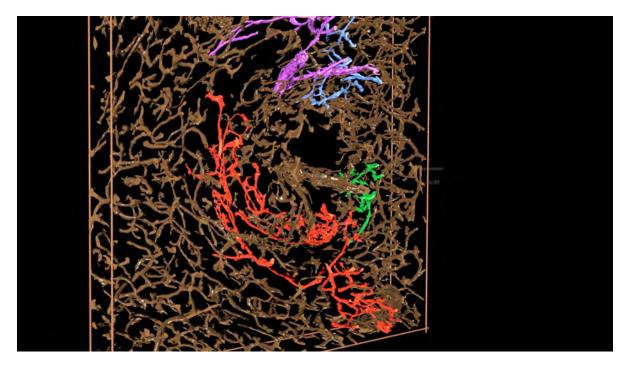
Figure D.2: A frame from Video 7.1c.
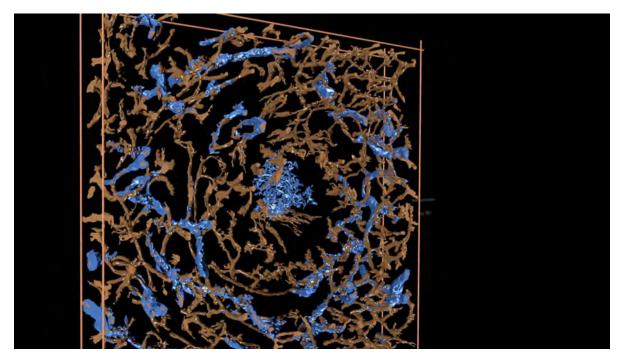
Figure D.3: A frame from Video 7.1d.
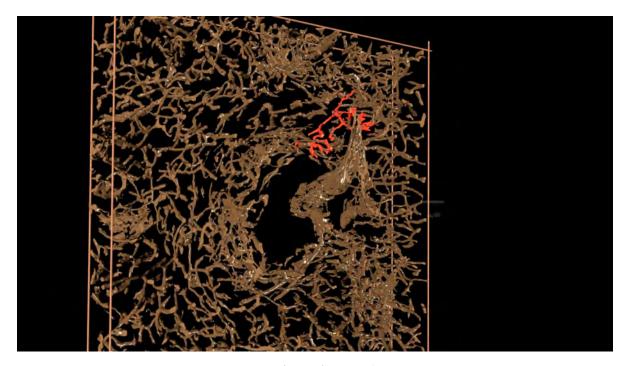
Figure D.4: A frame from Video 7.2c.
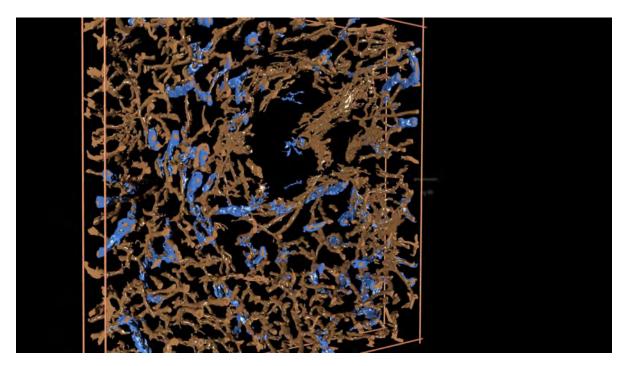
Figure D.5: A frame from Video 7.2d.
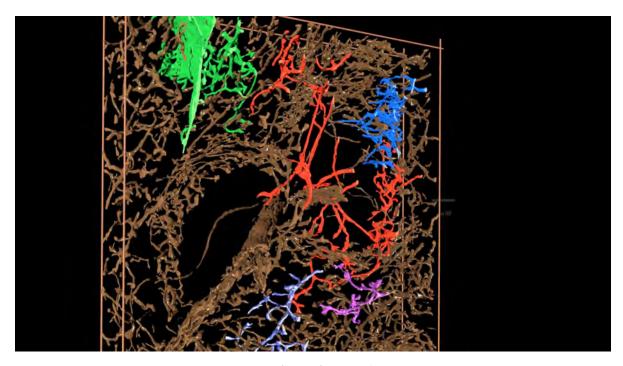
Figure D.6: A frame from Video 7.3c.
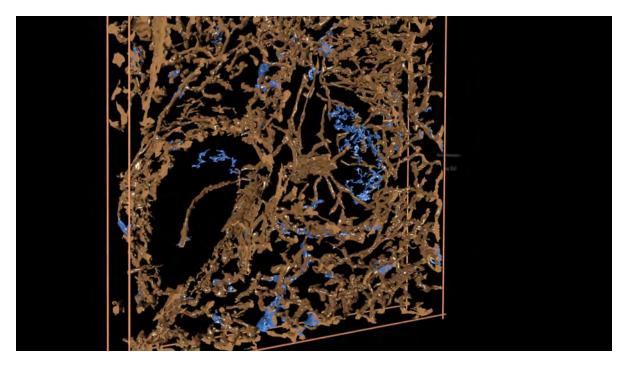
Figure D.7: A frame from Video 7.3d.

(a)

(b)
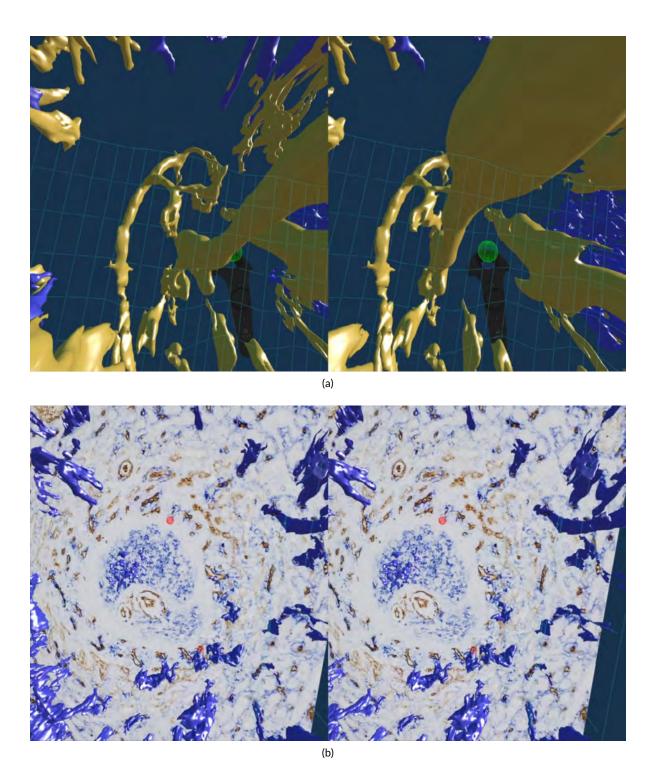
Figure D.8: Two frames from Video 7.4.

BIBLIOGRAPHY

> Archimedes will be remembered
> when Aeschylus is forgotten, be-
> cause languages die and mathemat-
> ical ideas do not.
>
> G. H. Hardy, *A Mathematician's*
> *Apology*

M. Acar, K. S. Kocherlakota, M. M. Murphy, J. G. Peyer, H. Oguro, C. N. Inra, C. Jaiyeola, Z. Zhao, K. Luby-Phelps, and S. J. Morrison. Deep imaging of bone marrow shows non-dividing stem cells are mainly perisinusoidal. *Nature*, 526(7571):126–130, 2015. 99

A. Alahi, R. Ortiz, and P. Vandergheynst. FREAK: Fast retina keypoint. In *IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '12, pages 510–517. IEEE, 2012. doi: 10.1109/CVPR.2012.6247715. 7, 25

P. F. Alcantarilla, A. Bartoli, and A. J. Davison. KAZE features. In A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *Computer Vision – ECCV 2012*, LNCS 7577, pages 214–227. Springer, 2012. ISBN 978-3-642-33782-6. doi: 10.1007/978-3-642-33783-3_16. 25

H. Alt and L. J. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. *Handbook of computational geometry*, 1:121–153, 1999. 121

B. Apolloni, C. Carvalho, and D. de Falco. Quantum stochastic optimization. *Stochastic Processes and their Applications*, 33(2):233–244, 1989. ISSN 0304-4149. doi: 10.1016/0304-4149(89)90040-9. URL http://www.sciencedirect.com/science/article/pii/0304414989900409. 151

T. I. Arnon, R. M. Horton, I. L. Grigorova, and J. G. Cyster. Visualization of splenic marginal zone B cell shuttling and follicular B cell egress. *Nature*, 493(7434):684–688, 01 2013. doi: 10.1038/nature11738. 134

D. Attali, D. Cohen-Steiner, and H. Edelsbrunner. Extraction and simplification of iso-surfaces in tandem. In *Proceedings of the 3rd Eurographics Symposium on Geometry Processing*, SGP '05. Eurographics, 2005. 68, 71, 72, 75, 79, 80, 81

M. Auer, P. Regitnig, and G. Holzapfel. An automatic nonrigid registration for stained histological sections. *IEEE Transactions on Image Processing*, 14(4):475–486, 2005. ISSN 1057-7149. doi: 10.1109/TIP.2005.843756. 7

J. Awni. Cryptographic key visualization, Mar. 2 2017. US Patent App. 14/837,652. Publication # US20170061199 A1. 168

231

U. Bagci and L. Bai. Automatic best reference slice selection for smooth volume reconstruction of a mouse brain from histological images. *IEEE Transactions on Medical Imaging*, 29(9):1688–1696, 2010. ISSN 0278-0062. doi: 10.1109/TMI.2010.2050594. 7, 26

U. Bagci, X. Chen, and J. Udupa. Hierarchical scale-based multiobject recognition of 3-D anatomical structures. *IEEE Transactions on Medical Imaging*, 31(3):777–789, 2012. ISSN 0278-0062. doi: 10.1109/TMI.2011.2180920. 6, 23, 26

D. Bagnara, M. Squillario, D. Kipling, T. Mora, A. M. Walczak, L. Da Silva, S. Weller, D. K. Dunn-Walters, J.-C. Weill, and C.-A. Reynaud. A reassessment of IgM memory subsets in humans. *The Journal of Immunology*, 195(8):3716–3724, 2015. ISSN 0022-1767. doi: 10.4049/jimmunol.1500753. URL http://www.jimmunol.org/content/195/8/3716. 126, 134, 136

R. Bajcsy and S. Kovačič. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46(1):1–21, 1989. ISSN 0734-189X. doi: 10.1016/S0734-189X(89)80014-3. 6

S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision*, 92(1):1–31, 2011. ISSN 1573-1405. doi: 10.1007/s11263-010-0390-2. 54

J. Balzer and S. Werling. Principles of shape from specular reflection. *Measurement*, 43(10): 1305–1317, 2010. ISSN 0263-2241. doi: 10.1016/j.measurement.2010.07.013. URL http://www.sciencedirect.com/science/article/pii/S0263224110001570. 151

J. Balzer, S. Höfer, and J. Beyerer. Multiview specular stereo reconstruction of large mirror surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, CVPR 2011, pages 2537–2544. IEEE, 2011. doi: 10.1109/CVPR.2011.5995346. 150, 151

M. Balázs, G. Horváth, L. Grama, and P. Balogh. Phenotypic identification and development of distinct microvascular compartments in the postnatal mouse spleen. *Cellular Immunology*, 212(2): 126–137, 2001. ISSN 0008-8749. doi: http://dx.doi.org/10.1006/cimm.2001.1847. 125

J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994. ISSN 1573-1405. doi: 10.1007/BF01420984. 56

U. Bağci and L. Bai. Registration of standardized histological images in feature space. In *Proc. SPIE*, volume 6914, pages 69142V–69142V–9, 2008. doi: 10.1117/12.770219. 27, 56

H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool. Speeded-up robust features (SURF). *Comput. Vision Image Understanding*, 110(3):346–359, 2008. ISSN 1077-3142. doi: 10.1016/j.cviu.2007.09.014. 7, 24, 29, 31

P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. ISSN 0162-8828. doi: 10.1109/34.121791. 6, 14, 25

E. A. Bier and K. R. Sloan. Two-part texture mappings. *IEEE Computer Graphics and Applications*, 6 (9):40–53, 1986. 155, 156

E. Boveri, F. Passamonti, E. Rumi, D. Pietra, C. Elena, L. Arcaini, C. Pascutto, A. Castello, M. Cazzola, U. Magrini, et al. Bone marrow microvessel density in chronic myeloproliferative disorders: a study of 115 patients with clinicopathological and molecular correlations. *British Journal of Haematology*, 140(2):162–168, 2008. 100

P. Bovet and S. Benhamou. Spatial analysis of animals' movements using a correlated random walk model. *Journal of Theoretical Biology*, 131(4):419–433, 1988. ISSN 0022-5193. doi: 10.1016/S0022-5193(88)80038-9. 151

G. Bradski and A. Kaehler. *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly, 2008. 13, 14, 57, 119, 142

F. P. Brooks Jr. What's real about virtual reality? *IEEE Computer Graphics and Applications*, 19(6): 16–27, 1999. ISSN 0272-1716. doi: 10.1109/38.799723. 84

L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992. ISSN 0360-0300. doi: 10.1145/146370.146374. 7, 23

T. Brox and J. Malik. Large displacement optical flow: Descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):500–513, 2011. ISSN 0162-8828. doi: 10.1109/TPAMI.2010.143. 54

T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European Conference on Computer Vision – ECCV '04*, LNCS 3024, pages 25–36. Springer, 2004. URL http://lmb.informatik.uni-freiburg.de//Publications/2004/Bro04a. 54

BSD General Commands Manual. *Manual page for ssh – OpenSSH SSH client*. Apple, 2017. 166, 168, 177

W. Burger and M. J. Burge. *Principles of digital image processing*. Springer, 2009. 54

M. Burns, M. Haidacher, W. Wein, I. Viola, and M. E. Gröller. Feature emphasis and contextual cutaways for multimodal medical visualization. In *Proceedings of the 9th Joint Eurographics / IEEE VGTC Conference on Visualization*, EUROVIS '07, pages 275–282. The Eurographics Association, 2007. ISBN 978-3-905673-45-6. doi: 10.2312/VisSym/EuroVis07/275-282. 96

A. Can, C. Stewart, B. Roysam, and H. Tanenbaum. A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):347–364, 2002. ISSN 0162-8828. doi: 10.1109/34.990136. 24, 26

L. Carata, D. Shao, M. Hadwiger, and E. Groeller. Improving the visualization of electron-microscopy data through optical flow interpolation. In *Proceedings of the 27th Spring Conference on Computer Graphics*, SCCG '11, pages 103–110. ACM, 2013. ISBN 978-1-4503-1978-2. doi: 10.1145/2461217.2461238. 56, 135

A. Cardona, S. Saalfeld, S. Preibisch, B. Schmid, A. Cheng, J. Pulokas, P. Tomancak, and V. Hartenstein. An integrated micro- and macroarchitectural analysis of the drosophila brain by computer-assisted serial section electron microscopy. *PLOS Biology*, 8(10):1–17, 10 2010. doi: 10.1371/journal.pbio.1000502. 55

A. Cardona, S. Saalfeld, J. Schindelin, I. Arganda-Carreras, S. Preibisch, M. Longair, P. Tomancak, V. Hartenstein, and R. J. Douglas. TrakEM2 software for neural circuit reconstruction. *PLOS ONE*, 7(6):1–8, 06 2012. doi: 10.1371/journal.pone.0038011. 135

A. Cariappa, I. B. Mazo, C. Chase, H. N. Shi, H. Liu, Q. Li, H. Rose, H. Leung, B. J. Cherayil, P. Russell, et al. Perisinusoidal B cells in the bone marrow participate in T-independent responses to blood-borne microbes. *Immunity*, 23(4):397–407, 2005. 99

P. Cassoni, L. Gaetano, R. Senetta, B. Bussolati, L. Molinaro, and G. Bussolati. Histology far away from flatland: 3D roller-coasting into grade-dependent angiogenetic patterns in oligodendrogliomas. *Journal of cellular and molecular medicine*, 12(2):564–568, 2008. ISSN 1582-4934. doi: 10.1111/j. 1582-4934.2007.00206.x. 112

C. L. Celso, H. E. Fleming, J. W. Wu, C. X. Zhao, S. Miake-Lye, J. Fujisaki, D. Côté, D. W. Rowe, C. P. Lin, and D. T. Scadden. Live-animal tracking of individual haematopoietic stem/progenitor cells in their niche. *Nature*, 457(7225):92–96, 2009. 99, 100

I. Cervesato, N. A. Durgin, P. D. Lincoln, J. C. Mitchell, and A. Scedrov. A meta-notation for protocol analysis. In *Proceedings of the 12th IEEE Computer Security Foundations Workshop*, pages 55–69, 1999. doi: 10.1109/CSFW.1999.779762. 167

S. Chan, F. Conti, K. Salisbury, and N. H. Blevins. Virtual reality simulation in neurosurgery: Technologies and evolution. *Neurosurgery*, 72(suppl_1):A154–A164, 2013. doi: 10.1227/NEU. 0b013e3182750d26. 84

S. Chandrasekhar. Stochastic problems in physics and astronomy. *Reviews of Modern Physics*, 15:1–89, 1943. doi: 10.1103/RevModPhys.15.1. URL http://link.aps.org/doi/10.1103/RevModPhys.15.1. 151

J. Chappelow, B. N. Bloch, N. Rofsky, E. Genega, R. Lenkinski, W. DeWolf, and A. Madabhushi. Elastic registration of multimodal prostate MRI and histology via multiattribute combined mutual information. *Medical Physics*, 38(4):2005–2018, 2011. doi: 10.1118/1.3560879. 6, 26, 56

X. Chen, L. Xu, Y. Wang, H. Wang, F. Wang, X. Zeng, Q. Wang, and J. Egger. Development of a surgical navigation system based on augmented reality using an optical see-through head-mounted display. *Journal of Biomedical Informatics*, 55:124–131, 2015. ISSN 1532-0464. doi: 10.1016/j.jbi.2015.04.003. 84

Y.-M. Cheng and C.-M. Wang. A high-capacity steganographic approach for 3D polygonal meshes. *The Visual Computer*, 22(9):845–855, Sep 2006. ISSN 1432-2315. doi: 10.1007/s00371-006-0069-4. 170

Y.-M. Cheng and C.-M. Wang. An adaptive steganographic algorithm for 3D polygonal meshes. *The Visual Computer*, 23(9):721–732, Sep 2007. ISSN 1432-2315. doi: 10.1007/s00371-007-0147-2. 170

E. V. Chernyaev. Marching cubes 33: Construction of topologically correct isosurfaces. Technical Report CN 95-17, CERN, 1995. URL http://cds.cern.ch/record/292771/files/cn-95-017.pdf. Retrieved 2017-08-29. 68

W. Cheung and G. Hamarneh. *n*-SIFT: *n*-dimensional scale invariant feature transform. *IEEE Transactions on Image Processing*, 18(9):2012–2021, 2009. ISSN 1057-7149. doi: 10.1109/TIP.2009. 2024578. 7, 25

K. S. Choi, S. T. Chan, C. H. M. Leung, and Y. P. Chui. Stereoscopic three-dimensional visualization for immersive and intuitive anatomy learning. In *IEEE 8ᵗʰ International Conference on Technology for Education*, T4E '16, pages 184–187. IEEE, 2016. doi: 10.1109/T4E.2016.046. 84

H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2–3):114–141, 2003. ISSN 1077-3142. doi: 10.1016/S1077-3142(03)00009-2. 6, 14

A. Cifor, L. Bai, and A. Pitiot. Smoothness-guided 3-D reconstruction of 2-D histological images. *NeuroImage*, 56(1):197–211, 2011. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2011.01.060. 6, 25, 56

A. Cifor, L. Risser, D. Chung, E. M. Anderson, and J. A. Schnabel. Hybrid feature-based diffeomorphic registration for tumor tracking in 2-D liver ultrasound images. *IEEE Transactions on Medical Imaging*, 32(9):1647–1656, 2013. ISSN 0278-0062. doi: 10.1109/TMI.2013.2262055. 27

P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia. Meshlab: an open-source mesh processing tool. In *Eurographics Italian Chapter Conference*, volume 2008, pages 129–136, 2008. xxvii, 120, 141

G. Cinamon, M. A. Zachariah, O. M. Lam, F. W. Foss, and J. G. Cyster. Follicular shuttling of marginal zone B cells facilitates antigen transport. *Nature immunology*, 9(1):54–62, 01 2008. doi: 10.1038/ni1542. 134

S. G. Cohen, T. Itkin, S. Chakrabarty, C. Graf, O. Kollet, A. Ludin, K. Golan, A. Kalinkovich, G. Ledergor, E. Wong, et al. PAR1 signaling regulates the retention and recruitment of EPCR-expressing bone marrow hematopoietic stem cells. *Nature Medicine*, 21(11):1307–1317, 2015. 99

G. Conti, J. Grizzard, M. Ahamad, and H. Owen. Visual exploration of malicious network objects using semantic zoom, interactive encoding and dynamic queries. In *IEEE Workshop on Visualization for Computer Security*, VizSEC '05, pages 83–90, 2005. doi: 10.1109/VIZSEC.2005.1532069. 169

I. Cox, M. Miller, J. Bloom, J. Fridrich, and T. Kalker. *Digital watermarking and steganography*. Morgan Kaufmann, 2007. 170

W. R. Crum, T. Hartkens, and D. L. G. Hill. Non-rigid image registration: theory and practice. *The British Journal of Radiology*, 77(suppl_2):S140–S153, 2004. doi: 10.1259/bjr/25329214. 25

G. Cuccuru, E. Gobbetti, F. Marton, R. Pajarola, and R. Pintus. Fast low-memory streaming MLS reconstruction of point-sampled surfaces. In *Graphics Interface*, pages 15–22, May 2009. URL http://vic.crs4.it/vic/cgi-bin/bib-page.cgi?id='Cuccuru:2009:FLM'. 71

P. M. Dammers, N. K. de Boer, G. J. Deenen, P. Nieuwenhuis, and F. G. M. Kroese. The origin of marginal zone B cells in the rat. *European Journal of Immunology*, 29(5):1522–1531, 1999. ISSN 1521-4141. doi: 10.1002/(SICI)1521-4141(199905)29:05<1522::AID-IMMU1522>3.0.CO;2-0. 135

C. DeCoro and N. Tatarchuk. Real-time mesh simplification using the GPU. In *Symposium on Interactive 3D Graphics and Games*, I3D '07, pages 161–166. ACM, 2007. 71

B. Denkena, H. Ahlers, F. Berg, T. Wolf, and H. K. Tönshoff. Fast inspection of larger sized curved surfaces by stripe projection. *CIRP Annals—Manufacturing Technology*, 51(1):499–502, 2002. 150

E. Derzapf and M. Guthe. Dependency-free parallel progressive meshes. *Computer Graphics Forum*, 31(8):2288–2302, 2012. ISSN 1467-8659. doi: 10.1111/j.1467-8659.2012.03154.x. 188

R. Dhamija and A. Perrig. Déjà vu: A user study. Using images for authentication. In *USENIX Security Symposium*, volume 9, pages 4–4, 2000. 169

C. D. Dijkstra, E. A. Döpp, P. Joling, and G. Kraal. The heterogeneity of mononuclear phagocytes in lymphoid organs: Distinct macrophage subpopulations in rat recognized by monoclonal antibodies ED1, ED2 and ED3. In G. G. B. Klaus, editor, *Microenvironments in the Lymphoid System*, pages 409–419. Springer, 1985. ISBN 978-1-4613-2463-8. doi: 10.1007/978-1-4613-2463-8_50. 126

C. E. Duchon. Lanczos filtering in one and two dimensions. *Journal of Applied Meteorology*, 18(8): 1016–1022, 1979. 54

G. Dupuy, B. Jobard, S. Guillon, N. Keskes, and D. Komatitsch. Parallel extraction and simplification of large isosurfaces using an extended tandem algorithm. *Computer-Aided Design*, 42(2):129–138, 2010. doi: 10.1016/j.cad.2009.04.016. 71, 80

C. Dyken, G. Ziegler, C. Theobalt, and H.-P. Seidel. High-speed marching cubes using histopyramids. *Computer Graphics Forum*, 27(8):2028–2039, 2008. ISSN 1467-8659. doi: 10.1111/j.1467-8659. 2008.01182.x. 70

J. Egger, M. Gall, J. Wallner, P. Boechat, A. Hann, X. Li, X. Chen, and D. Schmalstieg. HTC Vive MeVisLab integration via OpenVR for medical applications. *PLOS ONE*, 12(3):1–14, 03 2017. doi: 10.1371/journal.pone.0173972. 85, 135

J. Ehrhardt, R. Werner, D. Säring, T. Frenzel, W. Lu, D. Low, and H. Handels. An optical flow based method for improved reconstruction of 4D CT data sets acquired during free breathing. *Medical Physics*, 34(2):711–721, 2007. ISSN 2473-4209. doi: 10.1118/1.2431245. 56

K. W. Eliceiri, M. R. Berthold, I. G. Goldberg, L. Ibáñez, B. S. Manjunath, M. E. Martone, R. F. Murphy, H. Peng, A. L. Plant, B. Roysam, et al. Biological imaging software tools. *Nature Methods*, 9(7):697–710, 2012. 5, 24

M. Endler, O. Lobachev, and M. Guthe. Human-computer interaction using robust gesture recognition. In V. Skala, editor, *WSCG Posters proceedings*, WSCG '14, pages 45–48, 2014. ISBN 978-80-86943-72-5.

U. M. Erdem and S. Sclaroff. Automated camera layout to satisfy task-specific and floor plan-specific coverage requirements. *Computer Vision and Image Understanding*, 103(3):156–169, 2006. ISSN 1077-3142. doi: 10.1016/j.cviu.2006.06.005. Special issue on Omnidirectional Vision and Camera Networks. 151

G. Farnebäck. Two-frame motion estimation based on polynomial expansion. In J. Bigun and T. Gustavsson, editors, *Image Analysis*, LNCS 2749, pages 363–370. Springer, 2003. ISBN 978-3-540-40601-3. doi: 10.1007/3-540-45103-X_50. xxi, 39, 44, 54, 56, 57, 119, 135, 141

N. Fasnacht, H.-Y. Huang, U. Koch, S. Favre, F. Auderset, Q. Chai, L. Onder, S. Kallert, D. D. Pinschewer, H. R. MacDonald, F. Tacchini-Cottier, B. Ludewig, S. A. Luther, and F. Radtke. Specific fibroblastic niches in secondary lymphoid organs orchestrate distinct notch-regulated immune responses. *Journal of Experimental Medicine*, 211(11):2265–2279, 2014. ISSN 0022-1007. doi: 10.1084/jem.20132528. 139

O. D. Faugeras and R. Keriven. Complete dense stereovision using level set methods. In H. Burkhardt and B. Neumann, editors, *Computer Vision*, ECCV '98, pages 379–393. Springer, 1998. ISBN 978-3-540-69354-3. doi: 10.1007/BFb0055679. 150

A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, et al. 3D slicer as an image computing platform for the quantitative imaging network. *Magnetic resonance imaging*, 30(9):1323–1341, 2012. xxi, xxvii, 59, 120, 141

A. R. Ferguson, M. E. Youd, and R. B. Corley. Marginal zone B cells transport and deposit IgM-containing immune complexes onto follicular dendritic cells. *International Immunology*, 16 (10):1411–1422, 2004. doi: 10.1093/intimm/dxh142. 134

A. Finnila, M. Gomez, C. Sebenik, C. Stenson, and J. Doll. Quantum annealing: A new method for minimizing multidimensional functions. *Chemical Physics Letters*, 219(5):343–348, 1994. ISSN 0009-2614. doi: 10.1016/0009-2614(94)00117-0. 151

M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. ISSN 0001-0782. doi: 10.1145/358669.358692. 7, 9, 25, 29

S. Fleishman, D. Cohen-Or, and D. Lischinski. Automatic camera placement for image-based modelling. *Comput. Graph. Forum*, 5:231–239, 2000. 151

D. Fong and M. Saunders. LSMR: An iterative algorithm for sparse least-squares problems. *SIAM Journal on Scientific Computing*, 33(5):2950–2971, 2011. doi: 10.1137/10079687X. 28, 31, 39

A. S. Forsberg, D. H. Laidlaw, A. van Dam, R. M. Kirby, G. E. Karniadakis, and J. L. Elion. Immersive virtual reality for visualizing flow through an artery. In *Visualization*, VIS '00, pages 457–460. IEEE Computer Society Press, 2000. ISBN 1-58113-309-X. 84

N. Friedrich, O. Lobachev, and M. Guthe. Faking it: Simulating background blur in portrait photography using a coarse depth map estimation from a single image. In V. Skala, editor, *WSCG Short papers proceedings*, volume CSRN 2602 of *WSCG '16*, pages 15–24, 2016. ISBN 978-80-86943-58-9.

K. J. Friston, J. Ashburner, C. D. Frith, J.-B. Poline, J. D. Heather, and R. S. J. Frackowiak. Spatial registration and normalization of images. *Human Brain Mapping*, 3(3):165–189, 1995. ISSN 1097-0193. doi: 10.1002/hbm.460030303. 7, 23

M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 209–216, 1997. doi: 10.1145/258734.258849. 71, 72

M. Garland and P. S. Heckbert. Simplifying surfaces with color and texture using quadric error metrics. In *Proceedings of the Conference on Visualization '98*, pages 263–269, 1998. doi: 10.1109/VISUAL.1998.745312. 71

S. Gefen, O. Tretiak, and J. Nissanov. Elastic 3-D alignment of rat brain histological images. *IEEE Transactions on Medical Imaging*, 22(11):1480–1489, 2003. ISSN 0278-0062. doi: 10.1109/TMI.2003.819280. 6

J. M. Gijtenbeek, P. Wesseling, C. Maass, L. Burgers, and J. A. van der Laak. Three-dimensional reconstruction of tumor microvasculature: simultaneous visualization of multiple components in paraffin-embedded tissue. *Angiogenesis*, 8(4):297–305, 2005. 6, 25, 55, 112

H. J. Gilhuis, J. A. van der Laak, J. Pomp, A. C. Kappelle, J. M. Gijtenbeek, and P. Wesseling. Three-dimensional (3D) reconstruction and quantitative analysis of the microvasculature in medulloblastoma and ependymoma subtypes. *Angiogenesis*, 9(4):201–208, 2006. 6, 25, 112

GPG2. *Manual page for gpg2 – OpenPGP encryption and signing tool*. GNU Privacy Guard, 2016. 166

A. A. Goshtasby. *2-D and 3-D image registration: for medical, remote sensing, and industrial applications*. Wiley, 2005. 5

L. Grady. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, 2006. 151

D. Gray, I. C. M. MacLennan, H. Bazin, and M. Khan. Migrant $\mu^+\delta^+$ and static $\mu^+\delta^-$ B lymphocyte subsets. *European Journal of Immunology*, 12(7):564–569, 1982. ISSN 1521-4141. doi: 10.1002/eji.1830120707. 126, 134

D. Gray, D. S. Kumararatne, J. Lortan, M. Khan, and I. C. MacLennan. Relation of intra-splenic migration of marginal zone B cells to antigen localization on follicular dendritic cells. *Immunology*, 52(4):659–669, 08 1984. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1454654/. 134

N. Grund, E. Derzapf, and M. Guthe. Instant level-of-detail. In P. Eisert, J. Hornegger, and K. Polthier, editors, *Vision, Modeling, and Visualization*, VMV '11. EG, 2011. ISBN 978-3-905673-85-2. doi: 10.2312/PE/VMV/VMV11/293-299. 71, 75, 81

E. Guest, E. Berry, R. Baldock, M. Fidrich, and M. Smith. Robust point correspondence applied to two- and three-dimensional image registration. *IEEE T. Pattern Anal. Mach. Intell.*, 23(2):165–179, 2001. ISSN 0162-8828. doi: 10.1109/34.908967. 6, 25, 56

S. Guthe and M. Goesele. Variable length coding for GPU-based direct volume rendering. In M. Hullin, M. Stamminger, and T. Weinkauf, editors, *Vision, Modeling & Visualization*. EG, 2016. ISBN 978-3-03868-025-3. doi: 10.2312/vmv.20161345. 96

S. Guthe, M. Wand, J. Gonser, and W. Strasser. Interactive rendering of large volume data sets. In *IEEE Visualization*, VIS '02, pages 53–60, 2002. doi: 10.1109/VISUAL.2002.1183757. 96

J. V. Hajnal and D. L. Hill. *Medical image registration*. CRC, 2010. 5

N. Halper and P. Olivier. Camplan: A camera planning agent. In *Smart Graphics 2000 AAAI Spring Symposium*, pages 92–100, 2000. 151

H. Hammad, M. Vanderkerken, P. Pouliot, K. Deswarte, W. Toussaint, K. Vergote, L. Vandersarren, S. Janssens, I. Ramou, S. N. Savvides, J. J. Haigh, R. Hendriks, M. Kopf, K. Craessaerts, B. de Strooper, J. F. Kearney, D. H. Conrad, and B. N. Lambrecht. Transitional B cells commit to marginal zone B cell fate by Taok3-mediated surface expression of ADAM10. *Nature Immunology*, 18(3):313–320, 03 2017. doi: 10.1038/ni.3657. 126

X. Han. Feature-constrained nonlinear registration of lung CT images. *Medical image analysis for the clinic: a grand challenge*, pages 63–72, 2010. 24, 27

W. Heiden, T. Goetze, and J. Brickmann. Fast generation of molecular surfaces from 3D data fields with an enhanced "marching cubes" algorithm. *Journal of Computational Chemistry*, 14(2):246–250, 1993. 68

J. Heinly, E. Dunn, and J.-M. Frahm. Comparative evaluation of binary features. In A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *Computer Vision – ECCV 2012*, LNCS 7573, pages 759–773. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-33708-6. doi: 10.1007/978-3-642-33709-3_54. 7, 25

D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46(3):R1–R45, 2001. 5

J. L. Hintze and R. D. Nelson. Violin plots: A box plot-density trace synergism. *The American Statistician*, 52(2):181–184, 1998. doi: 10.1080/00031305.1998.10480559. xxi, xxii

D. Hong, H. Park, and H. Cho. Design of a multi-screen deflectometer for shape measurement of solder joints on a PCB. In *IEEE International Symposium on Industrial Electronics*, ISIE '09, pages 127–132. IEEE, 2009. doi: 10.1109/ISIE.2009.5213206. 150, 151

A. T. Hooper, J. M. Butler, D. J. Nolan, A. Kranz, K. Iida, M. Kobayashi, H.-G. Kopp, K. Shido, I. Petit, K. Yanger, et al. Engraftment and reconstitution of hematopoiesis is dependent on VEGFR2-mediated regeneration of sinusoidal endothelial cells. *Cell stem cell*, 4(3):263–274, 2009. 99, 100

H. Hoppe. Progressive meshes. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pages 99–108. ACM, 1996. ISBN 0-89791-746-4. doi: 10.1145/237170.237216. 188

B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1):185–203, 1981. ISSN 0004-3702. doi: 10.1016/0004-3702(81)90024-2. 54

Y.-C. Hou. Visual cryptography for color images. *Pattern Recognition*, 36(7):1619–1629, 2003. ISSN 0031-3203. doi: 10.1016/S0031-3203(02)00258-3. 169

I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. Transparent and specular object reconstruction. *Computer Graphics Forum*, 29(8):2400–2426, 2010. ISSN 1467-8659. doi: 10.1111/j.1467-8659.2010.01753.x. 150

A. Inoue, Y. Ikeda, K. Yatabe, and Y. Oikawa. Three-dimensional sound-field visualization system using head mounted display and stereo camera. *Proceedings of Meetings on Acoustics*, 29(1):025001–1–13, 2016. doi: 10.1121/2.0000381. 84

A. Iones, A. Krupkin, M. Sbert, and S. Zhukov. Fast, realistic lighting for video games. *IEEE Computer Graphics and Applications*, 23(3):54–64, 2003. 120

T. Itkin, S. Gur-Cohen, J. A. Spencer, A. Schajnovitz, S. K. Ramasamy, A. P. Kusumbe, G. Ledergor, Y. Jung, I. Milo, M. G. Poulos, et al. Distinct bone marrow blood vessels differentially regulate haematopoiesis. *Nature*, 2016. 99, 100

M. Jenkinson and S. Smith. A global optimisation method for robust affine registration of brain images. *Med. Image Anal.*, 5(2):143–156, 2001. ISSN 1361-8415. doi: http://dx.doi.org/10.1016/S1361-8415(01)00036-6. 6, 26

G. Johansson and H. Carr. Accelerating marching cubes with graphics hardware. In *Proceedings of the Conference of the Center for Advanced Studies on Collaborative research*, CASCON '06. ACM, 2006. 70

T. Ju. Robust repair of polygonal models. *ACM Transactions on Graphics*, 23(3):888–895, 2004. xxi, xxvii, 57, 60, 63, 120, 141

T. Ju, J. Warren, J. Carson, M. Bello, I. Kakadiaris, W. Chiu, C. Thaller, and G. Eichele. 3D volume reconstruction of a mouse brain from histological sections using warp filtering. *Journal of Neuroscience Methods*, 156(1–2):84–100, 2006. ISSN 0165-0270. doi: http://dx.doi.org/10.1016/j. jneumeth.2006.02.020. 6, 25

S. Kammel and F. P. León. Deflectometric measurement of specular surfaces. *IEEE Transactions on Instrumentation and Measurement*, 57(4):763–769, 2008. 150, 154

T. Katakai, H. Suto, M. Sugai, H. Gonda, A. Togawa, S. Suematsu, Y. Ebisuno, K. Katagiri, T. Kinashi, and A. Shimizu. Organizer-like reticular stromal cell layer common to adult secondary lymphoid organs. *The Journal of Immunology*, 181(9):6189–6200, 2008. ISSN 0022-1767. doi: 10.4049/ jimmunol.181.9.6189. 126

S. L. Keeling and W. Ring. Medical image registration and interpolation by optical flow with maximal rigidity. *Journal of Mathematical Imaging and Vision*, 23(1):47–65, 2005. ISSN 1573-7683. doi: 10.1007/s10851-005-4967-2. 54

R. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6):1153–1160, 1981. ISSN 0096-3518. doi: 10.1109/TASSP.1981. 1163711. 54

A. M. Khan, N. Rajpoot, D. Treanor, and D. Magee. A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution. *IEEE Transactions on Biomedical Engineering*, 61(6):1729–1738, 2014. ISSN 0018-9294. doi: 10.1109/ TBME.2014.2303294. xix, 27, 56, 141

M. J. Kiel and S. J. Morrison. Uncertainty in the niches that maintain haematopoietic stem cells. *Nature Reviews Immunology*, 8(4):290–301, 2008. 99

R. Kikinis, S. D. Pieper, and K. G. Vosburgh. 3D Slicer: A platform for subject-specific image analysis, visualization, and clinical support. In F. A. Jolesz, editor, *Intraoperative Imaging and Image-Guided Therapy*, pages 277–289. Springer, 2014. ISBN 978-1-4614-7657-3. doi: 10.1007/ 978-1-4614-7657-3_19. xxi

B. Kim, J. L. Boes, K. A. Frey, and C. R. Meyer. Mutual information for automated unwarping of rat brain autoradiographs. *NeuroImage*, 5(1):31–40, 1997. ISSN 1053-8119. doi: 10.1006/nimg. 1996.0251. 7

S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220 (4598):671–680, 1983. doi: 10.1126/science.220.4598.671. 151

S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. W. Pluim. elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, 2010. ISSN 0278-0062. doi: 10.1109/TMI.2009.2035616. xxii, 23, 25, 30, 37, 39, 40, 48, 50, 135

M. C. Knauer, J. Kaminski, and G. Häusler. Phase measuring deflectometry: A new approach to specular free-form surfaces. In *Optical Metrology in Production Engineering*, 2004. 150, 151

H.-G. Kopp, A. T. Hooper, S. T. Avecilla, and S. Rafii. Functional heterogeneity of the bone marrow vascular niche. *Annals of the New York Academy of Sciences*, 1176(1):47–54, 2009. 100

G. Kraal and R. Mebius. New insights into the cell biology of the marginal zone of the spleen. *International Review of Cytology*, 250:175–215, 2006. ISSN 0074-7696. doi: 10.1016/S0074-7696(06)50005-1. 126

Y. Kunisaki, I. Bruns, C. Scheiermann, J. Ahmed, S. Pinho, D. Zhang, T. Mizoguchi, Q. Wei, D. Lucas, K. Ito, et al. Arteriolar niches maintain haematopoietic stem cell quiescence. *Nature*, 502(7473):637–643, 2013. 100, 107, 112

A. P. Kusumbe, S. K. Ramasamy, and R. H. Adams. Coupling of angiogenesis and osteogenesis by a specific vessel subtype in bone. *Nature*, 507(7492):323–328, 2014. 100, 107

A. P. Kusumbe, S. K. Ramasamy, T. Itkin, M. A. Mäe, U. H. Langen, C. Betsholtz, T. Lapidot, and R. H. Adams. Age-dependent modulation of vascular niches for haematopoietic stem cells. *Nature*, 2016. 99, 100

S. Kusumi, D. Koga, T. Kanda, and T. Ushiki. Three-dimensional reconstruction of serial sections for analysis of the microvasculature of the white pulp and the marginal zone in the human spleen. *Biomedical Research (Tokyo)*, 36(3):195–203, 2015. doi: 10.2220/biomedres.36.195. 34, 91, 125, 126, 135

H. Kvasnicka and J. Thiele. Bone marrow angiogenesis: methods of quantification and changes evolving in chronic myeloproliferative disorders. *Histology and histopathology*, 19(4):1245—1260, 2004. ISSN 0213-3911. 100, 112

D. Le Gall. MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58, 1991. ISSN 0001-0782. doi: 10.1145/103085.103090. 25

T. M. Lehmann, C. Gonner, and K. Spitzer. Survey: interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging*, 18(11):1049–1075, 1999. ISSN 0278-0062. doi: 10.1109/42.816070. 56

H. Lester and S. R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149, 1999. ISSN 0031-3203. doi: 10.1016/S0031-3203(98)00095-8. 25

S. Leutenegger, M. Chli, and R. Siegwart. BRISK: Binary robust invariant scalable keypoints. In *IEEE International Conference on Computer Vision*, ICCV '11, pages 2548–2555. IEEE, 2011. doi: 10.1109/ICCV.2011.6126542. 7, 8

M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pages 131–144. ACM, 2000. ISBN 1-58113-208-5. doi: 10.1145/344779.344849. 148, 150

X.-M. Li, Z. Hu, M. L. Jorgenson, and W. B. Slayton. High levels of acetylated low-density lipoprotein uptake and low tyrosine kinase with immunoglobulin and epidermal growth factor homology domains-2 (Tie2) promoter activity distinguish sinusoids from other vessel types in murine bone marrow. *Circulation*, 120(19):1910–1918, 2009. 100

J. Liang and X.-J. Lai. Improved collision attack on hash function MD5. *Journal of Computer Science and Technology*, 22(1):79–87, 2007. ISSN 1860-4749. doi: 10.1007/s11390-007-9010-1. 169

B. Likar and F. Pernuš. A hierarchical approach to elastic registration based on mutual information. *Image and Vision Computing*, 19(1–2):33–44, 2001. ISSN 0262-8856. doi: 10.1016/S0262-8856(00)00053-6. 24, 25

P. Lindstrom. Out-of-core simplification of large polygonal models. In *Proceedings of the 27ᵗʰ Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pages 259–262. ACM, 2000. ISBN 1-58113-208-5. doi: 10.1145/344779.344912. 70, 71

Y.-J. Liu, S. Oldfield, and I. C. M. Maclennan. Memory B cells in T cell-dependent antibody responses colonize the splenic marginal zones. *European Journal of Immunology*, 18(3):355–362, 1988. ISSN 1521-4141. doi: 10.1002/eji.1830180306. 126, 134

Y.-J. Liu, J. Zhang, P. J. L. Lane, E. Y.-T. Chan, and I. C. M. Maclennan. Sites of specific B cell activation in primary and secondary responses to T cell-dependent and T cell-independent antigens. *European Journal of Immunology*, 21(12):2951–2962, 1991. ISSN 1521-4141. doi: 10.1002/eji.1830211209. URL http://dx.doi.org/10.1002/eji.1830211209. 134

Y. Livnat. *Noise, Wise and Sage: Algorithms for Rapid Isosurface Generation*. PhD thesis, University of Utah, 1999. 68

O. Lobachev. Direct visualization of cryptographic keys for enhanced security. *The Visual Computer*, 2017. ISSN 1432-2315. doi: 10.1007/s00371-017-1466-6. URL http://rdcu.be/Dtem. In press. ii, xxiv, 165

O. Lobachev, M. Schmidt, and M. Guthe. Optimizing multiple camera positions for the deflectometric measurement of multiple varying targets. *Journal of WSCG*, 21(2):145–152, 2013. ISSN 1213-6964. ii, xxiv, 147

O. Lobachev, C. Ulrich, B. S. Steiniger, V. Wilhelmi, V. Stachniss, and M. Guthe. Feature-based multi-resolution registration of immunostained serial sections. *Medical Image Analysis*, 35:288–302, 2017a. doi: 10.1016/j.media.2016.07.010. ii, xix, xxi, xxii, 21, 27, 30, 56, 59, 103, 106, 107, 110, 111, 119, 135, 141, 187

O. Lobachev, B. S. Steiniger, and M. Guthe. Compensating anisotropy in histological serial sections with optical flow-based interpolation. In I. Viola and J. Bittner, editors, *33ʳᵈ Spring Conference on Computer Graphics*, SCCG '17, 2017b. In press. ii, xix, xxi, 53, 135, 141

H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133–135, 1981. URL http://www.nature.com/doifinder/10.1038/293133a0. 150

W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987. ISSN 0097-8930. doi: 10.1145/37402.37422. 57, 68, 120, 141

K. S. Lorenz, P. Salama, K. W. Dunn, and E. J. Delp. A multi-resolution approach to non-rigid registration of microscopy images. In *9ᵗʰ IEEE International Symposium on Biomedical Imaging*, ISBI '12, pages 198–201. IEEE, 2012. doi: 10.1109/ISBI.2012.6235518. 25, 26

K.-L. Low and T.-S. Tan. Model simplification using vertex-clustering. In *Proceedings of the Symposium on Interactive 3D Graphics*, I3D '97, pages 75–81, 1997. 71

D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. ISSN 0920-5691. doi: 10.1023/B:VISI.0000029664.99615.94. 7, 24

B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7ᵗʰ International Joint Conference on Artificial Intelligence*, IJCAI '81, pages 674–679. Morgan Kaufmann, 1981. 7, 23, 54

D. P. Luebke. A developer's survey of polygonal simplification algorithms. *IEEE Computer Graphics and Applications*, 21, 2001. ISSN 0272-1716. 71

B. Ma, Z. Lin, S. Winkelbach, W. Lindenmaier, and K. E. Dittmar. Automatic registration of serial sections of mouse lymph node by using Image-Reg. *Micron*, 39(4):387–396, 2008a. ISSN 0968-4328. doi: 10.1016/j.micron.2007.03.005. 6, 25

B. Ma, L. Wang, R. von Wasielewski, W. Lindenmaier, and K. E. Dittmar. Serial sectioning and three-dimensional reconstruction of mouse Peyer's patch. *Micron*, 39(7):967–975, 2008b. ISSN 0968-4328. doi: http://dx.doi.org/10.1016/j.micron.2007.10.007. 55, 135

G. Malandain, E. Bardinet, K. Nelissen, and W. Vanduffel. Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations. *NeuroImage*, 23(1):111–127, 2004. ISSN 1053-8119. doi: http://dx.doi.org/10.1016/j.neuroimage.2004.04.038. 6

R. Martoňák, G. E. Santoro, and E. Tosatti. Quantum annealing of the traveling-salesman problem. *Physical Review E*, 70:057701, 2004. doi: 10.1103/PhysRevE.70.057701. 152

A. S. Mathur. Low cost virtual reality for medical training. In *IEEE Virtual Reality*, VR '15, pages 345–346, 2015. doi: 10.1109/VR.2015.7223437. 84

M. Meila and J. Shi. Learning segmentation with random walk. In *Advances in Neural Information Processing Systems*, NIPS '01, pages 873–879, 2000. URL http://papers.nips.cc/paper/1830-learning-segmentation-by-random-walks. 151

S. Meixlsperger, C. S. Leung, P. C. Rämer, M. Pack, L. D. Vanoaica, G. Breton, S. Pascolo, A. M. Salazar, A. Dzionek, J. Schmitz, R. M. Steinman, and C. Münz. CD141⁺ dendritic cells produce prominent amounts of IFN-α after dsRNA recognition and can be targeted via DEC-205 in humanized mice. *Blood*, 121(25):5034–5044, 2013. ISSN 0006-4971. doi: 10.1182/blood-2012-12-473413. 102

F. E. Mercier, C. Ragu, and D. T. Scadden. The bone marrow at the crossroads of blood and immunity. *Nature Reviews Immunology*, 12(1):49–60, 2012. 99

C. Montani, R. Scateni, and R. Scopigno. A modified look-up table for implicit disambiguation of marching cubes. *The Visual Computer*, 10(6):353–355, 1994. 68

H. Müller and M. Stark. Adaptive generation of surfaces in volume data. *The Visual Computer*, 9(4): 182–199, 1993. ISSN 0178-2789. doi: 10.1007/BF01901723. 68, 70

K. Murphy, B. van Ginneken, J. Reinhardt, S. Kabus, K. Ding, X. Deng, K. Cao, K. Du, G. Christensen, V. Garcia, T. Vercauteren, N. Ayache, O. Commowick, G. Malandain, B. Glocker, N. Paragios, N. Navab, V. Gorbunova, J. Sporring, M. de Bruijne, X. Han, M. Heinrich, J. Schnabel, M. Jenkinson, C. Lorenz, M. Modat, J. McClelland, S. Ourselin, S. Muenzing, M. Viergever, D. De Nigris, D. Collins, T. Arbel, M. Peroni, R. Li, G. Sharp, A. Schmidt-Richberg, J. Ehrhardt, R. Werner, D. Smeets, D. Loeckx, G. Song, N. Tustison, B. Avants, J. Gee, M. Staring, S. Klein, B. Stoel, M. Urschler, M. Werlberger, J. Vandemeulebroucke, S. Rit, D. Sarrut, and J. Pluim. Evaluation of registration methods on thoracic CT: The EMPIRE10 challenge. *IEEE Transactions on Medical Imaging*, 30(11):1901–1920, 2011. ISSN 0278-0062. doi: 10.1109/TMI.2011.2158349. 7, 27

A. T. Murray and K. Kim. Coverage optimization to support security monitoring. *Computers, Environment and Urban Systems*, 31:133–147, 2007. 151

K. Nagara, H. R. Roth, S. Nakamura, H. Oda, T. Moriya, M. Oda, and K. Mori. Micro-CT guided 3D reconstruction of histological images. In G. Wu, B. C. Munsell, Y. Zhan, W. Bai, G. Sanroma, and P. Coupé, editors, *Patch-Based Techniques in Medical Imaging*, LNCS 10530, pages 93–101. Springer, 2017. ISBN 978-3-319-67434-6. doi: 10.1007/978-3-319-67434-6_11. 26, 27, 195

M. Naor and A. Shamir. Visual cryptology. In A. De Santis, editor, *Advances in Cryptology*, EURO-CRYPT '94, pages 1–12. Springer, 1995. ISBN 978-3-540-44717-7. doi: 10.1007/BFb0053419. 169

L. Nataraj, S. Karthikeyan, G. Jacob, and B. S. Manjunath. Malware images: Visualization and automatic classification. In *Proceedings of the 8th International Symposium on Visualization for Cyber Security*, VizSec '11, pages 4:1–4:7. ACM, 2011. ISBN 978-1-4503-0679-9. doi: 10.1145/2016904. 2016908. 169, 170

H. Negaard, N. Iversen, I. Bowitz-Lothe, P. Sandset, B. Steinsvik, B. Østenstad, and P. Iversen. Increased bone marrow microvascular density in haematological malignancies is associated with differential regulation of angiogenic factors. *Leukemia*, 23(1):162–169, 2009. 100

T. S. Newman and H. Yi. A survey of the marching cubes algorithm. *Computers & Graphics*, 30(5): 854–879, 2006. ISSN 0097-8493. doi: 10.1016/j.cag.2006.07.021. 68

P. Nieuwenhuis and W. Ford. Comparative migration of B- and T-lymphocytes in the rat spleen and lymph nodes. *Cellular Immunology*, 23(2):254–267, 1976. ISSN 0008-8749. doi: 10.1016/ 0008-8749(76)90191-X. 125

C. Nikou, F. Heitz, A. Nehlig, I. J. Namer, and J.-P. Armspach. A robust statistics-based global energy function for the alignment of serially acquired autoradiographic sections. *Journal of Neuroscience Methods*, 124(1):93–102, 2003. ISSN 0165-0270. doi: 10.1016/S0165-0270(02)00369-2. 6, 24, 26

NIST. Federal information processing standards: Secure hash standard (SHS). Technical Report FIPS PUB 180-4, Information Technology Laboratory, National Institute of Standards and Technology, 2015. SHA2. 173, 176

M. A. Nolte, J. A. Beliën, I. Schadee-Eestermans, W. Jansen, W. W. Unger, N. van Rooijen, G. Kraal, and R. E. Mebius. A conduit system distributes chemokines and small blood-borne molecules through the splenic white pulp. *Journal of Experimental Medicine*, 198(3):505–512, 2003. ISSN 0022-1007. doi: 10.1084/jem.20021801. 136

C. Nombela-Arrieta, G. Pivarnik, B. Winkel, K. J. Canty, B. Harley, J. E. Mahoney, S.-Y. Park, J. Lu, A. Protopopov, and L. E. Silberstein. Quantitative imaging of haematopoietic stem and progenitor cell localization and hypoxic status in the bone marrow microenvironment. *Nature Cell Biology*, 15(5):533–543, 2013. 99, 100, 107

G. Olague and R. Mohr. Optimal camera placement for accurate reconstruction. *Pattern Recognition*, 35(4):927–944, 2002. ISSN 0031-3203. doi: 10.1016/S0031-3203(01)00076-0. 151

A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision*, 42(3):145–175, May 2001. ISSN 1573-1405. doi: 10.1023/A:1011139631724. 170

I. Oláh and B. Glick. Splenic white pulp and associated vascular channels in chicken spleen. *American Journal of Anatomy*, 165(4):445–480, 1982. ISSN 1553-0795. doi: 10.1002/aja.1001650408. 136

I. Oláh, B. Glick, and R. L. Taylor. Effect of soluble antigen on the ellipsoid-associated cells of the chicken's spleen. *Journal of Leukocyte Biology*, 35(5):501–510, 1984. URL http://www.jleukbio.org/content/35/5/501.abstract. 136

I. Oláh, N. Nagy, and L. Vervelde. *Structure of the Avian Lymphoid System*, pages 11–44. Elsevier, second edition, 2013. ISBN 9780123969651. doi: 10.1016/B978-0-12-396965-1.00002-9. 136

D. Onder, S. Zengin, and S. Sarioglu. A review on color normalization and color deconvolution methods in histopathology. *Applied Immunohistochemistry & Molecular Morphology*, 22(10), 2014. URL http://journals.lww.com/appliedimmunohist/Fulltext/2014/11000/A_Review_on_Color_Normalization_and_Color.1.aspx. 142

Open SSL. *Manual page for openssl – OpenSSL command line tool*. Open SSL, 2016. 177

S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache. Reconstructing a 3D structure from serial histological sections. *Image and Vision Computing*, 19:25–31, 2001. doi: 10.1016/S0262-8856(00)00052-4. 6, 25, 55

K. Panteli, N. Zagorianakou, N. Agnantis, K. Bourantas, and M. Bai. Clinical correlation of bone marrow microvessel density in essential thrombocythemia. *Acta haematologica*, 114(2):99–103, 2005. 100

H. Peng, F. Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, 2005. ISSN 0162-8828. doi: 10.1109/TPAMI.2005.159. 6

T. Peng, M. Yigitsoy, A. Eslami, C. Bayer, and N. Navab. 6[th] international workshop on biomedical image registration – wbir '14. LNCS 8545, pages 144–153. Springer, 2014. ISBN 978-3-319-08554-8. doi: 10.1007/978-3-319-08554-8_15. 25

A. Perrig and D. Song. Hash visualization: A new technique to improve real-world security. In *International Workshop on Cryptographic Techniques and E-Commerce*, CrypTEC '99, pages 131–138, 1999. 166, 168, 177

M. Piccardi. Background subtraction techniques: a review. In *IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3099–3104. IEEE, 2004. doi: 10.1109/ICSMC.2004. 1400815. 152

F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat. Comparing ICP variants on real-world data sets. *Auton. Robot*, 34(3):133–148, Feb. 2013. 14

J. Popović and H. Hoppe. Progressive simplicial complexes. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 217–224. ACM/Addison-Wesley, 1997. ISBN 0-89791-896-7. doi: 10.1145/258734.258852. 71

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, third edition, 2007. ISBN 0521880688, 9780521880688. 13, 39

N. Qian. Binocular disparity review and the perception of depth. *Neuron*, 18(3):359–368, 1997. 150

D. J. Quam, T. J. Gundert, L. Ellwein, C. E. Larkee, P. Hayden, R. Q. Migrino, H. Otake, and J. F. LaDisa, Jr. Immersive visualization for enhanced computational fluid dynamics analysis. *Journal of Biomechanical Engineering*, 137(3):031004–031004–12, 2015. doi: 10.1115/1.4029017. 84

F. Reck, C. Dachsbacher, R. Grosso, G. Greiner, and M. Stamminger. Realtime isosurface extraction with graphics hardware. In *Eurographics 2004, Short Presentations and Interactive Demos*, pages 33–36. EG, 2004. 70

A. Reger, J. Dumler, O. Lobachev, J. Neuberger, and R. Steinhilper. Machine allocation via pattern recognition in harmonic waves of manufacturing plants. In R. Teti, editor, *11th CIRP Conference on Intelligent Computation in Manufacturing Engineering*, ICME '17. CIRP, Elsevier, 2017.

E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21(5):34–41, 2001. ISSN 0272-1716. doi: 10.1109/38.946629. xix, 27, 56, 141

E. Rescorla. HTTP over TLS, 2000. Request for Comments: 2818. 175

J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. EpicFlow: Edge-preserving interpolation of correspondences for optical flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '15, pages 1164–1172, 2015. 54

D. Richardson and J. Lichtman. Clarifying tissue clearing. *Cell*, 162(2):246–257, 2015. ISSN 0092-8674. doi: 10.1016/j.cell.2015.06.067. 196

A. Roche, X. Pennec, G. Malandain, and N. Ayache. Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information. *IEEE Transactions on Medical Imaging*, 20(10):1038–1049, 2001. ISSN 0278-0062. doi: 10.1109/42.959301. 6, 26

S. Roettger. The volume library, 2013. URL http://www9.informatik.uni-erlangen.de/External/vollib/. Retrieved 2013-05-22. 77

J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering complex scenes. In B. Falcidieno and T. L. Kunii, editors, *Modeling in Computer Graphics*, IFIP TC 5, pages 455–465. Springer, 1993. ISBN 978-3-642-78116-2. doi: 10.1007/978-3-642-78114-8_29. 71

D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999. ISSN 0278-0062. doi: 10.1109/42.796284. 6, 25, 26, 56

A. C. Ruifrok and D. A. Johnston. Quantification of histochemical staining by color deconvolution. *Analytical and quantitative cytology and histology*, 23(4):291–299, 2001. ISSN 0884-6812. 142

A. Ruiz, M. Ujaldon, L. Cooper, and K. Huang. Non-rigid registration for large sets of microscopic images on graphics processors. *Journal of Signal Processing Systems*, 55(1–3):229–250, 2009. ISSN 1939-8018. doi: 10.1007/s11265-008-0208-4. 24, 26

S. Saalfeld, A. Cardona, V. Hartenstein, and P. Tomančák. As-rigid-as-possible mosaicking and serial section registration of large ssTEM datasets. *Bioinformatics*, 26(12):i57–i63, 2010. doi: 10.1093/bioinformatics/btq219. 24, 27, 187

S. Saalfeld, R. Fetter, A. Cardona, and P. Tomancak. Elastic volume reconstruction from series of ultra-thin microscopy sections. *Nature Methods*, 9(7):717–720, 2012. ISSN 15487091. doi: 10.1038/nmeth.2072. 6, 25, 56, 135

A. Sanderson, L. Weiss, and S. Nayar. Structured highlight inspection of specular surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(1):44–55, 1988. ISSN 0162-8828. doi: 10.1109/34.3866. 150

S. Schaefer and J. Warren. Adaptive vertex clustering using octrees. In *Proceedings of SIAM Geometric Design and Computing*, 2003. 71

S. Schaefer and J. Warren. Dual marching cubes: Primal contouring of dual grids. In *Processings of Pacific Conference on Computer Graphics*, PG '04, pages 70–76. IEEE Computer Society, 2004. ISBN 0-7695-2234-3. 68, 70

O. Scherzer, editor. *Mathematical models for registration and applications to medical imaging*. Springer, 2006. 5

J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682, 2012. doi: 10.1038/nmeth.2019. xxvii, 27, 55, 60, 135, 141, 142

M. Schmidt, O. Lobachev, and M. Guthe. Coherent Metropolis light transport using speculative mutations. *Journal of WSCG*, 24(1):1–8, 2016. ISSN 1213-6964.

R. Schmidt and M. Ratto. Design-to-fabricate: Maker hardware requires maker software. *IEEE Computer Graphics and Applications*, 33(6):26–34, 2013. 120

J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano-Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, et al. A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In W. J. Niessen and M. A. Viergever, editors,

Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001, LNCS 2208, pages 573–581. Springer, 2001. ISBN 978-3-540-42697-4. doi: 10.1007/3-540-45468-3_69. 6, 26, 56

B. Schneier. *Applied cryptography: protocols, algorithms, and source code in C*. John Wiley & Sons, 2007. 166

H. Schultz. Retrieving shape information from multiple images of a specular surface. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):195–201, Feb 1994. ISSN 0162-8828. doi: 10.1109/34.273732. 151

E. Shaffer and M. Garland. Efficient adaptive simplification of massive meshes. In *Proceedings of the Conference on Visualization '01*, VIS '01, pages 127–134. IEEE Computer Society, 2001. ISBN 0-7803-7200-X. 71

R. Shams, P. Sadeghi, R. Kennedy, and R. Hartley. A survey of medical image registration on multicore and the GPU. *IEEE Signal Processing Magazine*, 27(2):50–60, 2010. ISSN 1053-5888. doi: 10.1109/MSP.2009.935387. 5, 23

L. Shapira, A. Shamir, and D. Cohen-Or. Consistent mesh partitioning and skeletonisation using the shape diameter function. *The Visual Computer*, 24(4):249–259, 2008. 121

H.-W. Shen, C. D. Hansen, Y. Livnat, and C. R. Johnson. Isosurfacing in span space with utmost efficiency (ISSUE). In *Proceedings of the Conference on Visualization '96*, VIS '96, pages 287–295. IEEE Computer Society, 1996. ISBN 0-89791-864-9. doi: 10.1109/VISUAL.1996.568121. 68

R. Shen, P. Boulanger, and M. Noga. MedVis: A real-time immersive visualization environment for the exploration of medical volumetric data. In *Information Visualization in Medical and Biomedical Informatics*, MEDIVIS '08, pages 63–68, 2008. ISBN 978-0-7695-3284-4. doi: 10.1109/MediVis. 2008.10. 84

M. H. Shimabukuro and R. Minghim. Visualisation and reconstruction in dentistry. In *Information Visualization*, pages 25–31, 1998. ISBN 0-8186-8509-3. doi: 10.1109/IV.1998.694195. 84

D. A. Sipkins, X. Wei, J. W. Wu, J. M. Runnels, D. Côté, T. K. Means, A. D. Luster, D. T. Scadden, and C. P. Lin. In vivo imaging of specialized bone marrow endothelial microdomains for tumour engraftment. *Nature*, 435(7044):969–973, 2005. 100

T. Snook. Studies on the perifollicular region of the rat's spleen. *The Anatomical Record*, 148(2): 149–159, 1964. ISSN 1097-0185. doi: 10.1002/ar.1091480205. 125

Y. Song, D. Treanor, A. Bulpitt, and D. Magee. 3D reconstruction of multiple stained histology images. *Journal of Pathology Informatics*, 4(2):7, 2013. 7

Y. Song, D. Treanor, A. J. Bulpitt, N. Wijayathunga, N. Roberts, R. Wilcox, and D. R. Magee. Unsupervised content classification based nonrigid registration of differently stained histology images. *IEEE Transactions on Biomedical Engineering*, 61(1):96–108, 2014. ISSN 0018-9294. doi: 10.1109/TBME.2013.2277777. 25

A. Sotiras, C. Davatzikos, and N. Paragios. Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging*, 32(7):1153–1190, 2013. ISSN 0278-0062. doi: 10.1109/TMI.2013. 2265603. 24, 26

F. Spitzer. *Principles of random walk*. Springer-Verlag, 2001. 151

B. Steiniger and P. Barth. *Microanatomy and function of the spleen*. Springer, 1999. 136

B. Steiniger, E.-M. Timphus, R. Jacob, and P. J. Barth. CD27$^+$ B cells in human lymphatic organs: re-evaluating the splenic marginal zone. *Immunology*, 116(4):429–442, 2005. ISSN 1365-2567. doi: 10.1111/j.1365-2567.2005.02242.x. 126, 134, 140

B. Steiniger, V. Stachniss, H. Schwarzbach, and P. Barth. Phenotypic differences between red pulp capillary and sinusoidal endothelia help localizing the open splenic circulation in humans. *Histochemistry and Cell Biology*, 128(5):391–398, 2007. ISSN 0948-6143. doi: 10.1007/s00418-007-0320-8. 32, 127

B. Steiniger, M. Bette, and H. Schwarzbach. The open microcirculation in human spleens: a three-dimensional approach. *Journal of Histochemistry & Cytochemistry*, 59(6):639–648, 2011a. doi: 10.1369/0022155411408315. 125, 126, 127, 135, 140

B. S. Steiniger. Human spleen microanatomy: why mice do not suffice. *Immunology*, 145(3):334–346, 2015. ISSN 1365-2567. doi: 10.1111/imm.12469. 126, 127, 136, 138, 139

B. S. Steiniger, L. Rüttinger, and P. J. Barth. The three-dimensional structure of human splenic white pulp compartments. *Journal of Histochemistry & Cytochemistry*, 51(5):655–663, 2003. doi: 10.1177/002215540305100511. 6, 25, 28, 55, 126, 134, 135

B. S. Steiniger, M. Bette, and H. Schwarzbach. The open microcirculation in human spleens: a three-dimensional approach. *Journal of Histochemistry & Cytochemistry*, 59(6):639–648, 2011b. doi: 10.1369/0022155411408315. 6, 25

B. S. Steiniger, S. Bubel, W. Böckler, K. Lampp, A. Seiler, B. Jablonski, M. Guthe, and V. Stachniss. Immunostaining of pulpal nerve fibre bundle/arteriole associations in ground serial sections of whole human teeth embedded in Technovit®9100. *Cells Tissues Organs*, 198(1):57–65, 2013. doi: 10.1159/000351608. ii, 34, 56, 59, 109, 114

B. S. Steiniger, A. Seiler, K. Lampp, V. Wilhelmi, and V. Stachniss. B lymphocyte compartments in the human splenic red pulp: capillary sheaths and periarteriolar regions. *Histochemistry and Cell Biology*, 141(5):507–518, 2014a. ISSN 0948-6143. doi: 10.1007/s00418-013-1172-z. 32, 56, 59, 126, 127, 136

B. S. Steiniger, V. Wilhelmi, A. Seiler, K. Lampp, and V. Stachniss. Heterogeneity of stromal cells in the human splenic white pulp. Fibroblastic reticulum cells, follicular dendritic cells and a third superficial stromal cell type. *Immunology*, 143(3):462–477, 2014b. ISSN 1365-2567. doi: 10.1111/imm.12325. 34, 56, 59, 126, 127, 136, 138, 139, 140

B. S. Steiniger, V. Stachniss, V. Wilhelmi, A. Seiler, K. Lampp, A. Neff, M. Guthe, and O. Lobachev. Three-dimensional arrangement of human bone marrow microvessels revealed by immunohistology in undecalcified sections. *PLOS ONE*, 11(12):1–25, 12 2016. doi: 10.1371/journal.pone.0168173. ii, 34, 56, 59, 87, 93, 99, 122

B. S. Steiniger, C. Ulrich, M. Berthold, M. Guthe, and O. Lobachev. Capillary networks and follicular marginal zones in human spleens. three-dimensional models based on immunostained serial sections. *PLOS ONE*, 2018. doi: 10.1371/journal.pone.0191019. Accepted. ii, 125

M. Stevens. Fast collision attack on MD5. *IACR Cryptology ePrint Archive*, 2006. URL http://eprint.iacr.org/2006/104. 167

M. Stevens. Counter-cryptanalysis. In R. Canetti and J. A. Garay, editors, *Advances in Cryptology*, CRYPTO '13, pages 129–146. Springer, 2013. ISBN 978-3-642-40041-4. doi: 10.1007/978-3-642-40041-4_8. 169

M. Stevens, E. Bursztein, P. Karpman, A. Albertini, and Y. Markov. The first collision for full SHA-1, 2017. URL http://shattered.it/static/shattered.pdf. 167, 169, 177

J. E. Stone, A. Kohlmeyer, K. L. Vandivort, and K. Schulten. Immersive molecular visualization and interactive modeling with commodity hardware. In G. Bebis, R. Boyle, B. Parvin, D. Koracin, R. Chung, R. Hammound, M. Hussain, T. Kar-Han, R. Crawfis, D. Thalmann, D. Kao, and L. Avila, editors, *Advances in Visual Computing*, ISVC '10, pages 382–393. Springer, 2010. ISBN 978-3-642-17274-8. doi: 10.1007/978-3-642-17274-8_38. 84

M. S. Subhedar and V. H. Mankar. Current status and key issues in image steganography: A survey. *Computer Science Review*, 13:95–113, 2014. ISSN 1574-0137. doi: 10.1016/j.cosrev.2014.09.001. 170

D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *Computer Vision and Pattern Recognition*, CVPR '10, pages 2432–2439. IEEE Computer Society, 2010. doi: 10.1109/CVPR.2010.5539939. 54

X. Suo, Y. Zhu, and G. S. Owen. Graphical passwords: a survey. In *21st Annual Computer Security Applications Conference*, ACSAC '05. IEEE, 2005. doi: 10.1109/CSAC.2005.27. 169

J. S. Suru, D. L. Wilson, and S. Laximinarayan. *Handbook of biomedical image analysis*, volume 3. Springer, 2005. 5

R. Sviták and V. Skala. Surface reconstruction from orthogonal slices. In *International Conference on Computer Vision and Graphics*, pages 742–747. Silesian University of Technology, 2002. ISBN 8391768309. URL http://hdl.handle.net/11025/11697. 68

R. Sviták and V. Skala. A robust technique for surface reconstruction from orthogonal slices. *Machine Graphics & Vision International Journal*, 13(3):221–233, Jan. 2004. ISSN 1230-0535. URL http://dl.acm.org/citation.cfm?id=1073547.1073550. 68

T. Takaku, D. Malide, J. Chen, R. T. Calado, S. Kajigaya, and N. S. Young. Hematopoiesis in 3 dimensions: human and murine bone marrow architecture visualized by confocal microscopy. *Blood*, 116(15):e41–e55, 2010. 100, 112

A. Tanacs and Z. Kato. Fast linear registration of 3D objects segmented from medical images. In *4th International Conference on Biomedical Engineering and Informatics*, BMEI '11, pages 294–298. IEEE, 2011. doi: 10.1109/BMEI.2011.6098290. 6

N. Tatarchuk, J. Shopf, and C. DeCoro. Real-time isosurface extraction using the GPU programmable geometry pipeline. In *ACM SIGGRAPH 2007 Courses*, SIGGRAPH '07, pages 122–137. ACM, 2007. 70

G. Taubin. A signal processing approach to fair surface design. In *Proceedings of the 22ⁿᵈ Annual Conference on Computer Graphics and Interactive Techniques*, pages 351–358. ACM, 1995. 57, 120, 141

S. T. Teoh, T. Jankun-Kelly, K.-L. Ma, and S. F. Wu. Visual data analysis for detecting flaws and intruders in computer network systems. *IEEE/ACM Transactions on Networking*, 6(5):515–528, 1998. 169, 170

P. Thevenaz, T. Blu, and M. Unser. Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19 (7):739–758, 2000. ISSN 0278-0062. doi: 10.1109/42.875199. 56

J.-P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Med. Image Anal.*, 2(3):243–260, 1998. ISSN 1361-8415. doi: 10.1016/S1361-8415(98)80022-4. 6, 26, 56

K. Tokoyoda, S. Zehentmeier, A. N. Hegazy, I. Albrecht, J. R. Grün, M. Löhning, and A. Radbruch. Professional memory CD4$^+$ T lymphocytes preferentially reside and rest in the bone marrow. *Immunity*, 30(5):721–730, 2009. doi: 10.1016/j.immuni.2009.03.015. 99

K. Tokoyoda, A. E. Hauser, T. Nakayama, and A. Radbruch. Organization of immunological memory by bone marrow stroma. *Nature Reviews Immunology*, 10(3):193–200, 2010. doi: 10.1038/nri2727. 99, 100

G. M. Treece, R. W. Prager, and A. H. Gee. Regularised marching tetrahedra: improved iso-surface extraction. *Comput. Graph.*, 23(4):583–598, 1999. 68

V. M. Turner and N. A. Mabbott. Ageing adversely affects the migration and function of marginal zone B cells. *Immunology*, 151(3):349–362, 2017. ISSN 1365-2567. doi: 10.1111/imm.12737. 134

C. Ulrich, N. Grund, E. Derzapf, O. Lobachev, and M. Guthe. Parallel iso-surface extraction and simplification. In V. Skala, editor, *WSCG Communications proceedings*, WSCG '14, pages 361–368, 2014a. ISBN 978-80-86943-71-8. ii, xix, xxi, 67, 141

C. Ulrich, O. Lobachev, B. S. Steiniger, and M. Guthe. Imaging the vascular network of the human spleen from immunostained serial sections. In I. Viola, K. Buehler, and T. Ropinski, editors, *Visual Computing for Biology and Medicine*. The Eurographics Association, 2014b. ISBN 978-3-905674-62-0. doi: 10.2312/vcbm.20141185. ii, xix, xxi, xxii, 3, 22, 23, 24, 26, 28, 30, 31, 35, 37, 39, 40, 47, 48, 56, 59, 141

M. Urschler, J. Bauer, H. Ditt, and H. Bischof. SIFT and shape context for feature-based nonlinear registration of thoracic CT images. In R. R. Beichel and M. Sonka, editors, *Computer Vision Approaches to Medical Image Analysis*, LNCS 4241, pages 73–84. Springer, 2006. ISBN 978-3-540-46257-6. doi: 10.1007/11889762_7. 24, 27

A. van Dam, A. S. Forsberg, D. H. Laidlaw, J. J. LaViola, and R. M. Simpson. Immersive VR for scientific visualization: a progress report. *IEEE Computer Graphics and Applications*, 20(6):26–52, 2000. ISSN 0272-1716. doi: 10.1109/38.888006. 84

J. H. van Krieken, J. Te Velde, J. Hermans, and K. Welvaart. The splenic red pulp; a histomorpho-metrical study in splenectomy specimens embedded in methylmethacrylate. *Histopathology*, 9(4): 401–416, 1985. 6, 25

A. J. P. Veerman and W. van Ewijk. White pulp compartments in the spleen of rats and mice. *Cell and Tissue Research*, 156(4):417–441, 1975. ISSN 1432-0878. doi: 10.1007/BF00225103. 125

T. Wan, B. N. Bloch, S. Danish, and A. Madabhushi. A novel point-based nonrigid image registration scheme based on learning optimal landmark configurations. In *Medical Imaging 2013: Image Processing*, Proc. SPIE 8669, pages 866934–866934–12, 2013. doi: 10.1117/12.2007153. 7, 26, 56

X. Wang and H. Yu. How to break MD5 and other hash functions. In R. Cramer, editor, *Advances in Cryptology*, EUROCRYPT '05, pages 19–35. Springer, 2005. ISBN 978-3-540-32055-5. doi: 10.1007/11426639_2. 167, 169

X. Wang, Y. L. Yin, and H. Yu. Finding collisions in the full SHA-1. In V. Shoup, editor, *Advances in Cryptology*, CRYPTO '05, pages 17–36. Springer, 2005. ISBN 978-3-540-31870-5. doi: 10.1007/11535218_2. 167, 169, 177

Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. ISSN 1057-7149. doi: 10.1109/TIP.2003.819861. xxi, xxii, 39, 43

W. M. Wells, III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Med. Image Anal.*, 1(1):35—51, 1996. ISSN 1361-8415. doi: 10.1016/S1361-8415(01)80004-9. 6, 26

J. Wilhelms and A. V. Gelder. Octrees for faster isosurface generation. *ACM SIGGRAPH Computer Graphics*, 24(5):57–62, 1990. doi: 10.1145/99308.99321. 68

E. Willbold and F. Witte. Histology and research at the hard tissue–implant interface using technovit 9100 New embedding technique. *Acta biomaterialia*, 6(11):4447–4455, 2010. 109

S. Wirtz, N. Papenberg, B. Fischer, and O. Schmitt. Robust and staining-invariant elastic registration of a series of images from histologic slices. In *Medical Imaging 2005: Image Processing*, Proc. SPIE 5747, pages 1256–1262, 2005. doi: 10.1117/12.595246. 6, 25

G. Wyvill, C. McPheeters, and B. Wyvill. Data structure for soft objects. *The Visual Computer*, 2(4): 227–234, 1986. ISSN 0178-2789. doi: 10.1007/BF01900346. 68

Z. Xie and G. Farin. Image registration using hierarchical B-splines. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):85–94, 2004. ISSN 1077-2626. doi: 10.1109/TVCG.2004.1260760. 6, 26

Y. Xu, J. G. Pickering, Z. Nong, E. Gibson, J.-M. Arpino, H. Yin, and A. D. Ward. A method for 3D histopathology reconstruction supporting mouse microvasculature analysis. *PLOS ONE*, 10(5): 1–24, 05 2015. doi: 10.1371/journal.pone.0126817. 25, 55

R. Yang, C. M. Davies, C. W. Archer, and R. G. Richards. Immunohistochemistry of matrix markers in Technovit 9100 New-embedded undecalcified bone sections. *European Cells & Materials*, 6:57–71, 2003. 109

S. Zehentmeier, K. Roth, Z. Cseresnyes, Ö. Sercan, K. Horn, R. A. Niesner, H.-D. Chang, A. Radbruch, and A. E. Hauser. Static and dynamic components synergize to form a stable survival niche for bone marrow plasma cells. *European journal of immunology*, 44(8):2306–2317, 2014. 99

J. Zhao, S.-C. S. Cheung, and T. Nguyen. Optimal visual sensor network configuration. In H. Aghajan and A. Cavallaro, editors, *Multi-Camera Networks: Principles and Applications*, pages 139–141. Academic Press, 2009. 151

J. Y. Zheng and A. Murata. Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):913–920, 2000. ISSN 0162-8828. doi: 10.1109/34.868691. 151

P. R. Zimmermann. *The official PGP user's guide*. MIT Press, 1995. 166

C. W. Zitnick, N. Jojic, and S. B. Kang. Consistent segmentation for optical flow estimation. In *International Conference on Computer Vision*, ICCV '05, pages 1308–1315. IEEE, 2005. doi: 10.1109/ICCV.2005.61. 55

B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11): 977–1000, 2003. ISSN 0262-8856. doi: 10.1016/S0262-8856(03)00137-9. 7, 23