



FAKULTÄT FÜR **INFORMATIK**

Enhancement, Registration, and Visualization of High Resolution Episcopic Microscopy Data

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Computergraphik/Digitale Bildverarbeitung

eingereicht von

Clemens Brandorff

Matrikelnummer 9526005

an der

Fakultät für Informatik der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Betreuer: Univ.-Ass. Dipl.-Ing. Dr.techn. Peter Rautek

Wien, 20.07.2009


(Unterschrift Verfasser/in)

(Unterschrift Betreuer/in)

Clemens Brandorff

Enhancement, Registration, and Visualization of High Resolution Episcopic Microscopy Data

Master's Thesis



supervised by

Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Univ.-Ass. Dipl.-Ing. Dr.techn. Peter Rautek

Institute of Computer Graphics and Algorithms

Vienna University of Technology

Abstract

Weninger et al. [25] developed a novel methodology for rapid 2D and 3D computer analysis and visualization of gene expression patterns. The data is generated by staining a specimen followed by an iterating process of cutting thin slices and capturing them with an episcopic microscope. The result is a high resolution 3D dataset. One channel contains anatomical information and a second channel contains the gene expression patterns.

In this thesis we examine methods for enhancing, registering and visualizing this novel kind of data. We address the uneven illumination of slices that are introduced by the methodology. We developed an algorithm to fit a quadric surface through the background pixels to estimate the illumination situation over the whole slice. This estimate is used to correct the slices of one dataset.

Further, an extension of this methodology was researched. Recycling the already cut sections for staining them a second time allows the medical domain scientists to augment their technique with additional information. The result of the second data generation phase is a stack of unaligned slices. The manual processing of the sections introduces non-linear deformations. We explored several registration algorithms to align the two image stacks. We found a two step registration approach to yield the best results. In the first step a coarse affine registration is used to approximately align the datasets. The result of the first step is inspected and if necessary corrected by the user. In the second step a b-spline registration is used that compensates for the non-linear deformations of the 2D slices.

For the visual inspection of the registration results and to present an overview of the datasets we implemented two visualization approaches. A checkerboard view is used to compare 2D slices, and a three dimensional approach based on direct volume rendering incorporates surface enhancement by gradient magnitude opacity modulation to emphasize the alignment of tissue boundaries.

Kurzfassung

Weninger et al. [25] entwickelten eine neuartige Methode um schnelle 2D / 3D Computeranalysen und Computervisualisierungen von Genexpressionsmustern durchzuführen. Die Daten werden durch Einfärben, mit folgenden Schneiden und Aufnehmen mittels eines episkopischen Mikroskops erzeugt. Das Resultat ist ein hochaufgelöster 3D Datensatz. Ein Kanal enthält anatomische Information und ein zweiter Kanal enthält die Genexpressionsmuster.

In dieser Diplomarbeit untersuchen wir Ansätze um diese neue Art von Daten zu verbessern, zu registrieren und zu visualisieren. Wir befassen uns mit der unausgewogenen Beleuchtungssituation welche diese neue Methode mit sich bringt. Wir entwickelten einen Algorithmus um eine quadratische Fläche durch die Hintergrundpixel zu legen um die Beleuchtungssituation in den aufgenommenen Schnittbildern zu schätzen. Diese Schätzung wird zum Korrigieren des Datensatzes verwendet.

Weiters wurde eine Erweiterung dieser Methode erforscht. Dabei werden geschnittene Sektionen ein erneutes Mal gefärbt, um weitere Informationen in den Datensatz einzubringen. Das resultiert in einem zweiten Datensatz von unausgerichteten Schnittbildern. Da diese Erweiterung manuell ist, werden dadurch nicht lineare Deformationen eingeführt. Wir untersuchten verschiedene Registrierungsalgorithmen um die zwei Bildstapel zueinander auszurichten. Wir fanden heraus, dass ein Registrierungsansatz in zwei Schritten die besten Resultate liefert. Im ersten Schritt wird eine affine Registrierung benutzt um die Datensätze ungefähr auszurichten. Das Resultat dieser Registrierung wird vom Benutzer untersucht und gegebenenfalls korrigiert. Im zweiten Schritt wird eine b-spline Registrierung durchgeführt um die nicht linearen Deformationen in den Schnittbildern zu kompensieren.

Um die Registrierungsresultate visuell zu untersuchen und um einen Überblick über beide Datensätze zu bekommen präsentieren wir zwei Visualisierungsmethoden. "Checkerboard" Visualisierung wird benutzt um 2D Schnittbilder miteinander zu vergleichen. Eine dreidimensionale Methode basierend auf "direct volume rendering", bei der wir die Gewebsgrenzen mittels "gradient magnitude opacity modulation" betonen wird benutzt um die gesamten Datensätze miteinander zu vergleichen.

Contents

1	Introduction	1
1.1	Overview	1
1.2	Data Acquisition	2
1.2.1	Acquisition of the HREM data	3
1.2.2	Acquisition of the Histological data	3
1.3	Artifacts of HREM Data	8
1.3.1	Uneven Illumination	8
1.3.2	Shadow Artifacts	8
1.3.3	Sectioning Artifacts	11
1.4	Artifacts in the Histological Data	11
2	Background and Related Work	15
2.1	Classification of Registration Algorithms	15
2.1.1	Extrinsic Registration	16
2.1.2	Intrinsic Registration	16
2.1.3	Non Image Based Registration Techniques	19
2.1.4	Other Classifications	19
2.2	The Registration Framework	20
2.2.1	Transformations	20
2.2.2	Interpolators	23
2.2.3	Similarity Measures	27
2.2.4	Optimization	29
2.2.5	Extensions	31
3	Enhancement and Registration of HREM Data	32
3.1	Motivation	32
3.2	Illumination Correction	32
3.3	Registration Setup	36
3.4	Registration Results and Problems	43
3.4.1	Affine Registration	43
3.4.2	B-spline Registration	48

4	Visualization of the Registered Volumes	54
4.1	2D Visualization - Checkerboard View	54
4.2	3D Visualization - Direct Volume Visualization	57
4.2.1	Direct Volume Rendering	58
4.2.2	Implementation of the 3D Visualization	60
5	Summary	67
5.1	Introduction	67
5.2	Background and Related Work	68
5.3	Preprocessing of HREM Data	69
5.3.1	Illumination Correction	69
5.3.2	Registration	70
5.4	Visualization of the Registered Volumes	70
6	Conclusion	72

Chapter 1

Introduction

1.1 Overview

In this thesis we examine the preprocessing and visualization of High Resolution Episcopic Microscopy (HREM) data. Medical domain scientists use the HREM technique to explore the impact of genes on the cell growth of different tissue. This is done by staining the specimen and then generating the data by slicing and capturing the specimen. The slices are captured using two different filter sets. The result is a 3D dataset with two channels of different bands of the light spectrum. One channel shows the anatomy, while the other shows the stained regions. In Section 1.2 the data acquisition process is discussed in more detail. We provide a brief overview of the devices that are involved in data acquisition and show typical examples of HREM data.

The process of acquiring an HREM dataset involves mechanical slicing and indirect illumination of the specimen. These imperfections of the acquisition process introduce specific artifacts that differ significantly from other 3D imaging modalities. In Section 1.3 we give an overview of the artifacts that are specific to HREM data.

The generated HREM datasets show the general morphology of the specimen and the stained regions. The medical domain scientists experimentally extended the data acquisition process by staining the already cut slices a second time using standard histological staining techniques. These slices are manually stained and photographed using a standard light microscope. The histological staining generates an additional dataset, showing different features of the same specimen. The acquired dataset consists of unaligned slices that need to be registered for further processing. The histological dataset also lacks registration to the HREM dataset. The goal of our work was to enhance the data of the HREM modality as well as the registration of the experimental extension of the modality. In Chapter 2 we fo-

cus on the related work about registration and give a general overview of existing techniques.

in Chapter 3 we discuss the enhancement and registration of the datasets. One of the most hindering artifact is the uneven illumination of the slices. We developed a technique to correct the illumination. In Section 3.2 we give details about the illumination correction algorithm. Section 3.3 deals with the issues involved in registration. We explored different registration techniques to align histological slices with HREM slices. We found that an image based registration approach fits our needs best. Due to the manual processing of the histologically stained slices, the resulting images include additional artifacts. The two most prominent artifacts are a different field of view for HREM images, and non-linear deformations of the slices. The non-linear deformation happens because the slice curls up after cutting and needs to be manually embedded in a liquid. The differences in the field of view occur because each slice is handled separately without alignment. This leads to a linear transformation (because of the different field of view) and a non-linear transformation (because of the curling and the handling of the slice). To compensate these two transformation we investigated a two step registration process.

For the visual inspection the results of the preprocessing stage we developed 2D and 3D visualization methods. An image based comparison method was implemented for the inspection of the registration result [18]. We implemented a checkerboard method which is widely used for the inspection of registration results. Our second approach uses three dimensional direct volume rendering (DVR). We adopted DVR by incorporating a focus and context technique using a movable focus region. The context region shows the morphology of the original HREM volume, while the focus region shows the registration result. Further we improved the visibility of the tissue borders where registration errors are most easily detectable. We used a gradient magnitude opacity-modulation approach to enhance surface boundaries. In Chapter 4 the two visualization methods, are described including mathematical background and result images.

1.2 Data Acquisition

The data acquisition process is divided into two intertwined parts. The first part is to acquire High Resolution Episcopic Microscopy images which will be described in Section 1.2.1. The second part is to acquire corresponding photographs using histological colorings described in Section 1.2.2.

1.2.1 Acquisition of the HREM data

The first step of the HREM acquisition process is to stain the specimen using different labeling or staining techniques that lead to fluorescence of the stained regions. After the staining, the specimen is prepared for the sectioning. This is done by embedding the stained specimen in standard resin (i.e., JB4 developed by Polyscience [19]). After hardening, the block is mounted onto a rotary microtome (micro-Tec CUT 4060 E [17]). A block prepared for mounting is shown in Figure 1.1. The microtome was modified so that the stopping point of the block holder is kept constant after each rotation step with an accuracy of about one μm . For capturing one slice an episcopic microscope is used. Episcopic microscopes capture the reflected light in contrast to light microscopes where transmitted light is captured. The optical axis of a modified episcopic microscope (Leica DM LM [10]) was aligned with the stopping point. The microscope is equipped with fluorescence filter sets and a digital camera (Leica DC 480 [10]). The setup is depicted in Figure 1.2.

Each surface of the block is photographed after sectioning. Fluorescence filter sets are used to separate the light originating in the stained tissue from the light reflected from the unstained tissue. The result is an RGB image where each light source is stored in one channel. The image series resulting from this acquisition process consists of precisely aligned images of the specimen. The section thickness can be set from 0.5 to 10 μm . The channels of the RGB image are decomposed resulting in two gray scale images. One image shows the morphology, and the other one shows the signal of the stained tissue. An example of an image pair is shown in Figure 1.3. A closeup of each image is shown in Figure 1.4. A more detailed description of the HREM acquisition process can be found in the work of Weninger et al. [25]. The acquisition of 1000 sections takes about four hours. The process is fully automatic after the relatively short initial setup, consisting of mounting the block for sectioning, setting the field of view and focusing for capturing.

1.2.2 Acquisition of the Histological data

To acquire histologically colored images corresponding to the HREM sections an operator of the HREM setup has to stop the sectioning process manually, and save the actual section after cutting. This is only possible for slices that are thick enough (i.e., thicker than 2 μm). After the slice is saved it is put into a water quench. Slices which are curled after the sectioning, expand themselves in water. The expanded slice is taken out using a glass slide. The slice is colored using common histological colorings. An image of this colored slice is then acquired using a standard optical light microscope. Since this process is part of an experimental



Figure 1.1: A prepared and hardened block ready for the following sectioning on the microtome.






Figure 1.2: The setup for the High Resolution Episcopic Microscopy (HREM) data acquisition. (image taken from [25]): On the right side the rotary microtome can be seen. The episcopic microscope is mounted on the metal rod in the middle, and the digital camera can be seen on the left side.




(a) morphology image




(b) signal image

Figure 1.3: (a) shows the morphology image and (b) shows the signal image of a slice of a chicken heart. The actual signal of the stained tissue is visible in image (b) as dark regions.



(a) close-up of the morphology image



(b) close-up of the signal image

Figure 1.4: (a) shows a magnified region of the morphology image from Figure 1.3 and (b) shows the same region of the signal image of the corresponding slice.

setup it is performed manual. The manual intervention makes the acquisition of a histological dataset that corresponds to the HREM dataset very time consuming. Therefore the experimental setup typically allows the acquisition of a few slices. Nevertheless since the histological colorings are well established in medical research they are very useful to verify the quality of the new staining techniques employed in the HREM acquisition process. Another benefit is that, due to the possibility of staining the specimen a second, it is possible to incorporate new information into the HREM dataset.

1.3 Artifacts of HREM Data

The HREM data acquisition technique is still in an experimental stage. Therefore artifacts need be considered using this data. Some artifacts originate from the mechanical side of the data acquisition, others originate from the setup of the data acquisition. The different classes of artifacts are described in the following Sections.

1.3.1 Uneven Illumination

Due to the data acquisition setup, the images are not illuminated evenly. Figure 1.5 shows an image taken at a position of the block, where no actual specimen is present. The illumination situation considerably varies within one slice and also varies between datasets. It is possible that the lighting situation changes between consecutive slices for several reasons (e.g., the cleaning personal opens the door and turns on the light while the data acquisition is in progress). The uneven lighting causes different pixels to have different intensity values for the same signal. The illumination artifact hinders the use of simple segmentation techniques (e.g., threshold segmentation). Direct volume visualization of the HREM data is also more challenging since common transfer functions cannot be used to select features with constant intensity.

1.3.2 Shadow Artifacts

Stained regions from below the surface of the block shine through causing artifacts that appear as shadows on the image. An example of such an artifact is shown in Figure 1.6. The embedding medium is semi-transparent and therefore does not fully occlude structures below the topmost section. Structures below the current section appear blurred and with less intensity. The depth of the section which is shining through, and therefore adding wrong information to the image, depends on the type of embedding medium, the actual slicing thickness and the illumination



Figure 1.5: A slice taken at a position of the block where no actual specimen is present. It clearly shows the uneven illumination.




Figure 1.6: Example of a shadow artifact (circled in red): Shadow artifacts originate from structures that shine through the surface. The top of the block is photographed also capturing light that shines through the semi-transparent embedding medium.

situation. In the signal channel the shadow artifacts are more severe since more light from the underlying fluorescent stained structures shines through. Figure 1.7 shows a schematical view of shadow artifacts. The lower image of Figure 1.7 shows a cross section of the block with the embedded specimen. The upper image of Figure 1.7 shows the corresponding image that is captured from the top. The red part of the block shows the region that will be sliced after image acquisition. It is desirable that parts of the specimen embedded in this region contribute to the captured image. Stained parts of the specimen embedded in the yellow region of the block also contribute to the acquired image. Since these parts also contribute to subsequent images they are considered to be artifacts.

1.3.3 Sectioning Artifacts

Figure 1.8 shows the mechanical artifacts caused by the notches in the blade, which is used to cut the block. The artifacts appear as light gray stripes on the image. The width of these stripes is about eight pixels in the original data set, but can vary due to the dents in the blade. This artifact gets worse with a more used blade (i.e., towards the end of the block). Most of the embedded object is not as reflective as the embedding material causing the artifacts to vanish in the presence of the object. Therefore this class of artifacts is neglected in our work. Nevertheless when a new dataset is created it should be checked how grave the impact of these artifacts on the structures of interest is. The user also needs to be aware that slicing artifacts exist and can cause problems for further processing.

1.4 Artifacts in the Histological Data

The acquisition of the histological data is completely manual. Therefore the acquisition process introduces artifacts not found in the HREM dataset. Since each section is captured separately, the intensity can vary between consecutive sections. Because of the manual nature of this acquisition technique contaminations on the sections are common. Even if it is not a problem to capture the structures of interest within the sections, the acquired images are not aligned to each other (i.e., the field of view is different for each section). The most severe problem is that parts of the slice can flap over destroying the whole slice. All artifacts of histological data are shown in Figure 1.9.

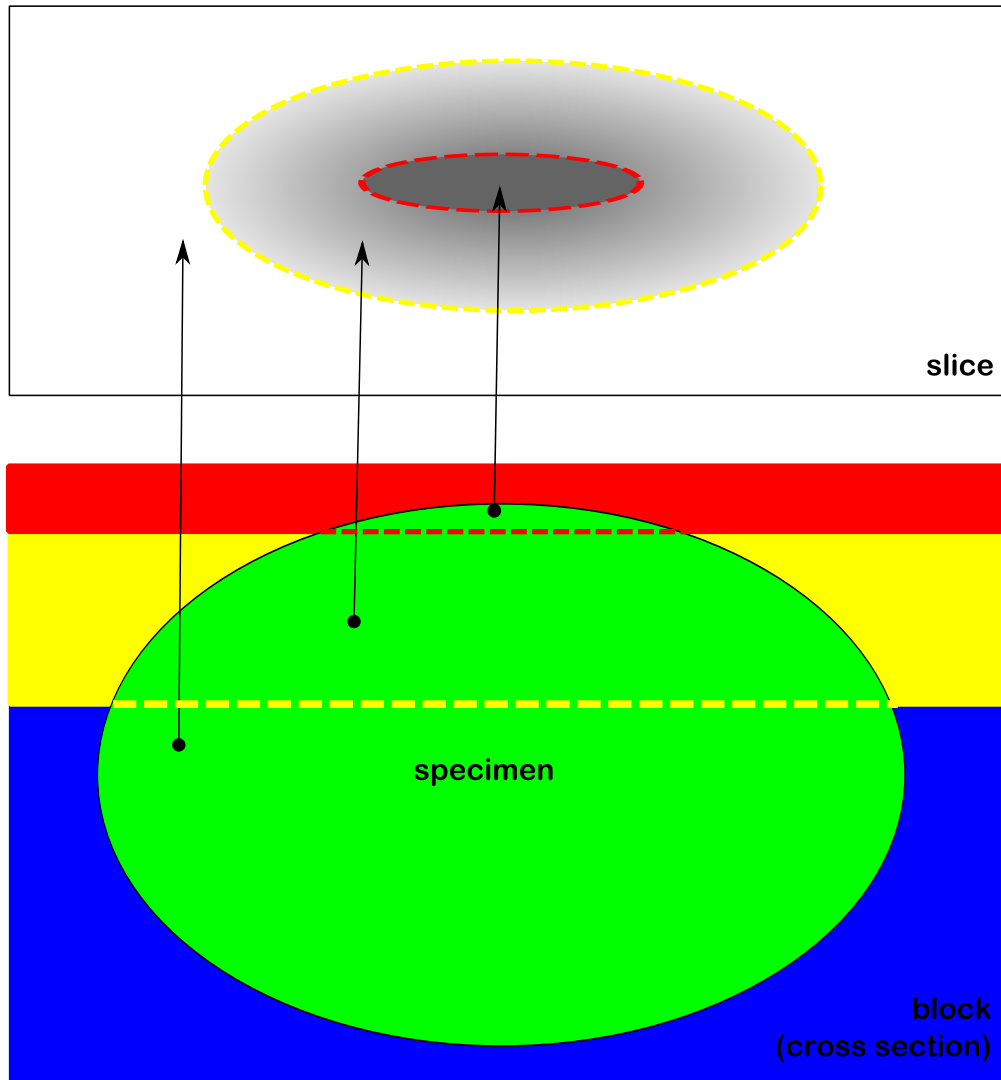


Figure 1.7: The square is representing the block with the embedded specimen (green). The block is divided into three sections. The red section represents the section which is photographed and then cut away. The yellow section represents the area which is shining through onto the current slice, contributing the shadow artifacts. The blue section represents the rest of the block.




Figure 1.8: The light stripes (indicated with the arrows) are the sectioning artifacts caused by dents on the blade.

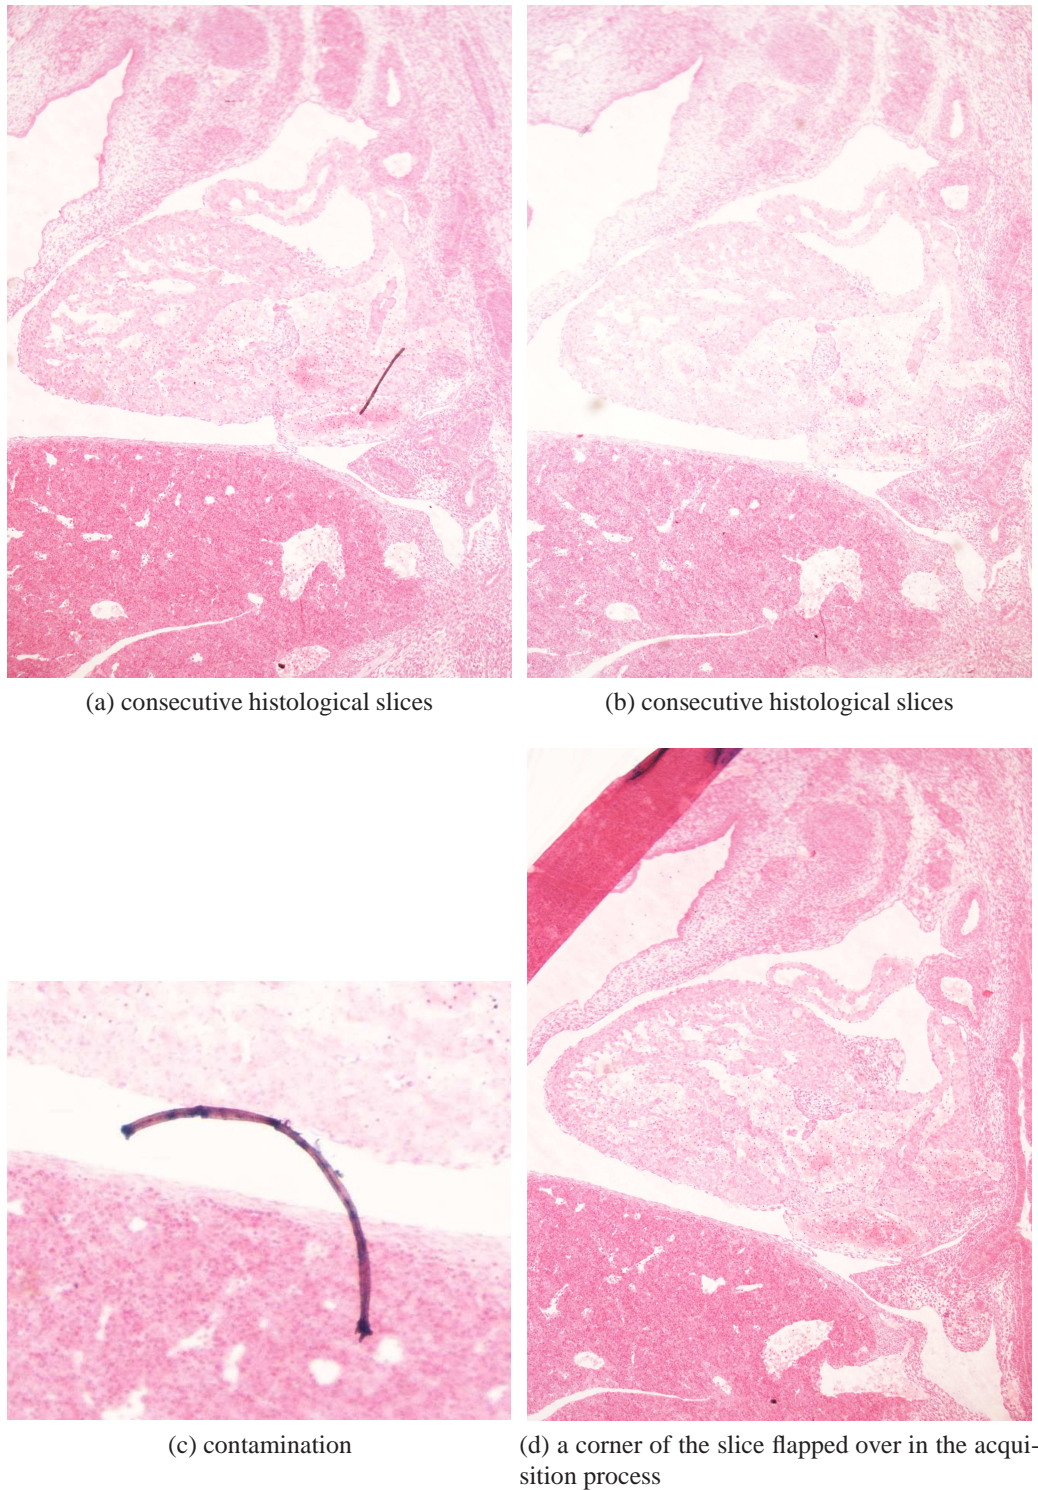


Figure 1.9: The different types of artifacts in histological images. Images (a) and (b) show the intensity variation and the varying field of view between consecutive slices. Image (c) shows a magnified contamination resulting from the acquisition process, and image (d) is a slice where a corner flapped over while preparing the slice.

Chapter 2

Background and Related Work

Registration is the process of aligning two or more datasets (e.g., 2D images, 3D grids, point clouds, etc.), that represent the same scene or parts of the same scene. The datasets might be measured from different viewpoints, at different times or with different modalities. Therefore registration is the process of finding a specific transformation which maps the points of one dataset to corresponding points of the other dataset. Three major areas where this task is frequently needed, are:

- **Computer Vision and Pattern Recognition:** Object recognition, shape reconstruction, motion tracking, stereo-mapping, etc. are dependent on specific forms (mostly specialized techniques) of registration.
- **Medical Image Analysis:** Diagnostics, surgery planning, analysis of microscopic images, comparative studies, etc. need the registration of datasets taken from different modalities or at different times.
- **Remote Sensing:** Satellite data processing requires the registration of many satellite images. Oceanography, geology, oil- and mineral exploration, pollution, disaster and urban studies require the registration of either satellite data or aerial photos.

In all three areas of research specific techniques have been developed to achieve this goal. The remainder of this chapter focuses mainly on the registration of medical data.

2.1 Classification of Registration Algorithms

In medical image analysis, registration is used to align two- or three dimensional datasets. The techniques developed to achieve this task can be classified in several ways. Although there are many classifications in literature, the classifications




Figure 2.1: *SIP Lab Innsbruck* frame as an example for a head mounted frame for extrinsic registration. Image taken from [3]

presented here are the most commonly used ones and are easily comprehensible. For more complete overviews refer to the image registration literature. Specifically Brown [2], Mainz et Viergever [15] and Zitova et Flusser [26] present extensive descriptions of registration algorithms. One way to classify registration is by which data of the dataset is used to achieve the registration:

2.1.1 Extrinsic Registration

Extrinsic Registration depends on artificial markers introduced in the image space. These markers are attached to the patient, and are easily detectable in all image modalities. The most common way to attach the markers is to let the patient wear a stereo tactic frame screwed to the outer skull. Another way is to use invasive markers such as screw mounted markers. In Figure 2.1 you can see the *SIP Lab Innsbruck frame* as an example for a head mounted frame for extrinsic registration. The actual registration is only performed on the markers to find a transformation that maps the markers in one image to the markers in the other image. Since only the position of the markers is used for the registration, this technique is typically very fast and can be automated. Often it can be modeled explicitly without the use of sophisticated optimization algorithms. The main drawback of this technique is that markers are not suitable for all applications (e.g., scans of inner organs). Another drawback of this method is that deformations that occur for example due to patient movement, cannot be modeled.

2.1.2 Intrinsic Registration

Intrinsic Registration makes only use of the image content of the generated images. The Intrinsic registration techniques in this Section can be further classified by the amount of image data that is used for the registration process.

- **Landmark Based Registration Techniques** use a number of manually selected salient points of both images. These selected points are usually accurately locatable points of the morphology which are manually selected by an expert. Since the selected landmarks are relatively sparse in contrast to the whole image content, this class of registration algorithms is typically very fast. In theory this technique can be used to find complex transformations. Landmark based registration methods can even model non-linear deformations, but their accuracy is highly dependent on the number of selected landmark points. The more landmarks are used the more complex transformations can be modeled reliably. The main drawback of this technique is that an expert is needed to select the landmarks. Errors made in the selection process can deteriorate the whole registration to the point where the result is useless. Figure 2.2 shows the process of landmark registration. First the salient points are chosen manually (first row), then a correspondence between the points of both images is made (second row). This information is then used to register the images. The result is shown in the third row.
- **Segmentation Based Registration Techniques** utilize segmented structures (mostly surfaces, but also curves and volumes) of the image content. Another variant of this technique uses predefined shapes for the registration process. Since only part of the image content is used, the algorithms of this class are fast, like the landmark based algorithms. Depending on the task at hand, any transformation between the selected areas can be used. The main drawbacks of this technique is similar to the drawbacks of the landmark based techniques. Therefore there is the need of an expert to segment the structures, which is more time consuming than the selection of landmarks and should be assisted with appropriate segmentation tools. Furthermore the accuracy of the segmentation is very important to the quality of the registration process. Another drawback is that the registration only models the transformation between the segmented areas, and therefore the registration of the rest of the image area can be faulty.
- **Pixel Based Techniques** use the content of the whole image for the registration task (this technique is also called voxel based for 3D registrations, but in this thesis it is referred to as pixel based for both 2D and 3D registrations). To compare the alignment of the images a similarity measure is used. In theory this technique is the most powerful and accurate since all available information is used in the registration process. It is possible to use any transformation with this approach. The main drawback is the tremendous computational cost with this technique. However, since computational




Figure 2.2: The process of landmark based registration. In the first two rows the salient points are selected and matched using invariant descriptors. The final row shows the result of the registration. (Figure taken from Zitova et Flusser [26])

power gets cheaper and more easily available, pixel based registration techniques are very promising in the future. Section 2.2 is focusing on this class of registration techniques.

2.1.3 Non Image Based Registration Techniques

Non Image Based Registration Techniques don't use any of the image content. This is only possible if the coordinate systems of the different acquisition techniques can be calibrated on each other. This calibration data is then used for the registration. Although this is not always possible, it is a very fast and easy means of registration. An example of such a scenario is the tracking of a hand held scanner (e.g. ultrasound) and a fixed scanning device. It is possible to use the tracking data to map both coordinate systems to each other.

2.1.4 Other Classifications

Another classification is based on the image acquisition process:

- **Multi-Modal Analysis** uses images of different modalities. This is done to incorporate the information acquired from different scanners (e.g, CT and MRI, SPECT and MRI, CT and Xray, etc.). This is often used for diagnostics and surgery planning.
- **Multi-Temporal Analysis** uses images from the same modality acquired at different times. The result of this registration is mostly used for diagnostics (e.g, progression of a disease) and treatment planning.

The last classification presented here is based on the subject that is examined:

- **Intra-Patient Registration** uses images of the same patient captured at different times or with different modalities. This is by far the most common class of registration in clinical usage.
- **Inter-Patient Registration** registers images taken from different patients with the same modality. This is used in comparative studies (for example in Alzheimer studies where scans of healthy and sick patients are compared).
- **Patient to Atlas Registration** is the type of registration where the image of a single patient is registered to an image information database constructed from images acquired of many subjects (i.e., an atlas).

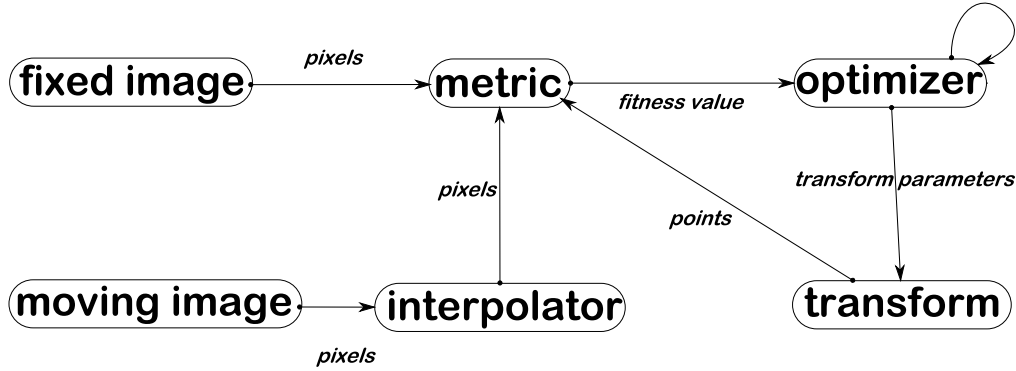


Figure 2.3: The Registration Framework

2.2 The Registration Framework

In the classical registration framework, shown in Figure 2.3, for pixel based registration two images are used. The fixed or reference image, will not be altered, while the moving or floating image, is transformed to match the fixed image. During transformation of the moving image we often need to evaluate positions between the pixel locations since both images are discrete an interpolator is needed. The similarity measure tells us how good two images are registered. The evaluated similarity measure using the actual transformation parameters is called fitness value and is an indication of the quality of the current fit. The optimization algorithm tries to find parameters of the transformation for which the fitness value reaches an optimum. If this optimum is a global optimum, the registration process was successful. It is possible for the optimization algorithm to get stuck in a local optimum which would result in a faulty registration. Therefore it is important to assure that the optimizer finds the global optimum. The following Sections explain the parts of the registration framework in more detail.

2.2.1 Transformations

Often the class of transformation between the fixed image and the moving image is known. Therefore the transformation is often the first part of the framework which is chosen. Figure 2.4 shows the result of different transformations. The first row shows the fixed and moving image. The second row shows transformation results for (c) rigid-, (d) affine- and (e) b-spline-transformations.

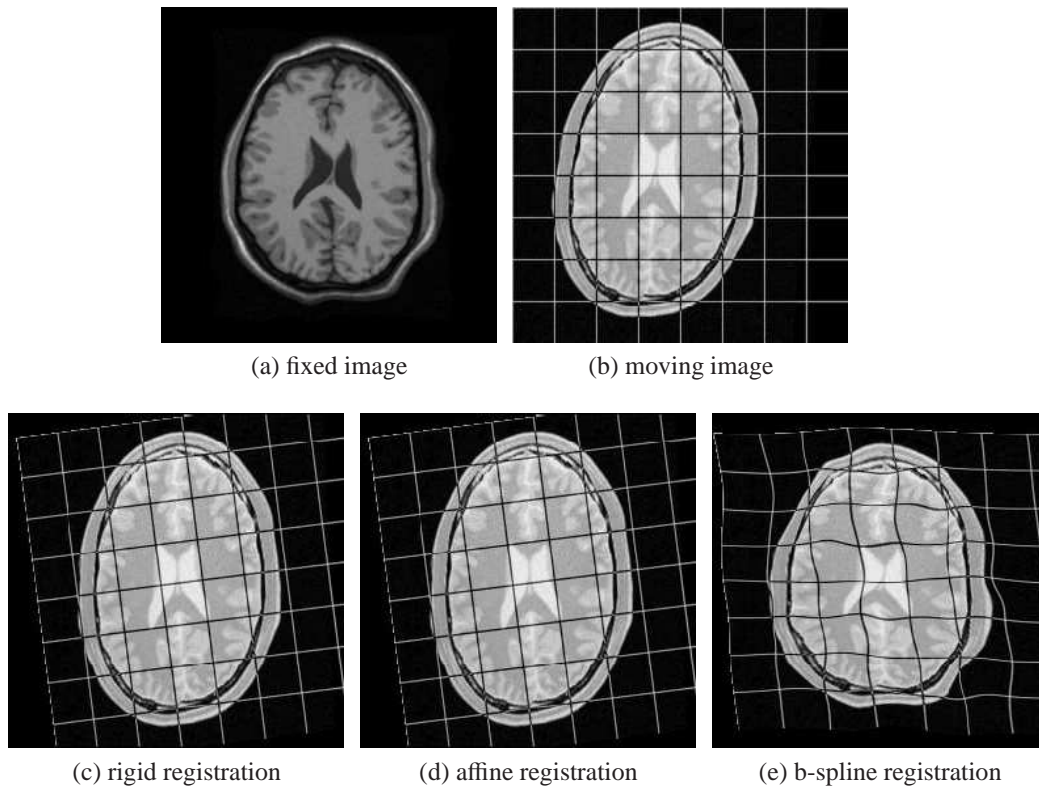


Figure 2.4: Comparison of the different transformations. First row: original fixed image (a), original moving image (b). Second row: result of a rigid registration (c), affine registration (d), b-spline registration (e). Images taken from the elastix manual [8]

Rigid Transformation

A rigid transformation consists of a translation and a rotation, therefore it has four parameters in 2D and six parameters in 3D. Because of the few parameters it is the fastest transformation commonly used for registration. The formal expression of the rigid transformation T_{rigid} of a location $\vec{l} = (l_x, l_y, l_z)^T$ in 3D is represented by a translation vector $\vec{t} = (t_x, t_y, t_z)^T$ and by a rotation matrix R :

$$T_{rigid}(\vec{l}) = R \cdot \vec{l} + \vec{t} \quad (2.1)$$

where the rotation matrix R is a 3x3 Matrix constructed from the Euler angles $\theta_x, \theta_y, \theta_z$ when rotated first around the x- then y- and z-axis and \cos_i is an abbreviation of $\cos(\theta_i)$ where $i \in \{x, y, z\}$:

$$R = \begin{pmatrix} \cos_y \cos_z & \sin_x \sin_z - \cos_x \cos_z \sin_y & \cos_z \sin_x \sin_y + \cos_x \sin_z \\ \sin_y & \cos_x \cos_y & -\cos_y \sin_x \\ -\cos_y \sin_z & \cos_z \sin_y + \cos_x \sin_y \sin_z & \cos_x \cos_z - \sin_x \sin_y \sin_z \end{pmatrix} \quad (2.2)$$

Using homogeneous coordinates the rigid registration can be expressed as:

$$T_{rigid}(\vec{l}) = \begin{pmatrix} R & \begin{matrix} t_x \\ t_y \\ t_z \end{matrix} \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} l_x \\ l_y \\ l_z \\ 1 \end{pmatrix} \quad (2.3)$$

Affine Transformation

A more powerful but still linear transformation is the affine transformation. It is able to handle translation, rotation, scaling, shearing and mirroring. It consists of six parameters in 2D and twelve parameters in 3D. Since this transformation is more general than the rigid transformation it can be written as a multiplication of the rigid transformation with a shear matrix and a scale matrix, using the scaling factors $\vec{s} = (s_x, s_y, s_z)^T$ and the shear vector $\vec{g} = (g_x, g_y, g_z)^T$. The affine transformation (T_{affine}) is given in homogeneous coordinates as:

$$T_{affine} = T_{rigid} \cdot G \cdot S \quad (2.4)$$

Where G is represents the shearing matrix:

$$G = \begin{pmatrix} 1 & g_x \cdot g_z & g_z & 0 \\ g_y & 1 & 0 & 0 \\ 0 & g_z & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (2.5)$$

and S is the scaling matrix :

$$S = \begin{pmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (2.6)$$

Spline Transformation

To model non-linear deformations it is common to use spline transformations described by a grid of control points. The number of parameters depends on the size of the grid. Each control point can be translated in two directions in 2D and three directions in 3D. Therefore the number of parameters for a spline transform is n^d where d is the dimensionality and n the number of control points. Typically n is very large and therefore the computational cost for spline-transformations is tremendous. Two examples of splines that are often used are b-splines or thin-plate-splines. Since this transformation has lots of independent parameters (i.e., the location of the control points) it has big problems with large translations, rotations and scalings. If translations, rotations and scalings need to be handled in addition of non-linear deformations, the registration task is usually split in two parts. The big translations, rotations and scalings are handled with an affine registration followed by a spline registration handling the non-linear transformations.

2.2.2 Interpolators

Since the images for the registration are discrete (that is we have intensities on discrete raster positions) a function to reconstruct data between the raster positions (see Figure 2.5) is needed. Commonly an interpolator is used to compute the values between raster positions. Since the interpolator is needed for each data value evaluation it has a huge impact on the computational cost of the registration algorithm. Furthermore, the quality of the interpolation has an impact on the similarity measure. Figure 2.6 shows the value of the similarity measure (mutual information in this case) of two images when a simple translation is done. Typically the value of the similarity measure decreases monotonically as the moves farther away from the original position. In presence of an imperfect interpolator the similarity measure has a jagged appearance. The spikes introduced by the interpolator make the task of finding a global optimum harder.




Figure 2.5: The rasters of the fixed image (black) and the moving image (blue). Interpolators are needed since positions within the raster have to be evaluated.

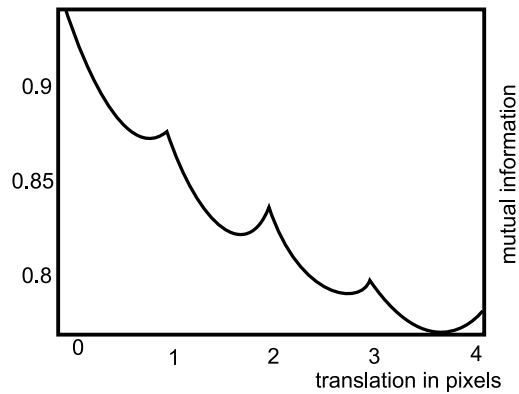


Figure 2.6: The value of the mutual information similarity measure of two images when a simple translation is performed. The horizontal axis shows the distance of the translation in pixels. The mutual information value should decrease monotonically as the image moves farther from the original position. With the presence of an imperfect interpolator the similarity measure has this jagged appearance.




Figure 2.7: The curve in red is called the Runge function. The blue curve is a interpolation using a 5th-order-polynomial and the green curve uses a 9th-order-polynomial on 10 equally spaced points. For higher interpolating polynomials Runge's phenomenon (i.e. the oscillating of the interpolated function near the border) gets worse. The error between the original and the interpolated function at the border therefore gets higher.




Figure 2.8: Three different types of interpolators (from left to right): nearest neighbor interpolator, linear interpolator, b-spline interpolator. While the nearest neighbor interpolator has very poor quality (i.e., it is very grainy), the bilinear interpolator has a better quality (i.e., it is more blurry and loses the fine details). The b-spline interpolator has good quality (i.e., the fine detail in white is still visible and not blurred).

Nearest Neighbor Interpolator

The nearest neighbor interpolator takes the nearest raster position as the interpolated value. An example of the nearest neighbor interpolator is shown in the left image of Figure 2.8. This operation is very fast since no floating-point operations are needed for the computation. On the other hand this interpolator has very low quality.

Linear Interpolator

Linear interpolation (bi-linear in 2D and tri-linear in 3D) assumes that the values between the raster positions vary linearly. An example is shown in the middle image of Figure 2.8. Linear interpolation is relatively fast although it uses floating point operations. It can easily be implemented in hardware to further speed up processing. Bi- and tri-linear interpolation is commonly available on nowadays graphics adapters.

Polynomial Interpolator

Polynomial interpolation is a generalized version of the linear interpolation, where the linear function is replaced by a polynomial of higher degree. Even as the computational complexity rises using higher order polynomials, the interpolated result is not exact. Especially at borders Runge's phenomenon occurs. Because of Runge's phenomenon the interpolated function tends to oscillate near the border

of the interpolating points negatively affecting the result at the borders. Figure 2.7 shows higher order interpolators for a given function. The curve in red is called the Runge function. The blue curve is a interpolation using a 5th-order-polynomial and the green curve uses a 9th-order-polynomial on 10 equally spaced points. For higher interpolating polynomials Runge's phenomenon (i.e. the oscillating of the interpolated function near the border) gets worse. The error between the original and the interpolated function at the border therefore gets higher.

Spline Interpolation

To overcome the drawbacks of the polynomial interpolation, spline interpolation can be used. An example can be found in the right image in Figure 2.8. Spline interpolation uses low-degree polynomials in each interpolation interval. The polynomials are chosen to fit smoothly together. Since lower-degree polynomials fit together piecewise the computational cost of this interpolation is lower than the cost of the polynomial interpolation while avoiding Runge's phenomenon.

2.2.3 Similarity Measures

To register two images we need means to quantify the goodness of the fit between the images. Considering pixel based registration methods this similarity measure has to be extracted directly from the pixel intensities. Usually pixel based registration makes use of the pixel intensities. However other approaches combine the similarity measure from pixel intensities with a penalty term to handle rigid structures within the image content. For more information about penalty terms refer to Maintz et al. [16], Rueckert et al. [21], Guimond et al. [5] and Hellieret al. [6]. To choose an adequate similarity measure the imaging modality needs to be considered. Corresponding structures are measured at different intensities when the images originate from different modalities. In the following four common similarity measures are explained.

Intensity difference based similarity measures

The sum of squared intensity differences (*SSD*) between both images is one of the easiest measures available. The SSD can be expressed as :

$$SSD = \frac{1}{N} \sum_{\vec{s}_A \in \Omega_{AB}^T} (A(\vec{s}_A) - B^T(\vec{s}_A))^2 \quad (2.7)$$

where A is the fixed image, B^T is the transformed moving image and N is the number of pixels from the overlapping region Ω_{AB}^T . The normalization makes the

similarity measure invariant to the number of pixels within the overlap of both images. Viola et Wells [24] state that this similarity measure is the best choice when registering two images from the same modality where only Gaussian noise makes the difference. However, even if the images originate from the same modality, they may contain an intensity bias, which lets the SSD fail. To overcome this problem, the sum of absolute differences (*SAD*) could be used as an alternative, which is expressed as:

$$SAD = \frac{1}{N} \sum_{\vec{s}_A \in \Omega_{AB}^T} |A(\vec{s}_A) - B^T(\vec{s}_A)| \quad (2.8)$$

Although *SAD* is better than *SSD* in presence of bias, *SAD* cannot overcome the problem with the bias completely. Therefore more stable similarity measures have been derived.

Cross Correlation based similarity measures

Cross correlation (*CC*) is formally expressed as:

$$CC = \frac{\sum_{\vec{s}_A \in \Omega_{AB}^T} (A(\vec{s}_A) - \bar{A}) \cdot (B^T(\vec{s}_A) - \bar{B})}{\sum_{\vec{s}_A \in \Omega_{AB}^T} (A(\vec{s}_A) - \bar{A})^2 \cdot (B^T(\vec{s}_A) - \bar{B})^2}^{\frac{1}{2}} \quad (2.9)$$

where \bar{A} and \bar{B} are the mean values of the pixel intensities within the overlap of A and B^T . This similarity measure is a good choice for mono-modal images, but it also works good for multi-modal images where the intensities are linear dependent. It is also able to overcome the problem of a bias in mono-modal images.

Joint entropy based similarity measures

Shannon [22] invented the concept of entropy as a measure of information. The entropy is formally expressed as:

$$H = - \sum p(s) \log p(s) \quad (2.10)$$

where $p(s)$ is the probability of the symbol s . This is also a measure of randomness where the entropy reaches its maximum if all symbols have the same probability, and it reaches a minimum when all but one symbols have the probability of zero. This property was the reason that joint entropy was introduced in registration. Studholm et al. [23] and Collingnon et al. [4] proposed to minimize

the joined entropy calculated from a histogram. The joint entropy of two images is given by:

$$H(A, B^T) = - \sum_{a \in \Omega_a^T} \sum_{b \in \Omega_b^T} p_{AB}^T(a, b) \log p_{AB}^T(a, b) \quad (2.11)$$

where $b \in \Omega_b^T$ represent the discrete intensity values if that are transformed over each of the intensities $a \in \Omega_a^T$ and $p_{AB}^T(a, b)$ represents the probability that a and b occur in the overlap region.

Mutual Information based similarity measures

As an extension of the joint entropy concept, the mutual information approach adds the individual information carried by the overlap Ω_{AB}^T . Expressing this information in terms of marginal entropies $H(A)$ and $H(B^T)$ the mutual information is formally expressed as:

$$I(A, B^T) = H(A) + H(B^T) - H(A, B^T) \quad (2.12)$$

where the marginal entropies are calculated as:

$$H(A) = - \sum_{a \in \Omega_a^T} p_A^T(a) \log p_A^T(a) \quad (2.13)$$

$$H(B^T) = - \sum_{b \in \Omega_b^T} p_B^T(b) \log p_B^T(b) \quad (2.14)$$

and p_A^T and p_B^T represent the marginal intensity distributions of the images A and B within the overlapping region Ω_{AB}^T . An important property is, that the marginal entropies and the joint entropy vary during the registration process. This variation is caused by the overlapping region which changes because of the estimated transformation. This means that the marginal entropies and the joint entropy has to be calculated for each new set of transformation parameters (i.e. in each iteration of the optimization process). Mutual information represents the information that one image contains about the other, and it reaches a maximum if both images are aligned correctly. Mutual information is the most promising approach when considering multi-modal registration where no a-priori assumption about the intensities can be made.

2.2.4 Optimization

All pixel based registration techniques use similarity measures S to quantify the goodness of the registration. Since no analytical solution for the registration task

exists, numerical optimization algorithms have to be used to find the optimal transformation $t(x)$ of the moving image I_M so that the moving image matches the fixed image I_F at every position x . As discussed in the subsection 2.2.1, each transformation is defined by a parameter vector μ . These parameters are the degrees of freedom of the optimization process. An optimization algorithm explores this parameter space in an iterative fashion using trial and error, beginning with an initial parameter estimation. While exploring the parameter space the fitness value is used as a function of the parameter set which is optimized. The current estimate of the transformation parameters is used to compute the actual value of the similarity measure. After that, the parameters are refined to maximize the fitness value. Usually the stopping criterion is defined as a minimum tolerance value for the changes of the similarity measure between consecutive iterations. This optimization is expressed mathematically as:

$$\hat{\mu} = \arg \min_{\mu} S(\mu, I_M, I_F) \quad (2.15)$$

Since there may be many local minimas, evaluation of the parameter vector $\hat{\mu}$ strongly depends on the optimization algorithm and the initial parameters. A way to suppress local minimas which should be avoided is to add a regularization term R to the optimization to penalize unwanted transformations. This is especially useful with non-linear transformations where some deformations might be undesired or just physically impossible (e.g., deformations of rigid structures like bones). The final expression including the regularization term is:

$$\hat{\mu} = \arg \min_{\mu} S(\mu, I_M, I_F) + \omega R(\mu) \quad (2.16)$$

Where ω is a weighting factor for the regularization term R , describing how strong the influence of the regularization term should be. Examples of regularization terms are curvature dependent penalty terms, volume preserving penalty terms and elastic energy penalty terms. The iterative fashion in which the parameter space is explored is described as:

$$\mu_{k+1} = \mu_k + a_k \mathbf{d}_k, k = 0, 1, 2, 3, \dots \quad (2.17)$$

Where μ_k are the parameters at iteration k , \mathbf{d}_k is the search direction at iteration k and a_k is a gain factor controlling the search direction at iteration k .

Optimization algorithms are classified according to their dependency on the derivative information of the similarity measure and how a_k and \mathbf{d}_k are computed. It is beneficial if the derivative information of the similarity measure can be easily computed. The biggest challenge considering the optimization algorithm is, not to stop at a local optimum caused by interpolation artifacts or by good local matches between the pixel intensities. The choice of the optimization algorithm

for a specific registration task depends on the remaining parts of the registration framework. Research on the behavior of optimization algorithms for specific registration tasks has been done by Maes et al. [14] Klein et al. [9]. For a detailed overview on different optimization algorithms please refer to the literature Press et al. [20] give a good overview over existing techniques.

2.2.5 Extensions

So far the components of a classical pixel based registration framework and common examples were described. Many improvements for the registration framework have been proposed.

One extension is the multi-resolution approach also called coarse-to-fine strategy. The multi-resolution approach uses an image pyramid that is constructed of the fixed and the moving image prior to the registration. The registration process is then started at the lowest resolution using an automated or manual initial estimation. The transformation parameters gained as a result from the pyramid at level i are then used as the initial parameters for the resolution at level $i - 1$. This is continued until the original resolution (i.e., level 0) is registered. The multi-resolution approach has two big advantages. Firstly, computation time is reduced because large scale translation, rotation and scaling are handled at lower resolutions. Computing the similarity measure at lower resolutions is also much cheaper due to the lower number of pixels. The second advantage is the smoothness of the similarity measure function at lower resolution levels, with smooth functions it is easier to find a global optimum for the optimization algorithm to be close to the global optimum (i.e., a good initialization for the next resolution level). In summary the multi-resolution approach leads to lower computational costs and more stability, since the optimization algorithm is more likely to find the global optimum and less likely to get stuck in a local optimum.

Another improvement is to use samples of the image content. This improves the computational cost of the similarity measure drastically. Klein et al. [9] state that it is important to use a new set of samples in each iteration (stochastic sampling), otherwise (deterministic sampling) a bias is introduced which results in approximation errors.

A third improvement is to construct the costly non-linear registration out of piecewise rigid registered sub-images of the original image. Therefore it is necessary to split up the original image. This approach is called hierarchical subdivision. After the subdivision, the image pieces are registered individually using a rigid or affine transformation. A dense deformation field is then reconstructed from these rigidly registered pieces. Examples for this improvement, can be found in Maintz et al. [16], Likar et Pernus [12] and Andronache et al. [1]

Chapter 3

Enhancement and Registration of HREM Data

3.1 Motivation

The HREM artifacts described in Section 1.3 vary in their severeness. Since segmentation of structures is often used on the HREM data, the most hindering artifact for automated segmentation algorithms (e.g., threshold segmentation) is the uneven illumination. Therefore we chose to correct the the uneven illumination. Our approach to correct the uneven illumination is described in Section 3.2. To incorporate the information gained by the histological staining, we have to register the manual generated histological dataset to the prior acquired HREM dataset. This is described in Section 3.3 and Section 3.4.

3.2 Illumination Correction

Since the medical researchers often need to segment structures within the HREM images based on gray values, the uneven illumination is a problem. One way to solve this problem is to take a blankfield-image (a slice without an embedded specimen - usually at the end of a block). This blankfield is used to correct the uneven illumination. Although this approach yields very good results, it is not always possible to take a blankfield for each dataset.

We also experimented with some standard image processing techniques like the TopHatFilter. Unfortunately the results were not pleasing.

Therefore we developed a more sophisticated method that is conceptually similar to the blankfield approach but also works in cases where blankfields are not available.

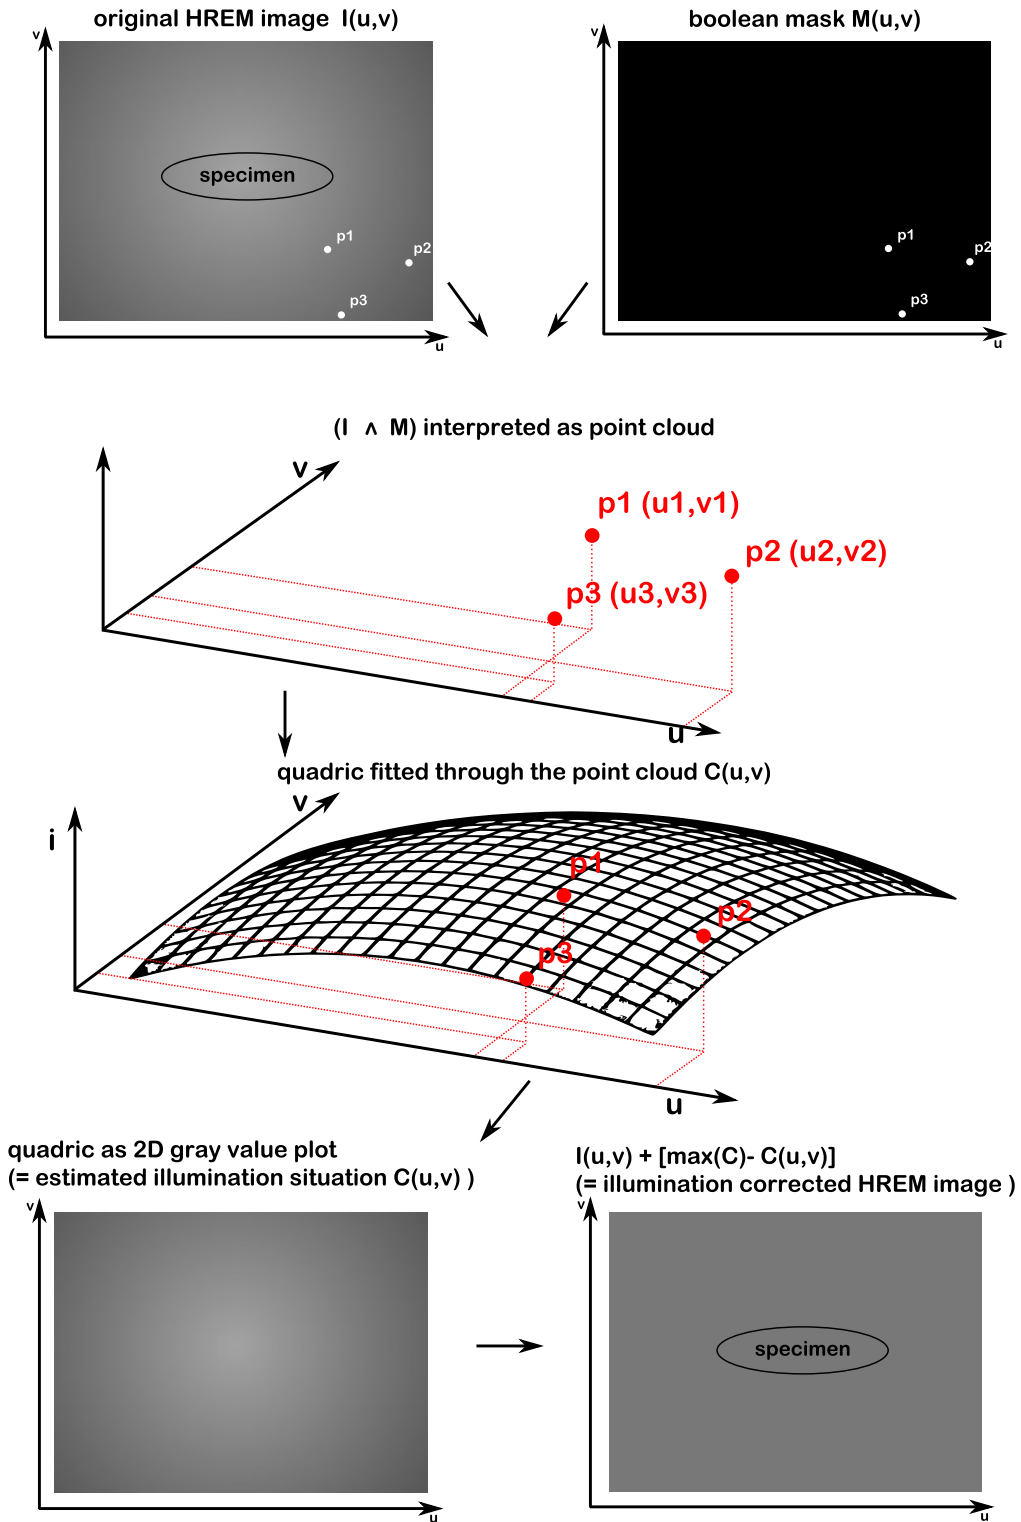


Figure 3.1: Overview of the illumination correction process. First background pixels are defined using a mask. These points are interpreted as a three dimensional point cloud. A quadric surface is fitted through this point cloud. Interpreting this quadric as a gray value image is an estimation for the illumination situation of the original HREM image. The illumination estimation is used to correct the HREM dataset.

An overview of our novel approach can be seen in Figure 3.1. First a mask is used to select pixels belonging to the background. Figure 3.1 illustrates the approach for three selected pixels. However, in reality a large subsection of the image is selected. These pixels are then interpreted as a point cloud in three dimensional space (with the intensity as z-value). The next step is to fit a quadric surface through the point cloud. Interpreting this quadric as a gray value image is an estimation for the illumination situation of the original HREM image. Using this newly computed illumination situation it is possible to correct the original HREM image resulting in an illumination corrected version of the original image.

In Figure 3.2 (a) an HREM slice is shown. The corresponding three dimensional plot is shown in Figure 3.2 (b). Since the background pixels have a light gray value they get a higher z-value. The background pixels are approximately lying on a quadric surface. The result of the fitted quadric surface creates an estimation of a blankfield-image. In Figure 3.3 the original HREM section (a) and the corresponding blankfield-image (b) is shown. The illumination corrected HREM slice is shown in 3.3 (c) and the computed artificial blankfield-image in (d). Also visible in Figure 3.3 (e) is the difference between the computed and the original blankfields. The difference is mostly in the range of the slicing artifacts and ranges between zero to fifteen.

As fitting algorithm we used the least squares technique (without weights or iterations). The least squares technique tries to minimize the squared distance of the data points and the fitted surface. This results in a highly overdetermined system which leads to an approximation when solved. The mathematical derivation for the least squares fitting follows:

We construct the overdetermined equation system

$$\sum_{j=0}^N a_j f_j(x_i, y_i) = z_i \quad (3.1)$$

where (x_i, y_i, z_i) represents the data points (i.e., the pixels selected by the mask) for $i = 1 \dots M$, a_j are the unknown coefficients and f_j are polynomials of the model function with the order $\leq N - 1$ for $j = 0 \dots N$ which should be fitted through the data points. This leads to a set of linear equations :

$$Aq = b \quad (3.2)$$

where

$$q = \begin{pmatrix} a_0 \\ \vdots \\ a_N \end{pmatrix} \quad (3.3)$$

,

$$b = \begin{pmatrix} z_1 \\ \vdots \\ z_M \end{pmatrix} \quad (3.4)$$

and

$$A = \begin{pmatrix} f_0(x_1, y_1) & \dots & f_N(x_1, y_1) \\ \vdots & & \vdots \\ f_0(x_M, y_M) & \dots & f_N(x_M, y_M) \end{pmatrix} \quad (3.5)$$

This set of equations is in general not solvable, therefore we look for a minimum of the squared distances between the data points and the model function :

$$f(q) := \arg \max_q \|b - Aq\|^2 \quad (3.6)$$

it holds that:

$$\nabla f = 2A^T Aq - 2A^T b \quad (3.7)$$

Therefore the approximation of q is:

$$q = (A^T A)^{-1} A^T b \quad (3.8)$$

It is possible to improve this approximation in an iterative process using weights for each data point and refine the weights in an iterative fashion. Since we chose to handle the data points manually, i.e., the researcher is responsible to spare out data points which don't belong to the background we do not need the iterative process to refine the approximation.

We tested the robustness of the least squares fitting algorithm concerning the number and location of pixels that are used for the fitting process. This was done using a blankfield-image. We used growing masks that increasingly spare out pixels, and measured the root mean square error. We used three different types of growing masks: a random mask sparing out an increasing number of random pixels and two square masks one starting to grow from the center pixels and the other starting to from the border pixels. The growing square masks are shown in Figure 3.4.

The random mask showed that the algorithm is quite invariant on the number of pixels, therefore the error started to rise when only ten percent of random pixels were used. Examining the behavior of the two square masks we concluded that the location of the pixels is important for the algorithm. The most important

pixels are at the border. A plot of the mean square error using the square mask growing from the inside is shown in Figure 3.5 (a) and the plot of the error from the mask growing from the border is shown in Figure 3.5 (b). When using this technique on images (with embedded specimen), a mask is applied to avoid the object's pixels. This mask is drawn by the user. A good mask spares out all the pixels from the specimens while preserving the background pixels, especially the border pixels that show background are important for a good fitting process. We found that it is usually sufficient to use a rectangle that masks out the specimen. With this approach the uneven illumination is corrected as long as all original values lie within the value range of the original image (in the case of HREM slices the range is 0 - 255). Therefore the illumination estimation gets worse if the original lighting situation lies outside the value range of the HREM image (i.e., the approximated quadric surface results in values lying outside the original HREM value range).


3.3 Registration Setup

To be able to compare the HREM staining techniques with well established histological coloring techniques the datasets need to be registered. We got two representative datasets from the medical researchers. Figure 3.6 shows an HREM morphology and a corresponding histologically colored image of the datasets. The resolution of a single slice is 2560 x 1920 pixels, and both datasets consist of 30 slices. The morphological HREM image was used as fixed image since it shows more features of the captured specimen than the signal HREM image. We chose a two dimensional registration approach, since the histological section images are not aligned to each other and sections might be missing (because they were destroyed in the acquisition process). An overview of this three step registration approach can be seen in Figure 3.7. The first step is to manually find an initial transformation (e.g., a coarse registration only using scale and translation). This initial transformation is used as the starting point for the affine registration, which handles the linear transformations of the different fields of view. The result of the affine registration is used as the starting point for the b-spline registration, which handles the non-linear deformations caused by the curling and further processing of the cut section.

To explore the possibility of registering the histological and HREM images we used elastix [8] which is a command-line tool based on the ITK (Insight Segmentation and Registration Toolkit) by Kitware [7]. Elastix implements the parts of the classical registration framework (described in Section 2.2) in the most commonly used ways, and therefore gives the possibility to tackle a registration problem trying out many different setups in a convenient way. Figure 3.8 shows the com-



(a) HREM slice



(b) HREM slice plotted in three dimensions plot

Figure 3.2: The plot in image (b) shows the data from image (a) plotted as a three dimensional plot in red. The fitted surface is shown in green

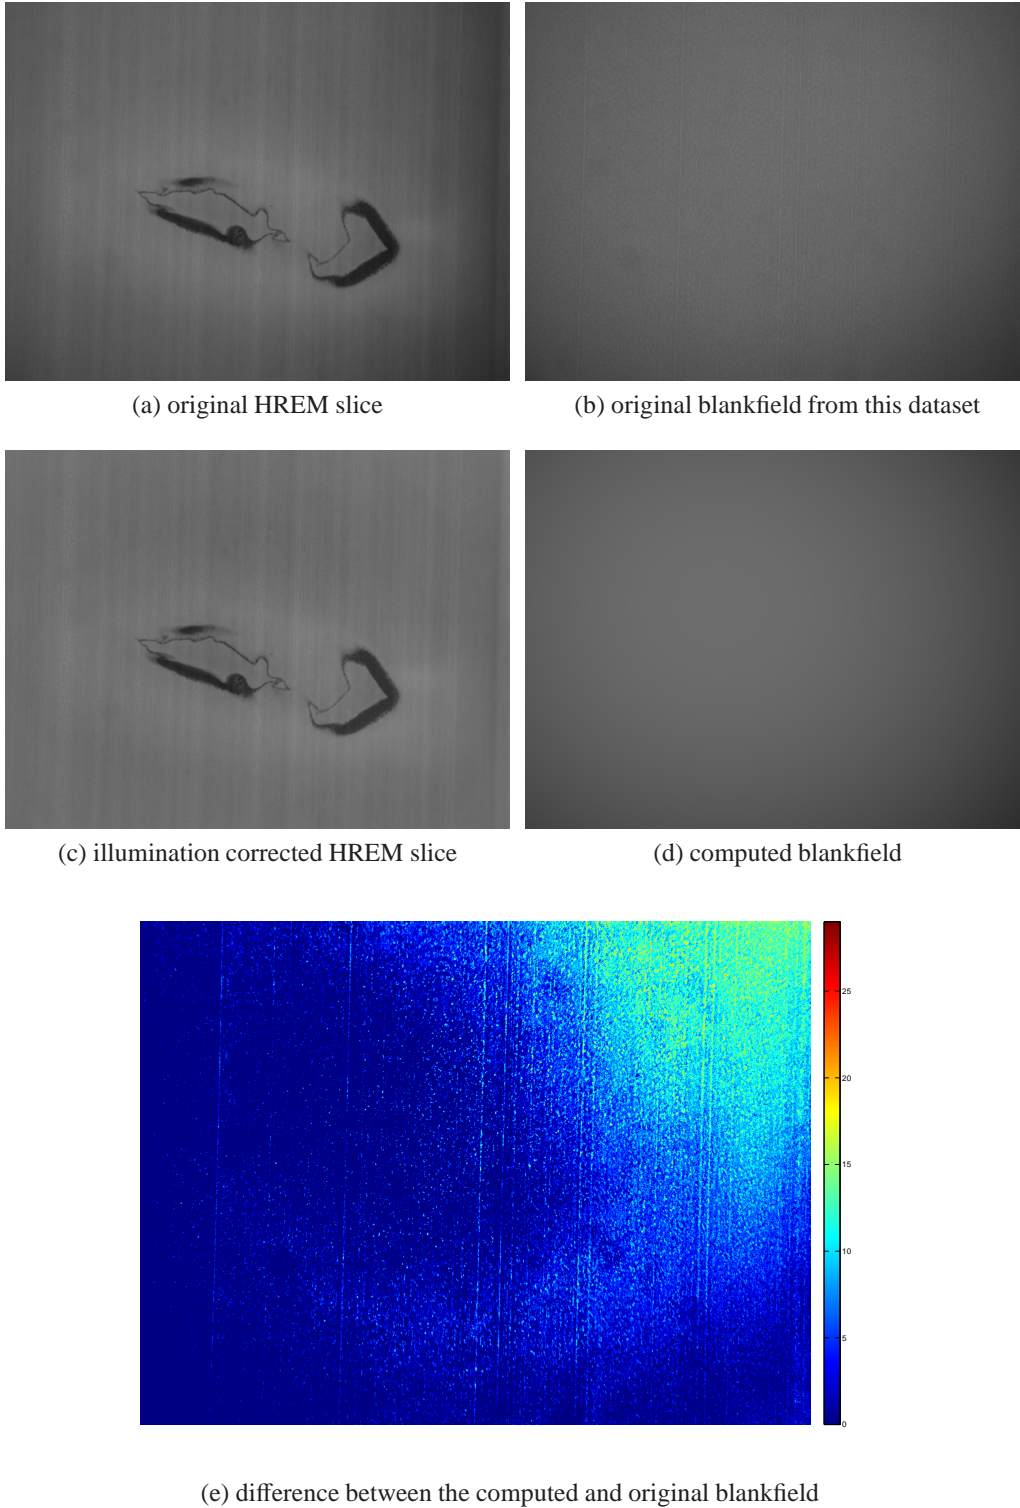


Figure 3.3: Image (a) shows an HREM image and image (b) shows a blankfield that was acquired from the same block. Image (c) shows the illumination corrected HREM slice using the computed blankfield image shown in (d). The difference between the computed and the original blankfield is shown in image (e)

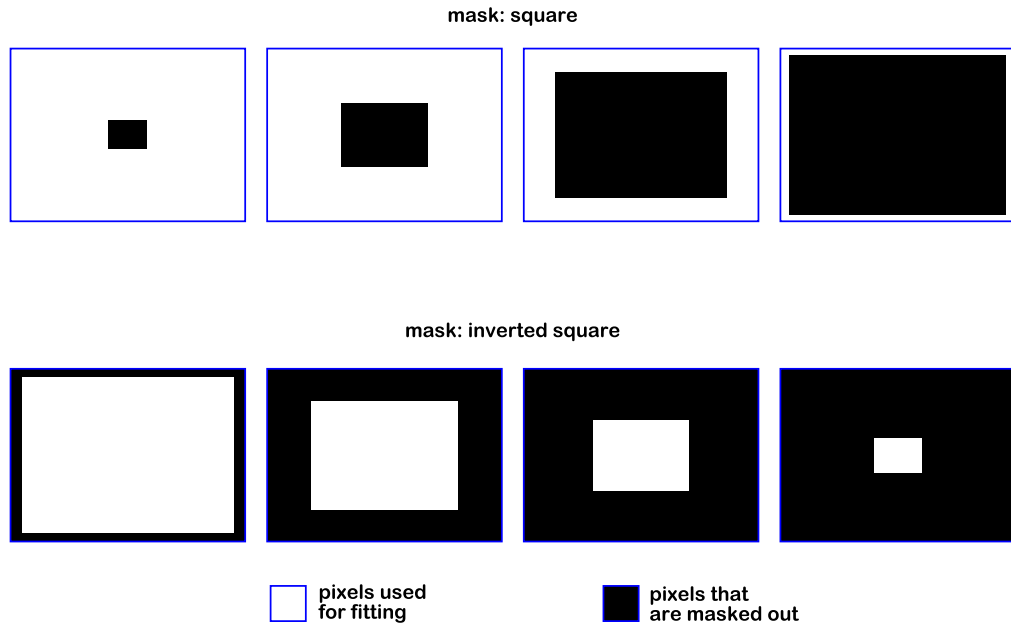


Figure 3.4: The square masks used for the examination of the robustness of the fitting algorithm.




Figure 3.5: Mean squared error plots, where pixels form different regions are favored for the computation of the illumination situation.




Figure 3.6: The two datasets. (a) and (b) show corresponding images of the dataset "specimen one", while (c) and (d) show corresponding images of the dataset "specimen two".

ponents used by elastix. Elastix incorporates the classical registration framework with the multi-resolution and sampling extensions described in Section 2.2.5.

These histologically colored images do not have the same field of view as the corresponding HREM images. Usually the histological images show the region of interest slightly magnified. Therefore we use an affine transformation (described in Section 2.2.1). For the similarity measure we tried out the normalized cross correlation and the advanced mutual information measures implemented in the elastix tool (they are variations of the similarity measures described in Section 2.2.3). As interpolator we used a b-spline interpolator (which is an instance of a spline interpolator described in Section 2.2.2). For the optimization we tried out the adaptive stochastic gradient descent optimizer and the quasi-newton optimizer. Both optimization algorithms compute the search direction \mathbf{d}_k using the derivative

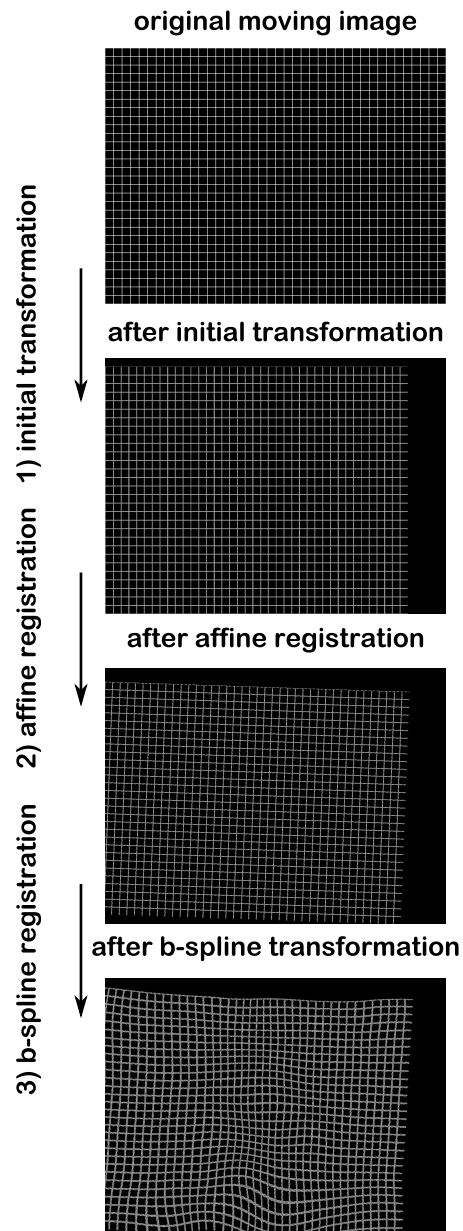


Figure 3.7: Overview of the registration process.

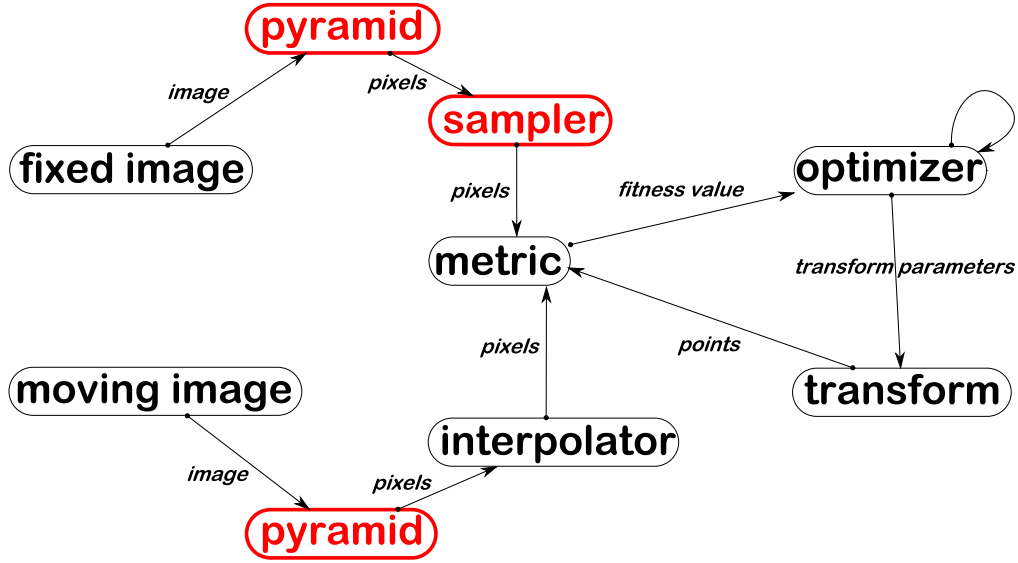


Figure 3.8: The components of the elastix registration tool. The components colored in red are extensions of the standard registration framework described in Section 2.2.5 and implemented in elastix.

information of the similarity measure $\delta S / \delta \mu$. It is assumed that the derivative information $\delta S / \delta \mu$ can be computed exactly. For the adaptive stochastic gradient descent optimizer the iterative search is described as:

$$\mu_{k+1} = \mu_k - a_k \mathbf{g}_k(\mu_k), k = 0, 1, 2, 3, \dots \quad (3.9)$$

where the gain factor a_k is defined as a decaying function of k :

$$a_k = a / (k + A)^\alpha \quad (3.10)$$

α , a and A are user specified parameters which have these properties: $a > 0$, $A \geq 1$, $0 \leq \alpha \leq 1$.

The iterative search of the quasi-newton algorithm is described as:

$$\mu_{k+1} = \mu_k - [H(\mu_k)]^{-1} \mathbf{g}(\mu_k), k = 0, 1, 2, 3, \dots \quad (3.11)$$

where $H(\mu_k)$ is the Hessian matrix of the similarity measure evaluated at μ_k . The evaluation of the inverse Hessian matrix is computationally expensive, especially for problems with high dimensionality (i.e., using non-linear transformations). Therefore the inverse Hessian matrix is approximated by $L_k \approx [H(\mu_k)]^{-1}$

where L_k is:

$$L_{k+1} = \left(I - \frac{\mathbf{s}\mathbf{y}^T}{\mathbf{s}^T\mathbf{y}} \right) L_k \left(I - \frac{\mathbf{s}^T\mathbf{y}}{\mathbf{s}^T\mathbf{y}} \right) + \frac{\mathbf{s}\mathbf{s}^T}{\mathbf{s}^T\mathbf{y}} \quad (3.12)$$

where I is the identity matrix, $\mathbf{s} = \boldsymbol{\mu}_{k+1} - \boldsymbol{\mu}_k$ and $\mathbf{y} = \mathbf{g}_{k+1} - \mathbf{g}_k$. L_0 is practically initialized with $L_0 = I$

For the extensions we used a gaussian pyramid in different resolution depths and a random sampler.

Since the histological images are non-linearly deformed due to the slicing and the preparation for the image acquisition, we used a b-spline registration using the result of the affine registration as input. The differences between the b-spline registration setup and the affine registration setup are the transformations (where we used the b-spline registration) and the sampler where we used the full image sampler (which means that all of the image content is used).

3.4 Registration Results and Problems

In this Section the problems, solutions and results of the two registration steps are described, starting with the affine registration and followed by the b-spline registration.

3.4.1 Affine Registration

Initial Transformation Parameters The starting point for the affine registration are the initial transformation parameters for the optimizer. In many cases it is sufficient to use the identity transform as a starting point. However, in other cases (e.g., in dataset "specimen two", where the aspect ratio differs after the histological images are rotated and mirrored to roughly match the orientation of the specimen in the HREM images), the identity transform is not sufficient if a big part of the moving image lies initially outside the fixed image region, a lot of information cannot be used by the optimizer. This leads to translations, that move the image out of the fixed image region after a few iterations.

For the registration it is beneficial to manually perform an initial transformation. A rough alignment of the images by scaling and translating the moving image is sufficient. Using a coarse manual registration as initial transform parameters leads to far better results than using the identity transform. Although this clearly leads to a semi-automatic approach, the manual intervention takes not more than a few minutes using an appropriate tool where the scale and translation can be easily adjusted. The researcher capturing the histological image usually tries to capture the section in a way that the image is aligned to the border of the section.

Therefore the rotational difference between the histological and the HREM image is usually less compared to the scale and the translation difference. For the coarse manual registration the rotational difference can usually be neglected. Another benefit of a coarse manual registration as starting point is that it saves a lot of computation time since the optimizer has to handle smaller transformations.

Similarity Measures To estimate the similarity we experimented with the normalized cross correlation and the advanced mutual information measures. In our experiments both measures led to similar results. Since mutual information is known to be more robust and the cross correlation metric needs either images from the same modality or that the intensities from both modalities only vary linearly, we concentrated on the mutual information measure.

in our experiments the mutual information metric was not as robust as the cross correlation measure i , and needed a carefully adjusted set of parameters to work appropriately. The most important parameter using the mutual information measure is the number of bins for the joint histogram. Using too many bins leads to a less stable detection of similarity between the images, using too few bins leads to detection of false similarities. The number of bins needs to be adjusted carefully for each dataset. In our experiments 64 bins turned out to work best for dataset of specimen one and 128 bins for dataset of specimen two. Nevertheless the mutual information measure is more powerful for this task since other histological colorings might be used in the future. We expect mutual information to work more robust with new colorings since it is based on information theory and only relies on the occurrence of similar structures in the image pairs.

Since the histological images are RGB color images, and the registration only works on gray scale images we selected a single channel of the color image. In our case the red channel led to the worst results since the color image has a very high red and white portion. Therefore the red channel has low contrast between visible structures and background. Even if the mutual information measure worked on the red channel, the results when using one of the other channels were far better (i.e., either the quality of the registration was better or the computation time was lower). The histograms of the the channels of a histologically stained section can be seen in Figure 3.9, where (a) shows the histogram of the red channel, (b) shows the histogram of the green channel and (c) shows the histogram of the blue channel of the original histologically stained section shown in Figure 1.9 (a). The histograms of the red and green channels show that the contrast of these channels is far better than the contrast of the red channel. Although it appears that the contrast of the green channel is even better than the contrast of the blue channel we could not find an improvement when using the green channel instead of the blue channel. This can be seen in Figure 3.10 where image (a) shows a checker board view of the

affine registration result of the red channel, while image (b) shows a checkerboard view of the affine registration result using the same parameters but uses the blue channel.

Optimizer The parameters of the optimizer have to be chosen very carefully, since a bad parameter setting leads to useless results. We experimented with two different optimizers, namely a gradient descent and a quasi newton optimizer. Both led to similar results as long as the parameters are chosen appropriately. The most important parameters are the initial transformation, the step size, the number of iterations, and the number of resolution-levels. Using more resolution-levels leads to a reduced computation time while improving the chances to find a global optimum. A multi-resolution approach works especially well for registration of high resolution images. The number of iterations within a resolution-level restricts the maximum computation time for the optimizer to find a solution. In our tests the optimizer almost never converged because of the non-linear deformations in the images. Therefore in our application a high number of iterations leads to high computation time without additional benefits. The number of iterations can be decreased once it is verified that the last iteration does not improve the result significantly. The step size defines the magnitude of the transformation at each iteration. The higher the step size, the harder it is to find a global optimum. With too small step sizes the optimizer needs much more iterations to find the optimum, which increases the computation time. In our experiments we found that five resolution levels worked well, and the especially important resolution levels were the first three (i.e., the smallest three). Furthermore we chose a step size of 16 which worked well in our experiments.

Results of the Affine Registration The computation time for the affine registration is about three minutes per slice. The quality of this registration strongly depends on the nonlinear deformation of the slice. With fewer non-linear deformations (which is the case when the slice thickness is smaller, as in the dataset of "specimen two", since it is expected that the tension in the cut section is less if the slice thickness is thin and therefore generating less non-linear deformations) the affine transformation leads to useful results. If the quality of the affine registration is not sufficient (i.e., severe non-linear deformations) a b-spline registration has to be done on the resulting images of the affine registration. Figure 3.12 (a) shows an HREM slice, (b) shows a sample histological slice before, and (c) shows the slice after the affine registration of the dataset of "specimen one". Figure 3.14 shows the same information of the dataset of "specimen two".

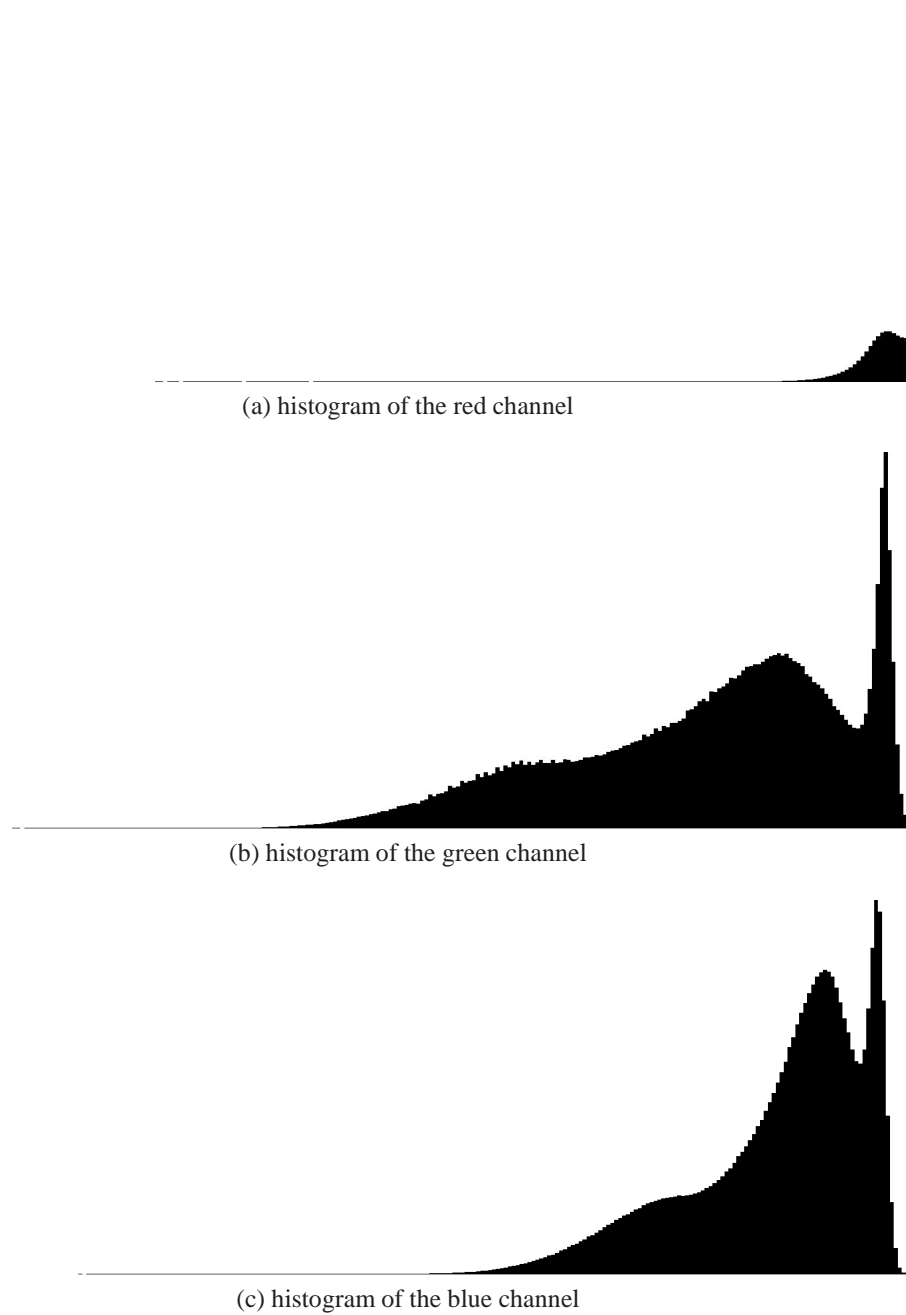




Figure 3.9: The histograms of the different color channels of Figure 1.9 (a). The red channel has the poorest contrast.



(a) registration result using the red channel of the original histologically colored image



(b) registration result using the blue channel of the original histologically colored image

Figure 3.10: An Example of the different affine registration results using the red channel of the original histologically colored image in image(a) and the blue channel in image (b) was used. All parameters were left unchanged.

3.4.2 B-spline Registration

Most of the parameters used for the affine registration can also be used for the b-spline registration. For the b-spline registration a random sampler provides not enough information for the optimizer therefore the whole image content is used for the registration. The affine registration was successful if the histologically colored images and HREM images differ only in non-linear deformations. A visual inspection is performed prior to the b-spline registration to assure sufficient alignment. The similarity measure is left unchanged, since a working similarity measure also works for the b-spline registration step. The only part of the framework where additional considerations have to be made is the b-spline transformation. The grid spacing of the transformation defines the number of parameters the optimizer has to evaluate. With too few parameters small deformations cannot be handled. Too many parameters resulting over fitting and increase the computation time tremendously. Figure 3.11, taken from the elastix manual [8], shows the impact of the grid spacing on the result of the registration. While image (a) and (b) show the fixed and moving images, image (c) shows the result of too wide grid spacing. Image (d) shows the result of a well chosen grid spacing. Image (e) shows the result of too fine grid spacing. The step size and the number of iterations also has to be adjusted for the b-spline registration step since we are dealing with new transformation parameters.

Results of the B-spline Registration We used b-spline registration to align all slices of the two datasets. Figure 3.12 (d) shows the result of the b-spline registration and in Figure 3.13 a combined view of the HREM slice and the b-spline registration result of the dataset of "specimen one" is shown. Figure 3.14 shows the same information of the dataset of "specimen two". Figure 3.15 (a) shows the whole volume of the dataset of "specimen one" before the registration and Figure 3.15 (b) after the registration. All results were generated on an Intel Core2-Duo E6750 CPU with 2 GB RAM, using a single core. Because of the large number of parameters used for the b-spline registration the computation time is approximately two hours per image.




Figure 3.11: The impact of the grid spacing of the b-spline transformation on the registration result. While image (a) and (b) show the fixed and moving images, the image (c) shows the result if the grid spacing is too wide, image (d) shows the result if the spacing is chosen well, and image (e) shows the result if the grid spacing is too fine. Images taken from the elastix manual[8]




Figure 3.12: Registration of a sample slice of the dataset of "specimen one": image (a) shows the HREM slice, image (b) shows the histological slice before the registration, and image (c) shows the histological slice after the affine registration, and image (d) shows the histological slice after the b-spline registration.




Figure 3.13: HREM slice combined with the final registration result from Figure 3.12.






Figure 3.14: Registration of a sample slice of the dataset of "specimen two": image (a) shows the HREM slice, image (b) shows the histological slice before the registration, and image (c) shows the histological slice after the affine registration, and image (d) shows the histological slice after the b-spline registration.



(a) before Registration



(b) after Registration

Figure 3.15: Volume rendering of dataset of "specimen one". Image (a) shows the volume before the registration and image (b) shows the volume after the b-spline registration. The moving dataset is shown in red.

Chapter 4

Visualization of the Registered Volumes

The main purpose of the visualization is to facilitate the visual inspection of the registration result and to provide a three dimensional overview of the registered volumes. The registration result found during optimization needs to be visually inspected to guarantee that the procedure found the global optimum. We describe a two dimensional checkerboard visualization for the comparison of registered slices in Section 4.1. Further we use a three dimensional visualization based on direct volume rendering. A focus region is used to show the borders of different structures. These borders can be used to visually compare the fixed and the moving dataset. The three dimensional approach is presented in Section 4.2.

4.1 2D Visualization - Checkerboard View

The checkerboard view is a simple two dimensional visualization technique where two corresponding images are displayed on the same canvas. Like on a checkerboard where white and black areas are alternating the checkerboard visualization alternately shows the two images. The resulting image provides a simple method for comparisons, especially at the borders. An example of a checkerboard view is shown in Figure 4.1 where two slices are compared before registration. In Figure 4.2 the same slices are shown after the registration.

The benefit of this visualization is that the user is able to inspect the correspondence of the two images simultaneously. Especially the connection between the two images (i.e., the border of the visualization between two adjacent checker areas) can be examined. The border regions present the information about the quality of the registration. A close-up of the checkerboard view is shown in Figure 4.3. Since the checkerboard structure can be moved over the images, it is



Figure 4.1: An example of a checkerboard view showing two slices before the registration.



Figure 4.2: An example of a checkerboard view showing two slices after the registration.




Figure 4.3: A close-up of the checkerboard view.

possible to follow structures of interest throughout the two images. Furthermore we implemented an interaction technique to adjust the size of the checkerboard areas. This technique provides a powerful tool for simple inspection of the registration quality of the structures of interest.

4.2 3D Visualization - Direct Volume Visualization

The two dimensional checkerboard technique is useful for inspecting a single slice, but it lacks the information about the registration in the third dimension and does not provide a conceptional overview. Therefore we implemented a three dimensional visualization technique. We make use of direct volume rendering (DVR), incorporating a focus region that is adjustable in size and position. This makes it possible to follow different structures of interest throughout the volume to make sure that the registration quality of those structures is sufficient. Section 4.2.1 explains the mathematical background of the direct volume rendering, while subsection 4.2.2 presents the adaption of DVR for our actual implementation.

4.2.1 Direct Volume Rendering

There are two major approaches to render three dimensional volumes. The first approach, called indirect volume rendering, extracts an iso-surface in a preprocessing step. A prominent example of indirect volume rendering techniques is the Marching Cubes algorithm [13]. The polygonalized result is then rendered using standard graphics hardware. The second approach renders the volume directly (therefore called direct volume visualization). The advantages of DVR are, that there is no preprocessing step for extraction of an iso-surface, and that different iso-surfaces can be shown by altering visualization parameters.

The direct volume rendering approach (proposed by Levoy [11]) uses a model of light going through semi-transparent material with different absorption coefficients. It works by shooting rays from the eye-point through the volume. Each sample along the ray is assigned a color and opacity. The colors of the samples are accumulated weighted by the opacity. The assignment of colors and opacity is usually done by a transfer function. A transfer function is a mapping from data values to color and opacity values. This transfer function is usually defined by the user in an interactive way.

The rays are either cast from back to front, or from front to back. In the case of front to back accumulation the opacity of the ray needs to be computed separately. If the accumulated opacity is fully opaque the ray is terminated. The resulting color of the pixel on the viewing plane is the accumulated color. Figure 4.4 shows a schematical view of the image plane, the volume, the samples and the ray through the volume.

The mathematical expression of the back to front accumulation is:

$$C_{\lambda}(u_i, v_j) = \sum_{k=0}^K \left[c_{\lambda}(x_i, y_j, z_k) \alpha(x_i, y_j, z_k) \prod_{m=k+1}^K (1 - \alpha(x_i, y_j, z_m)) \right] \quad (4.1)$$

where $C_{\lambda}(u_i, v_j)$ is the resulting color of the pixel (u_i, v_j) on the viewing plane, $c_{\lambda}(x_i, y_j, z_k)$ is the color and $\alpha(x_i, y_j, z_k)$ is the opacity of the voxel at location (x_i, y_j, z_k) . $c_{\lambda}(x_i, y_j, z_0) = c_{bkg, \lambda}$ and $\alpha(x_i, y_j, z_0) = 1$ are the color and opacity of the background.

The values, in a volume are given at discrete positions (i.e., the voxels). Since the evaluation involves sampling position between the voxels, the values need to be interpolated. To apply shading using a light source we need to estimate a normal vector for each sample. Since this information is usually not available in the dataset, the gradient of the voxel is commonly used to approximate the normal vector. This gradient usually has to be estimated using numerical derivation




Figure 4.4: Illustration of ray casting for DVR. A ray is shot from the eye-point through a pixel on the viewing plane and through the volume. The ray intersects the volume and is sampled inside. The resulting color is calculated by accumulating the opacity weighted colors of the samples.

approximation algorithms like the central difference:

$$\dot{f}(x, y, z) = \begin{pmatrix} \frac{f(x+h, y, z) - f(x-h, y, z)}{2h} \\ \frac{f(x, y+h, z) - f(x, y-h, z)}{2h} \\ \frac{f(x, y, z+h) - f(x, y, z-h)}{2h} \end{pmatrix} \quad (4.2)$$

Where the approximated derivative $\dot{f}(x, y, z)$ is defined by the central difference in each coordinate direction. h is the distance between the point under consideration and the samples taken on either side of this point (in our case h is the distance between two consecutive samples on the ray).


To enhance the visibility of surfaces (i.e., tissue boundaries) in DVR, gradient magnitude opacity modulation is used. The opacity of the sample is modulated by the gradient magnitude of the sample. Since the gradient is high at surface boundaries and low in homogeneous regions the modulation improves the visibility of surfaces within the volume and suppresses homogeneous regions. Figure 4.5 shows an HREM slice and its corresponding gradient magnitude. The influence of the gradient magnitude on the opacity, can be adjusted by the user with a transfer function.

Furthermore, if the result should exhibit shadows, a second ray from each sample to the light source needs to be evaluated.


4.2.2 Implementation of the 3D Visualization

We implemented direct volume visualization using OpenGL and gl shading language to achieve interactive frame rates. Both volumes (the illumination corrected HREM volume and the histologically colored and registered volume) were loaded as 3D textures onto the graphics hardware. 3D textures facilitate tri-linear interpolation of the samples implemented in hardware resulting in very high performance compared to software implementations. The main drawback of the graphics hardware is the limited amount of random access memory. Therefore the volumes are down-sampled to fit into the memory of the graphics adapter. The ray-casting is implemented in a shader program in a front to back manner incorporating early ray termination. To improve the visualization of tissue boundaries, which are important for the inspection of the registration results, we used gradient magnitude opacity modulation. The gradient was approximated by the central difference. A transfer function was used to modulate the opacity depending on the gradient magnitude of the sample. This leads to the following equation for the opacity of the sample:

$$\alpha(s) = \alpha_{tf}(s) \alpha_{tf,g}(\|g(s)\|) \quad (4.3)$$



(a) HREM slice



(b) Gradient magnitude image of the HREM slice

Figure 4.5: In image (a) a HREM slice is shown. In image (b) the corresponding gradient magnitude image is shown, where white denotes the highest and black the lowest gradient magnitude.

Where $\alpha(s)$ is the final opacity at the sample s , α_{tf} is the common opacity transfer function and $\alpha_{tf,g}$ is the gradient magnitude transfer function computed with the gradient magnitude $\|g(s)\|$.

The result of the gradient magnitude opacity modulation is shown in Figure 4.7 where in image (a) a DVR without gradient magnitude opacity modulation is shown, and in images (b) to (d) samples with low gradient magnitude are suppressed according to the transfer function.

Shading was not used, to allow visual inspection of the registration result without varying light intensities. To present both the fixed dataset and the registration result simultaneously we used a focus and context approach. The focus region shows the result of the registration (i.e., the moving images) and the HREM volume is used to provide context. The size and position of the focus region is adjustable. Since this focus region is movable the user is able to follow the boundaries of both volumes to inspect how good these boundaries match. We have chosen to use the actual sample color of the registration result to be displayed. However, as input of the opacity transfer function only the green channel is used. For the gradient estimation of the registration result we used the maximum gradient of all three channels since we left out the impact of color metrics (especially color differences needed for the gradient approximation), depending on the different color spaces, on the gradient magnitude result for further studies. To present the registration result within the focus region we switch to the registration result as soon as the ray enters the focus region and we switch back to the HREM volume as we enter the context region. The focus region of the registration result appears to be embedded within the HREM volume. A schematical view can be seen in Figure 4.6.

Examples of the focus and context DVR are shown in Figure 4.8. Examples of the focus and context DVR showing badly aligned regions are shown in Figure 4.9. We achieved interactive rendering frame rates on an Intel Core2-Duo E6750 CPU with 2GB RAM and an Nvidia 8800 GTS 512 graphics adapter. As stated the two volumes were down-sampled to fit into graphics memory.




Figure 4.6: Illustration of our focus and context approach. The black cuboid represents the context volume, and the red cuboid represents the focus region. The viewing ray is divided into three sections. We switch the volumes as the ray enters the focus region (indicated by the red part of the viewing ray), and use the context volume otherwise. The focus volume therefore appears to be embedded within the context volume.

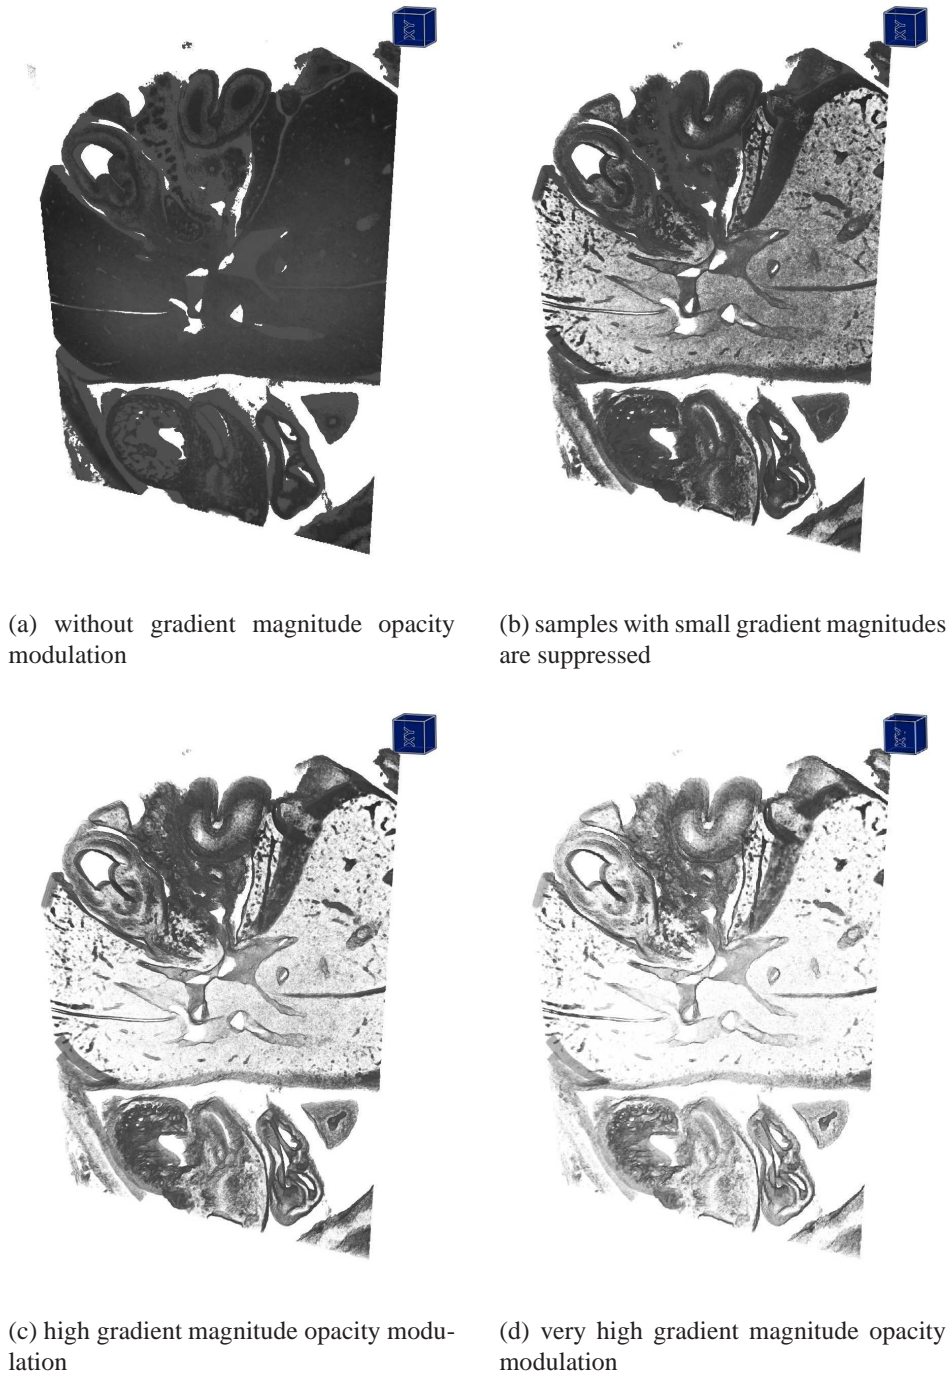




Figure 4.7: The impact of gradient magnitude opacity modulation. Image (a) shows DVR without gradient magnitude opacity modulation. Images (b) to (d) show DVR with regions of low gradient magnitude increasingly suppressed.



(a) 3D view of the registered volumes



(b) top view

Figure 4.8: Exemplary results of our DVR approach. The focus volume (in red) shows the registration result. The registration result can easily be inspected at surface boundaries regions.




Figure 4.9: Exemplary result of our DVR approach showing a badly aligned region.

Chapter 5

Summary

5.1 Introduction

High Resolution Episcopic Microscopy (HREM) is a technique medical domain scientists use to explore the impact of genes on different cell growth. For this task the scientists stain a specimen using fluorescent markers. Next the specimen is embedded in a resin to fix the specimen for slicing. The embedded specimen is then cut using an modified microtome. To make this technique fully automatic the microtome was modified by aligning the stopping point of the cutting with the optical axis of the episcopic microscope.

This results in images which are already aligned to each other. Fluorescent filter sets are used to separate the the light reflected from the specimen and the light, which is emitted form the fluorescent stained regions. Therefore the HREM acquisition process results in two datasets, one dataset shows the general morphological context, while the other dataset shows gene expression patterns as a result of the staining. Both datasets have a resolution of 2560 x 1920, while the number of slice images depends on the cutting thickness, which can be chosen between one and ten μm . Due to the mechanical slicing of the specimen there are many artifacts. The most hindering artifact for further processing is the uneven illumination. Therefore we developed a simple technique to correct the illumination.

The scientists experimentally extended the HREM modality by staining the already cut sections of the specimen another time using standard histological staining techniques. This second staining is fully manual. The operator of the microtome has to save the cut section, put it in a water quench to uncurl it, and capture it using a standard optical light microscope. This extension results in an image series of unaligned images. The new image series is also not aligned to the original HREM dataset. The only information about the images of the newly generated series is the correspondence to the HREM slice. Due to the manual

nature of this technique there are many artifacts in the newly generated images. The most prominent artifacts are a different field of view as the HREM images and non-linear deformations. To register the datasets we explored different registration techniques.

For the visual inspection of the registration results we implemented a two dimensional checkerboard view. This two dimensional visualization is widely used to inspect registration results. We also implemented a 3D visualization based on direct volume rendering. We adopted the standard DVR by incorporating a focus and context and a boundary enhancement technique.

5.2 Background and Related Work

Registration is the process of aligning two or more datasets (e.g., 2D images, 3D grids, point clouds, etc.), that represent the same scene or parts of the same scene. There are three major fields in which registration is needed: computer vision and pattern recognition, remote sensing and medical image analysis. There are many ways to classify registration algorithms, while the most meaningful classification is on the data which is used for the registration.

For instance markers can be used for the registration task. This is called **extrinsic registration**. Since only few data points are used for the registration, this approach is usually very fast. Although this approach is very simple and fast there are many applications where external markers can not be used. Furthermore deformations (e.g., of soft tissue) can not be captured.

Another method, called **non image based registration** does not take any information of the captured data into account. Instead it relies on information about the coordinate systems of the acquisition modalities. Therefore if the transformation between the coordinate systems of the acquisition modalities is known (e.g., by tracking the scanners), the datasets can be registered. Because the only task is the actual transformation of the datasets this technique is extremely fast. On the other hand it can not model deformations, and not all acquisition devices can be tracked easily.

The most general registration approach is the **intrinsic registration**, where the image content is used. The intrinsic registration is based on points selected by a user (e.g., so called landmarks), segmented or predefined structures or the whole image content. The approach where the whole image is used is called pixel based technique. This is the most flexible approach.

The input of a pixel based registration are two datasets (e.g., images). The fixed dataset represents the original to which the moving dataset has to be transformed. The moving image has to be evaluated for each transformation. Therefore locations between the data points are accessed. An interpolator is used to

resample the discrete data. A transformation that aligns the two datasets has to be found by the optimizer. The class of transformations that are applied depends on the registration task. In general the simplest class of transformations (i.e., with the fewest parameters) that transforms the moving dataset onto the fixed dataset should be chosen. The optimizer samples the parameter space of the given class of transformations.

A similarity metric measures the quality of the current fit. The fitness value which is maximized by the optimization algorithm to find the transformation that best aligns the fixed and moving image.

5.3 Preprocessing of HREM Data

The preprocessing of the HREM data described in this thesis is split in two parts. The first part deals with the uneven illumination correction, while the second part describes the registration of the images of the second staining process.

5.3.1 Illumination Correction

For the illumination correction we developed a novel technique specifically tailored to HREM data. First the user has to mask out regions of the image that do not contain background pixels. Pixels that are considered to be background are interpreted as a point cloud in three dimensional space using the intensity as z-value. A quadric surface is fitted through this point cloud. This quadric surface is interpreted as an image where the z value is used as intensity. This image is an estimation of the illumination situation without a specimen. Using the illumination estimation we are able to correct the original images.

We evaluated the quality of the estimated illumination situation depending on the number and location of selected pixels. For this task we used an original image without a specimen embedded and measured the root mean squared error between the original image and the estimated illumination situation. We used different masks to analyze the robustness. One mask was used to first delete the center pixels and growing towards the outside. The other mask deleted the border pixels first and grew towards the inside. The results led us to the conclusion that the border pixels of the original image are more important than the center pixels. We measured the quality of our estimated illumination correction approach using an original image without a specimen embedded. The error lies in the range of the remaining artifacts.

5.3.2 Registration

For the registration we proposed an approach in three steps. We used the HREM image showing the morphological context as a fixed image. The histologically colored images were used as the moving images. For the registration we used only a single channel of the RGB images. We identified the red channel as the channel with the lowest contrast, and therefore chose the blue channel with much higher contrast. The first step is a coarse manual registration using translation and scale. This first step can often be neglected when the scale and translation between the fixed and the moving images is small. It is however important for large image regions outside the fixed image at the start. The second step is an affine registration. The affine registration is done because the field of view is not the same for the fixed and moving image. It uses mutual information as a similarity measure. Mutual information is used because it does not rely on the absolute intensity values of the involved images. The most important parameter for the optimization algorithm is the step size. Too small step size results in a high computation time, a too large step size yields undesired results. Since the affine registration cannot handle the non-linear deformations an additional b-spline registration step is used. In the b-spline registration step all parameters from the affine registration except the transformation parameters and the step size can be kept. The transformation parameters define the grid of the b-spline transformation. A good choice for the size of the grid is important, since a too wide grid cannot handle small deformations, and a too fine grid is very sensible to noise. While the affine registration takes in the order of minutes per slice, the b-spline registration takes in the order of hours. The registration steps were performed on an Intel Core2-Duo E6750 CPU (using a single core) with 2GB RAM.

5.4 Visualization of the Registered Volumes

The main task of the visualization is to facilitate the visual inspection of the registration results. We implemented the widely used technique called checkerboard technique for 2D comparisons. Like on a checkerboard alternately areas from the two images are shown. This results in an visualization where the quality of the registration can be inspected at the border of adjacent areas. Since the checkerboard is movable the user is able to track different structures throughout the images.

Since this visualization lacks the information of the third dimension, we proposed another visualization based on direct volume rendering. DVR was proposed as an alternative to older approaches where surfaces were extracted before visualization using algorithms like marching cubes. The DVR approach uses a model of light going through transparent materials with different absorption coefficients.

For the visualization of different surfaces within the volume a transfer function is used. This transfer function is a user defined mapping between the data values of the volume and opacity and color values. To render an image rays are cast from the eye-point through the volume. Each ray is sampled inside the volume. The transfer function maps the sample values to color and opacity values. The resulting color and opacity are then accumulated either in a front-to-back or back-to-front manner.

We implemented our DVR approach using OpenGL and the gl shading language to achieve interactive frame rates. To improve the visibility of surfaces within the volume we used gradient magnitude opacity modulation. This technique uses the gradient, which is high at tissue borders and low in homogeneous regions, to modulate the opacity. It suppresses homogeneous areas in the volume and enhances tissue boundaries. We approximated the gradient by using central differences. For the registered volume (which has RGB channels), we used the maximum gradient of all three channels. Furthermore we incorporated a focus region in our DVR approach, which shows the registration result. Since the focus region can be changed in size and position, the user is able to follow tissue borders to inspect the registration result while gaining an overview on both datasets at the same time. To render the focus region we simply switch to the registered volume as soon as the ray enters the focus region, and we switch back to the HREM volume as soon as we leave the focus region. This results in an image where the focus region appears to be embedded in the context volume.

Chapter 6

Conclusion

The High Resolution Episcopic Microscopy (HREM) developed by medical domain scientists is able to acquire high resolution volumes of small specimen. This technique is a mechanical and destructive approach where the specimen is cut and photographed. This process yields slice images of the specimen with the thickness of about one to ten μm . Different genes of the specimen can be deactivated or activated. Staining of the specimen results in gene expression patterns. Exploring these gene expression patterns gives information about the different genes.

Different artifacts occur because of the mechanical nature of the acquisition technique. In this work we investigated approaches that improve the quality of HREM images. The most hindering artifact of the HREM images is the uneven illumination. This uneven illumination poses a big problem when trying to segment structures of interest (either with standard segmentation techniques or using a transfer function for direct volume visualization). This artifact was eliminated by computing an image that estimates the uneven illumination. This image is used to correct the whole image stack.

For experimental extension of the HREM modality slices were used for further medical staining, resulting in a second image series. We examined the possibility of registering the HREM images with the second image series. Because of the misalignment within the newly generated image stack we employed a two dimensional registration approach. The histologically colored images exhibit a combination of linear (i.e., different fields of view) and non-linear (i.e., deformations caused by the manual intervention of the already cut slices) transformations. For the linear part of the transformations an affine registration step was used. For the non-linear part of the transformations a b-spline registration step was used. With these two steps in the registration phase the two volumes can be registered generating visually matching images. The quality of the registration highly depends on the parameter settings of the registration framework. While the affine registration takes several minutes per slice, the b-spline registration takes in the

order of hours due to the huge number of parameters involved. In this thesis we examined different registration algorithms and found a parameter set that works for the two representative datasets.

Further, visualization tools are presented to examine the registration results. We used a two dimensional checkerboard visualization approach, and a three dimensional visualization approach based on direct volume rendering. To improve the visibility of the tissue borders the DVR was adapted for this task by incorporating gradient magnitude opacity modulation. To examine the quality of the registration we used a movable focus region which facilitates the task of following different tissue boundaries.

Acknowledgements

My thanks goes to Peter Rautek my supervisor for the best supervision I could think of. Next I want to thank Stefan Geyer for his time and patience to explain me the medical issues involved in my thesis. Furthermore I want to thank Eduard Gröller, our beloved master, as well as Wolfgang Weninger for their input at our meetings. Last but not least I want to thank the whole visualization group of the Institute of Computer Graphics for listening to my talks and making suggestions about my work.

Bibliography

- [1] A. Andronache, M. von Siebenthal, G. Székely, and Ph Cattin. Non-rigid registration of multi-modal images using both mutual information and cross-correlation. *Medical Image Analysis*, 12(1):3–15, 2008.
- [2] L. Gottesfeld Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.
- [3] M. Capek, R. Wegenkittl, A. König, W. Jaschke, R. Sweeny, and R. J. Bale. Multimodal medical volume registration based on spherical markers. In *WSCG*, pages 17–24, 2001.
- [4] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multi-modality image registration based on information theory. *Information Processing in Medical Imaging*, pages 263–274, 1995.
- [5] A. Guimond, A. Roche, N. Ayache, and J. Meunier. Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *Medical Imaging, IEEE Transactions on*, 20(1):58–69, 2001.
- [6] P. Hellier, C. Barillot, E. Memin, and P. Perez. Hierarchical estimation of a dense deformation field for 3-d robust registration. *Medical Imaging, IEEE Transactions on*, 20(5):388–402, 2001.
- [7] Kitware. <http://www.itk.org/>, retrieved on 20. 07. 2009.
- [8] S. Klein and M. Staring. *elastix the manual*. Image Sciences Institute, University Medical Center Utrecht, http://elastix.isi.uu.nl/download/elastix_manual_v4.0.pdf, retrieved on 20. 07. 2009.
- [9] S. Klein, M. Staring, and J. P. Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines. *IEEE transactions on Image Processing*, 16(12):2879–2890, 2007.

- [10] Leica. <http://www.leica-microsystems.com/>, retrieved on 20. 07. 2009.
- [11] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, 1988.
- [12] B. Likar and F. Pernus. A hierarchical approach to elastic registration based on mutual information. *Image and Vision Computing*, 19(1-2):33–44, 2001.
- [13] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, volume 21, pages 163–169, July 1987.
- [14] F. Maes, D. Vandermeulen, and P. Suetens. Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information. *Medical Image Analysis*, 3(4):373–386, 1999.
- [15] J. Maintz and M. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [16] J. B. A. Maintz, E. H. W. Meijering, and M. A. Viergever. General multimodal elastic registration based on mutual information. In *Image Processing*, pages 144–154. SPIE Press, 1998.
- [17] micro Tec. <http://www.micro-tec.de/>, retrieved on 20. 07. 2009.
- [18] James R. Miller. Attribute blocks: Visualizing multiple continuously defined attributes. *IEEE Computer Graphics and Applications*, 27(3):57–69, 2007.
- [19] Polyscience. <http://www.polysciences.com/>, retrieved on 20. 07. 2009.
- [20] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 2007.
- [21] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
- [22] C. E. Shannon. A mathematical theory of communication. *Bell system technical journal*, 27, 1948.

- [23] C. Studholme, D. L. G. Hill, and D. J. Hawkes. Multiresolution voxel similarity measures for mr-pet registration. In *Proceedings of the 14th International Conference on Information Processing in Medical Imaging*, pages 287–298, 1995.
- [24] P. Viola and W. M. Wells, III. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997.
- [25] W. J. Weninger, S. H. Geyer, T. J. Mohun, D. Rasskin-Gutman, T. Matsui, I. Ribeiro, L. da F. Costa, J. C. Izpisúa-Belmonte, and G. B. Müller. High-resolution episcopic microscopy: a rapid technique for high detailed 3d analysis of gene activity in the context of tissue architecture and morphology. *Anatomy and Embryology*, 211(3):213–221, 2006.
- [26] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, October 2003.